



HAL
open science

Piecewise affine dynamical systems applied to the performance evaluation of emergency call centers

Marin Boyet

► **To cite this version:**

Marin Boyet. Piecewise affine dynamical systems applied to the performance evaluation of emergency call centers. Optimization and Control [math.OC]. Institut Polytechnique de Paris, 2022. English. NNT : 2022IPPAX031 . tel-04086980

HAL Id: tel-04086980

<https://theses.hal.science/tel-04086980>

Submitted on 2 May 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT
POLYTECHNIQUE
DE PARIS

NNT : 2022IPPAX031

Thèse de doctorat



Piecewise Affine Dynamical Systems applied to the Performance Evaluation of Emergency Call Centers

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École polytechnique

École doctorale n°574, Mathématiques Hadamard (EDMH)
Spécialité de doctorat : Mathématiques appliquées

Thèse présentée et soutenue à Palaiseau, le 25 mai 2022, par

Marin Boyet

Composition du jury :

Xavier Allamigeon Chargé de recherche, Inria et CMAP, École polytechnique	Co-directeur de thèse
Pierre Carli Professeur des universités - Praticien hospitalier Chef de service au SAMU de Paris, Hôpital Universitaire Necker-Enfants Malades, Université Paris Cité	Examineur
Pierre L'Ecuyer Professeur, Département d'informatique et de recherche opérationnelle, Université de Montréal	Examineur
Stéphane Gaubert Directeur de recherche, Inria et CMAP, École polytechnique	Directeur de thèse
Bruno Gaujal Directeur de recherche, Inria et LIG, Université Grenoble Alpes	Rapporteur
Alessandro Giua Professeur, Département de génie électrique et électronique, Université de Cagliari	Rapporteur
Éric Goubault Professeur, Département d'informatique, École polytechnique	Examineur
Régis Reboul Directeur de programme, Préfecture de Police de Paris	Invité
Amandine Véber Directrice de recherche, CNRS et MAP5, Université Paris Cité	Présidente

Inria



AGENCE
INNOVATION
DÉFENSE

Piecewise Affine Dynamical Systems applied to the Performance Evaluation of Emergency Call Centers

by Marin Boyet



This work is licensed under a **Creative Commons Attribution-ShareAlike 3.0 Unported License**, whose terms can be found at the link <https://creativecommons.org/licenses/by-sa/3.0/>.

This license allows reusers to distribute, remix, adapt, and build upon the material in any medium or format, so long as attribution is given to the creator. The license allows for commercial use. If you remix, adapt, or build upon the material, you must license the modified material under identical terms.

This document was typesetted using Lua \TeX and the fonts Linux Libertine and Brandon Text. The template was inspired by the PhD manuscripts of Pierre Lairez and Jean Auriol, and was composed with the help of Xavier Allamigeon.

Cette thèse est dédiée à tous les personnels de l'urgence en France.

ABSTRACT

We develop in this thesis mathematical methods for the performance analysis and dimensioning of emergency call centers. To this purpose, we use tools from the field of discrete-event dynamical systems, in particular the formalism of timed Petri nets with preselection and priority rules, to describe the handling of emergency calls by dedicated platforms. These models are characterized by piecewise affine and recursive dynamical equations which are a subclass of controlled switched systems. We show that the continuous relaxation (or fluid) approximation of these dynamics is asymptotically precise under the application of a scaling factor.

We establish formal correspondence results between the dynamical analysis of monotonic timed Petri nets and the study of the value function of semi-Markov decision processes with discount factors as well as terminal and stopping costs. Leveraging on this “backwards” correspondence, we obtain several characterizations of the throughput vector of Petri nets, which indicate the handling rate of the organization that is modeled. We deduce practical staffing recommendations for emergency call centers in terms of the calls features. We compute analytical upper bounds on the time needed for a class of monotonic and hierarchical timed Petri nets to absorb a perturbation affecting an input by exhibiting a correspondence with stochastic shortest path problems. In the context of emergency call centers, this indicates how quickly an extra peak of calls can be treated. We also study Petri nets describing organizations with priority routings and associated with non-monotonic dynamics, whose congestion phases can be represented by polyedral complexes. This allows us to compute the minimum staffing of more involved call center layouts. In addition, we focus on tropical posynomial systems, which are an algebraical abstraction and static counterpart of our dynamical systems. In particular, these systems generalize Markov decision processes. We show that the resolution of these equations via linear programming features geometric problems of independent interest, involving separation and tangency properties of several convex bodies.

We challenge our theoretical contributions on two real-life case studies carried out in collaboration with emergency call centers in the Paris area: the four SAMU of AP-HP for medical needs and the PFAU for rescue and police requests. Our comparison is based on numerical simulations that reproduce call center operations with as few simplifications as possible, in line with the “digital twin” logic. The arrival processes, service times and patience levels of callers are randomly generated on the basis of histograms derived from the data analysis of millions of calls. We show that our analytical formulas that neglect abandonment phenomena yield good polyhedral approximations of the real throughput. We also evaluate the accuracy of queueing theory estimates which are frequently applied to the sizing problem of call centers. The simulation approach allows us to compare the single-tier and the two-tier architectures of emergency call centers. We show that the latter causes fewer losses of urgent calls than the former provided that the first-level instruction is quick. The two-tier layout is also more resilient to burst of calls. We finally quantify the benefit in practice of merging or bringing together several pools of agents.

RÉSUMÉ EN FRANÇAIS

Nous développons dans cette thèse des méthodes mathématiques pour l'évaluation de performance et le dimensionnement de centres d'appels d'urgence. Nous utilisons pour cela des outils issus de la théorie des systèmes dynamiques à événements discrets, en particulier le formalisme des réseaux de Petri avec règles de préselection ou priorité, afin de décrire le traitement d'appels d'urgence par des plateformes dédiées. Ces modèles sont gouvernés par des équations dynamiques récursives et affines par morceaux, une sous-classe des systèmes contrôlés à commutation. Nous montrons que l'approximation relâchée-continue (ou fluide) de cette dynamique est asymptotiquement précise sous une limite d'échelle.

Nous établissons une correspondance entre l'analyse de la dynamique des réseaux de Petri temporisés monotones et l'étude de la fonction valeur des processus de décision semi-Markoviens escomptés avec coûts finaux et coûts d'arrêt. Cette correspondance "en sens rétrograde" nous permet d'obtenir plusieurs caractérisations du débit des réseaux de Petri, et de déduire des recommandations pratiques de dimensionnement en agents pour les centres d'appels d'urgence en fonctions des caractéristiques des appels. Nous calculons des bornes explicites sur le temps de retour à la normale pour une classe de réseaux de Petri temporisés monotones et hiérarchiques, afin d'absorber la perturbation d'une consigne. Nous exploitons pour cela une correspondance avec les problèmes de plus court chemin stochastique. Dans le monde des centres d'appels, cela fournit une borne supérieure du temps de traitement d'un pic d'appels. Nous étudions aussi des réseaux de Petri d'organisations présentant des mécanismes de priorité et associées à des dynamiques non-monotones, dont les phases de congestion peuvent être décrites par des complexes polyédraux. Cela nous permet de calculer le dimensionnement minimum de centres d'appels plus complexes. De plus, nous nous intéressons aux systèmes posynomiaux tropicaux, qui sont une abstraction algébrique et le pendant statique de nos systèmes dynamiques. Ces systèmes généralisent notamment les processus de décisions Markoviens. La résolution de ces équations par programmation linéaire fait apparaître des problèmes géométriques d'un intérêt indépendant, impliquant des notions de tangence et de séparation entre ensembles convexes.

Nous évaluons nos contributions théoriques à travers deux études de cas réalisées en collaboration avec des centres d'appels d'urgence de la région parisienne, les quatre SAMU de l'AP-HP pour les appels santé, et la PFAU pour les requêtes orientées police ou secours à personne. Nous avons pour cela recours à la simulation numérique, reproduisant le fonctionnement de centres d'appels avec aussi peu d'hypothèses simplificatrices que possible, selon le principe des jumeaux numériques. Les lois des arrivées d'appels, des temps de conversation ou des niveaux de patience des requérants sont générés aléatoirement sur la base d'histogrammes provenant de l'analyse de plusieurs millions d'appels, et tenant compte de divers phénomènes de corrélations entre les variables. Nous montrons que nos formules analytiques négligeant les abandons fournissent de bonnes estimations polyédrales du débit réel. Nous étudions aussi la précision des formules de performance prédites par la théorie des files d'attente et le calcul d'Erlang, souvent utilisés dans les problèmes de dimensionnement de centres d'appels. L'approche par simulation nous permet de comparer des architectures de centres d'appels mononiveau ou biniveau devant traiter les mêmes flux d'appels. Nous montrons que le mode biniveau cause moins d'abandons que le mononiveau pourvu que le premier niveau d'instruction soit court. L'architecture biniveau est également plus robuste aux pics d'appels. Nous quantifions enfin l'intérêt opérationnel de mutualiser partiellement ou intégralement des plateformes distinctes, ou encore de faire travailler ensemble différents groupes d'agents.

REMERCIEMENTS

Ce travail de thèse, long de quasi quatre années, n'aurait tout simplement pas été possible sans le concours de nombreuses personnes que je souhaite remercier ici.

Mes premiers mots vont évidemment à mes directeurs de thèse, Stéphane Gaubert et Xavier Allamigeon. Stéphane, je te remercie pour la confiance que tu m'as accordée depuis le début, pour toutes les portes que tu m'as ouvertes et en particulier pour m'avoir offert ce sujet unique. Ton érudition, ta capacité à faire des liens entre de nombreux domaines scientifiques et ta pédagogie se retrouvent je crois beaucoup dans ce travail. Xavier, tu as été d'un soutien sans faille tout au long de ces années et travailler avec toi a été un plaisir de tous les instants. Cette thèse doit autant à ta finesse mathématique qu'à ton intelligence humaine, et je souhaite à de nombreux élèves de pouvoir en bénéficier dans le futur. Sur le plan scientifique, je veux aussi remercier Marianne Akian et Philippe Robert qui ont pris de leur temps pour m'aider à traiter plusieurs points techniques de ce travail.

Je suis très reconnaissant à Bruno Gaujal et Alessandro Giua d'avoir accepté de relire minutieusement ce manuscrit et rapporter mon travail de thèse, dont ils ont su percevoir tous les aspects. Je remercie également Pierre L'Ecuyer, Éric Goubault et Amandine Veber d'avoir bien voulu faire partie du jury.

Rien de ce travail n'aurait été possible sans l'investissement de la Préfecture de Police (PP), à travers la direction de programme PFAU. À ce titre, je souhaite en premier lieu remercier le LCL Stéphane Raclot dont l'expertise sur les systèmes d'urgence, la connaissance métier, les contacts et parfois le flair ont tant contribué à donner à cette thèse sa dimension appliquée. Je remercie beaucoup Régis Reboul pour la confiance et le soutien qu'il a apportés à notre équipe, et je veux saluer la patience et l'abnégation dont il a toujours fait preuve. À travers lui, je veux aussi remercier la Direction de l'Innovation, de la Logistique et des Technologies de la PP, en particulier son directeur Arnaud Mazier mais aussi Christophe Guillonnet pour son aide sur les questions fonctionnelles. Je remercie la Brigade des Sapeurs Pompiers de Paris, pour avoir renouvelé son intérêt dans nos travaux et nous avoir régulièrement ouvert ses portes. Je suis notamment très reconnaissant au SCH Benjamin Berhault pour son aide dans l'extraction des données de téléphonies de la PFAU, et au CNE Denis Daviau qui nous a aidé à les interpréter.

Initialement permise par l'intermédiaire de la PP, la collaboration que nous avons pu développer avec les SAMU a été un tournant majeur de cette thèse. J'exprime toute ma gratitude au DR Éric Lecarpentier, chef de service du SAMU 94, pour le rôle moteur qu'il n'a cessé de jouer dans ce projet, tant sur les aspects organisationnels que techniques, avec le soutien de Thomas Pérennou. À l'hôpital Necker, le SAMU de Paris nous a réservé un accueil exceptionnel et je tiens à remercier chaleureusement le PR Pierre Carli pour nous avoir hébergé et intégré à ses équipes durant les premières semaines de la crise Covid, et maintenant pour avoir accepté de faire partie de mon jury de thèse. Ces remerciements s'étendent au DR Caroline Telion, au DR Christelle Dagon et évidemment au DR Jean-Sébastien Marx dont les facéties et l'infatigable imagination n'auront eu de cesse d'égayer ces années de thèse. Je remercie aussi beaucoup Jean-Michel Gourgues et Isabelle Lhomme pour leur aide dans le traitement des données du SAMU. J'ai eu grand plaisir à travailler avec les équipes du SAMU 92, où le DR Thomas Loeb nous a toujours accueilli avec

beaucoup d'amitié, tout comme le DR Jérémie Boutet, le DR Anna Ozguler et le DR Margot Cassuto avec qui nous avons épluché de nombreux tableaux de données de téléphonie... Je remercie le DR Érick Chanzy du SAMU 93 qui avec le DR Laurent Goix et Frédéric Linval a toujours pris de son temps pour nous recevoir et répondre à nos questions. À l'AP-HP, je souhaite remercier le DR Christophe Leroy qui nous a aussi soutenu avec beaucoup d'enthousiasme, ainsi que Stéphane Crezé et François Planeix pour nos nombreux échanges sur les questions techniques. Enfin, je veux remercier les équipes du SAMU 69 pour l'intérêt qu'ils ont manifesté pour nos travaux et leur accueil à Lyon, en particulier le DR Christian Di Filippo, le DR David Schiavo, le DR Yann-Franck Lourcy et le PR Pierre-Yves Gueugniaud.

De retour au CMAP, il me faut bien sûr remercier Baptiste Colin, le Bill de "Boule et Bill", tantôt coloc' et tantôt collègue, pour le temps passé à rire au Magnan comme pour celui passé à développer le simulateur de réseaux de Petri. Cette tâche a été reprise avec succès par Benjamin Nguyen Van Yen à qui je souhaite une très bonne continuation. Je tiens à saluer chaleureusement mes grands et petits frères de thèse, avec qui j'ai eu plaisir à échanger : Nikolas Stott, Mateusz Skomra, Jean-Bernard Eytard, Paulin Jacquot, Maxime Grangereau et Omar Saadi, puis Maël Forcier, Nicolas Vandame et Quentin Canu. Je souhaite particulièrement remercier mon prédécesseur Vianney Bœuf : tu as non seulement été de bon conseil avant mon arrivée à l'X, mais surtout, s'il y a bien une thèse dont s'inspire la mienne, c'est la tienne ! Un grand merci également à Mathilde Boissier et Céline Bonnet pour leur accueil au laboratoire. À l'École polytechnique ou à l'école doctorale, j'ai pu compter sur l'aide de tous les personnels et je veux particulièrement remercier Pierre Straebler, Alexandra Noiret, Nasséra Naar, Hanadi Dib, et Clotilde d'Épenoux. Je veux aussi saluer le rôle joué par Inria Saclay : l'institut a toujours su témoigner son soutien moral et matériel à cette thèse. Sur un tout autre plan, je veux remercier et saluer Alexandra Elbakyan pour ses actions en faveur de la science ouverte.

À l'École des Mines, je remercie très amicalement Frédéric Fontane et Matthieu Mazière, mais aussi Philippe Mustar et Sandrine Kletz, qui malgré mes promesses de moins fréquenter le boulevard Saint-Michel m'ont à tour de rôle trouvé de nombreuses occupations en parallèle de la thèse...

Un immense merci à tous mes amis pour m'avoir soutenu *et* supporté durant toutes ces années : Baptiste à nouveau, comme Charles (pour cette inoubliable et débile année de colocation, #709forever), Arnaud (pour nos soirées *VS Code – Nuggets* ©), Dog (pour nos nombreux "RDV passerelle"), Rudy et Robin (pour nos fraternelles virées dans le 13^{ème}), Aurélien (et tous ces saumons finis en galettes), Laurie (pour nos collaborations artistiques mais aussi pour avoir toléré les soirées *VS Code – Nuggets* ©), Corentin et Valentin (en souvenir de nos nuits sacrifiées pour une minable feuille de chou), Fanny (pour nos vaustrages parisiens et aixois), Lucie, Béatrice et François (pour avoir accueilli un passager clandestin en plein Covid), Barrou, Vincent, Romain et Océane (pour ces moments partagés en altitude avenue d'Italie), Juliette (pour notre colocation éphémère mais nos commérages éternels), Haowen (en souvenir de soirées piano et raviolis), Chloé et Agathe (pour m'avoir accueilli en leurs terrasses fleuries), Goutal (pour ce hackathon où nous avons tous laissé beaucoup de points de QI), Rémi et Edward (pour nos repas de petits gourmands à Lille) et Laurène (pour notre amour partagé de la mise en page et de la typographie).

Je veux pour terminer remercier mes parents, qui ont peu contribué à cette thèse mais ont bien travaillé pendant les vingt-quatre années précédentes ! Enfin, je remercie Gwendal, pour tout; j'ai sans doute passé à cause de toi un peu moins de temps que prévu sur ma thèse pendant ces deux dernières années... mais c'était tellement mieux comme ça ! Garde ton sourire, ton intelligence, ton énergie et ton humour, bref, ne change pas, car il faut bien que l'un de nous deux continue à être payé une misère pour écrire des articles !

CONTENTS

Abstract	iv
Résumé en français	v
Remerciements	vi
Introduction	1
1 Context of the thesis and related work	1
2 Organization of the emergency calls reception in France	5
3 Contributions and organization of the manuscript	7
4 Publications related to the thesis	11
I Piecewise-affine dynamics of fluid Timed Petri Nets	15
1 Discrete dynamics of timed Petri nets	17
1 Basics of untimed Petri nets	18
2 Addressing underdetermined situations	23
3 Timed Petri nets	31
2 Continuous relaxation and large-scale limit of Petri nets	47
1 Semantics and tools for fluid Petri nets	49
2 Control aspects of monotonic discrete and fluid dynamics	62
3 Fluid Petri nets as large-scale limits of discrete models	69
3 Semi-Markov decision processes under weak non-Zeno conditions	85
1 Basic terminology and recalls on Markov chains	87
2 The finite-horizon value of semi-Markov decision processes	92
3 Long-run properties of solutions of semi-Markov dynamics	121
4 Stochastic Shortest Path configurations	126
4 The correspondence between monotonic fluid Petri nets and semi-Markov decision processes	133
1 A running example: an emergency call center for medical needs	136
2 The finite-horizon correspondence theorems	140
3 Studying the throughput of fluid Petri nets through SMDP tools	148
4 Deviation to a congestion-free regime	159
5 Convergence time of the value iteration algorithm for semi-Markov SSP problems	168
6 Application to a staffing problem of the SAMU	185

5	Non-monotonic Petri net models of emergency call centers and their congestion diagrams	189
1	Stationary regimes of Petri nets with priority rules	190
2	A bilevel emergency call center with a priority mechanism	194
3	An emergency call center for medical needs with a monitored reservoir	201
4	Variations of the model EMS-B	208
6	Posynomial systems and underlying geometric problems	219
1	Real and tropical posynomial systems	220
2	Properties of the colorful interior of convex sets	232
II A comprehensive study of a real-life and large-scale emergency call-center		239
Foreword and context of the study		241
1	Object of the study	241
2	Presentation of the PFAU emergency platform	242
3	Summary of the data analysis step	244
4	Summary of the comparative study of organizations	245
7	Data analysis of phone logs	251
1	Description of used databases and global counts	252
2	Throughputs and proportions of calls	255
3	Identification of logged agents	261
4	Service times of calls	264
5	Patience of callers	275
6	Some specific days	281
8	Performance analysis using numerical simulation	289
1	Methodology	290
2	Performance analysis of a bilevel and single-job system, comparison with a single-tier layout	296
3	Validation of the analytical results by simulations	314
4	Study of the PFAU architecture and its possible evolutions	326
5	Behaviour of the bilevel layout under non-standard arrivals	338
Conclusion and perspectives		345
Appendix: Summary of our work during the Covid-19 pandemic		353
1	Adaptation of the organizations	353
2	Contributions to the dimensioning of response systems	354
3	Contributions to the forecast of the epidemic	357
4	Study of the Interactive Voice Response of SAMU 92	358
Index		360
Bibliography		363

INTRODUCTION

1 Context of the thesis and related work

- 1.1 THE PERFORMANCE EVALUATION ENDEAVOUR.** — Paraphrasing the definition given by Le Boudec in his book [LB10], performance evaluation can be described as the process of quantifying the service delivered by a system, whether it is a digital protocol (communication networks, computer software...) or a physical and human-managed organization (manufacturing systems, traffic infrastructure...). When one wants to address a problem that involves meeting a demand, having objectified performance grounds is essential to shape a solution; either for the step of system dimensioning (the more resources we put in it, the better the performance is likely to be, but what is a satisfactory level?) or to nail down the system architecture itself (there may be several design alternatives, is one of them any better?). Performance evaluation typically is about determining if a system can withstand an incoming load, finding if it has an intrinsic capacity and what are the bottlenecks. In an input/output approach, one is interested in relating the service level to the load and the system parameters.

Assessing performance of a system necessarily starts with a formalization phase, where one needs to rigorously establish what are the key parameters and factors that affect the performance and what are the metrics of interest. Performing real-life trials is tempting, but these are often expensive to carry out, little reproducible, and can rarely pretend to comprehensiveness. Therefore, one generally has to develop a model that mimics the system but is easier to manipulate. This whole process comes with some abstraction cost and requires a very good understanding of the original system's purposes (as well as its users) in order to accurately reflect its behaviour. As customary in the field of model-building or engineering, it is crucial to carefully break down the whole system into simpler pieces. This should be done without oversimplifications and by retaining the fundamental features. At each step, one must control the potential approximation errors and be aware of any blind spots of the model.

Performance evaluation sits at the confluence of several mathematical fields. First, the systems we focus on frequently have to cope with a form of uncertainty (if the incoming demand is not totally foreseeable, or random failure of some service units occur...). Hence it is natural to resort to some tools of probabilities or stochastic processes in order to account for nondeterminism. In this view, efficiency criteria may bring some statistical quantities into play, such as mean performance and corresponding variance, quantiles, best and worst-case scenarios, etc., allowing one to perform either average or robust system tuning. Moreover, when time intervenes (because the demand is spread over some duration or because the system itself introduces delays or buffer steps), it is often insufficient to study static algebraic properties. Instead, one has to tackle the problem through the prism of dynamical systems. By further including the notion of feedback or system memory, (i.e., the fact that some protocols or organizations can adjust their behaviour depending upon the load they receive or external decisions), one ends up looking at central questions of control theory.

Finally, and contrary to multiple problems arising in physics, man-made and service-oriented systems often involve phenomena of an integer nature (computer operations, limited number of resources, channels or system components) and subject to on/off evolution (at the abrupt onset of an exterior event, or when some threshold becomes exceeded). These observations motivated in the past century the development of the field of discrete event systems (DES), operations research, and the rise of suitable frameworks to deal with switched (or piecewise) dynamics.

1.2 MATHEMATICAL TOOLS AND APPROACHES. — A popular method to gain quantitative insight on the performance of a discrete event system, which also encompasses stochastic and control aspects, is simulation. In [Ban05], Banks recalls that it is a field of its own and lists many dedicated software. The main interest of simulation is its ability capture a wide spectrum of features with only minor simplifications needed to be made on the real system. It allows one to test and compare many hypotheses and tweak varied parameters. On the other hand, even implementing a simulation model, let alone developing an *ad hoc* software, may come with an important entry cost. One also needs to ensure that the implemented program correctly simulates the desired system. In addition, the results of the experiments are sometimes hard to interpret, with many entangled effects. Finally, an important number of simulation runs may be needed, which can turn out to be expensive and time-consuming. For these reasons, it is highly desirable to also have analytical performance estimates, taking the form of closed formulas in which the system parameters intervene.

In their introduction to discrete event systems [CLo8], Cassandras and Lafortune review the main theoretical tools used in the field of performance evaluation. Among these, the concept of queue has been used early by Erlang in [Erl17] to describe the situation of a customer waiting for a server (or in general for a resource that is in limited number) in a stochastic setting. Assuming some primitive distributions for the arrivals and the service times, queueing theory aims at computing stability conditions (in order for queues not to build up indefinitely) and performance results such as the distribution of the waiting times, the queue length, or the probability that some request shall be rejected or lost [Gro08]. Queues are considered to be one of the main building blocks of discrete event systems, and have received broad attention since, with many generalizations of server and customer classes or tasks types [BCMP75]. Burke’s Theorem [Bur56] paved the way for the study of queueing networks in which multiple queues are cascaded and where one can derive product-form solutions, in particular Jackson networks and Kelly networks (see [CY01, Rob13]). However, the theoretical assumptions made on the network structures or primitive distributions (typically exponential) are often strong and may not always be realistic. Cassandras and Lafortune further precise that queueing theory is rather “descriptive” in nature, and one often resorts to the study of controlled Markov processes to make room for some decision-making and “prescriptive” aspects, see [KV15, Sti85]. Unfortunately, such techniques require additional assumptions. From the numerical point of view, very large Markov chains are often introduced to compute steady-state distributions and performance criteria of queueing networks, and specific techniques need to be used to cope with the curse of dimensionality [Buc99, PF91].

In order to capture the event-to-event evolution of discrete event systems, many authors have used the framework of automata (see the book of Hopcroft et al. [HMU01] for background and [CLo8] again for the specific use in DES). Among these tools, the expressive framework of Petri nets has been praised for its modularity and its ease in graphically representing somewhat complex systems. Introduced by Petri in a work centered on conceptualizing communicating systems [Pet62], Petri nets have been designed to offer a general formalism covering synchronization of resources (in a more general way than queues) and model concurrent processes (running

in parallel and in an asynchronous way). They have later been extended to take timed phenomena into account, as detailed in [Pet81] or in the historical perspective [GS18]. Reachability problems and deadlock detection have been active research topics in the Petri net literature (see for instance [Reu89, EN94]), in relation with the field of verification and performance analysis. Petri nets also come with flexible ways to arbitrate conflicts and resource allocation, with the possibility to implement priority or inhibition rules. Studying the dynamics of timed Petri nets allows to compute performance estimates, like the throughput of the modeled system, or to identify its bottlenecks. Algebraic techniques like those developed in min-plus and max-plus algebras have shown to be particularly fruitful to derive analytical results of some subclasses of timed Petri nets. This includes the class of timed event graphs, for instance used by Carrier and Chretienne for application in scheduling [CC88] and also thoroughly studied by Baccelli, Cohen, Olsder and Quadrat in [BCOQ92] with deterministic or stochastic time aspects. In the subsequent work [BGF94, BFG96, BM98], in-depth analysis of Petri nets was extended to the class of free-choice nets, building on fundamental monotonicity properties of the dynamics — see in particular the book of Glasserman and Yao [GY94]. However, for more general models, mixing the Petri net formalism with stochastic processes generally yields instances out of analytical reach, and one often has recourse to fluidification methods or scaling limits to derive mean ergodic results, as done in [CGQ95, GGo4b, RS01]. Hybrid systems, featuring both discrete-event and fluid parts, are also considered [BGM00] and remain active nowadays.

For the system dimensioning purposes, obtaining lower or upper bounds on the performance can already be very satisfactory. This is the core motivation of the so-called network calculus, introduced by Cruz in [Cru91], where one seeks to lean on an algebraic framework and symbolic propagation techniques to compute envelopes on the system dynamics. Just as Petri nets, the theoretical tools of idempotent structures play an important role in network calculus. The latter can be seen as a filtering theory in the min-plus algebra, in particular building on the min-plus convolution operation. Customarily, the arrival and service curves are bounded by piecewise affine profiles, which makes computation more handy, but the theory also covers more general envelope shapes. As shown by Chang in his book [Cha00], network calculus has proven advantageous to guarantee bounds on backlog and delays in communication networks — other applications include scheduling and buffer dimensioning (see Le Boudec and Thiran, [LBT01]). It was noted that although the network calculus approach is well-adapted to worst-cases studies, it can be over-pessimistic for average scenarios. This motivated further developments on the theory with stochastic counterparts. We also refer to Bouillard who covered recent advances in [Bou14].

It is worth mentioning that queueing theory, Petri net analysis, and network calculus share several common features and yield first-order estimates that are consistent. For instance, the stability conditions of simple queues or the celebrated Little's law (see [Lit61] or [CLo8, p. 457]) express intuitive “rule-of-three” equilibria, that are retrieved by computation of throughput functions in Petri nets. Similarly, evaluating transience bounds in the latter give affine formulas that are reminiscent of network calculus, as we shall see in this thesis. Finally, all these analytical methods are complementary with computer simulation, and it is often advantageous to carry them out in parallel.

- ### 1.3 THE CASE OF CALL CENTERS AND EMERGENCY CALL CENTERS. —
- Call centers are nowadays a well-known industry, employing millions of people in the world. We usually distinguish the concept of outbound and inbound call centers (respectively emitting and receiving calls). Both types of platform fit into the setting of discrete event systems (the call sender initiates the call, the receiver picks it up, the call is hung up after some time, etc.). Inbound call centers can

fully be thought of as systems delivering a service to customers, for which we want to evaluate the delay levels and the congestion situations. Queueing theory was actually developed for this very application, giving rise to the pioneering Erlang calculus [Erl17] and some decades later to the Palm calculus [Pal57]. Erlang formulas are still used for some practical dimensioning of call centers nowadays. The Erlang-C formula predicts the average waiting time or the probability for a customer to wait depending on the arrival rate, the mean service time and the number of servers, provided the queue is stable. The Erlang-B and Erlang-A formulas refine these estimates by considering that calls are blocked if all the servers are busy or that customers can hang up themselves if the wait is too long.

More generally, call centers exhibit several performance evaluation or operations research problems. In [ALo5], the latter are reviewed by Avramidis and L'Ecuyer, who also cover the main associated analytical models and simulation techniques. We also refer to the overview of staffing and routing algorithms of Koole and Pot [KP06]. Beyond Erlang formulas, integer programming and linear programming techniques can be implemented to solve scheduling problems in the call center, like those of [ICWC10] and [AAC⁺12]. Other models include continuous-time Markov chains to approximate service levels and find a satisfactory staffing level, as in [DLP⁺07], or approximate dynamic programming to improve routing of calls [RB12]. Simulation tools are also widely used in the call center domain, with some dedicated libraries (for instance [BLo5]). Much effort has also been made to precisely understand the primitive parameters of call centers, like the arrival rate of calls, the conversation times, the patience of the customers, etc., for both prediction purposes but also to feed the analysis — see for instance the thorough study [BGM⁺05b] of Brown et al. based on more than a million of calls over a year to a bank call center.

Emergency call centers are dedicated to the handling of medical distress or civil security matters, with the ability to dispatch immediate assistance and response teams. Dealing with these call centers has several specificities. First, owing to the very nature of the requests, any congestion of the system resulting in long treatment delays or call abandonments can have a dramatic fallout. Therefore, it is paramount that the calls are handled as soon as possible with minimum delay. This has an influence on the staffing policy, that will typically be more conservative than commercial call centers in which a trade-off between cost (number of agents) and efficiency (quality of service) is tolerated, as discussed for instance in [KM02]. Second, there can be highly skilled agents in emergency call centers (such as emergency physicians), hence the architecture and the call handling procedures must be well thought through to correctly allocate these resources. In addition, emergency call centers are more exposed to bursts in demand than commercial call centers, due to unexpected events that have health or civil security repercussions, and have to resist these peaks of calls while maintaining their performance level. As a critical public service, one cannot afford to experiment with too many modifications or solutions directly on the real-life emergency platforms, hence the recourse to analytical models and simulation is of great help in practice.

In the emergency response field, most of the literature has been focused on the dispatching chain, i.e., dealing with the issue of choosing the number, the location and the composition of the response teams and vehicle fleet. In comparison, there are fewer works on emergency call centers themselves (see for instance the recent review of Petidmange in her dissertation [Pet20]). Building on the work [AEHo4], Ta, L'Ecuyer and Bastin study the staffing problem in emergency call centers with a stochastic average approximation problem approach [TLB16], with application to a 911 call center in Montreal, Canada. In [dQLMSG14], Lima et al. mix reliability block diagrams and stochastic Petri nets to evaluate the performance and failure risks for an emergency call center in Brazil. Discrete event system simulation was used by Gustavsson in his licenciate

thesis [Gus18] with application to a Swedish emergency platform. In [vBKvdMB15], van Buuren et al. rely on queueing models to simulate the performance of a medical emergency call center as a function of the number of agents and choose the staffing accordingly. In the subsequent work [vBKvdMB17], they compare different models with varied agent tasks, taking both the call taking and the dispatching domains into account. In [ABG15], Allamigeon, Bœuf and Gaubert proposed a modelization of an emergency call center located in Paris, France, using timed Petri nets with priority rules, and derived analytical formulas of the system throughput as a function of the staffing, hence deducing a minimum number of agents for all the calls to be handled with no delay.

2 Organization of the emergency calls reception in France

The models of emergency call centers studied in this thesis have been developed in partnership with French emergency services which we briefly present below. To begin with, there are four main emergency phone numbers that the population can dial in France:

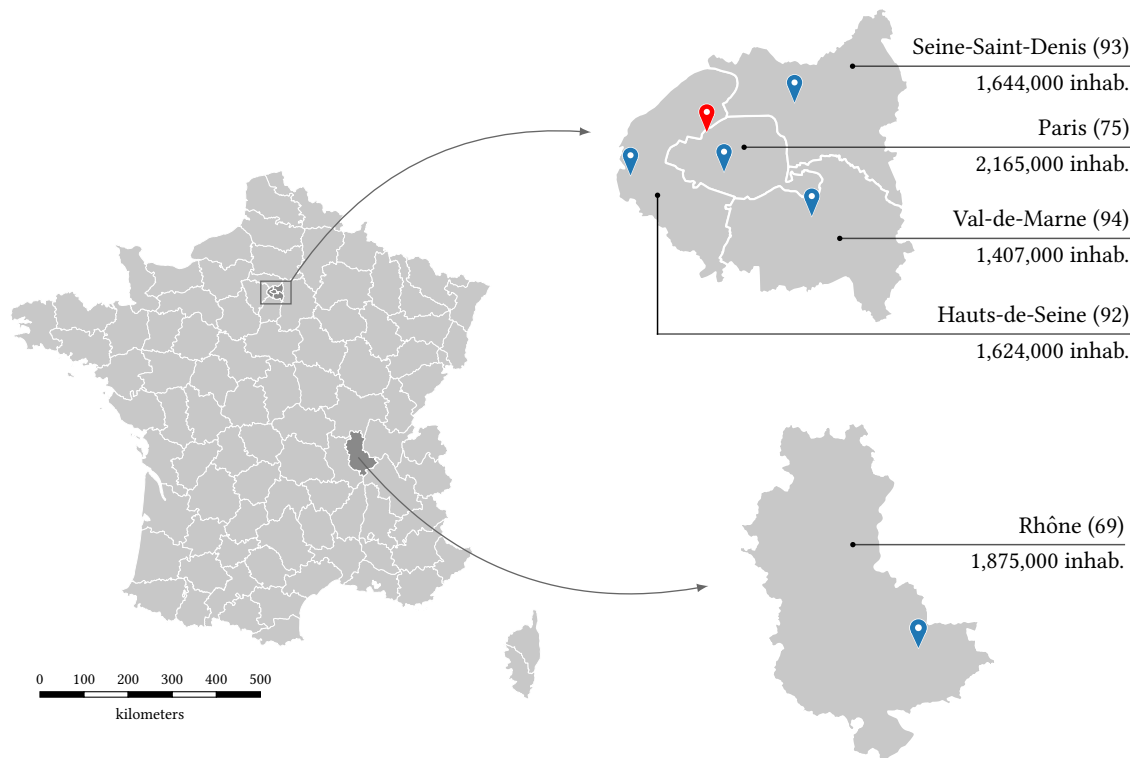
- the national number 15, dedicated to medical emergencies. It is operated by the *Service d'Aide Médicale Urgente* (SAMU), and organizes the prehospital response, with ambulance or mobile intensive care unit (MICU) dispatch.
- the national number 17, to contact the police when witnessing a crime, violence, traffic accident, public disorder etc. The calls are picked up either by the national police or the *gendarmerie* (local law enforcement), able to send teams on scene.
- the national number 18, to reach the firefighters and report fires, gas leaks, explosions etc. or situations involving rescuing a victim (public road accident, injuries...). It is operated by the *Service d'Incendie et Secours* (SIS), which handles fire trucks and rescue vehicles.
- the European number 112, which is a hub to contact all of the above (in several European countries, the 112 is the only emergency number to dial regardless of the matter). In most regions of France, the 112 calls are routed directly to the SIS and one often talks about the 18–112.

The emergency calls are typically centralized at the departmental level (France is divided into 101 “départements”), so that there are about 450 emergency call centers throughout the country, involving more than 8000 responding agents. In 2018, 29 millions of calls were received by the SAMU, 23 millions by the police or the *gendarmerie*, and 22 millions by the SIS (see [MP19]), thus representing a total of 74 millions of emergency calls for a population of 67 millions inhabitants.

As far as medical emergencies are concerned, a characteristic of the French system is that an emergency physician is present all along the intervention, from the call regulation to the on-site intervention. This contributes to avoid “overtriage”, that is to say sending unneeded ambulances and overcrowding emergency rooms. French emergency services tend to follow the “stay and play” logic (see [SC09]), meaning that MICU can be dispatched in the most severe situations to initiate on-site primary treatments and provide advanced life support (abb. ALS, which includes invasive techniques). Other countries (such as the United States of America via the 911) adopt the “scoop and run” model, in which only basic life support methods (abb. BLS, involving non invasive techniques of immobilization, oxygen administration, bandaging...) are used, often by paramedics, in order to bring the patient to an hospital facility as fast as possible. We refer the reader to [NPJ98, AL04, JPM20] for more in-depth description of the French SAMU system.

The health, police and rescue services work hand in hand to fulfill their missions of assistance to the public. Due to their training in BLS techniques, the French firefighters are often called in support of police operations (extraction and evacuation of victims) but also of SAMU teams (patient transportation, cardiac massage...). For life threatening situations requiring immediate care, the requests received by the SAMU and the firefighters sometimes overlap; in case of a heart attack, the first witness dials the number 15 in one out of two cases, and the number 18 in the other case [MP19, p. 39].

In the Paris area (i.e., the city and its inner suburbs, highlighted in the map below), the emergency medical response is carried out the four SAMU of *Assistance Publique – Hôpitaux de Paris* (AP-HP). The high number of hospital facilities in the region is an extra challenge for the answering physicians who seek to send the patients to most appropriate units. Another regional specificity is the military status of the Parisian firefighters, who are members of the *Brigade des Sapeurs-Pompiers de Paris* (BSPP) and have specific means (see [BL19]). Likewise, police forces of Paris area are unified under the interdepartmental unit referred to as the *Direction de la sécurité de proximité de l'agglomération parisienne* (DSPAP). The BSPP and the DSPAP work under the authority of the *Préfecture de Police de Paris*, which led to the inauguration in 2016 of a common emergency calls platform with firefighters and police forces, the *Plateforme d'Appels d'Urgence* (PFAU), or also the 17–18–112 call center. This platform receives all the law enforcement and



The 96 departments of metropolitan France (left), and a focus on the Paris “inner ring” region and Rhône department (right). The location of the PFAU (17–18–112 call center) is shown in red, and locations of the five SAMU call centers we worked with in blue.

rescue emergency calls of the Paris area, making it the biggest emergency call center of the country.

During this thesis, we had the opportunity to collaborate with the BSPP and DSPAP to carry out the performance evaluation of the PFAU call center, and with the SAMU of Paris, Hauts-de-Seine, Seine-Saint-Denis and Val-de-Marne to study each one of their regulation platforms as well. During the Covid-19 crisis, we had the chance to also work with the SAMU of Rhône, located in the city of Lyon, to adapt our models to their organizations and needs.

3 Contributions and organization of the manuscript

The thesis is divided in two parts, reflecting two complementary approaches to analyze real-life organizations such as emergency call centers that lend themselves well to mathematical modeling and numerical simulation. We present our contributions chapter by chapter.

Part I deals with analytical models of timed Petri nets and semi-Markov decision processes, and it addresses performance evaluation issues of organizations modeled by these means. The central object that we actually study is a set of dynamical recursive piecewise-affine equations with delays, arising in these two families of models.

- Chapter 1 presents the formalism of Petri nets and timed Petri nets, with exogenous prescribed inputs and subject to preselection or priority routings, and when the holding times are constant and attached to places. We endeavour to show the well-posedness and compatibility of the firing rules. We resort in particular to the notion of simultaneous firing of multisets of transitions in the untimed setting, which allows one (Theorem 1.13) to avoid the situations of so-called “confusions” brought up in [AMBC⁺98]. We recall the dynamical equations verified by the integer-valued counter functions of the net, already featured in [ABG15]. In comparison, we work in a slightly more expressive context where arbitrary degrees of priority rules are permitted.

In the priority-free framework, we show that this system of equations provide well-posed dynamics even if the holding times of certain places are null, provided that no circuit of the net is composed of only such places (Theorem 1.23). These sufficient “non-Zeno” condition allow us to safely model real-life processes where some steps are instantaneous, with no need to add artificial positive waits. We further show that the counter dynamics involving priority rules also determine unambiguous trajectories but we have to this purpose to rule out some pathological patterns.

- Chapter 2 focuses on continuous-relaxed (or fluid) timed Petri nets, in which integrality constraints are discarded. The dynamical equations governing these models can be postulated by analogy with their discrete counterpart. However, we show that the natural fluidification of the delayed dynamics can be obtained by amending the usual ODE system based on exponential firings of transitions considered by [AD98].

We propose a single operator-based formalism to treat both the discrete and continuous-relaxed dynamics using a control theory approach. We establish structural properties of the underlying evolution operators or evolution semigroups – in particular the fact that they are order-preserving (or monotonic). The main contribution of the chapter is the quantitative analysis of the deviation between the trajectories of the discrete and continuous-relaxed dynamics. Provided the inputs at stake and the preselection rules have bounded differences, we prove that the fluid approximation is asymptotically exact under a scaling of both dynamics, that is to say when more and more tokens navigate in the nets (Theorem 2.19). This conclusion is still valid when the discrete

dynamics is subject to multiple nondeterministic events; we prove under nonhomogeneous Poisson arrivals, Bernoulli routings and arbitrarily-distributed holding times that after conjugation by a scaling parameter, the stochastic discrete dynamics admit the continuous-relaxed trajectory as a deterministic fluid limit (Theorem 2.23). This “strong law of large numbers for Petri nets” answers a question raised in [CGQ95] and is an additional argument to justify the well-foundedness of the fluid approximation.

- Chapter 3 introduces the setting of semi-Markov decision processes (SMDPs) with finite state and action spaces and constant sojourn times. Unlike the majority of the SMDP literature, we allow most sojourn times of the model to be null. We exhibit minimal non-Zeno assumptions under which the expected number of moves achieved in the game in a given time-horizon stays finite (Theorems 3.15 and 3.16). We establish under these weak non-Zeno conditions the finite-horizon dynamic programming equations of several cost criteria, involving discount factors, terminal costs and stopping costs (Theorems 3.18 and 3.22). Our whole setting requires tools from non-linear Perron–Frobenius theory (see [AGN11]) to control the contraction properties of the semi-Markov dynamics, in particular the Zeno phenomena that a priori induce implicit systems of equations. It leads us to state an induction principle that we use to show multiple structural properties of the trajectories via propagation in the dynamical equations, and we still use the high-level semigroup formalism introduced in the Chapter 2.

Standard results on the SMDP average-cost problem are recalled, to characterize the long-run behaviour of trajectories. In the last section of the chapter, we extend to the semi-Markov setting the analysis of the stochastic shortest path problem introduced in [BT91] as a subclass of MDPs. We prove the existence of the value and that all trajectories converge towards it (Theorem 3.41).

- Chapter 4 contains our main contributions to the analysis of timed Petri net dynamics. We enounce a formal correspondence between fluid timed Petri nets and semi-Markov decision processes based on their respective evolution equations (Correspondence Theorem 1), which extends ideas sketched in [CGQ95]. We show that advanced modeling aspects in SMDP such as the terminal and the stopping costs are essential to account for exogenous inputs in Petri nets. Relying on the notion of stoichiometric invariants, we reduce the analysis of timed Petri nets to the one of *undiscounted* SMDPs (Correspondence Theorem 2).

This allows us to consider the throughput vector of Petri nets as the average-cost vector of an SMDP (Correspondence Theorem 3). We leverage on classical results of the SMDP literature and on our non-Zeno extensions to retrieve or refine several characterizations of the throughput of Petri nets. We show that the latter arises as the solution of a set of lexicographic piecewise affine optimality equations and that it can also be explicitly expressed using the initial Petri net parameters. This formula proves that the throughput vector of a Petri net is a piecewise-affine, nondecreasing and concave mapping of the resources of the net. The corresponding polyhedral complex provides a *congestion diagram*, which is an operational tool from which staffing recommendations can be drawn for real-life organizations. Exploiting linear programming techniques for SMDPs (see [DF68]), we also retrieve that the throughput vector of this class of Petri nets can be computed as the solution of a linear program, which was first shown in [GG04b].

We build further on our correspondence theorems to study the transience regimes of single-input Petri nets satisfying congestion-free conditions. We establish that the problem of evaluating the time needed to catch-up the input after a perturbation reduces to computing the convergence time of the value iteration procedure for semi-Markov decision processes in stochastic shortest path configuration (Correspondence Theorem 4 and Theorem 4.11). We prove that the value iteration terminates in finite-time if and only if the transition matrices of the optimal policies are nilpotent (Theorem 4.13), which is also new in the setting of MDPs. We exhibit a class

of hierarchical SMDPs for which this finite-time convergence conditions are satisfied, and we compute analytical upper bounds on the length of the transient regimes (Theorem 4.19). We thus obtain analytical maximal convergence times to catch-up the input in congestion-free Petri nets, from which we can derive security margins on the staffing of real-life organizations in order to recover quickly from a burst in the demand.

The entirety of the Chapter 4 is illustrated on the running example (EMS-A) of a medical emergency call center that is drawn from our collaboration with the SAMU of Paris area. Examples of practical recommendations are presented in the last section.

- Chapter 5 discusses how the use of priority rules in Petri nets impacts the results obtained for the monotonic models of Chapter 4. We use the equations of [ABG15] in which the throughput of non-monotonic Petri nets are looked for based on min-plus systems on germs of affine functions. Our approach is inspired by the work [FGQ11] applied to the modeling of road traffic networks.

We illustrate on two examples of emergency call centers implementing priority rules that we can still compute congestion diagrams in this setting. The first case study models the (BiLvl) system of the PFAU, i.e., the 17–18–112 police and firefighters call center of Paris area. This call center was already thoroughly studied in [ABG15, Bœu17]. Our second example, the system (EMS-B), is an enhanced and complexified version of the medical emergency call center (EMS-A) introduced in Chapter 4. It is also designed with the help of the SAMU, and its purpose is to study an innovative reservoir mechanism.

The analytical expressions of the throughput that we derive for these two practical cases allow us to demonstrate the interest of priority rules for certain call center architectures (protection of the most urgent calls, desynchronization of limiting resources, ...) and on the other hand, it brings to light some pathological behaviours that can arise if the system is not correctly tuned. For instance, we exhibit a situation where increasing the number of agents in the system (EMS-B) decreases the throughput of the task with the highest priority of the call center. The advantage of our approach is that it is modular and systematic, with a characterization of the system throughput even in the degraded or unanticipated phases. We explore further three alternative models to (EMS-B), and using the congestion diagram tool, we show that nontrivial comparison criteria between models or extra security margins can be computed analytically.

- Chapter 6 focuses on posynomial systems, that are multivariate polynomials on positive variables but with arbitrary exponents. We consider either classical posynomials or tropical posynomials. Our motivation in this chapter is that tropical posynomial systems are induced by the fixed-point equations arising in the dynamics of both continuous-relaxed timed Petri nets or semi-Markov decision processes. As a result, the understanding of such systems is a way to gain insight on those dynamics from the algebraical point of view.

We prove that solving classical or tropical posynomial systems is in general NP-hard (Proposition 6.3 and Theorem 6.4). Under two sufficient geometric conditions on the exponent vectors of the posynomials – the *pointed* exponents property and the existence of a *colorful* vector – we show that a solution of tropical posynomial systems can be found by resolution of a linear program (Theorem 6.7). The same conditions imply the existence of solutions to classical posynomial systems, that can be determined by solving a geometric program (Theorem 6.15). We establish that the tropical posynomial systems that are pointed and that admit a colorful vector are equivalent to Markov decision processes up to a change of variables (Theorem 6.12). We build on this result to reduce the effective discount factor of discounted MDPs. We also study the properties of the set of colorful vectors associated with convex bodies. We recall a result of [LS09] showing that this set is the interior of a simplex if nonempty. We conjecture a necessary and sufficient condition for this set to be nonempty, that we prove up to dimension 3 (Proposition 6.24).

Part II of the thesis is an in-depth case study of the PFAU, the 17–18–112 police and firefighters emergency call center of Paris area. The originality of this platform is threefold: first it is the largest emergency call center in France, second it gathers under the same roof two distinct emergency missions (security and rescue) most of the time physically separated in the rest of the country, and finally implements a bilevel (or two-tier) procedure to handle inbound calls. In this case study, the analytical results of Part I are challenged by a simulation approach with a very high degree of fidelity with respect to the real organization. This work is part of a collaboration with the *Préfecture de Police de Paris* initiated in 2015, and was funded by the *Institut des Hautes Études du Ministère de l'Intérieur* (IHEMI). It is here presented in English, while a French version can be found in [ABCG21].

- Chapter 7 presents the analysis of the call logs for more than 5 millions of emergency calls received by the PFAU in 2019, in the spirit of the study [BGM⁺05b] carried out for a commercial call center. We detail the arrival rates of the calls received through the three emergency numbers 17, 18 and 112, the distribution of conversation times at both levels of instruction depending upon the calls severity degree, and the average staffing levels over the time. The scale of the study and the fine granularity of the analysis seem unprecedented in the academic literature on emergency call centers. We are able to detect phenomena on the arrivals that are invisible to the answering operators, and we highlight unforeseen correlations between several primitive variables such as conversation times at the different levels of instruction.

Building on statistical estimators used in survival analysis, we reconstruct the patience of the callers at both levels of the instruction. We prove that the bilevel system increases the patience of the callers between the two steps of the instruction, and so does the escort mechanism for very urgent calls. We also present unaveraged data corresponding to particular events.

- Chapter 8 uses a dedicated simulation software to evaluate the performance of the PFAU for different staffing policies. This is done by simulating the execution of very realistic random scenarios which implement the arrivals, conversation times and patience levels computed in the Chapter 7. A first interest of the simulation technique is to compare the precision of the analytical formulas derived in the Part I using continuous-relaxed Petri nets without abandonment. We show in the case of the PFAU that our polyhedral throughput expressions are good estimates of the real simulated throughput, in particular deep inside the linearity regions. We also show that the Erlang formulas – that are very often used for the sizing of call centers – perform reasonably well as far as the proportion of abandonment is concerned, but that they do not precisely predict the average waiting time of callers.

A central contribution of this chapter is the comparison of the current bilevel system adopted by the PFAU and an equivalent single-tier architecture, privileged in other firefighters or police French local call centers. The main conclusion of this aspect of the study is that with the same number of operators, the bilevel system does not alter the global performance compared to the single-tier system, provided the instruction at the first level is performed quickly enough relatively to the second level – this also requires a “critical size” of the call center with a significant number of inbound calls. We confirm that the two-tier organization offers a much better service to the life-threatening calls when there are too many calls or not enough operators, which is why this designed was brought up in Paris.

Finally, we use the simulation techniques to prospectively estimate the performance of some modifications of the organization considered by the PFAU, in particular the merger of some pools of operators. We are in addition able to quantify the interest of having two different forces (police and firefighters) collaborating in the same platform.

We quickly summarize our contributions by themes in the [Conclusion](#), and discuss there some perspectives and open problems that are left for future work. We present in a final [Appendix](#) the work that we have done with the four SAMU of AP-HP during the beginning of the Covid-19 pandemic. This was a very concrete situation in which we applied our call center staffing methodology with immediate use.

4 Publications related to the thesis

4.1 JOURNAL AND CONFERENCE ARTICLES. — We mention below the articles we have published throughout this work, and what part of the manuscript are drawn from them.

- [\[ABG21\]](#) was published in the proceedings of the Petri nets 2020 conference, and later invited in the journal *Fundamenta Informaticae* for an extended version. Basic material from the Chapter 1 come from this article, as well as the three first correspondence theorems and the long-run behaviour analysis of the Chapter 4.
- [\[AABG20\]](#) was published in the proceedings of the ICMS 2020 conference, it is resumed in our Chapter 6 with additional figures and results on MDPs.
- [\[ABG22\]](#) was accepted in the Petri nets 2022 conference. Most results of the Sections 3 and 4 in Chapter 4 on the computation of transient regimes of hierarchical Petri nets are drawn from this article.
- [\[GAA⁺20\]](#) was released in the French *Comptes-Rendus de l'Académie des Sciences*. It arises from a crisis work that occurred in the Covid-19 pandemic, when we worked with the four SAMU of AP-HP to adapt our call center models and used the data at our disposal to model the epidemic evolution. It is jointly authored by a team of mathematicians, computer scientists and physicians from the SAMU.
- [\[Cov20\]](#) was published in the PLOS One medical journal, at the initiative of the crisis medical director of AP-HP and intended to a medical audience. It provides a broader perspective on epidemic indicators, featuring some material of [\[GAA⁺20\]](#).

4.2 TECHNICAL REPORTS. — In addition, our collaboration with French emergency services led to the production of two reports:

- [\[ABG20b\]](#) refers to four reports delivered to SAMU 75 (Hôpital Necker in Paris), SAMU 92 (Hôpital Raymond Poincaré in Garches), SAMU 93 (Hôpital Avicenne in Bobigny) and SAMU 94 (Hôpital Henri Mondor in Créteil), where we analyze one month of emergency call logs from their call centers and model the SAMU organization before computing practical staffing recommendations. This work gave rise to the (EMS-A) model extensively discussed in our Chapter 4. The Section 6 of the latter is drawn from our report to SAMU 94.
- [\[ABCG21\]](#) is the final report resulting from a work ordered and funded by the IHEMI, centered around the performance study of the PFAU using data analysis and numerical Petri net simulation. It constitutes the Part II of the manuscript.

Based on the study presented in the studies of [\[ABG20b\]](#), and on the methods presented in Chapters 4 and 5 involving the analytic computation of the congestion-fluid regime thresholds, we provided a prototype of emergency call center dimensioning tool to the four SAMU of AP-HP.

During the Covid-19 crisis, we had the opportunity to develop contacts with the SAMU 69 (Hôpital Édouard Herriot at Lyon). We developed a specific model of their call center, and adapted our dimensioning tool for them, also based on the present modeling and analytic methodology.

Notation

We briefly list below some notational conventions used in the manuscript:

- In all the thesis except in Chapter 6, we use the “hat” superscript to refer to random variables, like \widehat{U} . In this case, no hat (for instance U) means that we are talking about a realization of the random variable \widehat{U} .
- In any vector space, we use the notation $\mathbf{0}$ to stand for the zero vector, except for function spaces where we add a tilde and denote the null function by $\tilde{\mathbf{0}}$. Likewise, the notation $\mathbf{1}$ and $\tilde{\mathbf{1}}$ respectively refer to the unit vector of a vector space (with all components set to one) and the constant function equal to 1.
- Whenever S is a finite set, we denote by \mathbb{R}^S the vector space over \mathbb{R} with dimension $|S|$, isomorphic to the space of functions from S to \mathbb{R} . In addition if x is a vector of \mathbb{R}^S , we may use the notation x_s to refer to the coordinate of x associated with the element s of S . As customary in Markov chain analysis, we may as well write $x(s)$ (in particular in Chapters 3 and 4).
- If x is a vector of \mathbb{R}^S and S' is a subset of S , we denote by $x|_{S'}$ the vector of $\mathbb{R}^{S'}$ that is the restriction of x to the elements of S' . The same goes for matrices, with the notation $M|_{A' \times B'}$ if M is a matrix of $\mathbb{R}^{A \times B}$, $A' \subset A$ and $B' \subset B$.
- When working with square matrices, the notation I always denotes the identity matrix, regardless of the dimension.
- We use the symbol $\mathbb{1}$ to denote the indicator function of any set (equal to one over this set, zero elsewhere). It is sometimes identified with a vector.
- The floor and ceiling functions are respectively denoted by $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$.
- The dot product between two vectors is denoted by $\langle \cdot, \cdot \rangle$.
- The symbols \odot and \oslash respectively denote the Hadamard (componentwise) products and division of vectors.
- The conic and convex hulls of a set X are denoted by $\text{cone}(X)$ and $\text{conv}(X)$.

Acknowledgement

This doctoral thesis was supported by a joint PhD grant from Inria and AID (*Agence Innovation Défense*, part of the French *Direction Générale de l'Armement*).

Part I

**PIECEWISE-AFFINE DYNAMICS
OF FLUID TIMED PETRI NETS**

1

DISCRETE DYNAMICS OF TIMED PETRI NETS

Contents

1	Basics of untimed Petri nets	18
1.1	Topology of a Petri net	18
1.2	Markings and transition firings	19
1.3	Elementary incidence patterns	21
1.4	Petri nets with multipliers	22
1.5	Petri nets controlled by external inputs	23
2	Addressing underdetermined situations	23
2.1	Preselection routing	24
2.2	Priority routing	26
2.3	Simultaneous firing of non-conflicting transitions	27
3	Timed Petri nets	31
3.1	Framework and semantics of timed Petri nets	31
3.2	Counter functions of the net	34
3.3	Bringing stochasticity to the nets	41

There are several reasons why we choose to work with Petri nets in some parts of this thesis. As detailed in the [Introduction](#), they offer a great modeling power to describe discrete-event systems. The notions of synchronization and concurrency were prime motivations of the work of Carl Adam Petri [[Pet62](#)]. Various enhancements of the formalism – bringing time into play or allowing fine resource allocation schemes – subsequently made Petri nets a very popular modeling tool in applied mathematics. We refer for instance to the following reviews in the field of biology [[HR04](#)], workflow management [[VdA98](#)], production scheduling [[TB07](#)], large computing infrastructures and web services [[CEHMR12](#)] or urban traffic [[NRA13](#)] (and this list is not exhaustive). Petri nets also come with simple graphical representations, which make them easy to present in the layman’s term. Our main interest in Petri nets is however the dynamical systems they give rise to

and that govern their evolution. These switched and controlled dynamics are actually the core object that we study in this thesis.

In Section 1, we focus on the so-called “untimed” Petri nets, that is to say the ground models introduced in the 1960s. These were widely used to study structural properties of systems, in particular to detect possible blockages. The purpose of Section 2 is to somehow “rigidify” the setting of untimed Petri nets to guarantee that they behave in a univocal manner, without resorting to the intervention of an external observer. Thus, they give rise to a single and well-posed evolution sequence. We finally introduce time in Section 3, that enriches further the modeling power of Petri nets, and that makes it quite natural to study their behaviour under the perspective of controlled dynamical systems. In particular, we focus on the evolution of the counter functions of the net. We also discuss several ways in which nondeterminism can be brought to models of timed Petri nets.

Some original material on the uniqueness of Petri net trajectories is developed in Sections 2 and 3. For the rest of the chapter, that is a synthesis of basic definitions and main properties of Petri nets, we rely on the monographs of Baccelli, Cohen, Olsder and Quadrat [BCOQ92], of Heidergott, Olsder and van der Woude [HOvdW05], of Ajmone Marsan, Balbo, Conte, Donatelli and Franceschinis [AMBC⁺98] and of Reisig [Rei13]. We also refer the reader to the surveys of David and Alla [AD94] and of Bowden [Bow00].

1 Basics of untimed Petri nets

This section gives the formal definition of a Petri net and its evolution semantics, as well as the graphical representation conventions that we shall use throughout the thesis.

1.1 TOPOLOGY OF A PETRI NET. — The fundamental and common layer to all the models of Petri nets that we shall encounter involves specifying the connections between objects of two kinds: places and transitions.

► **DEFINITION 1.1.** *A **Petri net** is a triple $(\mathcal{P}, \mathcal{Q}, \mathcal{E})$, where \mathcal{P} is a finite set whose elements are called **places**, \mathcal{Q} is a finite set whose elements are called **transitions**, and \mathcal{E} is a subset of $(\mathcal{P} \times \mathcal{Q}) \cup (\mathcal{Q} \times \mathcal{P})$ representing the **incidence relation** between places and transitions.*

It comes natural in practice to assume that $\mathcal{P} \cap \mathcal{Q} = \emptyset$ and $\mathcal{P} \cup \mathcal{Q} \neq \emptyset$. A Petri net can then be seen as a **directed bipartite graph**, with vertices set $\mathcal{P} \cup \mathcal{Q}$ and arcs set \mathcal{E} . In particular, if p is a place of \mathcal{P} and q a transition of \mathcal{Q} , there is an arc from p to q (resp. q to p) if and only if (p, q) (resp. (q, p)) belongs to \mathcal{E} .

If x and y are two vertices of $\mathcal{P} \cup \mathcal{Q}$ such that the arc $x \rightarrow y$ is in the Petri net (i.e., the pair (x, y) is in \mathcal{E}), we say that x is an **upstream** vertex of y , and that y is a **downstream** vertex of x . The set of upstream (resp. downstream) vertices of x is denoted by x^{in} (resp. x^{out}). Due to the bipartite character of the incidence relation, upstream or downstream vertices of a place (resp. a transition) are transitions (resp. places).

We say that two vertices x and y create a self-loop in the Petri net if both pairs (x, y) and (y, x) are in \mathcal{E} . Petri nets with no self-loops are called **pure**, and constitute the most widely used setting in the literature. The incidence relation of pure Petri nets can be represented by a $|\mathcal{P}| \times |\mathcal{Q}|$ matrix with elements in $\{-1, 0, 1\}$, denoted by C and referred to as the **incidence matrix**. This is

done by setting for all places p in \mathcal{P} and transitions q in \mathcal{Q}

$$C_{pq} := \begin{cases} 1 & \text{if } (q, p) \in \mathcal{E} \\ -1 & \text{if } (p, q) \in \mathcal{E} \\ 0 & \text{otherwise} \end{cases} . \quad (1.1)$$

The incidence matrix C is customarily decomposed under the form $C = C^+ - C^-$, where C^+ and C^- are two 0–1-matrices with same dimension than C , respectively indicating the places to transitions and transitions to places incidences. For a pure Petri net, the specification of the set \mathcal{E} is equivalent to the one of the matrix C , or to the knowledge of the matrices C^+ and C^- .

Example 1.2. We depict below in Figure 1.1 a pure Petri net with places set $\mathcal{P} = \{p_1, p_2, p_3, p_4\}$ and transitions set $\mathcal{Q} = \{q_1, q_2, q_3\}$. By convention, places are depicted by circles, while transitions are depicted by thick segment lines. Just like usual (bipartite) directed graphs, incidence between a place-transition pair is represented by one-way arrows.

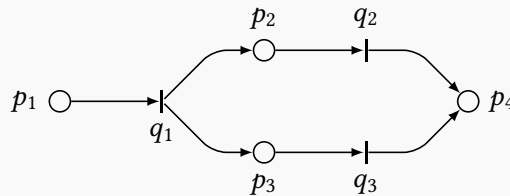


Figure 1.1: A pure Petri net with four places and three transitions

The incidence relation associated with the Petri net of Figure 1.1 amounts to seven pairs, namely $\mathcal{E} = \{(p_1, q_1), (q_1, p_2), (q_1, p_3), (p_2, q_2), (p_3, q_3), (q_2, p_4), (q_3, p_4)\}$, but is more conveniently represented by the equivalent incidence matrix

$$C = \begin{pmatrix} -1 & 0 & 0 \\ 1 & -1 & 0 \\ 1 & 0 & -1 \\ 0 & 1 & 1 \end{pmatrix} \begin{matrix} p_1 \\ p_2 \\ p_3 \\ p_4 \end{matrix} . \\ \begin{matrix} q_1 & q_2 & q_3 \end{matrix}$$

1.2 MARKINGS AND TRANSITION FIRINGS. — A Petri net starts to come alive once some of its places are populated with tokens, that we keep count of using the notion of marking.

► **DEFINITION 1.3.** Given a Petri net $(\mathcal{P}, \mathcal{Q}, \mathcal{E})$, we call *marking* (or sometimes *state* of the Petri net) an integer-valued vector of $\mathbb{N}^{\mathcal{P}}$. If m is a marking and p is a place of \mathcal{P} , the integer m_p gives the number of *tokens* present in p .

Although introduced in an abstract way, tokens shall be used in practice to represent multiple kind of information: pending job, available resource, operation in progress, etc. Only places are equipped with tokens, and may be thought of as “tokens reservoirs”. By definition, a Petri net undergoes a change of state when tokens appear or vanish in some of its places. This can only happen if a transition is fireable.

- **DEFINITION 1.4.** For a Petri net $(\mathcal{P}, \mathcal{Q}, \mathcal{E})$ with marking m in $\mathbb{N}^{\mathcal{P}}$ and associated incidence matrix decomposition (C^+, C^-) in $\mathbb{N}^{\mathcal{P} \times \mathcal{Q}}$, we say that the transition q in \mathcal{Q} is **fireable** if $m \geq [C^-]_q$, where $[C^-]_q$ denotes the q -column of the matrix C^- and the inequality is taken entrywise.

If the transition q is fireable for the marking m , we further say that q **fires** if the Petri net evolves in state m' , such that $m' = m + [C^+]_q - [C^-]_q = m + [C]_q$.

In other words, the firing of a transition q is possible only if there are tokens present in each upstream place of q (i.e., each place of q^{in}). When performed, the firing of q *consumes* one token in each place of q^{in} and *produces* one token in each place of q^{out} (i.e., each downstream place of q). The following principle epitomizes this fundamental evolution rule:

In a Petri net, tokens pop up in places or are cleared out from them as transitions fire.

We point out that it would be slightly misleading to say that “tokens go from places to places by crossing transitions”. Indeed, depending on the incidence relation \mathcal{E} , the firing of a transition may not preserve the total amount of tokens at stake. In addition, the tokens of two different places may represent objects of a very heterogeneous nature. As a consequence, the idea of identical tokens traveling through the net should be considered with caution.

The firing condition of a transition q may still be satisfied after one or multiple firings of q . Alternatively, an initially non fireable transition q' may become fireable after the firing of q . To deal with successive firings of one or multiple transitions, the notion of firing sequence comes in handy. Given an **initial marking** m^0 in $\mathbb{N}^{\mathcal{P}}$ and ℓ in $\mathbb{N} \cup \{\infty\}$, we say that a (possibly finite) sequence of transitions $\sigma = (q^{(1)}, q^{(2)}, \dots, q^{(k)}, \dots)$ of \mathcal{Q}^ℓ is an **admissible firing sequence** for m^0 if for all $1 \leq k \leq \ell$, the transition $q^{(k)}$ is fireable for the marking $m^0 + \sum_{i=1}^{k-1} [C]_{q^{(i)}}$. For such a finite sequence σ and denoting by $|\sigma|$ the vector of $\mathbb{N}^{\mathcal{Q}}$ whose q -coordinate gives the number of occurrences of transition q in σ , the final marking of the Petri net starting from initial marking m^0 is equal to $m^0 + C|\sigma|$. This equation is sometimes referred to as the fundamental equation of Petri nets.

Example 1.5. We illustrate the notion of transition firing in Figure 1.2, which depicts a Petri net with four places and with marking $m^0 = (3 \ 0 \ 1 \ 0)^\top$. The transition q_4 is not fireable due to the absence of tokens in p_2 , however both q_2 and q_3 are fireable owing to the presence of tokens in p_1 .

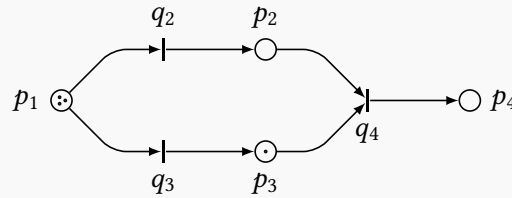


Figure 1.2: A Petri net with two marked places

Once q_2 is fired (consuming a token in p_1 and producing a token in p_2), the transition q_4 becomes fireable. Firing q_4 has the effect of consuming a token in both p_2 and p_3 and producing a token in p_4 . We let the reader check that both $(q_2 \ q_2 \ q_2 \ q_4)$ and $(q_2 \ q_4 \ q_3 \ q_2 \ q_4)$ are admissible firing sequences from m^0 , and that they respectively lead to final markings $(0 \ 2 \ 0 \ 1)^\top$ and $(0 \ 0 \ 0 \ 2)^\top$.

A lot of attention in the Petri net literature has been received by the **reachability** problem, i.e., determining for two markings m^0 and m in $\mathbb{N}^{\mathcal{P}}$ if there is a firing sequence σ such that $m = m^0 + C|\sigma|$. In this case, we say that m is **reachable** from m^0 . In various applications, reaching a specific marking can model the termination of a process of interest, or the triggering of some key step, etc. It was further observed that numerous other problems on Petri nets reduce to the reachability question, see for instance [Hac76a]. In particular, it is often considered that a sound behaviour for a Petri net is that the latter stays **bounded** from a given initial marking m^0 , which means that no reachable marking from m^0 can have an arbitrary large amount of tokens in some place. We shall also be interested in Petri nets that are **live** for an initial marking m^0 , meaning that for every reachable marking m from m^0 and transition q of \mathcal{Q} , there exists a firing sequence σ such that q is fireable for the marking $m + C|\sigma|$. In other words, it is always possible to make any transition fireable indefinitely many times. Conversely, we say that the net admits a **deadlock** from m^0 if there exists a reachable marking from m^0 from which no transition is fireable anymore.

The elements of the left and right nullspaces of the matrix C are often used to characterize structural properties of Petri nets. We shall use later on the following notion of T-invariants, that expresses balance equations for produced and consumed tokens in places (in the manner of a Kirchhoff's nodal rule):

- **DEFINITION 1.6.** Let $(\mathcal{P}, \mathcal{Q}, \mathcal{E})$ be a pure Petri net with associated incidence matrix C . We say that a vector y in $\mathbb{N}^{\mathcal{Q}}$ is a **T-invariant** if $Cy = \mathbf{0}$.

It is known in particular that if a bounded and deadlock-free Petri net admits an *elementary* positive T-invariant (i.e., which is not a positive combination of T-invariants with smaller supports), then it is live [LR96]. We also refer to [Reu89, EN94] for other results on structural analysis of Petri nets.

- 1.3 ELEMENTARY INCIDENCE PATTERNS.** — Topology-wise, the incidence patterns involving a place or a transition of a Petri net are interpreted differently depending on the number of upstream and downstream vertices at stake. The Figure 1.3 below depicts four elementary arrangements of places and transitions, that all incidence patterns are extensions of.

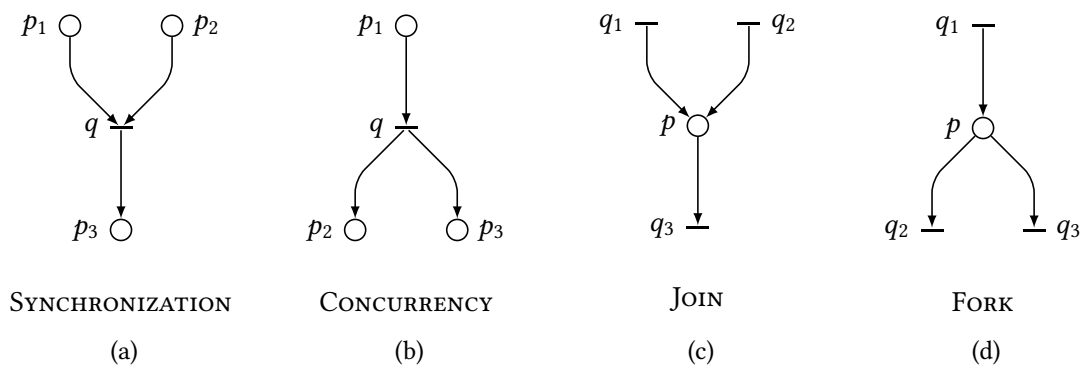


Figure 1.3: Four elementary incidence patterns

The first situation (1.3a) occurs when a transition q has multiple upstream places. As outlined previously, this transition cannot be fired until a token is present in each place of q^{in} . This lends itself well to modeling the need for a synchronization of resources. If q admits multiple

downstream places (1.3b), the firing of q produces one token in each place of q^{out} and allows to consider concurrent (or parallel) lanes. When a place p has multiple upstream transitions (1.3c), the firing of any transition of p^{in} produces tokens in p , and one can talk about tokens meeting or joining – this may be used to model the merger of several lanes into a single one. Finally, a somehow more delicate pattern arises when a place p can admit multiple downstream transitions (1.3d). Indeed, if one token lies in p , a single transition of p^{out} may be fired (provided these transitions do not have other upstream places with no tokens), and such a firing makes all the transitions of p^{out} no longer fireable. As a result, there is here room for a choice, and as illustrated by the Section 1.2 associated with the Figure 1.2, choosing which transitions to fire highly influences the possible firing sequences and resulting markings to come in the future. We shall discuss these cases in detail when introducing routing rules in Section 2.

In the case where only synchronization and concurrency patterns are involved, the class of systems at stake reduces to **event graphs**, that are defined in [BCOQ92] as Petri nets where all places have at most one upstream and one downstream transition. It comes with little surprise that the structural and the dynamical analyses of event graphs are less difficult than those of Petri nets. It is in fact customary in the Petri net literature to focus on several subclasses with extra restrictions on the net topologies, to be able to derive analytical results.

1.4 PETRI NETS WITH MULTIPLIERS. – The modeling power of Petri nets is sometimes extended by taking into account **multipliers** (or **arc weights**). Formally, one may consider that a Petri net with multipliers is defined by a triple $(\mathcal{P}, \mathcal{Q}, \mathcal{E}')$, where \mathcal{P} and \mathcal{Q} still denote finite sets of places and transitions, but \mathcal{E}' is a subset of $(\mathcal{P} \times \mathcal{Q} \times \mathbb{N}^*) \cup (\mathcal{P} \times \mathcal{Q} \times \mathbb{N}^*)$. An element (x, y, α) of \mathcal{E}' stands for an incidence relation between vertices x and y (i.e., an arc from x to y) with weight α in \mathbb{N}^* . If a Petri net with multipliers is pure (in the sense that the underlying Petri net with no multipliers is pure), the weighted incidence relation \mathcal{E}' is conveniently represented by a $|\mathcal{P}| \times |\mathcal{Q}|$ matrix, similarly to the matrix C introduced in (1.1), with this time coefficients in \mathbb{N} . If p is a place of \mathcal{P} and q a transition of \mathcal{Q} , we indeed let $C_{pq} = \alpha$ if (q, p, α) belongs to \mathcal{E}' , $C_{pq} = -\alpha$ if (p, q, α) belongs to \mathcal{E}' and $C_{pq} = 0$ otherwise. Note that the resulting matrix C still admits a unique decomposition into the form $C = C^+ - C^-$, where the matrices C^+ and C^- have nonnegative entries. Following the convention of [CGQ98], if a Petri net with multipliers features a $p \rightarrow q$ arc from place p to transition q , we shall refer to its weight by α_{qp} , while the weight of a $q \rightarrow p$ arc from a transition q to a place p shall be denoted by α_{pq} .

When working with Petri nets with multipliers, the condition for a transition q of \mathcal{Q} to be fireable from a marking m in $\mathbb{N}^{\mathcal{P}}$ that was introduced in Definition 1.4 is unchanged. It remains $m \geq [C^-]_q$, and the new marking obtained after the firing of q is still given $m + [C]_q$. This means that a *single* firing of transition q consumes α_{qp} tokens in place p for each upstream place p in q^{in} , and produces α_{pq} tokens in place p for each downstream place p in q^{out} . Allowing weights in Petri nets can be used in practice to represent the fact that some operations require multiple resource of the same kind to be performed, or conversely provide multiple same elements when achieved once (e.g., buying *one* box of eggs actually gives you *six* eggs).

Note that the definition of an admissible firing sequence σ for a marking m^0 introduced in Section 1.2 extends to Petri nets with multipliers, and the resulting marking m after applying σ still verifies $m = m^0 + C|\sigma|$. Furthermore, Definition 1.6 also carries over and we say that a vector y in $\mathbb{N}^{\mathcal{Q}}$ is a T-invariant of a pure Petri net with multipliers if $Cy = \mathbf{0}$.

Example 1.7. Consider the pure Petri net with multipliers depicted in Figure 1.4. The arc weights are represented along the arcs if not equal to one.

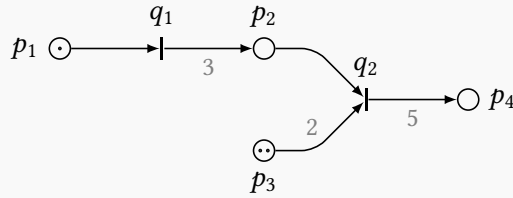


Figure 1.4: A Petri net with multipliers

The weighted incidence matrix C associated with this Petri net is given by

$$C = \begin{array}{cc} \begin{pmatrix} -1 & 0 \\ 3 & -1 \\ 0 & -2 \\ 0 & 5 \end{pmatrix} & \begin{matrix} p_1 \\ p_2 \\ p_3 \\ p_4 \end{matrix} \\ \begin{matrix} q_1 & q_2 \end{matrix} & \end{array} .$$

An initial marking $m^0 = (1 \ 0 \ 2 \ 0)^T$ is also represented in Figure 1.4. Observe that $(q_1 \ q_2)$ is an admissible firing sequence, and it leads to the new marking $(0 \ 2 \ 0 \ 5)^T$.

- 1.5 PETRI NETS CONTROLLED BY EXTERNAL INPUTS.** — So far, in the class of Petri nets that we have described, the knowledge of a topology and an initial marking is sufficient to determine the admissible firing sequences of the net, and thus its reachable states. Alternatively, one may want to have leverage on the fireable character of certain transitions, for instance temporarily halt a part of a chain of treatment, regulate the amount of tokens present in the net, etc. To this purpose, it is useful to introduce **controlled Petri nets**, that are Petri nets of the form $(\mathcal{P}, \mathcal{Q}, \mathcal{E})$ in which a subset $\mathcal{Q}_{\text{input}}$ of \mathcal{Q} is distinguished, with elements called **input transitions**. By convention, these transitions have no upstream places, and their firings are prescribed externally, hence producing tokens in all places they are located upstream of. In fact, if q_{input} is an input transition, we can always suppose without loss of generality that q_{input} admits a single downstream place p_{input} , that the latter admits a single downstream transition q_{internal} in $\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$, and conversely that $p_{\text{input}}^{\text{in}} = \{q_{\text{input}}\}$.

Allowing input transitions raises questions that are customary in control theory. In particular, one would be interested in the *controllability* of a Petri net, that is to say the possibility to enforce specific markings or behaviours of the net (liveness, boundedness, deadlock avoidance, etc.) by choosing adequate input firings. More broadly, one can also study the *stability* of the system, looking for equilibria and characterizing their properties.

2 Addressing underdetermined situations

The Petri nets that are not “controlled” in the sense of Section 1.5 and where $\mathcal{Q}_{\text{input}}$ is empty are sometimes called “autonomous” Petri nets, for they do not rely on any external commands. This designation is somewhat misleading, since whether there are input transitions or not, the

associated Petri nets may still require the intervention of an observer to make some *choices*. For instance, the “fork” pattern featured in the Figure 1.3 of Section 1.3 is fundamentally undetermined; if several fireable transitions can consume the same resource, one needs to decide which one of them should be actually fired. This situation is often referred to as a structural conflict:

- **DEFINITION 1.8.** We say that two transitions q and q' of Q are in *structural conflict* if they have a common upstream place, or in other words $q^{\text{in}} \cap q'^{\text{in}} \neq \emptyset$.

We denote by $\mathcal{P}_{\text{conflict}}$ the subset of \mathcal{P} consisting of the places involved in structural conflicts, i.e., $\mathcal{P}_{\text{conflict}} := \{p \in \mathcal{P} : |p^{\text{out}}| > 1\}$. The term “structural” is used to precise that the conflict is considered regardless of the marking of the net. Some authors convene that these conflicts are not “effective” if enough tokens are present in the common upstream places mentioned in Definition 1.8 to fire all the downstream transitions (see [AMBC⁺98, Chapter 2]).

Nonetheless, for a Petri net to behave autonomously, it is desirable to be able to *arbitrate* any such conflict. This is the purpose of *routing rules*, or routing policies, that will in effect restrict admissible firing sequences. In the following two subsections, we present two deterministic routing rules that we shall use in practice in our applications. The triple $(\mathcal{P}, Q, \mathcal{E})$ will denote a pure Petri net with multipliers, and m^0 an initial marking of the latter.

- 2.1 PRESELECTION ROUTING.** — We say that a place p of \mathcal{P} is ruled by a *preselection* (or pre-allocation) routing rule if a transition of p^{out} is associated to each token of p upon its production. This token is then *reserved* to fire its related transition, and cannot be consumed by any other transition of p^{out} , even if they are fireable. We denote by $\mathcal{P}_{\text{psel}}$ the set of places ruled by preselection, and by Q_{psel} the set of transitions that are downstream of a place ruled by preselection, i.e., $Q_{\text{psel}} := \{q \in Q \mid q^{\text{in}} \cap \mathcal{P}_{\text{psel}} \neq \emptyset\}$.

Formally, the preselection routing for a place p is described by a collection of nondecreasing mappings $(\Pi_q^p)_{q \in p^{\text{out}}}$ from \mathbb{N}^* to \mathbb{N}^* , and nonnegative integers $(m_{p,q})_{q \in p^{\text{out}}}$, satisfying the property:

$$\forall n \in \mathbb{N}^*, \quad \sum_{q \in p^{\text{out}}} \Pi_q^p(n) = n, \quad \text{and} \quad \sum_{q \in p^{\text{out}}} m_{p,q} = m_p^0. \quad (1.2)$$

For q in p^{out} , $\Pi_q^p(n)$ represents the number of tokens which are reserved to fire transition q , amongst the n first tokens *produced* in place p (i.e., not including the initial marking), while $m_{p,q}$ stands for the number of tokens reserved amongst the initial marking m_p^0 to fire q . Observe that such a routing rule prescribes the fate on any token that will visit p , with no dependence on the global topology of the net nor on the markings of other places.

A trivial example of preselection routing occurs when for some q in p^{out} , we have $\Pi_q^p(n) = n$ for all n in \mathbb{N}^* , that is, all the tokens produced in p are routed towards q , and we talk about a 0–1 allocation. Another more general instance is given by the *periodic routings*. If $p^{\text{out}} = \{q_1, q_2, \dots, q_k\}$, consider a positive integer L , a partition (J_1, J_2, \dots, J_k) of $\{1, 2, \dots, L\}$ and define $\Pi_{q_k}^p(n) = \text{card}(\{1, 2, \dots, n\} \cap (J_k + L\mathbb{N}))$. Such maps are indeed nondecreasing, periodic (with period L), and satisfy the requirements of (1.2). Observe that for large values of n , we have $\Pi_{q_k}^p(n) \sim n \cdot \text{card}(J_k)/L$. More general and abstract preselection routings that are not periodic can be built from Sturmian words, see for instance [Lot02] and [MV98].

Example 1.9. Consider the pure Petri net with three places and four transitions depicted in Figure 1.5. It has initial marking $m^0 = (1 \ 0 \ 0)^T$.

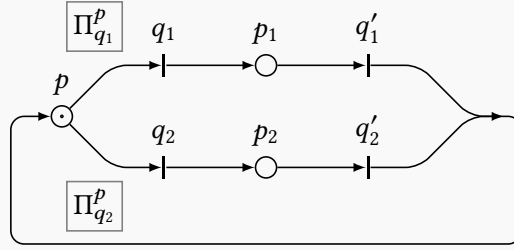


Figure 1.5: A Petri net with a place ruled by periodic preselection routing

The transitions q_1 and q_2 are in structural conflict for they share the common upstream place p . We here assume that p is subject to a periodic preselection routing, associated with the two mappings $\Pi_{q_1}^p$ and $\Pi_{q_2}^p$ such that for all n in \mathbb{N}^* , we have

$$\Pi_{q_1}^p(n) := \text{card}(\{1, \dots, n\} \cap 3\mathbb{N}) \quad \text{and} \quad \Pi_{q_2}^p(n) := \text{card}(\{1, \dots, n\} \cap (3\mathbb{N} + \{1, 2\})).$$

We use the “boxed” notation in Figure 1.5 to indicate the preselection mappings at stake. We also let $m_{p,q_1} := 1$ and $m_{p,q_2} := 0$. This means that the first token that is in p , present in the initial marking, shall be routed towards transition q_1 . Afterwards, the only fireable transition is q_1' and its firing produces a new token in p , that will be directed towards q_2 . The same goes for the second token produced in p . However, the third token to pop up in p shall be routed towards q_1 . As a result, denoting by $m^1 = (0 \ 1 \ 0)^\top$ and $m^2 = (0 \ 0 \ 1)^\top$, observe that the visited states form the sequence $(m^0, m^1, m^0, m^2, m^0, m^2, m^0, m^1, \dots)$. In particular, this Petri net is live and bounded.

To prevent any ambiguity in the future models to be introduced, we shall assume that preselection routings are only allowed for places whose downstream transitions do not admit other upstream places, i.e., $\mathcal{P}_{\text{psel}} \subset \{p \in \mathcal{P} \mid \forall q \in p^{\text{out}}, q^{\text{in}} = \{p\}\}$, or equivalently, $|q^{\text{in}}| = 1$ for all q in $\mathcal{Q}_{\text{psel}}$. This entails no loss of generality; indeed for a place p in $\mathcal{P}_{\text{psel}}$ and a transition q in p^{out} , if we want to take into account another constraint on q and therefore another upstream place p' in q^{in} , we may as well consider a *modified* Petri net with an extra place p^* and an extra transition q^* , such that $p^{*\text{in}} = \{q\}$, $p^{*\text{out}} = \{q^*\}$, and where the arc $p' \rightarrow q$ is replaced by a $p' \rightarrow q^*$ arc. This scheme, depicted in Figure 1.6, consists in introducing the “buffer” place p^* so that q indeed admits p as its single upstream place and q^* still fulfills the synchronization role with p' . Note that if p' itself is involved in another preselection rule, this modification may be applied twice (and in general, repeated for all transitions of p^{out}). We shall actually use the configuration depicted in Figure 1.6a (not allowed in our setting) to stand as a shorthand notation for Figure 1.6b.

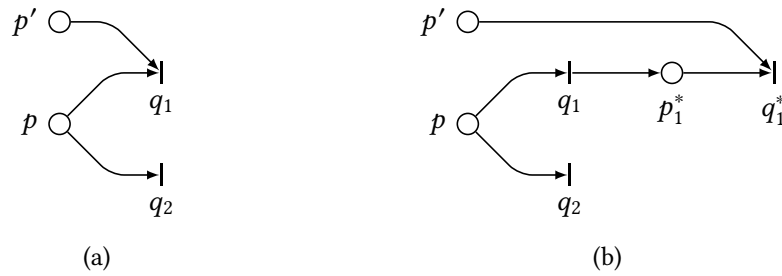


Figure 1.6: Transformation of a Petri net to ensure that transitions located downstream of a place ruled by preselection do not have other upstream places

2.2 PRIORITY ROUTING. — Another way to arbitrate a structural conflict between two transitions q and q' is to give **priority** to one of them for the consumption of tokens present in a place of $q^{\text{in}} \cap q'^{\text{in}}$. More generally, we may need to provide a priority order on multiple transitions to obtain an unequivocal firing sequence.

We say that a place p in \mathcal{P} is ruled by priority routing, and we denote by $\mathcal{P}_{\text{prio}}$ the set of such places, if the set p^{out} of downstream transitions of p is equipped with a total order ($<_p$). The principle of this routing rule is that a transition q in p^{out} can be fired only if it is fireable and there is no other fireable transition q' in p^{out} with a higher priority, i.e., $q' <_p q$, or equivalently $q >_p q'$ (we convene that the minimal element of p^{out} with respect to $<_p$ is the highest priority transition). In graphical representations (see the examples below), we will use increasing number of arrow tips on arcs to account for priorities, the transition with highest priority level being pointed by the highest number of tips. As before, we denote by $\mathcal{Q}_{\text{prio}}$ the set of transitions downstream of a place ruled by priority routing, i.e., $\mathcal{Q}_{\text{prio}} := \{q \in \mathcal{Q} \mid q^{\text{in}} \cap \mathcal{P}_{\text{prio}} \neq \emptyset\}$.

Priority routings are particularly interesting when for a place p in $\mathcal{P}_{\text{prio}}$, some of the downstream transitions of p admit other upstream places (otherwise it reduces to a 0–1 allocation such as introduced in Section 2.1). It can then model subtle conditional behaviours involving the constraints associated with different places. Hence, and contrary to the preselection routing, the priority routing from a place is a dynamical allocation that may require the knowledge of the marking in other places of the net. The following example (such as the Example 1.22 of Section 3 to come) illustrates well this observation.

Example 1.10. We show in Figure 1.7 a simple setting with three places and two transitions where a priority rule is involved.

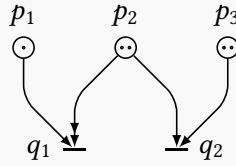


Figure 1.7: A Petri net with a place ruled by priority routing

The place p_2 is ruled by priority and according to the graphical convention above, we have $q_1 <_{p_2} q_2$, which means that whenever q_1 and q_2 are both fireable, we should fire q_1 . This is here the case since the initial marking is $(1 \ 2 \ 2)^T$. Observe although that after firing q_1 once, there is no more tokens in p_1 , as a consequence q_1 is no longer fireable, but q_2 still is. The only admissible firing sequence here is thus $(q_1 \ q_2)$. On the other hand, should the priority levels be flipped, the only admissible firing sequence would have been $(q_2 \ q_2)$.

Although we allow transitions in $\mathcal{Q}_{\text{prio}}$ to admit multiple upstream places, a question of consistency arises if several of these upstream places are ruled by priority. We shall allow this setting as long as the following compatibility condition is met.

► **DEFINITION 1.11.** We say that the rules $(<_p)_{p \in \mathcal{P}_{\text{prio}}}$ are **compatible** if their union as binary relations is **acyclic**.

Acyclicity means that the transitive closure of the union of the local total orders $(<_p)_{p \in \mathcal{P}_{\text{prio}}}$ forms a global partial order on the set \mathcal{Q} of all transitions. The Figure 1.8a shows an example of two compatible priority rules. Indeed, we have $q_1 <_{p_1} q_3 <_{p_1} q_2$ and $q_3 <_{p_2} q_2 <_{p_2} q_4$, so

that the orders (\prec_{p_1}) and (\prec_{p_2}) are compatible with the global linear extension (\prec) for which $q_1 \prec q_3 \prec q_2 \prec q_4$. However, no such global ordering can be found in the setting of Figure 1.8b, due to the cyclic pattern $q_3 \prec_{p_1} q_2 \prec_{p_2} q_3$.

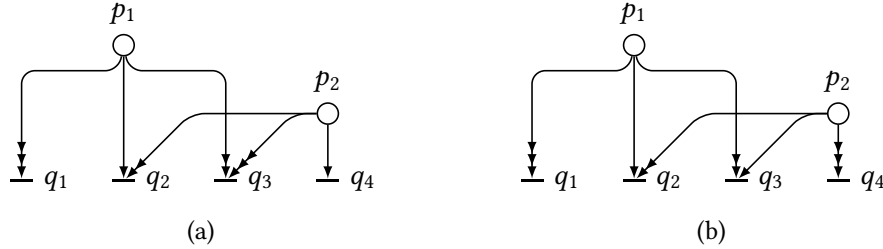


Figure 1.8: Two Petri nets with two places ruled by compatible (left) and incompatible (right) priority rules

Remark that the priority routing – and once again contrary to the preselection routing – is essentially “non-monotonic”. This means that it is possible that a token is produced in a place p earlier than another token is produced in a place p' , and that this second token in p' is consumed *before* the first one in p , owing to the priorities rules of the net. This simple observation shall have huge consequences on our analysis of Petri nets. In fact, such non-monotonic allocations are not totally standard in most models of Petri nets. As a result, we shall always explicitly mention in this thesis when a studied Petri net uses priority rules; otherwise it is by default “priority-free”.

It is also worth mentioning that many authors do not use local priority rules as introduced here, but adopt **inhibitor arcs** instead. By definition, if there is a $p \rightarrow q$ inhibitor arc for some p in \mathcal{P} and some q in \mathcal{Q} , then q cannot fire as long as tokens are present in p , see for instance [AMBC87]. As shown by Hack in [Hac76b] (see also [CDF91]), the use of two-level priorities can be equivalently replaced by inhibitors arcs. In particular, building on the result of Peterson [Pet81], we can state that allowing priority rules confer Petri nets the modeling power of Turing machines.

Observe that by construction, we have that $\mathcal{Q}_{\text{psel}} \cap \mathcal{Q}_{\text{prio}} = \emptyset$. We denote by $\mathcal{Q}_{\text{sync}}$ the set $\mathcal{Q} \setminus (\mathcal{Q}_{\text{input}} \cup \mathcal{Q}_{\text{psel}} \cup \mathcal{Q}_{\text{prio}})$ of transitions only subject to synchronization of upstream places (possibly a single one) and that admit at most one downstream place. The whole set of transitions \mathcal{Q} is therefore partitioned into input transitions, synchronization transitions, transitions that are downstream of places ruled by preselection and finally transitions that are downstream of places ruled by priority.

2.3 SIMULTANEOUS FIRING OF NON-CONFLICTING TRANSITIONS. – Arbitrating situations of structural conflict using routing rules is necessary yet not quite sufficient to eliminate equivocal behaviours of Petri nets. Consider indeed the following example, which is drawn from the book of Ajmone Marsan et al. [AMBC⁺98, Figure 2.11].

Example 1.12. The Figure 1.9 represents a net involving at the same time a synchronization pattern (we have arcs $p_2 \rightarrow q_3$ and $p_3 \rightarrow q_3$) and a conflict situation (due to the two arcs $p_2 \rightarrow q_2$ and $p_2 \rightarrow q_3$). The places p_1 and p_2 have an initial marking of one token.

In this Petri net, both the transitions q_1 and q_2 are fireable at the beginning of the execution, and three different (maximal) firing sequences are admissible. The two sequences $(q_1 q_2)$ and $(q_2 q_1)$ lead to the same final marking; so we might want to say that the firings of q_1 and q_2 “commute”. However, firing q_1 first without firing q_2 then allows to fire q_3 (i.e., the firing

sequence $(q_1 q_3)$ is admissible), while the opposite is not true. As a result, even though they are not in structural conflict, the order of firing of q_1 and q_2 does matter.

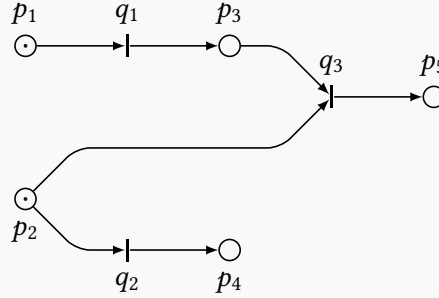


Figure 1.9: A “confusion” pattern

The situation described in Example 1.12 is called a **confusion** in [AMBC⁺98], and the issue partly comes from the fact that p_2 has multiple downstream transitions, that are therefore in structural conflict. As we have seen before, the routing rules are designed to resolve such situations. Equipping the place p_2 with a preselection routing rule as defined in Section 2.1 indeed removes underdetermination since depending on the destination transition assigned to the token present in p_2 , we will either apply the firing sequence $(q_1 q_3)$, or one of the two equivalent sequences $(q_1 q_2)$ and $(q_2 q_1)$ (under this setting, the execution of the Petri net is to be understood in the sense of the aforementioned reduction, based on the Figure 1.6). On the other hand, ruling p_2 with a priority routing may not bring more clarity. Indeed, if q_3 has a higher priority level over q_2 (meaning that $q_3 <_{p_2} q_2$), both sequences $(q_1 q_3)$ and $(q_2 q_1)$ remain admissible. Since they lead to different markings, this shows that we must rigidify the semantics further.

It is common to avoid confusion patterns by restricting to **free-choice** nets, denoting the Petri nets in which for a place p and a transition q , there can be a $p \rightarrow q$ arc (or a couple (p, q) in the incidence relation \mathcal{E}) only if $q^{\text{in}} = \{p\}$ or $p^{\text{out}} = \{q\}$ – we refer for instance to the book of Desel and Esparza [DE95]. Note that our choice in Section 2.1 to restrict preselection routings so that $|q^{\text{in}}| = 1$ for all q in \mathcal{Q}_{pse} is compatible with the free-choice setting. In free-choice nets, specifying routing rules to all the places having multiple downstream transitions suffices to obtain a unique maximal admissible firing sequence. However, this makes priority rules trivial as we pointed it out in Section 2.2. Another solution to end up with a single firing sequence without resolving to free-choice nets is to adopt another model of priorities, imposing a global firing order on all the transitions of the net, even those that are not downstream of places in $\mathcal{P}_{\text{conflict}}$ – this is the approach studied by Ajmone Marsan et al.

In this thesis, we choose to address the confusion configurations by sticking to our model of local priority rules introduced in Section 2.2 but by refining the notion of firing sequence. Indeed, the definition provided in Section 1.2 and widely used in the Petri net literature supposes that transitions in a firing sequence are fired *sequentially*, in a step-by-step fashion. Instead, we consider sensible that some transitions must fire *simultaneously*, when they are not in effective conflict. Sometimes, the *same* transition may also simultaneously fire several times, if enough resources are available for it. Formally, given a marking m in $\mathbb{N}^{\mathcal{P}}$ of the net, whose incidence matrix C still writes under the form $C = C^+ - C^-$ (see Sections 1.1 and 1.4), we say that a **multiset** $X := (\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}, \nu)$ of transitions of $\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$ (i.e., a subset of $\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$ with possibly repeated elements, the vector ν in $\mathbb{N}^{\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}}$ giving the *multiplicity* of each element) is **simultaneously fireable** if $\sum_{q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}} \nu_q [C^-]_q \leq m$. When a simultaneously fireable multiset X of transitions

with multiplicity ν fires, the vector $\sum_{q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}} \nu_q [C^-]_q$ gives the number of tokens *consumed* in each place of \mathcal{P} , while the vector $\sum_{q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}} \nu_q [C^+]_q$ gives the number of tokens *produced* in each place.

An admissible **generalized firing sequence** for a given initial marking m^0 in $\mathbb{N}^{\mathcal{P}}$ is (as in Section 1.2) a sequence $(X_1, X_2, \dots, X_k, \dots)$ of multisets of transitions of $\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$, where for all $k \geq 1$, $X_k := (\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}, \nu^{(k)})$ is simultaneously fireable for the marking $m^k := m^{k-1} + \sum_{q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}} \nu_q^{(k)} [C]_q$. Observe that the definition of a firing sequence is recovered if each multiset happens to be a singleton. In the simultaneous firing setting, the firings are performed so that at each step, the transitions of a multiset X do not have access to the tokens produced by the firings of other transitions in X . Hence, the notion of simultaneously fireable multiset is the right one to consider transitions whose firings “commute”. Naturally, when working with routing rules such as preselection or priorities, the conditions for a multiset to be simultaneously fireable are restricted in the same way as in Sections 2.1 and 2.2.

Example 1.12 (continued). Arbitrating the conflict occurring at place p_2 by the priority routing such that $q_3 \prec_{p_2} q_2$, there is a single maximal multiset that is simultaneously fireable from the initial marking of Figure 1.9, and this is $\{q_1, q_2\}$. After firing it, no more transition is fireable. Using the notion of simultaneously fireable multisets, we can thus see here that the transition q_3 can *never* be fired.

The logic behind the idea of simultaneously fireable multisets of transitions is for every marking to look for *maximal* such multisets (in the sense of the inclusion of multisets, i.e., if X and X' have multiplicity vectors ν and ν' , then $X \subset X'$ if and only if $\nu'_q \leq \nu_q$ for every q in $\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$). Due to the partially ordered structure of multisets and also to the conflict patterns, showing that such a maximum exists is not entirely trivial. The next theorem, that is constructive, addresses this question and confirms that a maximal simultaneously fireable multiset exists. Actually, the proof shows that transitions in $\mathcal{Q}_{\text{psel}}$ do not cause much problem since by construction, they admit only one upstream place that reserves tokens for them, and similarly to transitions of $\mathcal{Q}_{\text{sync}}$, the number of times that they can simultaneously fire can be maximized regardless of the behaviour of the rest of the net. On the contrary, dealing with priority rules that are sometimes mixed requires more care, because tokens need to be allocated in accordance with the local orderings.

- **THEOREM 1.13.** *If all the conflict places of a Petri net with multipliers are ruled by preselection or priority, and that the priority rules $(\prec_p)_{p \in \mathcal{P}_{\text{prio}}}$ are compatible in the sense of Definition 1.11, then for any marking of the net, there exists a maximal (and therefore unique) simultaneously fireable multiset of transitions that comply with the routing rules.*

Proof. Let $m \in \mathbb{N}^{\mathcal{P}}$ be a marking of the net. We first build a simultaneously fireable multiset of transitions X , and show in a second phase that it is maximal and unique.

◆ Our construction is based on the following algorithmic procedure:

- (0) Define $X := \emptyset$.
- (1) For all $q \in \mathcal{Q}_{\text{sync}}$, add q in X with multiplicity $\nu_q := \min_{p \in q^{\text{in}}} \lfloor \alpha_{qp}^{-1} m_p \rfloor$.
- (2) For all $q \in \mathcal{Q}_{\text{psel}}$ with upstream place p , add q in X with multiplicity $\nu_q := \lfloor \alpha_{qp}^{-1} m_{p,q} \rfloor$.
- (3) Define $m^{(1)} := m$, and set $i := 1$. Choose a global order (\prec) compatible with the orders $(\prec_p)_{p \in \mathcal{P}_{\text{prio}}}$. Label the transitions in $\mathcal{Q}_{\text{prio}}$ according to this order: $q_1 \prec q_2 \prec \dots \prec q_{|\mathcal{Q}_{\text{prio}}|}$.
- (4) Add q_i to X with multiplicity $\nu_{q_i} := \min_{p \in q_i^{\text{in}}} \lfloor \alpha_{qp}^{-1} m_p^{(i)} \rfloor$.
- (5) If $i = |\mathcal{Q}_{\text{prio}}|$, stop and return X , otherwise, define $m^{(i+1)} := m^{(i)} - \nu_{q_i} [C^-]_{q_i}$, let $i := i + 1$ and apply step (4).

- ◇ We first prove that X is indeed simultaneously fireable, that is, $\sum_{q \in X} v_q [C^-]_q \leq m$. For $\mathcal{Q}_{\text{sync}}$, $\mathcal{Q}_{\text{psel}}$ and $\mathcal{Q}_{\text{prio}}$ are pairwise disjoint, we may verify the inequality on each of these subsets. If $p \notin \mathcal{P}_{\text{conflict}}$, then $|p^{\text{out}}| \leq 1$, so we should only prove the inequality when p admits a single downstream transition $q \in \mathcal{Q}_{\text{sync}}$. We have

$$\sum_{q' \in X} v_{q'} C_{p,q'}^- = v_q C_{p,q}^- = \alpha_{qp} \lfloor \alpha_{qp}^{-1} m_p \rfloor \leq m_p .$$

If $p \in \mathcal{P}_{\text{psel}}$, we have similarly

$$\sum_{q \in X} v_q C_{p,q}^- = \sum_{q \in p^{\text{out}}} v_q C_{p,q}^- = \sum_{q \in p^{\text{out}}} \alpha_{qp} \lfloor \alpha_{qp}^{-1} m_{p,q} \rfloor \leq \sum_{q \in p^{\text{out}}} m_{p,q} = m_p .$$

Finally if $p \in \mathcal{P}_{\text{prio}}$, we prove that $\sum_{q \in X} v_q C_{p,q}^- = \sum_{i=1}^{|\mathcal{Q}_{\text{prio}}|} v_{q_i} C_{p,q_i}^- \leq m_p$ by finite induction. We rely on the invariant given by the above procedure, that stipulates $v_{q_j} [C^-]_{q_j} \leq m^{(j)}$ for all j in $\{1, \dots, |\mathcal{Q}_{\text{prio}}|\}$ (it is obtained by inequalities similar as above). This invariant specified for $j = 1$ gives the initialization step $v_{q_1} C_{p,q_1}^- \leq m_p^{(1)} = m_p$. Suppose now that for some $1 \leq j < |\mathcal{Q}_{\text{prio}}|$, we have $\sum_{i=1}^j v_{q_i} C_{p,q_i}^- \leq m_p$. Our invariant provides $v_{q_{j+1}} C_{p,q_{j+1}}^- \leq m^{(j+1)}$, but we can also derive by immediate induction that $m^{(j+1)} = m^{(1)} - \sum_{i=1}^j v_{q_i} [C^-]_{q_i}$, therefore by induction hypothesis we obtain $\sum_{i=1}^{j+1} v_{q_i} C_{p,q_i}^- \leq m_p^{(1)} = m_p$.

- ◇ We now prove that X is maximal, in the sense that if X' is another simultaneously fireable set of transitions with multiplicities $(v'_q)_q$ verifying the routing rules, therefore $v' \leq v$. Suppose for contradiction that there is some q in \mathcal{Q} such that $v'_q > v_q$. A first situation occurs if $q \in \mathcal{Q}_{\text{sync}}$, in which case for each p in q^{in} , the v'_q firings of q consume $v'_q \alpha_{qp}$ tokens. However, we have $p^* \in q^{\text{in}}$ such that $v_q = \lfloor \alpha_{qp^*}^{-1} m_{p^*} \rfloor$, as a result we can write $v'_q \alpha_{qp^*} \geq (v_q + 1) \alpha_{qp^*} = (\lfloor \alpha_{qp^*}^{-1} m_{p^*} \rfloor + 1) \alpha_{qp^*} > m_{p^*}$, using a standard property of the floor function. This would mean that X' consumes more tokens in p^* than its marking, hence the contradiction. If $q \in \mathcal{Q}_{\text{psel}}$ and p is its unique upstream place, we obtain by the exact same argument that X' consumes more than $m_{p,q}$ tokens in p to fire q , which is not possible.

We now focus on the case $q \in \mathcal{Q}_{\text{prio}}$, still assuming by contradiction that $v'_q > v_q$. We know that we have j^* between 1 and $|\mathcal{Q}_{\text{prio}}|$ such that $q = q_{j^*}$. Actually, up to replacing q_{j^*} by another transition with higher priority degree, we can suppose that j^* is the minimum integer such that q_{j^*} has a strictly bigger multiplicity in X' than in X , i.e.,

$$j^* = \operatorname{argmin} \{ j \in \{1, \dots, |\mathcal{Q}_{\text{prio}}|\} \mid v'_{q_j} > v_{q_j} \} . \quad (1.3)$$

In other words, for all i such that $1 \leq i < j^*$, we have $v'_{q_i} - v_{q_i} \leq 0$. We now state a fundamental family of equations that are derived from our algorithm in the beginning of the proof, expressing the maximality of the $(v_q)_{q \in \mathcal{Q}_{\text{prio}}}$:

$$\forall j \in \{1, \dots, |\mathcal{Q}_{\text{prio}}|\}, \quad m_p - \sum_{i=1}^j v_{q_i} \alpha_{q_i p} < \alpha_{q_j p} . \quad (1.4)$$

Indeed, the $(v_q)_{q \in \mathcal{Q}_{\text{prio}}}$ have been constructed so that none of them can be incremented; when considering any transition q of $\mathcal{Q}_{\text{prio}}$, after firing the transitions $(q')_{q' < q}$ with higher priority than q , the tokens remaining in places of q^{in} only permit v_q firings of q .

We now state the simultaneous fireability condition of the multiset X' , that is $\sum_{i=1}^{|\mathcal{Q}_{\text{prio}}|} v'_{q_i} \alpha_{q_i p} \leq m_p$. Because the $(v'_{q_i} \alpha_{q_i p})_{1 \leq i \leq |\mathcal{Q}_{\text{prio}}|}$ are all nonnegative, we can also state that

$$\forall j \in \{1, \dots, |\mathcal{Q}_{\text{prio}}|\}, \quad \sum_{i=1}^j v'_{q_i} \alpha_{q_i p} \leq m_p .$$

We can particularize this last equation to $j = j^*$, given by (1.3), and we obtain

$$\sum_{i=1}^{j^*} v'_{q_i} \alpha_{q_i p} = \sum_{i=1}^{j^*} (v_{q_i} + (v'_{q_i} - v_{q_i})) \alpha_{q_i p} \leq m_p, \quad \text{or also,} \quad m_p - \sum_{i=1}^{j^*} v_{q_i} \alpha_{q_i p} \geq \sum_{i=j^*}^{j^*} (v'_{q_i} - v_{q_i}) \alpha_{q_i p} .$$

However, we know by (1.3) that for all i such that $1 \leq i < j^*$, we have $v'_{q_i} - v_{q_i} \leq 0$. If all these nonpositive integers are actually null, i.e., for all $1 \leq i < j^*$, $v'_{q_i} - v_{q_i} = 0$, then maximality of v at step j^* (expressed in (1.4)) is contradicted due to $v'_{q_{j^*}} > v_{q_{j^*}}$ (meaning that q_{j^*} could have been used at least one more time in X). Otherwise, it means that there exists $i^* < j^*$ such that $v'_{q_{i^*}} < v_{q_{i^*}}$, but this is also a contradiction, because in this case, X' does not respect the priority rule. Indeed, q_{i^*} has higher priority level over q_{j^*} , so the number of firings of the former cannot be decreased to increase the one of the latter! \clubsuit

The following result is a direct consequence of Theorem 1.13, when the latter is applied by induction. It claims that the routing rules and the framework of simultaneous firings of multisets of transitions suffice to make the evolution of a Petri net univocal. Note that this corollary is here stated in an uncontrolled setting, as we reckon that these controls find their place better in the next section where time is also introduced.

- **COROLLARY 1.14.** *Let be an untimed Petri net with multipliers and no input transitions, where conflicts are arbitrated by preselection or priority. For any initial marking, there exists a unique admissible generalized firing sequence of the net made of maximal multisets.*

We point out that the unique firing sequence mentioned in Corollary 1.14 can have finite or infinite length. Actually, this sequence is finite if and only if the net admits a deadlock for the considered initial marking. Conversely, in a live Petri net, the unique firing sequence must be infinite.

3 Timed Petri nets

The models of Petri nets that we have presented so far are very useful to explore the structural properties of a system implementing synchronization and concurrency patterns, and in which the firings of the transitions (possibly simultaneous) only involve a **logical time**. Even though we talk about *sequential steps*, these could very well happen instantaneously, *in a flash*.

For many applications though, it is natural to augment these models to take into account a **physical time**, mainly to describe the time needed for some tasks to be completed. Such an extension was first introduced by Ramchandani in his doctoral thesis [Ram73], and a great variety of models of timed Petri nets arose afterwards – see for instance the survey of Bowden [Bow00]. In this section, we present a model of timed Petri nets along with tools to study the dynamical systems that they give birth to.

- 3.1 FRAMEWORK AND SEMANTICS OF TIMED PETRI NETS.** – We begin by formally introducing our model. Although the most popular setting in the Petri net literature is to attach duration and holding times only to transitions, we choose to rather have timed places and untimed transitions, as done for instance in [Sif79]. These two frameworks are however equivalent.

- **DEFINITION 1.15.** *A **timed Petri net** is a 4-tuple $(\mathcal{P}, \mathcal{Q}, \mathcal{E}, \tau)$, where the triple $(\mathcal{P}, \mathcal{Q}, \mathcal{E})$ is a Petri net with multipliers and τ is a vector of $\mathbb{R}_{\geq 0}^{\mathcal{P}}$. For p in \mathcal{P} , the nonnegative real number τ_p is called the **holding time** of the place p .*

Compared to an untimed Petri net, the semantics of a timed Petri net are modified as follows. When a token is produced in a place p of \mathcal{P} , a fictional clock is triggered to retain its *age*. This token cannot be consumed by any transition of p^{out} before it is **mature**, i.e., before its age exceeds τ_p , the holding time of p . Conversely, a transition is fireable if and only if all its upstream places contain enough tokens with an age greater than their respective holding times (“enough” here corresponding to the amount of tokens given by the places-to-transitions multipliers). The firing of a transition is *instantaneous*, and when a fireable transition fires, it consumes tokens in each of its upstream places and produces tokens in each of its downstream places, with null ages. Between two successive firing instants (say t_1 and t_2) of two transitions of the net (possibly the same), the marking of the net is left unchanged, but the age of every token in every place is incremented by $t_2 - t_1$. This last point prompts us to distinguish the notions of marking and state for timed Petri nets, a state now being a given marking (the number of tokens in every place) *and* the age of

every token in every place. If a timed Petri net is equipped with an initial marking, we convene that the related tokens have an infinite age at the beginning of the execution of the net, so that they are immediately consumable.

Since the holding times attached to places introduce additional constraints on the fireable character of transitions, there will typically be less admissible firing sequences in a timed Petri net than in the corresponding untimed Petri net (with the holding times ignored) for the same marking and at a given instant. Still, the notion of simultaneously fireable multiset of transitions discussed in Section 2.3 carries over, being understood that it is now defined for a given state and not a given marking, and therefore that only mature tokens can be used to fire transitions.

To avoid confusion patterns, we require that anytime a transition is fired, we actually fire (with multiplicity) all the transitions of the maximal simultaneously fireable multiset exhibited in Theorem 1.13. Note a subtlety here, the notion of simultaneously fireable transitions essentially refers to the *logical time* (measurable with a click counter, governing the step-by-step evolution of untimed Petri nets, from markings to markings), and not the *physical time* (measurable with a clock, comparable with the holding times of places or the ages of tokens). In particular, two transitions of a simultaneously fireable multiset of transitions shall be fired at the *same* physical time, while the converse may be false – see the Example 1.17 below.

The newly added requirement that tokens produced in a place must await for a duration greater than the holding time of this place only provides a lower bound on the firing times. Indeed, nothing in the semantics prevents that transitions are fired long after becoming fireable. To the extent of tightening the evolution and reducing the degrees of freedom of the Petri nets we study, we shall always place ourselves in the **earliest behaviour** setting, meaning that transitions are fired as soon as possible.¹ In other words, a token can have an age strictly greater than the holding time of the place it lies in if and only if no downstream transition of this place can be fired.

One talks about **controlled timed Petri nets** if as set out in Section 1.5, there are some transitions of the net whose firings are prescribed externally. More precisely, all the (possibly repeated) firing (physical) instants of input transitions are known in advance. As for other transitions of the net, the firings of input transitions cost no time, and produce tokens in all places that are downstream of them.

The following theorem extends Corollary 1.14 to the case of controlled timed Petri nets. It guarantees that the setting described above is rigid enough to yield a single and univocal evolution of the net.

- **THEOREM 1.16.** *Let be a controlled timed Petri net with multipliers, where conflicts are arbitrated by preselection or priority, and subject to the earliest behaviour rule. For any initial marking, there exists a unique trajectory of the net.*

Proof. The difference with the direct proof of Corollary 1.14 by induction is that time needs to be taken into account. At the beginning of the execution of the net (say at physical time $t = 0$), we denote by $m_{t=0}^0 \in \mathbb{N}^P$ the marking vector composed of the tokens of the initial marking and the tokens possibly produced by the firing of input transitions at time $t = 0$ in places with null holding times, i.e., all the tokens that are mature. According to Theorem 1.13, there exists a unique maximal simultaneously fireable multiset $X(m_{t=0}^0)$ of transitions in $Q \setminus Q_{\text{input}}$, whose firing provides a marking $m_{t=0}^1$, still at physical time $t = 0$. If in turn $X(m_{t=0}^1) \neq \emptyset$, the earliest behaviour rule imposes to fire this new multiset, giving a new marking $m_{t=0}^2$. This continues until for some $k \in \mathbb{N}^*$, we have $X(m_{t=0}^k) = \emptyset$ (otherwise the evolution of the net stays uniquely determined but does not exceed $t = 0$). Then, we consider the minimal physical instant $t_1 > 0$ such that the marking of mature tokens evolves (i.e., some young token becomes old) or an input transition fires, giving a marking of mature tokens denoted by $m_{t=t_1}^0$. The earliest behaviour rule imposes

¹ No doubt that this assumption will be quite licit in the systems later considered in this thesis, as they originate from real-life emergency call centers.

to reapply Theorem 1.13 to this marking, and the previous scheme goes on. If at some step k , we have $t_k = \infty$, i.e., no more external token arrives while all tokens are mature, then the execution of the net stops.

Example 1.17. The Figure 1.10 depicts a timed Petri net basically consisting in two lanes, synchronizing in transition q_3 . We provide the holding times of the places p_2, p'_2, p_3 and p'_3 (respectively 1, 0, 3 and 2 units of time), the other ones being not relevant in this example.

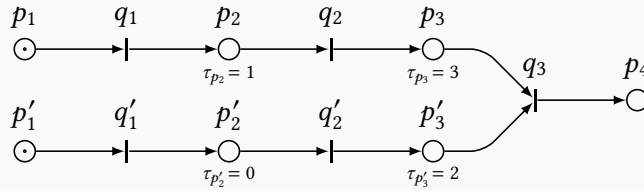


Figure 1.10: A two-lane timed Petri net

Enumerating places in the order $p_1, p'_1, p_2, p'_2, p_3, p'_3$ and p_4 , this Petri net is equipped with an initial marking $(1\ 1\ 0\ 0\ 0\ 0\ 0)^T$. When its execution starts, say at $t = 0$, both tokens present in p_1 and p'_1 have infinite age and thus q_1 and q'_1 are immediately and simultaneously fireable. The earliest behaviour then imposes to fire them, producing tokens with null ages in p_2 and p'_2 . For $\tau_{p'_2} = 0$, the transition q'_2 is immediately fireable, and its firing produces a token in p'_3 , still at $t = 0$. No more transition is then fireable, and we must wait for 1 unit of time before the token in p_2 is old enough to fire q_2 . At $t = 1$, q_2 fires and produces a token in p_3 . At $t = 2$, the age of the token in p'_3 exceeds $\tau_{p'_3}$, but due to the synchronization pattern, we must wait until the instant $t = 4$ to fire q_3 – since the token present in p_3 is not old enough beforehand. Finally, at $t = 4$, a token is produced in p_4 and execution stops.

To summarize, five different markings are encountered during the execution of the net, associated with five (distinct) logical instants:

Logical time	Physical time	Marking
0	0	$(1\ 1\ 0\ 0\ 0\ 0\ 0)^T$
1	0	$(0\ 0\ 1\ 1\ 0\ 0\ 0)^T$
2	0	$(0\ 0\ 1\ 0\ 0\ 1\ 0)^T$
3	1	$(0\ 0\ 0\ 0\ 1\ 1\ 0)^T$
4	4	$(0\ 0\ 0\ 0\ 0\ 0\ 1)^T$

Observe in particular that due to the null time associated with p'_2 , the firing of q'_2 occurs at the same physical time than those of q_1 and q'_1 , although in this example, q'_1 and q'_2 are never simultaneously fireable.

Our choice to allow zero holding times for some places is not very usual in the literature. Indeed, most authors prefer to use a vector τ with positive entries, which spares many technicalities. Instantaneous steps are then either ruled out, or modeled using very small holding times compared with the rest of the system. In order to work in a somewhat purified framework, we shall embrace in this thesis the rough ride that mercilessly comes with null times. In particular, a restriction is in order to avoid the **Zeno behaviour** of a Petri net, which designates the phenomenon of counting indefinitely many firings of a transition over a zero timespan. Such a predicament can only arise

if a transition belongs to a cycle of the graph where all places have zero holding times. As a result, we will choose to work only with non-Zeno Petri net.

- **DEFINITION 1.18.** *We say that a timed Petri net is **non-Zeno** if its underlying bipartite graph of places and transitions does not contain any cycle in which all places have null holding time.*

Throughout the coming chapters and models, we shall come back to the notion of “non-Zenoness” and refine it when possible. When working with **controlled timed Petri nets**, i.e., in which there are input transitions whose firings are prescribed externally (see Section 1.5), we will always assume that these controls are non-Zeno as well, in the sense that not only the input transitions cannot fire indefinitely many times in a null time (Zeno effect on the count of produced tokens), but also that they cannot fire indefinitely many times over a finite timespan (Zeno effect on the dates of token productions), which is more restrictive. See below our discussion on piecewise-constant functions.

- 3.2 COUNTER FUNCTIONS OF THE NET.** — Now that the time comes into play, it is natural to study Petri nets from a **dynamical** perspective. One typically wants to keep track of the net markings or the firings of the transitions over time. The evolution summary provided in Example 1.17 already gives such a flavour. Instead of focusing on instantaneous markings, we consider an “integral” (or cumulated) version of it, expressed by counter variables. We will always convene that the execution of the timed Petri nets we study starts at the (physical) instant $t = 0$.

- **DEFINITION 1.19.** *Let $(\mathcal{P}, Q, \mathcal{E}, \tau)$ be a non-Zeno timed Petri net equipped with an initial marking. We call **counter functions** of the net the mappings $(X_p)_{p \in \mathcal{P}}$, $(Y_p)_{p \in \mathcal{P}}$ and $(Z_q)_{q \in Q}$ from $\mathbb{R}_{\geq 0}$ to \mathbb{N} such that for all $t \geq 0$:*

$X_p(t)$ denotes the number of tokens that have been present in place p up to time t included (i.e., either the tokens produced in p or those of the initial marking in p);

$Y_p(t)$ denotes the number of tokens that have been present and mature in place p up to time t included;

$Z_q(t)$ denotes the number of firings of transition q up to time t included.

Dually to counter functions, one can also introduce “dater” functions from \mathbb{N} to $\mathbb{R}_{\geq 0}$, compiling the times at which (cumulated) events occur. Counters and daters – as tools for the study of systems involving concurrency and synchronization mechanisms – were already used by Baccelli, Cohen, Olsder and Quadrat in [BCOQ92]. The following proposition recalls the form of the counter functions. We use the term “**càdlàg** function” to denote a function that is right-continuous with left limits.²

- **PROPOSITION 1.20.** *The counter functions of a non-Zeno timed Petri net are nondecreasing piecewise-constant and càdlàg functions.*

⋮ *Proof.* The well definition of counter functions comes from the fact that owing to the non-Zeno character of the net (and of the possible input transitions), starting from a finite marking and using finite multipliers, the number of firings of a transition over the time period $[0, t]$ stays finite. Counters are nondecreasing by definition, and stay constant between two successive events (again, this notion is well-posed thanks to the non-Zeno property). Finally, the càdlàg character is ensured by the fact that the Definition 1.19 is based on a number of events up to a certain instant *included*. ♠

² The abbreviation comes from the French “continue à droite, limite à gauche”

- **REMARK 1.21.** Note that in all the thesis, we use the term “piecewise” to refer to a property verified over a countable (hence not necessarily finite) collection of disjoint intervals; however we require that there are finitely many such pieces over any compact set. For instance, the usual floor function $x \mapsto \lfloor x \rfloor$ is piecewise-constant over \mathbb{R} , but if we denote $b_n := \sum_{k=1}^n 2^{-k}$ for all n in \mathbb{N}^* , then the function $x \mapsto \sum_{n=1}^{\infty} b_n \mathbb{1}_{[b_n, \infty)}(x)$ is not piecewise-constant. This is due to the fact that the discontinuity points accumulate at $x = 1$, thus there are infinitely many pieces over $[0, 1]$. This function is still càdlàg though.

It will be useful to extend the definition of counter functions from $\mathbb{R}_{\geq 0}$ to \mathbb{R} , by simply setting $X_p(t) = Y_p(t) = m_p$ and $Z_q(t) = 0$ for all $t < 0$ and all places p or transitions q , where the vector m in $\mathbb{N}^{\mathcal{P}}$ denotes the initial marking. When working with a càdlàg function f , we shall also denote by $f(t^-)$ the left limit of f at point t in \mathbb{R} .

Example 1.22. We depict in Figure 1.11 a timed Petri net with four places and four transitions. The net is equipped with an initial marking of six tokens in p_0 , two tokens in p_1 and one token in p_2 , or equivalently $m = (6 \ 2 \ 1 \ 0)^T$. The holding times of places are set to $\tau_{p_1} = 4$, $\tau_{p_2} = 3$ and $\tau_{p_3} = 2$ (the one of p_0 does not play a role here).

As an external control, we choose the behaviour of the transition q_0 and we decide that it will fire once at $t = 4$ and twice at $t = 6$, so that we already know that the counter function associated with q_0 will be $Z_{q_0} = \mathbb{1}_{[4, \infty)} + 2 \cdot \mathbb{1}_{[6, \infty)}$.

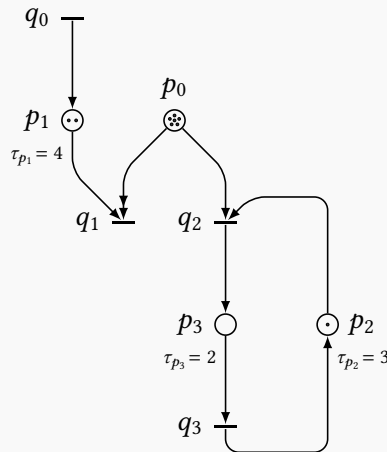


Figure 1.11: A non-Zeno timed Petri net associated with the counters of Figure 1.12

We are interested in plotting the counter functions corresponding to this Petri net. First, we describe its complete execution.

At $t = 0$, the tokens of the initial marking in p_1 and p_2 can be used to fire q_1 twice and q_2 once, this also consumes three tokens in p_0 , and produces one token in p_3 . At $t = 2$, the latter is old enough to trigger the firing of q_3 and a new token enters p_2 . At $t = 4$, q_0 fires and a token is produced in p_1 . At $t = 5$, the token present in p_2 is old enough for q_2 to be fired, consuming the fourth token of p_0 . At $t = 6$, accordingly to our input, q_0 fires twice, adding two tokens in p_1 . At $t = 7$, q_3 fires and a new token pops up in p_2 . At $t = 8$, the first token originating from the command is old enough to fire q_1 , which consumes the fifth resource initially present in p_0 . At $t = 10$, both q_1 and q_2 are fireable, the former can use mature tokens generated by the input four unit of times before, the latter can use the mature token in p_3 .

However, only one last token is available in p_0 , the priority rule applies and q_1 fires only once. At the end of the execution, we are left with a marking of $(0 \ 1 \ 1 \ 0)^T$.

We show the counter functions in the Figure 1.12 below, representing the families $(X_p)_{p \in \mathcal{P}}$ and $(Z_q)_{q \in \mathcal{Q}}$. To simplify the notation, we write X_0 and Z_0 instead of X_{p_0} and Z_{q_0} , etc.

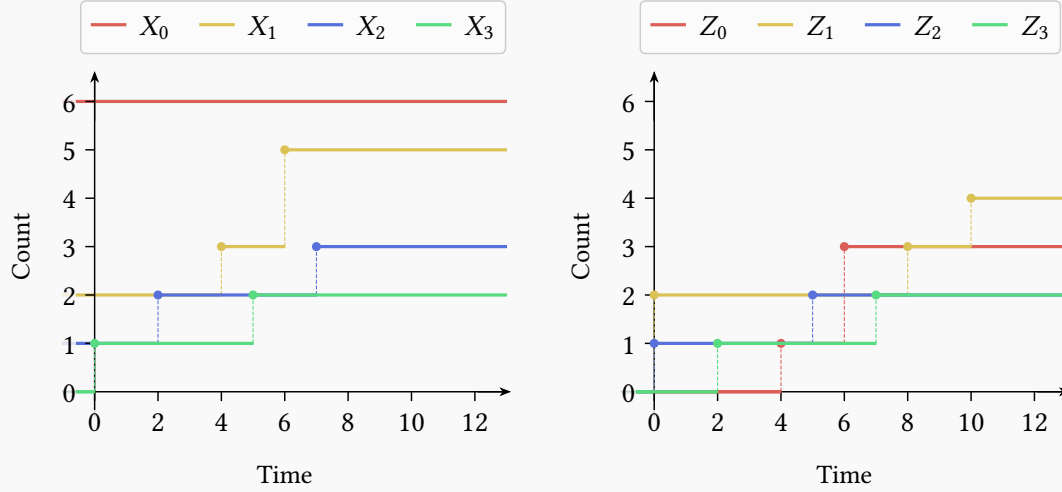


Figure 1.12: Counter functions corresponding to the net of Figure 1.11

We verify in particular the claim of Proposition 1.20. Here, since none of the transitions are live and the execution stops in finite time, the counters are stationary functions.

We point out that although the transitions q_1 and q_2 fired six times in total, the priority rule was actually only used once to arbitrate an effective conflict. In all the other situations, either one of the transitions was not fireable, either there was enough tokens to fire both of them as many times as possible.

The three families of counter variables defined in Definition 1.19 are linked together by the evolution rules of the net described previously, and we now detail how to put them into equations. First and foremost, a token can be present in a place p in \mathcal{P} if and only if it is part of the initial marking of p or if it has been produced by the firing of a transition in p^{in} . Denoting by m the initial marking vector of the net and recalling that a single firing of q in p^{in} produces α_{pq} tokens in p , we have:

$$\forall p \in \mathcal{P}, \quad X_p(t) = m_p + \sum_{q \in p^{\text{in}}} \alpha_{pq} Z_q(t). \quad (\text{P1})$$

Since the two definitions of X_p and Y_p differ from the condition that tokens are old enough to fire the downstream transitions of the place p , the latter is obtained by delaying the former by the time τ_p :

$$\forall p \in \mathcal{P}, \quad Y_p(t) = X_p(t - \tau_p). \quad (\text{P2})$$

We know using the earliest behaviour rule that a transition not involved in a structural conflict must fire as soon as possible, which occurs when all its upstream places contain enough mature tokens. More precisely, if q is in $\mathcal{Q}_{\text{sync}}$ and p is a place of q^{in} , it takes α_{qp} mature tokens in p to fire q once, and therefore k mature tokens in p can fire q exactly $\lfloor \alpha_{qp}^{-1} k \rfloor$ times, where k is an integer and $\lfloor \cdot \rfloor$ denotes the floor function. Because the upstream place of q with the least number

of old enough tokens constrains the number of firings of q , if for all p in q^{in} , p contains k_p mature tokens, then q can fire $\min_{p \in q^{\text{in}}} \lfloor \alpha_{qp}^{-1} k_p \rfloor$. In terms of counter functions and cumulated number of firings, we obtain

$$\forall q \in \mathcal{Q}_{\text{sync}}, \quad Z_q(t) = \min_{p \in q^{\text{in}}} \left\lfloor \alpha_{qp}^{-1} Y_p(t) \right\rfloor. \quad (\text{P3})$$

If a transition is involved in a structural conflict, meaning that one of its upstream places admits several downstream transitions, the equation (P3) must be amended to take into account the routing rules. We first address the case of transitions in $\mathcal{Q}_{\text{psel}}$, that by convention admit a single upstream place (see Section 2.1). If q is in $\mathcal{Q}_{\text{psel}}$ with upstream place p , then among the initial marking of m_p tokens (with infinite age), $m_{p,q}$ tokens are reserved to fire q , and among the n first tokens produced in p , $\Pi_q^p(n)$ tokens are reserved to fire q . Since the total number of tokens produced in p (i.e., excluding the initial marking) that are mature at time t is $Y_p(t) - m_p$, we have

$$\forall q \in \mathcal{Q}_{\text{psel}}, \quad Z_q(t) = \left\lfloor \alpha_{qp}^{-1} \left(m_{p,q} + \Pi_q^p(Y_p(t) - m_p) \right) \right\rfloor \quad (q^{\text{in}} = \{p\}). \quad (\text{P4})$$

The equations for transitions that are downstream of a place ruled by priority are a little more subtle. If q is in $\mathcal{Q}_{\text{prio}}$ and t is in \mathbb{R} , we focus on $Z_q(t) - Z_q(t^-)$, which precisely gives the number of firings of q at time t . The amount of tokens available in a place p of q^{in} at the physical time t before actually firing any transition of p^{out} is $Y(t) - \sum_{q' \in p^{\text{out}}} \alpha_{q'p} Z_{q'}(t^-)$. These tokens can be used to fire the fireable transitions of p^{out} , by allocating them in order of priority. In details, the transition q_1 in p^{out} with the highest priority degree relatively to p can fire all these tokens provided it is not constrained by another upstream place. The $Z_{q_1}(t) - Z_{q_1}(t^-)$ firings of q_1 at instant t leave an amount of $Y(t) - \alpha_{q_1p} Z_{q_1}(t) - \sum_{q' \in p^{\text{out}} \setminus \{q_1\}} \alpha_{q'p} Z_{q'}(t^-)$ tokens in p to fire q_2 , the transition of p^{out} with second higher priority relatively to p . This scheme goes on until hitting q , which can use the $Y(t) - \sum_{q' <_p q} \alpha_{q'p} Z_{q'}(t) - \sum_{q' \geq_p q} \alpha_{q'p} Z_{q'}(t^-)$ remaining mature tokens in p . As a result, we obtain $Z_q(t) - Z_q(t^-)$ by packing these tokens by groups α_{qp} and taking the minimum of resulting terms over q^{in} to account for the synchronization, which by canceling the terms $Z_q(t^-)$ in both sides provides:

$$\forall q \in \mathcal{Q}_{\text{prio}}, \quad Z_q(t) = \min_{p \in q^{\text{in}}} \left\lfloor \alpha_{qp}^{-1} \left(Y_p(t) - \sum_{q' <_p q} \alpha_{q'p} Z_{q'}(t) - \sum_{q' \geq_p q} \alpha_{q'p} Z_{q'}(t^-) \right) \right\rfloor. \quad (\text{P5})$$

The equations (P1)–(P5) can be combined together to form a system of equations that features only the transition counters $(Z_q)_{q \in \mathcal{Q}}$ based on their “type”, and we write them in Table 1.13. Observe that in case the preselection or priority routings are trivial, we recover for transitions in $\mathcal{Q}_{\text{psel}}$ and $\mathcal{Q}_{\text{prio}}$ the equation of transitions in $\mathcal{Q}_{\text{sync}}$. We let the reader check that these relations are verified for the counters depicted in Figure 1.12.

In [ABG15, Theorem 1], Allamigeon, Bœuf and Gaubert prove that any admissible firing sequence of transitions of a timed Petri net yield counter functions that verify equations (P1)–(P5) (in a slightly lighter setting without arc multipliers and only two priority levels, but their demonstration extends to our framework). Conversely, Bœuf shows that if the equations (P1)–(P5) admit nondecreasing càdlàg solutions, then the latter correspond to a firing sequence of the Petri net, see [Bœu17, Section 2.3.6]. We close the loop by studying well-posedness properties of the counter equations, that is to say in which cases they determine unique solutions. We begin with the standard priority-free setting.

Type	Counter equation in the discrete model
$q \in \mathcal{Q}_{\text{input}}$	$Z_q(t)$ is prescribed externally
$q \in \mathcal{Q}_{\text{sync}}$	$Z_q(t) = \min_{p \in q^{\text{in}}} \left[\alpha_{qp}^{-1} \left(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} Z_{q'}(t - \tau_p) \right) \right]$
$q \in \mathcal{Q}_{\text{psel}}$	$Z_q(t) = \left[\alpha_{qp}^{-1} \left(m_{p,q} + \Pi_q^p \left(\sum_{q' \in p^{\text{in}}} \alpha_{pq'} Z_{q'}(t - \tau_p) \right) \right) \right]$
$q \in \mathcal{Q}_{\text{prio}}$	$Z_q(t) = \min_{p \in q^{\text{in}}} \left[\alpha_{qp}^{-1} \left(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} Z_{q'}(t - \tau_p) - \sum_{q' <_p q} \alpha_{q'p} Z_{q'}(t) - \sum_{q' >_p q} \alpha_{q'p} Z_{q'}(t^-) \right) \right]$

Table 1.13: Dynamic equations followed by transitions counter functions

► **THEOREM 1.23.** *In a non-Zeno and priority-free controlled timed Petri net, the knowledge of the input profiles $(Z_q)_{q \in \mathcal{Q}_{\text{input}}}$ and the system of equations of Table 1.13 determine unique nondecreasing càdlàg functions $(Z_q)_{q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}}$ over \mathbb{R} .*

Proof. The proof is based on an inductive scheme exploiting the delays arising in the equations of the Table 1.13. The most important step is first to reduce to a set of equations featuring only positive delays.

◇ We introduce the weighted and directed graph $\mathcal{G} := (\mathcal{Q}, \mathcal{E}_{\mathcal{G}}, w)$ with nodes set \mathcal{Q} , arcs set $\mathcal{E}_{\mathcal{G}}$ and arc weights function w , so that for q_1 and q_2 in \mathcal{Q} , the arc $q_1 \rightarrow q_2$ is in $\mathcal{E}_{\mathcal{G}}$ if and only if $q_1^{\text{in}} \cap q_2^{\text{out}} \neq \emptyset$, and it has weight $w(q_1 \rightarrow q_2) := \min \{ \tau_p \mid p \in q_1^{\text{in}} \cap q_2^{\text{out}} \}$. Based on \mathcal{G} , we introduce the binary relation \succrightarrow over \mathcal{Q} , where if $q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$ and $q' \in \mathcal{Q}$, we say that $q \succrightarrow q'$ if and only if there exists $k \in \mathbb{N}^*$ and transitions q_1, q_2, \dots, q_k in \mathcal{Q} such that either the conditions (i)+(ii)+(iii)+(iv) or the conditions (i)+(ii)+(iii)+(v) below are verified (note that (iv) implies (iii) due to the non-Zeno assumption):

- (i) $q_1 = q$ and $q_k = q'$,
- (ii) $q_1 \rightarrow q_2 \rightarrow \dots \rightarrow q_k$ is a directed walk in \mathcal{G} ,
- (iii) $q_1 \rightarrow q_2 \rightarrow \dots \rightarrow q_{k-1}$ is a walk in \mathcal{G} that does not contain any cycle,
- (iv) $w(q_i \rightarrow q_{i+1}) = 0$ for all $1 \leq i < k-1$ and $w(q_{k-1} \rightarrow q_k) > 0$,
- (v) $q' \in \mathcal{Q}_{\text{input}}$.

If the conditions (i)+(ii)+(iii)+(iv) hold, we say that the walk $q_1 \rightarrow q_2 \rightarrow \dots \rightarrow q_k$ is a *maximal non-Zeno walk* from q to q' . This means that the counter equation of q will causally depend on the counter function $Z_{q'}$, delayed at least by $w(q_{k-1} \rightarrow q_k)$ units of time (that is positive). Note that Z_q might depend on itself, although delayed. Similarly, if the conditions (i)+(ii)+(iii)+(v) hold, we talk about a *maximal input-headed walk* from q to q' , and the counter equation of q will causally depend on input function $Z_{q'}$ (possibly with no positive delay). In both cases, we talk about *maximal walks* from q to q' .

For all $q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$, we denote by $C(q) := \{q' \in \mathcal{Q} \mid q \succrightarrow q'\}$ the set of transitions that are in relation with q in the previous sense. For $q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$ and $q' \in C(q)$, we denote by $\tau_{q,q'}$ the set

$$\tau_{q,q'} := \{w(q_{k-1} \rightarrow q_k) \mid q_1 \rightarrow q_2 \rightarrow \dots \rightarrow q_{k-1} \rightarrow q_k \text{ is a maximal walk from } q \text{ to } q'\}$$

To lighten a little bit the notation, we denote $C_1(q) := C(q) \cap \mathcal{Q}_{\text{input}}$ and $C_2(q) := C(q) \setminus \mathcal{Q}_{\text{input}}$.

◇ Exploiting the Table 1.13, for all $q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$, we can apply a finite number of substitutions in the equation on Z_q to obtain an equality of the form

$$Z_q(t) = F_q \left((Z_{q'}(t - \tau))_{q' \in C_1(q), \tau \in \tau_{q,q'}}, (Z_{q'}(t - \tau))_{q' \in C_2(q), \tau \in \tau_{q,q'}} \right), \quad (\text{M})$$

where F_q only involves affine combinations of counters, preselection operators, minimum and floor functions operations. In other words, $Z_q(t)$ depends only on input counters $(Z_{q'}(t - \tau))_{q' \in C_1(q), \tau \in \tau_{q,q'}}$ possibly delayed by a null time, and on other counters $(Z_{q'}(t - \tau))_{q' \in C_2(q), \tau \in \tau_{q,q'}}$ necessarily delayed by a positive time. Indeed, if $q' \in C_2(q)$, then the condition (iv) is verified for any maximal walk from q to q' , therefore $\tau > 0$ for all $\tau \in \tau_{q,q'} > 0$ by definition.

Hence, denoting by $\tau_{\min} := \min_{p \in \mathcal{P}} \{\tau_p \mid \tau_p > 0\}$ and $\tau_{\max} := \max_{p \in \mathcal{P}} \{\tau_p \mid \tau_p > 0\}$, we can see that for all $q \in Q \setminus Q_{\text{input}}$, the number $Z_q(t)$ is fully determined by the knowledge of the restrictions $(Z_{q'}|_{[t-t_{\max}, t]})_{q' \in Q_{\text{input}}}$ and $(Z_{q'}|_{[t-t_{\max}, t-t_{\min}]})_{q' \in Q \setminus Q_{\text{input}}}$. We can thus apply the following inductive scheme; for all $k \in \mathbb{N}$, compute the restrictions $(Z_q|_{[k\tau_{\min}, (k+1)\tau_{\min}])}_{q \in Q \setminus Q_{\text{input}}}$, each step based on the inputs and on the restrictions computed at the previous steps. In particular, for all $t \geq 0$, the value of $Z_q(t)$ can be computed by induction in $\lceil t/\tau_{\min} \rceil$ steps.

- ◇ The fact that the resulting counters are nondecreasing and càdlàg follows from the fact that the input profiles are such functions and the form of the equations of Table 1.13 (to check this part, we may for the sake of simplicity suppose that all the holding times are positive to avoid the above substitutions). \clubsuit

The induction technique employed in the proof of Theorem 1.23 to propagate the dynamics is fundamental, and we shall encounter it again multiple times in our work. To stress its importance and easily refer to it, we frame it as a principle.

- **INDUCTION PRINCIPLE (I).** *In a non-Zeno and priority-free controlled timed Petri net, denoting by $\tau_{\min} := \min_{p \in \mathcal{P}} \{\tau_p \mid \tau_p > 0\}$ and $\tau_{\max} := \max_{p \in \mathcal{P}} \{\tau_p \mid \tau_p > 0\}$, the trajectory of the net over $[0, \infty)$ can be determined by induction, since for all k in \mathbb{N} , the counter functions over $[k\tau_{\min}, (k+1)\tau_{\min})$ only depend on their values over $[k\tau_{\min} - \tau_{\max}, k\tau_{\min})$ and on the input profiles.*

When working with priority rules and zero holding times, we need to pay extra care for the dynamics associated with the equations of Table 1.13 to remain well-posed. Observe indeed the Petri net of the Figure 1.14, where the place p_0 is ruled by priority, with $q_2 <_{p_0} q_1$, and the holding times of both places p_0 and p_1 are null.

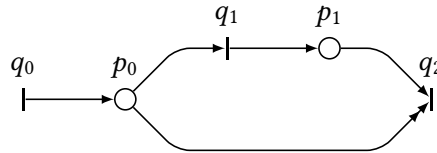


Figure 1.14: A “diamond” pattern leading to ambiguous dynamics

We have according to our dynamical equations that for all $t \geq 0$, $Z_2(t) = \min(Z_1(t), Z_0(t) - Z_1(t^-))$ and $Z_1(t) = Z_0(t) - Z_2(t)$. Even though Z_0 is known, the value of Z_2 (or Z_1) at any instant depends upon itself, and we may find several solutions to our dynamics. These ambiguous equations arise despite the fact that there are no “Zeno circuit” in the net, and the behaviour of the latter seems intuitive; here q_2 shall fire exactly once every two firings of q_0 (while it would never fire if the priority degrees were swapped).

We think it is desirable in general to prevent such configurations and work only with timed Petri nets in which substitution schemes in the dynamical equations always suffice for these dynamics to determine unique trajectories. To this purpose, we make use of an additional topological condition in the Theorem 1.24 below.

Reusing the maximal walks in zero time introduced in the proof of Theorem 1.23, we denote for all q in $\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$ by $\mathcal{D}(q)$ the subset of transitions defined by

$$\mathcal{D}(q) := \bigcup_{q' \in \mathcal{C}(q)} \{q_i \mid 1 < i < k \text{ and } q_1 \rightarrow \cdots \rightarrow q_{k-1} \rightarrow q_k \text{ is a maximal walk from } q \text{ to } q'\}.$$

The transitions of $\mathcal{D}(q)$ precisely indicate which equations are used and substituted to go from the initial equation giving Z_q in the Table 1.13 to the relation (M) featured in the proof of Theorem 1.23. Note that $\mathcal{D}(q)$ is empty if all the places of q^{in} have positive holding times.

- **THEOREM 1.24.** *In a non-Zeno and controlled timed Petri net, such that for all q in $\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$ we have $\mathcal{D}(q) \cap \mathcal{Q}_{\text{prio}} = \emptyset$, the knowledge of the input profiles $(Z_q)_{q \in \mathcal{Q}_{\text{input}}}$ and the system of equations of Table 1.13 determine unique nondecreasing càdlàg functions $(Z_q)_{q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}}$ over \mathbb{R} .*

Proof. We reuse the notation of the proof of Theorem 1.23. In addition to the sets $\mathcal{C}_1(q)$ and $\mathcal{C}_2(q)$, we introduce for all $q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$ the sets $\mathcal{C}_3(q)$ and $\mathcal{C}_4(q)$ (empty whenever $q \notin \mathcal{Q}_{\text{prio}}$) defined by

$$\mathcal{C}_3(q) := \bigcup_{p \in \mathcal{P}_{\text{prio}} \cap q^{\text{in}}} \{q' \in p^{\text{out}} \mid q' <_p q\} \quad \text{and} \quad \mathcal{C}_4(q) := \bigcup_{p \in \mathcal{P}_{\text{prio}} \cap q^{\text{in}}} \{q' \in p^{\text{out}} \mid q' >_p q\}.$$

For transitions in $\mathcal{Q}_{\text{sync}} \cup \mathcal{Q}_{\text{psel}}$, the equation (M) is unchanged. However, if $q \in \mathcal{Q}_{\text{prio}}$, the master equation governing the evolution of Z_q is after finitely many substitutions of the form

$$Z_q(t) = F_q \left((Z_{q'}(t - \tau))_{q' \in \mathcal{C}_1(q), \tau \in \tau_{q,q'}}, (Z_{q'}(t - \tau))_{q' \in \mathcal{C}_2(q), \tau \in \tau_{q,q'}}, (Z_{q'}(t))_{q' \in \mathcal{C}_3(q)}, (Z_{q'}(t^-))_{q' \in \mathcal{C}_4(q)} \right). \quad (1.5)$$

Indeed, according to the assumption that $\mathcal{D}(q) \cap \mathcal{Q}_{\text{prio}} = \emptyset$ and the equations of Table 1.13, only the siblings of q can appear with no delay in the equation. This would not be true if some $q' \in \mathcal{Q}_{\text{prio}}$ lied in $\mathcal{D}(q)$, because we would have used the equation of $Z_{q'}$ in the substitution scheme and would have introduced with no positive delay the transitions of $\mathcal{Q}_{\text{prio}}$ in structural conflict with q' .

- ◇ For all $q \in \mathcal{Q}_{\text{input}}$, we denote by \mathcal{T}_q the set of firing instants of q , which are also the discontinuity points of Z_q , since $Z_q = \sum_{t \in \mathcal{T}_q} c_t \mathbb{1}_{[t, \infty)}$ (where c_t is the number of firings at instant $t \in \mathcal{T}_q$). By assumption on the input profiles, the set \mathcal{T}_q is countable, with no accumulation point (this entails that the sequence of the successive firing times of q tends to ∞ if \mathcal{T}_q is countably infinite). We also define the subset \mathcal{T} of $\mathbb{R}_{\geq 0}$ by

$$\mathcal{T} := \left(\bigcup_{q \in \mathcal{Q}_{\text{input}}} \mathcal{T}_q \right) + \left(\sum_{p \in \mathcal{P}} \tau_p \mathbb{N} \right).$$

The set \mathcal{T} is countably infinite (it could be finite in the trivial case where the Petri net has a tree structure with only zero holding times), and most importantly, for all $t \geq 0$, the set $\mathcal{T} \cap [0, t]$ is finite, even though \mathcal{T} is ultimately dense. Therefore, the elements of \mathcal{T} can be written as an increasing sequence $(t_n)_{n \in \mathbb{N}}$, with $\lim_{n \rightarrow \infty} t_n = \infty$.

- ◇ We now show by induction that the predicate “The functions $(Z_q|_{(-\infty, t_n)})_{q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}}$ are entirely determined by the $(Z_q)_{q \in \mathcal{Q}_{\text{input}}}$ and the equations of Table 1.13. In addition, their discontinuity points lie in \mathcal{T} .” holds for all $n \in \mathbb{N}$. Since $t_0 = 0$, the property for $n = 0$ follows by definition of the counter functions that are taken null over $\mathbb{R}_{< 0}$.

Let us now suppose that the above predicate is verified for some $n \in \mathbb{N}$. We first show that the functions $(Z_q|_{[t_n, t_{n+1})})_{q \in \mathcal{Q}_{\text{prio}}}$ are uniquely determined, using the equation (1.5). We reason cluster by cluster, where for q in $\mathcal{Q}_{\text{prio}}$, the cluster of q is the transitive closure of $\{q\}$ by the binary relation \sim defined by $q \sim q'$ if and only if $(q^{\text{in}} \cap q'^{\text{in}}) \cap \mathcal{P}_{\text{prio}} \neq \emptyset$, that is to say q and q' are siblings under some priority rules. The clusters form a partition of $\mathcal{Q}_{\text{prio}}$, as shown by Boëuf in [Boëu17]. Let $\{q_1, q_2, \dots, q_k\}$ be a cluster of the Petri net of cardinal $k \in \mathbb{N} \setminus \{0, 1\}$. Due to the compatibility of the priority rules, we have for all $1 \leq i < j \leq k$ that $p \in \mathcal{P}_{\text{prio}}$ and $p \in q_i^{\text{in}} \cap q_j^{\text{in}}$ implies $q_i <_p q_j$ (up to possible relabeling of the elements of the cluster).

This means that in the equation of type (1.5) for q_1 , the transition with highest priority degree in the cluster, the set $\mathcal{C}_3(q_1)$ is empty, and $Z_{q_1}(t_n)$ only depends on the $(Z_{q'}(t_n - \tau))_{q' \in \mathcal{C}_1(q), \tau \in \tau_{q,q'}}$ (known because

coming from input profiles), the $(Z_{q'}(t_n - \tau))_{q' \in C_2(q), \tau \in \tau_{q,q'}}$ (known because in this case $t_n - \tau < t_n$) and on the $Z_{q'}(t_n^-)_{q' \in C_4(q)}$. The latter are also known by our induction hypothesis, hence $Z_{q_1}(t_n)$ is indeed fully determined by the $(Z_q|_{(-\infty, t_n)})_{q \in Q}$. We can then focus on q_2 , for which $C_3(q_2) \subset \{q_1\}$, and still deduce $Z_{q_2}(t_n)$. More generally, since $C_3(q_i) \subset \{q_1, \dots, q_{i-1}\}$ and $C_4(q_i) \subset \{q_{i+1}, \dots, q_k\}$, all the $(Z_{q_i}(t_n))_{1 \leq i \leq k}$ are determined using the past values. We know use the fact that derived from the canonical equation for $q \in Q_{\text{prio}}$ in the Table 1.13 that $Z_q(t)$ can jump only if some $Z_{q'}(t - \tau_p)$ jumps for some $p \in q^{\text{in}}$ and $q' \in p^{\text{in}}$. Because the functions $t \mapsto (Z_{q'}(t - \tau))_{q' \in C_1(q) \cup C_2(q), \tau \in \tau_{q,q'}}$ are constant for $t \in [t_n, t_{n+1})$ by construction of \mathcal{T} , we deduce that so are the $(Z_{q_i})_{1 \leq i \leq k}$, in other words $Z_{q_i}(t) = Z_{q_i}(t_n)$ for all $1 \leq i \leq k$ and $t \in [t_n, t_{n+1})$. This proves our predicate for the transitions in Q_{prio} .

Dealing with transitions in $Q_{\text{sync}} \cup Q_{\text{psel}}$ is easier because we just need to reuse equation (M) from the proof of Theorem 1.24. The fact that the discontinuity points of the $(Z_q)_{q \in Q_{\text{sync}} \cup Q_{\text{psel}}}$ is immediate using the piecewise-constant character of the counters. \clubsuit

Observe that the extra assumption required by Theorem 1.24 was not verified for the Petri net in the Figure 1.14, in which we had $\mathcal{D}(q_2) = \{q_1\}$ and q_1 was indeed an element of Q_{prio} . We think in general that requiring $\mathcal{D}(q) \cap Q_{\text{prio}} = \emptyset$ for all q in Q_{prio} is too strong, and that finer necessary conditions could be exhibited. This is however left for future work.

3.3 BRINGING STOCHASTICITY TO THE NETS. — The models of controlled timed Petri nets that we have presented so far are **deterministic**, in the sense that multiple executions of the same Petri net with same controls will provide the same evolution and final states. However, one may also want to study systems in which there is room for nondeterminism at certain steps. This will typically make the counter functions of the net random variables, and we shall write $(\widehat{X}_p, \widehat{Y}_p, \widehat{Z}_q)_{p \in \mathcal{P}, q \in Q}$ with the “hat” notation to emphasize the stochastic character. More specifically, we work in a probability space $(\Omega, \mathcal{A}, \mathbb{P})$, these random variables take values in the sample space Ω , and we denote by $X_p(t; \omega)$, $Y_p(t; \omega)$ and $Z_q(t; \omega)$ the non-random values taken by the counters at the instant t for the event ω in Ω .

The term “stochastic Petri net” does exist in the literature and mostly denotes the Petri nets in which the transitions have firing duration given as realizations of nonnegative random variables. Most authors consider exponentially distributed firing times, as introduced independently by Molloy [Mol82] and Florin and Natkin [FN85]. This model allows to consider a new routing rule to arbitrate conflict, known as the “race policy”. When several transitions have a common upstream place and are all fireable at the same instant, each one of them starts a virtual firing with exponentially distributed duration; the first one to finish the race (that is to say the transition with the minimum firing time) actually fires and consumes the resources involved in the conflict. Marsan et. al. talk about “generalized stochastic Petri nets” (see [AMCB84]) when firing duration can be both null or exponentially distributed.

In our model, the firings of the transitions are always immediate and the race policy does not apply. However, we may want to allow stochastic holding times to account for the fact that a same abstract task can in practice be performed with a non-constant duration. This amounts to adopting holding times that are no longer *place-dependent* but *token-dependent*. More formally, we say that a timed Petri net has **stochastic holding times** if there are nonnegative random variables $(\widehat{\tau}_p)_{p \in \mathcal{P}}$, such that every time a token is produced in the place p of \mathcal{P} , it is given a nonnegative holding time that is a realization of $\widehat{\tau}_p$. This token cannot be used to fire any transition of p^{out} before its age exceeds its holding time. Note that this formalism still supposes that the *distributions* of the holding times are place-dependent. However, the $(\widehat{\tau}_p)_{p \in \mathcal{P}}$ can follow any distribution, not necessarily exponential, since the latter is unrealistic for many applications. In particular some places may still use Dirac distributions if we want to model a task that lasts for a deterministic and constant duration.

By using stochastic holding times, the instants at which tokens are produced in places or at which transitions fire become random variables, and so do the counter functions. Nevertheless, equations (P1), (P3), (P4) and (P5) of Section 3.2 are still valid. Only the equation (P2) needs to be modified, since for a place p in \mathcal{P} the relationship between \widehat{Y}_p and \widehat{X}_p is no longer a simple delay. If p is in \mathcal{P} and k in \mathbb{N}^* , we denote by $\widehat{T}_k^{(p)}$ the k -th instant at which a token is produced in p (we convene that $\widehat{T}_k^{(p)} = \infty$ if no transition of p^{in} is live and that $X_p(t) < k$ for all $t \geq 0$), so that

$$\widehat{X}_p(t) = m_p + \sum_{k=1}^{\infty} \mathbb{1}_{[\widehat{T}_k^{(p)}, \infty)}(t).$$

Now denoting by $\widehat{\tau}_k^{(p)}$ the holding time of the k -th token to be produced in p , we have

$$\widehat{Y}_p(t) = m_p + \sum_{k=1}^{\infty} \mathbb{1}_{[\widehat{T}_k^{(p)} + \widehat{\tau}_k^{(p)}, \infty)}(t). \quad (\text{P2}')$$

The above expression actually defines an operation between the random variable \widehat{X}_p and the random variable $\widehat{\tau}_p$, and we shall denote $\widehat{Y}_p = \widehat{X}_p * \widehat{\tau}_p$, where

$$\forall \omega \in \Omega, \quad \forall t \in \mathbb{R}, \quad (\widehat{X}_p * \widehat{\tau}_p)(t; \omega) := m_p + \sum_{k=1}^{\infty} \mathbb{1}_{[T_k^{(p)}(\omega) + \tau_k^{(p)}(\omega), \infty)}(t). \quad (1.6)$$

The equation (P2') surely is more involved and harder to work with than (P2). It illustrates that the non-uniform holding times are likely to operate a “shuffle” of the tokens, in the sense that some young tokens may become mature before their older siblings. We let the reader check that if $\widehat{\tau}_p$ is in fact deterministically equal to the value τ_p , we retrieve $\widehat{Y}_p(t) = \widehat{X}_p(t - \tau_p)$, that is to say (P2).

A second possible source of nondeterminism may be implemented for preselection routings. In Section 2.1, the allocation of the initial marking (via the integer collections $(m_{p,q})_{q \in p^{\text{out}}}$) and of the produced tokens (via the mappings $(\Pi_q^p)_{q \in p^{\text{out}}}$) were known in advance for all p in $\mathcal{P}_{\text{psel}}$. Alternatively, each $m_{p,q}$ or $\Pi_q^p(n)$ with n in \mathbb{N} could be a random variable, for example to model systems with multiple treatment lanes where the allocation to one of these is unknown. In this case, we talk about **stochastic preselection routings**.

The most natural stochastic preselection routing is the so-called “Bernoulli routing” (see for instance [BGMo5a]), in which we consider that the preselection of a token entering the place p in $\mathcal{P}_{\text{psel}}$ towards the transition q in p^{out} is given by the realization of a boolean random variable \widehat{R}_q . The $(\widehat{R}_q)_{q \in p^{\text{out}}}$ are of course not independent, and we will suppose that their joint distribution is such that

$$\mathbb{P}(\widehat{R}_q = 1 \text{ and } \forall q' \in p^{\text{out}} \setminus \{q\}, \widehat{R}_{q'} = 0) = \pi_{qp}, \quad (1.7)$$

where the $(\pi_{qp})_{q \in p^{\text{out}}}$ are nonnegative real numbers verifying $\sum_{q \in p^{\text{out}}} \pi_{qp} = 1$. Observe that the “preselection mappings” $(\Pi_q^p)_{q \in p^{\text{out}}}$ used so far can be recovered as follows. It indeed suffices to consider independent and identically distributed sequences $((\widehat{R}_q^{(k)})_{q \in p^{\text{out}}})_{k \in \mathbb{N}^*}$ of random variables verifying (1.7), that allocate each token entering p to a single one of its downstream transitions. For $q \in p^{\text{out}}$ and $n \in \mathbb{N}^*$, the cumulative number of tokens routed from p to q among the n first is given by

$$\widehat{\Pi}_q^p(n) := \sum_{k=1}^n \widehat{R}_q^{(k)}, \quad \text{still satisfying } \sum_{q \in p^{\text{out}}} \widehat{\Pi}_q^p(n) = n. \quad (1.8)$$

In particular, the random variables $(\widehat{\Pi}_q^p(n))_{q \in p^{\text{out}}}$ follow a multinomial law with probabilities $(\pi_{qp})_{q \in p^{\text{out}}}$. The same reasoning could also be applied to the partition of the initial marking of each place ruled by preselection, and using the definition (1.8), the equation (P4) governing the evolution of counters for transitions in $\mathcal{Q}_{\text{pse1}}$ still holds. It is interesting to note that with Bernoulli routings, the fraction $\widehat{\Pi}_q^p(n)/n$ of tokens routed towards the transition q of p^{out} almost surely tends to π_{qp} when n approaches ∞ . Recall that this asymptotic property was (deterministically) true as well for periodic routings.

Finally, since we are interested in controlled Petri nets and that the evolution of the system highly depends on the prescribed firing instants of transitions in $\mathcal{Q}_{\text{input}}$, it is also worth considering a randomization of these entries. This can model systems having to cope with a demand that is not totally foreseeable. In this case, we talk about **stochastic inputs**.

Such stochastic inputs could be any random process on the nonnegative half-line. However, Poisson processes stand out from this class, for both their mathematical ease of use and the realistic representation they provide in many applications (we refer to the Figure 7.6 in Chapter 7 to justify how well Poisson processes are adapted to the models we study). Recall that a **Poisson process** \mathcal{P} with rate parameter $\lambda > 0$ is given by a sequence of (random) nonnegative physical times $(\widehat{t}_n)_{n \in \mathbb{N}}$, such that the *interarrival times* $(\widehat{t}_{n+1} - \widehat{t}_n)_{n \in \mathbb{N}}$ are independent realizations of an exponential random variable with parameter λ . It is customary to write a Poisson process under the càdlàg function form giving the cumulated number of arrivals over time, denoted by \widehat{G} :

$$\widehat{G}(t) := \sum_{n=1}^{\infty} \mathbb{1}_{[\widehat{t}_n, \infty)}(t). \quad (1.9)$$

Recall also that the number of arrivals of a Poisson process with parameter λ in a time-window of length Δt follows a Poisson law with parameter $\lambda \Delta t$, i.e., for all $t \geq 0$ and k in \mathbb{N} :

$$\mathbb{P}(\widehat{G}(t + \Delta t) - \widehat{G}(t^-) = k) = \frac{(\lambda \Delta t)^k}{k!} e^{-\lambda \Delta t}, \quad (1.10)$$

and in particular if we are interested in the average number of arrivals since the beginning, we have $\mathbb{E}(\widehat{G}(t)) = \lambda t$.

Poisson processes can be extended by **nonhomogeneous Poisson processes**, in which the arrival rate λ is no longer a constant but a measurable nonnegative function over $\mathbb{R}_{\geq 0}$. This rate function can be used for instance to take into account hourly or daily effects. It is customary in this context to introduce the integrated rate Λ defined for all $t \geq 0$ by

$$\Lambda(t) := \int_0^t \lambda(s) ds.$$

In nonhomogeneous Poisson processes, the interarrival times are no longer independent (if $t, s \geq 0$, we have $\mathbb{P}(\widehat{t}_{n+1} - \widehat{t}_n \geq t \mid \widehat{t}_n = s) = \exp(\Lambda(s+t) - \Lambda(s))$), but the function \widehat{G} still defined by (1.9) has *independent increments*, i.e., the random variables $(\widehat{G}(\widehat{t}_{n+1}) - \widehat{G}(\widehat{t}_n))_{n \in \mathbb{N}^*}$ are mutually independent. The equation (1.10) admits the nonhomogeneous counterpart

$$\mathbb{P}(\widehat{G}(t + \Delta t) - \widehat{G}(t^-) = k) = \frac{(\Lambda(t + \Delta t) - \Lambda(t))^k}{k!} e^{-(\Lambda(t + \Delta t) - \Lambda(t))}, \quad (1.11)$$

and we have in average $\mathbb{E}(\widehat{G}(t)) = \Lambda(t)$. We refer the reader to [Kal97] for advanced discussion on Poisson processes.

Example 1.25. We show in Figure 1.15 a timed Petri net that is very simple in terms of topology, but that is stochastic in the three different ways introduced before. The input transition q_1 fires according to a Poisson process with parameter 0.2 s^{-1} and produces tokens in p_1 , the holding time of a token in p_1 follows a Gamma law with shape parameter 6 s and scale parameter 2 , to which we add a constant delay of 1 s . Mature tokens in p_1 are reserved to fire q_2 and q_3 according to a Bernoulli preselection scheme, with proportions $3/10$ and $7/10$.

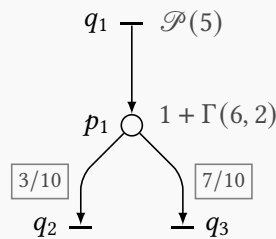


Figure 1.15: A timed Petri net with stochastic inputs, holding times and preselection routings

This system could model many simple real-life organizations. Think for instance of a polling station where voters cast their ballots in q_1 , that it takes a varying time to open and count, and finally the function Z_2 gives the number of votes over time received by a candidate expected to obtain 30 % of total votes...

We depict in Figure 1.16 the counter functions Z_1 , X_1 and Z_2 for two different executions of the Petri net.

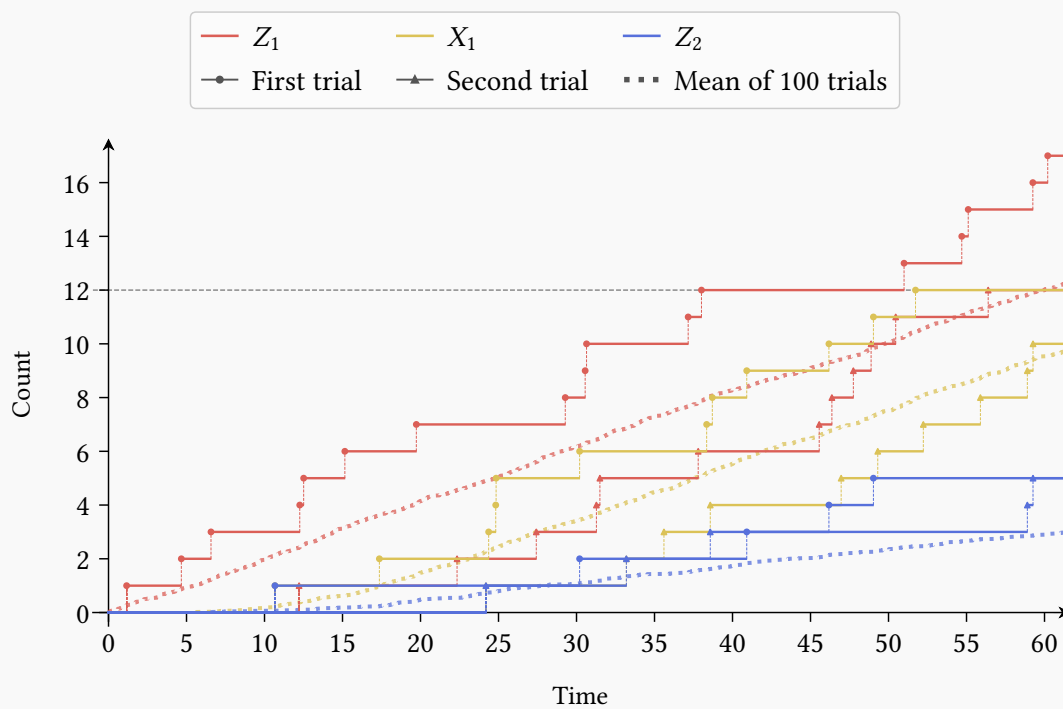


Figure 1.16: Stochastic counters functions associated with the Petri net of Figure 1.15

We can observe that these two trials provide very different outputs, for instance after 40 s, the transition q_1 has fired twelve times in the first execution, and only six times in the second. Observe that the horizontal shift between the curves of Z_1 and X_1 for a same trial varies importantly, for the holding times are no longer uniform. The jump times of Z_2 are contained in those of X_1 but with a three in ten chance, possibly leading to quite different empirical results. For instance still at 40 s, two tokens out of eight have been routed to q_2 in the first experience, versus three out of four in the second one.

We show in dotted lines the empirical mean obtained by averaging the counter functions of a hundred executions of the net. We can then see deterministic expectations arising from this operation, and we retrieve the fact that the expectation of Z_1 is linear with slope $1/5 \text{ s}^{-1}$ (and in particular in 60 s, q_1 has fired in average twelve times). We note that the average counter of X_1 follows the one of Z_1 up to a seemingly constant delay, except in the beginning of the execution where the slope is significantly lower. Finally, the average counter of Z_2 is obtained by multiplying the one of X_1 by $3/10$.

We shall discuss further in the Section 3 of Chapter 2 these models of Petri nets in which stochastic behaviours are allowed. In particular, we will explain how they induce deterministic systems through averaging techniques.

2

CONTINUOUS RELAXATION AND LARGE-SCALE LIMIT OF PETRI NETS

Contents

1	Semantics and tools for fluid Petri nets	49
1.1	Fluid Petri nets and their firing rules	49
1.2	Firing-rate semantics and fluid counter functions	52
1.3	Conflict-free transformation of preselection routings	61
2	Control aspects of monotonic discrete and fluid dynamics	62
2.1	Abstract form of the systems	62
2.2	The operator approach	63
3	Fluid Petri nets as large-scale limits of discrete models	69
3.1	The bounded deviations model	69
3.2	The threefold stochastic model	74

The Petri net paradigm introduced in Chapter 1 can represent accurately a huge variety of real-life organizations that fit in the discrete-event systems formalism. As mentioned in the [Introduction](#), whether it is in logistics, telecommunication networks or hardware design, problems are very often expressed by considering tasks or resources of an *atomic* or at least integer nature. This is well captured by the notion of tokens in a net, and also by counter functions that by definition are integer-valued and experience whole jumps at precise instants.

It is however known that these aspects generally come with issues of computational complexity. Although working with Petri nets or timed event graphs subject to integrality constraints is possible through an appropriate and advanced formalism (see for instance [[BCOQ92](#), [CHB13](#), [TCHR20](#)]), there are in this case many properties associated with structural analysis (reachability, liveness, boundedness, ...) that involve exploring structures whose size exponentially grows with the number of tokens at stake (known as the *state explosion* problem). In order to cope with these difficulties, [continuous relaxation](#) techniques have been proposed, in the work of Trivedi,

Kulkarni, Horton and Nicol [TK93, HKNT98] and in the one of Alla and David [AD98] – see also the work of Recalde and Silva [RS01]. They consist in deriving “continuous-state counterparts” of discrete models, in which integrality constraints are relaxed, but where the evolution semantics stay close from those of original Petri nets. These continuous-relaxed Petri nets are sometimes called “fluid” since everything happens in their evolution as if atomic tokens were replaced by fluids, that can be divided in arbitrary small parts. We shall be careful with such denomination to avoid confusion with the notion of fluidity of a system, that is to say for which the inputs are the only limiting factors in the evolution.

From another perspective, the integrality constraints of the models in Chapter 1 generally make dynamical aspects of timed Petri nets more intricate and less predictable, without resorting to explicit computation of the trajectory of the system. It thus turns desirable to simplify the recursive system of equations characterizing the evolution of the counter functions in Table 1.13, and to propose a **continuous-relaxed dynamics** in which all integrality constraints are discarded. These simplified dynamics have been studied by Cohen, Gaubert and Quadrat in [CGQ95, CGQ98] and by Gaujal and Giua [GG04b], in particular to derive analytical evaluations of firing throughputs in the net. These objectives were also pursued in [ABG15]. Among other things, the present part of the thesis aims at extending this work.

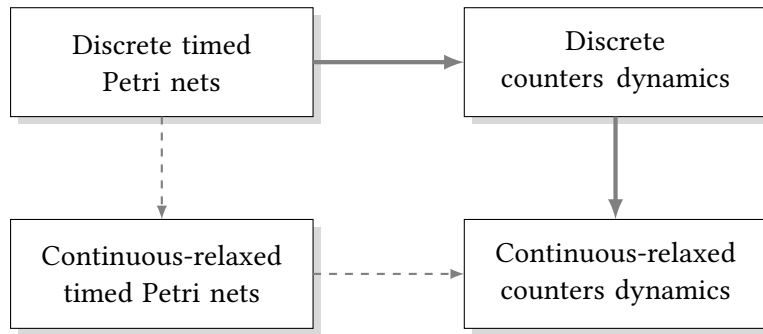


Figure 2.1: Two different routes leading to continuous-relaxed dynamics of Petri nets

We insist on the fact that our main motivation in the coming chapters is not to study the dynamics associated with continuous-relaxed Petri nets, but rather to focus on the continuous-relaxed dynamics of classical discrete Petri nets, as shown by the solid arrows in the Figure 2.1. The two notions may indeed not coincide, unless we adopt a suitable semantics for the former. Still, building a bridge between the commonly used models of “fluid” Petri nets and the set of dynamical equations we are going to study is valuable (dashed arrows in Figure 2.1), if only to get interpretation of these dynamics in terms of standard fluid models. That is the reason why we discuss fluid Petri nets in Section 1 and explain why relaxing integrality constraints in equations of Table 1.13 is in a sense well-founded. The material of Section 1.2 is purposefully introduced in a heuristic spirit, as establishing formal equivalences between different choices of semantics are beyond the scope of this discussion.

Section 2 takes a more abstract view on both classical and continuous-relaxed dynamics of Petri nets. It introduces evolution operators and evolution semigroups, that are standard in control theory for infinite dimensional time-delay systems (see for instance [Ber12a] for a case where dynamics are given by partial differential equations). In Section 3, we demonstrate that the approximation of a discrete timed Petri net by its continuous relaxation is asymptotically exact when applying a scaling parameter on the initial marking and inputs firings, even in presence of many non-deterministic events, which legitimizes the use of fluid models.

1 Semantics and tools for fluid Petri nets

1.1 FLUID PETRI NETS AND THEIR FIRING RULES. — The core idea of continuous-relaxed Petri nets is to allow non-integer and infinitely divisible tokens, along with non-integer firings of transitions. To this purpose, we first adapt the definitions of the previous chapter to authorize arbitrary markings and arc weights:

➤ **DEFINITION 2.1.** *We say that a 4-tuple $(\mathcal{P}, \mathcal{Q}, \mathcal{E}, \tau)$ is a **continuous-relaxed** (or **fluid**) **timed Petri net with multipliers** if contrary to the requirements of Definition 1.3 and Section 1.4 of Chapter 1, the marking vectors and the arc weights take **nonnegative real values**.*

Because the markings of the places in a continuous-relaxed Petri net are not necessarily integer, we may talk about **fluid-tokens** to refer to these “fractional” tokens. They have the key property to be divisible at will, since a real quantity can be split into an arbitrary number of arbitrary smaller real quantities. The term “fluid Petri net” is very often used in lieu of “continuous-relaxed Petri net” to convey the idea that fractional and divisible tokens are better understood if thought of as particles of fluids, that flow and pour into places, crossing transitions by being mixed (or unmixed) according to the multipliers stoichiometry.

The definition of the incidence matrix C in presence of multipliers given in Section 1.4 of Chapter 1 is unchanged, except that we now have that C is in $\mathbb{R}^{\mathcal{P} \times \mathcal{Q}}$. As before, C can be decomposed under the form $C = C^+ - C^-$, with C^+ and C^- in $\mathbb{R}_{\geq 0}^{\mathcal{P} \times \mathcal{Q}}$. The fundamental principle that transitions of the net fire by consuming resource in their upstream places and produce resource in their downstream places does survive to the continuous-relaxed case, but we now authorize **non-integer firings**. If q is a transition of \mathcal{Q} and ν is a positive real number, ν firings of q consume $\alpha_{qp}^{-1}\nu$ fluid-tokens in each place p of q^{in} , and produce $\alpha_{pq}\nu$ fluid-tokens in each place p of q^{out} . In the analogy with fluids circulating in the net, non-integer firings of transitions can be thought of as adjusting the opening of a tap to regulate the fluid outflow. Remark that non-integer firings preserve the idea of *packing* and *unpacking* tokens according to the arc multipliers, introduced in Section 1.4 of Chapter 1 – indeed, if a place p in q^{in} contains m_p fluid-tokens (with $m_p > 0$), their infinitely divisible character allows us to create α_{qp} “packs” of $\alpha_{qp}^{-1}m_p$ tokens. The notion of T-invariant introduced in Definition 1.6 of Chapter 1 also carries over, but these are now elements of $\mathbb{R}_{\geq 0}^{\mathcal{Q}}$ instead of $\mathbb{N}^{\mathcal{Q}}$.

As in discrete models, we shall also consider routing rules (resp. age constraints) to address possible conflicts (resp. to represent the progress of the physical time) that will restrict the possibility to fire some transitions. Apart from these, it can be seen that a major difference with the discrete framework is that a transition of a continuous-relaxed Petri net can be fired as soon as each of its upstream places contain any positive quantity of fluid-tokens, with no minimum threshold.

The study of structural properties of discrete Petri nets also have its continuous-relaxed counterpart; it is even made richer in the sense that some executions of fluid Petri nets can result in convergent marking sequences. The set of reachable or *lim-reachable* markings is therefore often studied, as well as associated notions of liveness, boundedness, etc. We refer to this purpose to the work of Fraca and Haddad [FH15] in which they show that several of these problems are decidable in polynomial time in the continuous-relaxed framework, which is not the case when dealing with discrete models of Petri nets (see also [VMJS13]). The following example illustrates the notions of non-integer firings.

Example 2.2. Figure 1.1 depicts a pure Petri net with three places and three transitions. As in Section 1.4 of Chapter 1, arc weights, that are not necessarily integer anymore, are given along the arcs when not equal to one. In this example, the net is equipped with an initial marking of $(V \ V \ 0)^T$, where V is a positive number (we do not represent atomic tokens anymore, but just in this example use pictures of fluids...)

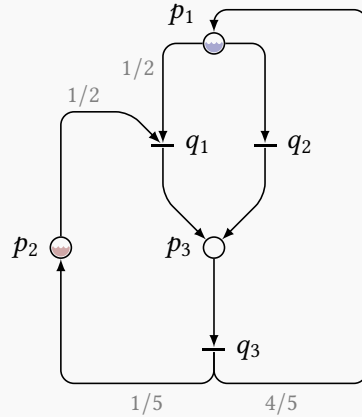


Figure 2.2: A continuous-relaxed Petri net

This continuous-relaxed Petri net could model the water circuit of a house, where p_1 is a cold water tank, and p_2 a hot water tank, both filled with a volume of V liters. Firing q_1 consists in taking a lukewarm shower, that is a fifty-fifty mix of hot and cold water, while firing q_2 uses only cold water. Firing the transition q_3 amounts to activate a recycling circuit that empties the bath tub and fills back the tanks, but this operation produces only 20 % of hot water again!

If a lukewarm shower consuming a total volume $2V$ is taken (i.e., q_1 fires $2V$ times) the marking vector becomes $(0 \ 0 \ 2V)^T$, and after recycling water, we obtain the marking $(8V/5 \ 2V/5 \ 0)^T$, meaning that the next person who wants to take the least cold and longest shower possible can only use a volume of $2V/5$ liters of hot water. Repeating this scheme indefinitely many times makes the marking of p_2 decrease geometrically with a rate of $2/5$, and the net reaches a *limit marking* of $(2V \ 0 \ 0)^T$, so that only cold showers can ultimately be taken (or fewer and fewer infinitesimal amounts of hot water are used, resulting in using no water)... Conversely, note that in this toy model, taking only cold showers indefinitely many times (i.e., using only transitions q_1 and q_3) also results in cold water shortage due to the little production of hot water by the recycling circuit.

It is also possible to combine non-integer firings of both q_1 and q_2 . For example, if q_2 is fired V times and q_1 is fired $V/2$ times from the initial situation, resulting in a water mix colder than lukewarm but warmer than cold, we obtain the marking $(12V/10 \ 8V/10 \ 0)^T$ after the recycling circuit is triggered, which is closer from the starting volumes than when using only q_1 . An equilibrium exists and can be found by looking at the incidence matrix associated with the Petri net of Figure 2.2, given by

$$C = \begin{pmatrix} -1/2 & -1 & 4/5 \\ -1/2 & 0 & 1/5 \\ 1 & 1 & -1 \end{pmatrix} \begin{matrix} p_1 \\ p_2 \\ p_3 \end{matrix} \cdot \begin{matrix} q_1 \\ q_2 \\ q_3 \end{matrix} .$$

The vector $y := (3/5 \ 2/5 \ 1)^\top$ verifies $Cy = \mathbf{0}$ and therefore is a T-invariant of the net. We deduce that using q_1 and q_2 in respective 60 % and 40 % proportions guarantees that both tanks are always refilled to their initial capacity. Alternatively, observe that the vector $e := (1 \ 1 \ 1)^\top$ is a so-called P-invariant of the incidence matrix ($eC = \mathbf{0}^\top$), which translates the conservation of the total fluid quantity in this particular example.

In the same manner as explained in Section 2 of Chapter 1, we are interested in Petri nets models where it is not up to an observer to make the choices of which fireable transitions should actually be fired, and in which proportions. The first underdetermined case to address is the one of conflicts, in case the same fluid-tokens present in a place p can be used to trigger several transitions in p^{out} . In this scenario, we still implement routing rules, either preselection or priorities. The notation $\mathcal{P}_{\text{psel}}$, $\mathcal{P}_{\text{prio}}$, $\mathcal{Q}_{\text{psel}}$, $\mathcal{Q}_{\text{prio}}$ and $\mathcal{Q}_{\text{sync}}$ from Chapter 1 are hereafter reused with the same meaning.

In the Section 2.1 of Chapter 1, the preselection routing for a place p was explicitly introduced using nondecreasing mappings $(\Pi_q^p)_{q \in p^{\text{out}}}$ from \mathbb{N}^* to itself such that $\sum_{q \in p^{\text{out}}} \Pi_q^p(n) = n$ for all n in \mathbb{N}^* . An identical construction could be brought up for continuous-relaxed models, replacing \mathbb{N}^* by $\mathbb{R}_{\geq 0}$ and imposing that $\sum_{q \in p^{\text{out}}} \Pi_q^p(x) = x$ for all $x \geq 0$. For instance, in a two-lane fork situation where $p^{\text{out}} = \{q_1, q_2\}$, one could define $\Pi_{q_1}^p: x \mapsto \int_0^x \sum_{k \in \mathbb{N}} \mathbb{1}_{[2kV, (2k+1)V)}(u) du$ and $\Pi_{q_2}^p: x \mapsto \int_0^x \sum_{k \in \mathbb{N}} \mathbb{1}_{[(2k+1)V, (2k+2)V)}(u) du$, where $V > 0$. This would model an allocation of fluid-tokens produced in p between q_1 and q_2 in a “one after another” fashion, by portions of V units of fluid each time. Although this formulation is the most general, we will in practice always consider much simpler preselection routings, that we may call **proportional preselection routings**, in which a vector of probabilities $(\pi_{qp})_{q \in p^{\text{out}}}$ is given, i.e., we have $\sum_{q \in p^{\text{out}}} \pi_{qp} = 1$, and for all q in p^{out} and $x \geq 0$ we have $\Pi_q^p(x) = \pi_{qp}x$. This amounts to saying that *any* quantity V of fluid-tokens produced in p is automatically split into the portions $(\pi_{qp}V)_{q \in p^{\text{out}}}$, and each one of these portions is *reserved* for the firing of one transition in p^{out} . In particular in this type of routing, the ratio $\Pi_q^p(x)/x$ is constant for all $x > 0$ and q in p^{out} . This property was only asymptotically reached for many preselection routings of interest for discrete models. Note that to obtain a fully determined routing, one also needs to precise the allocation of the initial marking V^0 of p . In the proportional setting, this can be done by specifying a probability vector $(\pi_{qp}^0)_{q \in p^{\text{out}}}$, either identical or different from the one governing the fluid-tokens produced in p , such that for all $q \in p^{\text{out}}$, the fraction $\pi_{qp}^0 V^0$ of the initial marking is reserved for the firing of q .

The formalism of priority routing introduced in the Section 2.2 of Chapter 1 does not need to be changed and applies also to continuous-relaxed nets. Given a place p equipped with a marking V and ruled by the local priority order (\prec_p) , a maximum amount of fluid-tokens must be used to fire the transition q_1 of p^{out} with highest priority degree (this quantity is necessarily V if $q_1^{\text{in}} = \{p\}$), then the second transition q_2 in priority order is fired a maximum number of times using the remaining tokens in p , and so on. We extend to fluid models the convention introduced on discrete models that transitions in $\mathcal{Q}_{\text{psel}}$ have only one upstream place (recall that it is not restrictive, cf. Figure 1.6), while those in $\mathcal{Q}_{\text{prio}}$ may have multiple upstream places.

Finally, we intend to avoid situations of “confusion”, in the same way as done in Section 2.3 of Chapter 1. To this purpose, we still rely on a notion of simultaneous firings. Since the transitions can fire a real number of times, we do not talk about simultaneously fireable multisets of transitions anymore, but we can keep the idea of a weight vector indicating if simultaneous and multiple firings of transitions are possible. Formally, given a vector v in $\mathbb{R}_{\geq 0}^{\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}}$ and a marking m in $\mathbb{R}_{\geq 0}^{\mathcal{P}}$, we say that the transitions of the net (or rather those for which v is positive) are **v -simultaneously fireable** if $\sum_{q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}} v_q [C^-]_q \geq m$. The simultaneous firings of these transitions consume (resp.

produces) fluid-tokens in the places of \mathcal{P} in amounts given by the vector $\sum_{q \in Q \setminus Q_{\text{input}}} v_q [C^-]_q$ (resp. $\sum_{q \in Q \setminus Q_{\text{input}}} v_q [C^+]_q$).

We extend to the continuous-relaxed setting the concept of **holding times** introduced in the Section 3.1 of Chapter 1, and impose similar semantic constraints. Any fluid-token produced in a place p in \mathcal{P} must wait for a nonnegative (so possibly null) physical time τ_p before being used or split to fire any transition of p^{out} . We may again say that a fluid-token with an age exceeding the holding time of its place is **mature**, and by convention the tokens of the initial marking have an infinite age. Like in the Definition 1.18 of Chapter 1, a continuous-relaxed Petri net is said to be non-Zeno if none of its cycles contain only places with zero holding times.

Finally, we point out that continuous-relaxed Petri nets may also admit input transitions, and to prevent Zeno behaviours, either on firing instants or quantities of produced fluid-tokens, we require as in the Chapter 1 that only a finite amount of fluid-tokens can be fired by these input transitions over any finite time-window.

1.2 FIRING-RATE SEMANTICS AND FLUID COUNTER FUNCTIONS. — Studying the behaviour of continuous-relaxed Petri nets over time first requires to carefully specify their evolution semantics. This is what we have done in the Chapter 1, by setting more and more evolution rules until there was no more underdetermination. We recalled in addition the results of [ABG15, Boeu17] who showed that these evolution rules were equivalent to the system of equations of Table 1.13, on discrete counter functions of the net.

So far, the evolution rules that we have described for continuous-relaxed Petri nets are contrary to their discrete peers not tight enough, due to a new behaviour made possible by the fluid context. Indeed, while a transition in a discrete Petri net could maybe fire one token (and not less!) every second, it is totally licit for its continuous-relaxed counterpart to fire a tenth of the fluid token every tenth of second, or even a thousandth fluid-token every thousandth of second. Pushing this reasoning to the limit, one can sense that the rightful notion to consider is that transitions of fluid Petri nets can have **firing rates**, also referred to as continuous firings.

In particular, it would be of little use in practice (although it is true) to state in the spirit of Theorem 1.13 that there exists a maximal vector v in $\mathbb{R}_{\geq 0}^{Q \setminus Q_{\text{input}}}$ such that the transitions of a continuous-relaxed net are v -simultaneously fireable. Indeed, this vector can be null and yet the transitions of the net can fire with positive firing rates. We need in fact to authorize both instantaneous and continuous firings in the semantics. Formally, if $v^{(0)}$ and $v^{(1)}$ are two nonnegative vectors of $\mathbb{R}_{\geq 0}^{Q \setminus Q_{\text{input}}}$, we should say that a continuous-relaxed Petri net with marking m in $\mathbb{R}_{\geq 0}^{\mathcal{P}}$ is $v^{(0)}$ -simultaneously and $v^{(1)}$ -continuously fireable if we have $\sum_{q \in Q \setminus Q_{\text{input}}} v_q^{(0)} [C^-]_q \geq m$, and $\sum_{q \in Q \setminus Q_{\text{input}}} v_q^{(1)} [C]_q \geq 0$. In this case, for $dt > 0$ small enough, the firing of transitions of the net with immediate multiplicity $v^{(0)}$ and rate $v^{(1)}$ at time $t = 0$ would produce a marking $m(t) = m + \sum_{q \in Q \setminus Q_{\text{input}}} (v^{(0)} + v^{(1)} dt) [C]_q + o(dt)$ (note that the transitions can have time-dependent firing rates). Proceeding as in Chapter 1 would therefore prompt us to show that there exists at each instant two maximal vectors $v^{(0)}$ and $v^{(1)}$ such that the net is $v^{(0)}$ -simultaneously and $v^{(1)}$ -continuously fireable.

Instead, we develop below an approach that is converse of the one applied in the Chapter 1. We postulate that the evolution of the net is dictated by some equations on the counter functions, which keep track of the evolution of fluid-tokens in the net. The following definition formally introduces them, and to insist on the fact that their dynamics is different from the one arising in the discrete context, we denote them by using lowercase letters.

- **DEFINITION 2.3.** Let $(\mathcal{P}, \mathcal{Q}, \mathcal{E}, \tau)$ be a non-Zeno continuous-relaxed timed Petri net equipped with an initial marking. We call **counter functions** of the net the mappings $(x_p)_{p \in \mathcal{P}}$, $(y_p)_{p \in \mathcal{P}}$ and $(z_q)_{q \in \mathcal{Q}}$ from $\mathbb{R}_{\geq 0}$ to $\mathbb{R}_{\geq 0}$ such that for all $t \geq 0$:

$x_p(t)$ denotes the (nonnegative real) number of fluid-tokens that have been present in place p up to time t included (i.e., either the fluid-tokens produced in p or those of the initial marking in p);

$y_p(t)$ denotes the (nonnegative real) number of fluid-tokens that have been present and mature in place p up to time t included;

$z_q(t)$ denotes the (nonnegative real) number of firings of transition q up to time t included.

Denoting by m the initial marking of the net, we extend the domain of these counters to \mathbb{R} by setting $x_p(t) = y_p(t) = m_p$ and $z_q(t) = 0$ for all $t < 0$, p in \mathcal{P} and q in \mathcal{Q} . The main difference with discrete counters is that the functions $(x_p)_{p \in \mathcal{P}}$, $(y_p)_{p \in \mathcal{P}}$ and $(z_q)_{q \in \mathcal{Q}}$ do no longer take values in \mathbb{N} but in $\mathbb{R}_{\geq 0}$, hence are not piecewise-constant anymore.

- **PROPOSITION 2.4.** The counter functions of a continuous-relaxed non-Zeno timed Petri net are nondecreasing piecewise-continuous càdlàg functions.

∴ *Proof.* The proof is the same as the Proposition 1.20 of Chapter 1. ♣

Recall that monotone functions are almost everywhere continuous (see for instance [Rud76, Theorem 4.30, p. 90]), and even almost everywhere differentiable (see [WZ77, Theorem 7.22, p. 177]); it is then licit to talk about the derivatives of the counter functions. Actually, we shall soon only work with Petri nets whose input profiles are such that the counters are piecewise-continuously-differentiable – we will come back on this aspect after the semantics are tightened. Following Proposition 2.4, observe that any jump of a counter z_q at the instant $t \geq 0$ (characterized by $z_q(t) > z_q(t^-)$) translates that $z_q(t) - z_q(t^-)$ instantaneous firings of q occurred at t , while continuity of z_q over $[t_1, t_2)$ means that q fires with a rate $\dot{z}_q(t)$ for all t in $[t_1, t_2)$ (wherever this derivative exists).

We now discuss the form of the equations that we choose to govern the evolution of our continuous-relaxed Petri nets (we will not prove that this yields identical semantics than imposing maximum vectors $\nu^{(0)}$ and $\nu^{(1)}$ at each instant). The fact that firing rates are involved and that counters functions (and therefore the markings of places as well) admit derivatives does explain why various authors introduce fluid semantics of Petri nets as systems of ordinary differential equations (ODEs). Since $\dot{z}_q(t)$ gives the firing rate of transition q in \mathcal{Q} , the equation governing the instantaneous variations of the amount $m_p(t)$ of fluid-tokens in p of \mathcal{P} at instant t is

$$\dot{m}_p(t) = \sum_{q' \in p^{\text{in}}} \alpha_{pq'} \dot{z}_{q'}(t) - \sum_{q \in p^{\text{out}}} \alpha_{qp} \dot{z}_q(t). \quad (2.1)$$

Remark that contrary to the $(\dot{z}_q)_{q \in \mathcal{Q}}$ that are nonnegative, the $(\dot{m}_p)_{p \in \mathcal{P}}$ are *a priori* allowed to be negative (if more tokens are consumed than produced). The equation (2.1) is nothing else than the derivative of the fundamental equation $m(t) = m(0^-) + Cz(t)$ for the place p of \mathcal{P} . However, since only mature tokens of p can be used to fire transitions of p^{out} , it is desirable to decompose the term $m_p(t)$ into $m_p^{\text{YOUNG}}(t) + m_p^{\text{OLD}}(t)$, where m_p^{YOUNG} (resp. m_p^{OLD}) is the number of fluid-tokens present in p at instant t that are not yet mature (resp. that are mature). We can then explicitly introduce a **conversion rate** $\phi_p^{\text{YOUNG} \rightarrow \text{OLD}}(t)$ indicating for all p in \mathcal{P} the amount of non-mature

tokens of p becoming mature at instant t , by writing

$$\begin{aligned} \dot{m}_p^{\text{YOUNG}}(t) &= \sum_{q' \in p^{\text{in}}} \alpha_{pq'} \dot{z}_{q'}(t) - \phi_p^{\text{YOUNG} \rightarrow \text{OLD}}(t), \\ \text{for all } p \in \mathcal{P}, \end{aligned} \quad (2.2)$$

$$\dot{m}_p^{\text{OLD}}(t) = - \sum_{q \in p^{\text{out}}} \alpha_{qp} \dot{z}_q(t) + \phi_p^{\text{YOUNG} \rightarrow \text{OLD}}(t).$$

We naturally retrieve (2.1) by taking the sum of both equations in (2.2). To characterize the firing rates $(\dot{z}_q)_{q \in Q \setminus Q_{\text{input}}}$, we remark that in the discrete framework of Chapter 1, each transition admitted at each instant an upstream place with no idle mature token, in accordance with the requirement to fire as many tokens as possible. We shall later refer to these places (i.e., the set $\{p \in \mathcal{P} \mid m_p^{\text{OLD}}(t) = 0\}$) as **bottleneck places**, since they have no fluid-tokens to offer to the firing of their downstream transitions. In addition, by the **earliest behaviour** rule that we still want verified in the continuous-relaxed setting, if there are several bottleneck upstream places of a transition q , then the firing rate \dot{z}_q is limited by the creation rate of mature tokens $\phi_p^{\text{YOUNG} \rightarrow \text{OLD}}$ of one of these bottleneck places. Except for transitions involved in structural conflicts (and whom we will focus on afterwards), these two conditions can both be coded by the equation

$$\text{for all } q \text{ in } Q_{\text{sync}}, \quad \min_{\substack{p \in q^{\text{in}} \\ m_p^{\text{OLD}}(t) = 0}} (\dot{m}_p^{\text{OLD}}(t)) = 0. \quad (2.3)$$

Note indeed that for all q in Q_{sync} , $\{p \in q^{\text{in}} \mid m_p^{\text{OLD}}(t) = 0\}$ must be nonempty for a minimum taken on this set to be non-infinite. The semantics are completed by specifying the form of the conversion rates $(\phi_p^{\text{YOUNG} \rightarrow \text{OLD}})_{p \in \mathcal{P}}$. A standard choice is that a constant fraction of non-mature tokens does become mature at each instant,¹ so that $\phi_p^{\text{YOUNG} \rightarrow \text{OLD}}(t)$ is proportional to $m_p^{\text{YOUNG}}(t)$:

$$\text{for all } p \in \mathcal{P}, \quad \phi_p^{\text{YOUNG} \rightarrow \text{OLD}}(t) := \frac{m_p^{\text{YOUNG}}(t)}{\tau_p} \quad (\text{exponential conversion rate}), \quad (2.4)$$

where $\tau_p > 0$. Re-using (2.2), one can see that (2.4) leads to a proper ODE system, and this is the model studied by Allamigeon, Bœuf and Gaubert in [ABG17]. It is equivalent to the system obtained by David and Alla in [AD98] and revisited by Recalde and Silva in [RS01]. In this setting, if a unit of fluid is produced in a place p at time $t = 0$, with $p^{\text{out}} = \{q\}$ and $q^{\text{in}} = \{p\}$, then q will fire this fluid-token according to an exponential law with parameter $\tau_p > 0$, i.e., a total proportion of $(1/\tau_p) \int_0^t e^{-u/\tau_p} du$ units of fluid originating from the initial fluid-token will be consumed in p between 0 and t . This indeed corresponds to a firing rate that is proportional to the marking of p itself (remark that it takes an infinite time to fully consume the initial fluid-token).

Although it yields non-trivial and rich dynamics, the postulate (2.4) does not match our requirement that fluid-tokens become mature after a constant holding time. To enforce this behaviour, we shall instead work with conversion rates such that

$$\text{for all } p \in \mathcal{P}, \quad \phi_p^{\text{YOUNG} \rightarrow \text{OLD}}(t) := \sum_{q' \in p^{\text{in}}} \alpha_{pq'} \dot{z}_{q'}(t - \tau_p) \quad (\text{delayed conversion rate}), \quad (2.5)$$

which means that the conversion rate between non-mature and mature fluid-tokens in place p is precisely given by p 's inflow of tokens, delayed by τ_p . Contrary to the form (2.4), the choice (2.5)

¹ Similarly to the usual decay model of radioactive nuclei, each nucleus having a probability $dP = dt/\tau$ to break down over an infinitesimal duration dt

leads to a somehow “factitious” ODE system, since the two relations of (2.2) become easily integrable and we have $m_p^{\text{OLD}}(t) = m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}(t - \tau_p) - \sum_{q \in p^{\text{out}}} \alpha_{qp} z_q(t)$ for all p in \mathcal{P} . The fact given by (2.3) that any transition q of $\mathcal{Q}_{\text{sync}}$ admits bottleneck upstream places (and that these upstream places only admit q as downstream transition due to the absence of conflict) then provides

$$\text{for all } q \text{ in } \mathcal{Q}_{\text{sync}}, \quad z_q(t) = \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}(t - \tau_p) \right).$$

We recognize the equation obtained for transitions in $\mathcal{Q}_{\text{sync}}$ for discrete timed Petri nets (see Table 1.13), except that the integer-parts have been dropped, and the arc multipliers are not compelled to be integer. In this sense, continuous-relaxed dynamics equations on counter functions arise from the semantics of continuous-relaxed timed Petri nets under the “delayed conversion rate” rule. Similarly to the discrete framework, the previous equation also writes $z_q(t) = \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} y_p(t)$ if q is in $\mathcal{Q}_{\text{sync}}$, where for all p in \mathcal{P} , we have $x_p(t) = m_p + \sum_{q \in p^{\text{in}}} \alpha_{pq} z_q(t)$ and $y_p(t) = x_p(t - \tau_p)$. These three relations are the fluid counterparts of equations (P1), (P2), and (P3) from Chapter 1.

One could also have obtained a fluid equation for transitions in $\mathcal{Q}_{\text{psel}}$ that resembles the discrete one. Observe that if p is a place ruled by proportional preselection routing with the probabilities $(\pi_{qp})_{q \in p^{\text{out}}}$, the second equation of (2.2) also writes $\dot{m}_p^{\text{OLD}}(t) = -\sum_{q \in p^{\text{out}}} (\alpha_{qp} \dot{z}_q(t) + \pi_{qp} \phi_p^{\text{YOUNG} \rightarrow \text{OLD}}(t))$. In accordance with the preselection logic, we shall then require that a fraction π_{qp} of the non-mature fluid-tokens becoming mature are allocated to the firing of the transition q in p^{out} , and introduce the number of tokens $(m_{p,q}^{\text{OLD}}(t))_{q \in p^{\text{out}}, t \in \mathbb{R}}$ that are mature in p and reserved for q at time t , so that $\dot{m}_{p,q}^{\text{OLD}}(t) = -\alpha_{qp} \dot{z}_q(t) + \pi_{qp} \phi_p^{\text{YOUNG} \rightarrow \text{OLD}}(t)$ for all q in p^{out} . The equation (2.3) becomes $m_{p,q}^{\text{OLD}}(t) = \dot{m}_{p,q}^{\text{OLD}}(t) = 0$ since no transition q of p^{out} is involved in synchronization patterns and should fire mature tokens of p as soon as possible. Equivalently, this means using the initial condition $m_p^{\text{OLD}}(0^-) = \pi_{qp}^0 m_p$ that

$$\text{for all } q \text{ in } \mathcal{Q}_{\text{psel}} \text{ such that } q^{\text{in}} = \{p\}, \quad z_q(t) = \alpha_{qp}^{-1} \left(\pi_{qp}^0 m_p + \pi_{qp} \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}(t - \tau_p) \right), \quad (2.6)$$

which is indeed a continuous-relaxed counterpart of preselection equation of Table 1.13, since discrete initial markings and preselection mappings degenerate into real fractions.

Our point is that starting from natural fluid semantics where evolution rules are stated on firing rates, one recovers under the “delayed conversion rate” convention the dynamical equations of the Table 1.13 for discrete Petri nets, with integrality constraints removed. Hence, echoing the Figure 2.1, we highlight the following principle for priority-free nets:

The continuous-relaxation of dynamical equations on counter functions of Table 1.13 for discrete timed Petri nets also describes the dynamics of fluid counter functions for continuous-relaxed timed Petri nets, under the *delayed conversion rate* semantics.

This observation legitimizes the fluid equations that we shall use in the thesis, and shows that they are not only brought up for computation convenience. Instead, they characterize the behaviour of systems that are idealized, but still obey very reasonable evolution rules.

Note that the case of priority routings is once again a little more subtle when firing rates are involved. It is addressed in [ABG17] (using exponential conversion rate, but it could be adapted to the delayed conversation rate as well) for local priority patterns involving two transitions, and in

Type	Counter equation in the continuous-relaxed model
$q \in \mathcal{Q}_{\text{input}}$	$z_q(t)$ is prescribed externally
$q \in \mathcal{Q}_{\text{sync}}$	$z_q(t) = \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}(t - \tau_p) \right)$
$q \in \mathcal{Q}_{\text{psel}}$	$z_q(t) = \alpha_{qp}^{-1} \left(\pi_{qp}^0 m_p + \pi_{qp} \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}(t - \tau_p) \right)$
$q \in \mathcal{Q}_{\text{prio}}$	$z_q(t) = \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}(t - \tau_p) - \sum_{q' <_p q} \alpha_{q'p} z_{q'}(t) - \sum_{q' >_p q} \alpha_{q'p} z_{q'}(t^-) \right)$

Table 2.3: Dynamic equations followed by transitions counter functions

this case we need to provide additional semantic equations. We do not develop this aspect further, and when priority rules will be featured in continuous-relaxed timed Petri nets, we shall use the same dynamical counter equations than those stated in Table 1.13 for discrete nets, but with integer-parts dropped, in accordance with the above principle. This approach is justified by the fact shown in [ABG17, Bœu17] that both settings provide identical evaluations of the long-run behaviour of the system.

We gather in the Table 2.3 above the system of dynamical equations that governs the evolution of fluid counter functions of transitions of the net based on their type. Similarly to the discrete case, as soon as the Petri net is non-Zeno, we can use the induction principle (\mathcal{I}) to study the effect of these equations by chunks of τ_{\min} units of time. In particular, we obtain that the continuous-relaxed dynamics define unique trajectories:

- **THEOREM 2.5.** *In a non-Zeno and priority-free continuous-relaxed controlled timed Petri net, the knowledge of $(z_q)_{q \in \mathcal{Q}_{\text{input}}}$ and the system of equations of Table 2.3 determine unique nondecreasing càdlàg functions $(z_q)_{q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}}$ over \mathbb{R} .*

⋮ *Proof.* The proof is identical to the one of Theorem 1.23, making use of non-Zeno substitutions in the equations and the induction scheme (\mathcal{I}) to propagate the dynamics. ✍

We can use the same technique to show that the unique solution of the dynamics inherits the regularity of the input profiles, hence if the latter are piecewise-continuously-differentiable, so is the whole trajectory.

Note that we did state Theorem 2.5 only in the “monotonic” framework, that is to say without priority rules. The fact that pathological input profiles can yield accumulating sequences of instants for which minimizing terms change in the equations Table 2.3 makes it more difficult to extend Theorem 1.24. In Chapter 5, we shall use the result of Theorem 2.5 also under priority rules, but with much simpler inputs.

In the following example, we review some of the similarities and differences between the two above models of conversion rates, for a very simple Petri net with no structural conflict. It is the only place of the thesis where we consider the exponential conversion rate; all the future computations on continuous-relaxed Petri nets will systematically use the equations of Table 2.3, hence supposing delayed conversion rates.

Example 2.6. We show in the Figure 2.4 a simple timed Petri net representing a so-called “single-tier” inbound call-center, which means there is only one pool of identical agents picking up the inbound calls, and carrying out the instruction in its entirety. The calls arrive via the uppermost transition q_0 and cumulated arrivals are counted by z_0 . A characteristic time τ_0 must elapse before an arrived call is presented to one of the N agents, who can then pick up the calls by firing transition q_1 . The instruction will take a typical time τ_1 , and once the conversation is over, q_2 fires, hence sending back the agent in his or her pool. Note that we suppose that a characteristic time τ_2 must elapse before this agent is available again before picking up new calls (it can be thought of either as a recovery time or a back-office task). Since we are working in the continuous-relaxed framework, firings of transitions can be non-integer. In the coming illustrations, we have used $N = 3$, $\tau_0 = 10$ s, $\tau_1 = 120$ s and $\tau_2 = 5$ s.

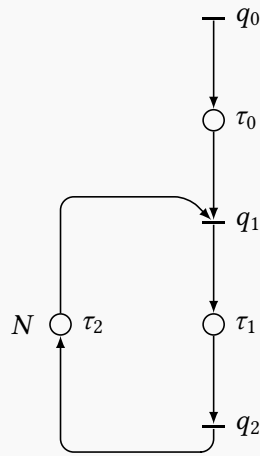


Figure 2.4: A timed Petri net representing a single-tier inbound call-center

On one hand, the response of the call-center depends on the input profile $t \mapsto z_0(t)$ of the inbound demand. We study below three simple cases for this profile: low-slope ramp, high-slope ramp, and high step. Yet, the semantics rule we choose also play a key role. In the *delayed conversion rate* setting (later abbreviated DCR), the dynamics equations write

$$\begin{aligned} z_1(t) &= \min(z_0(t - \tau_0), N + z_2(t - \tau_2)) \\ z_2(t) &= z_1(t - \tau_1) \end{aligned} \quad (\text{ST-DCR})$$

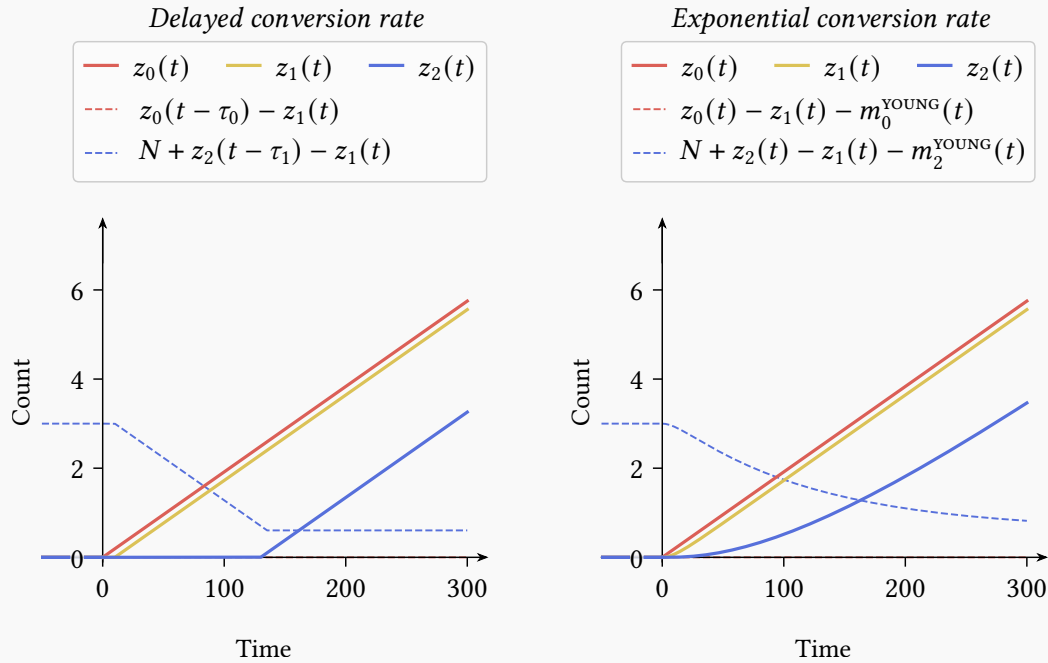
In the contrary, if we place ourselves in the *exponential conversion rate* setting (abb. ECR), we have to solve the following piecewise-linear differential system:

$$\begin{aligned} \dot{m}_0^{\text{YOUNG}}(t) &= \dot{z}_0(t) - m_0^{\text{YOUNG}}(t)/\tau_0 \\ \dot{z}_1(t) &= \begin{cases} m_0^{\text{YOUNG}}(t)/\tau_0 & \text{if } z_0(t) - z_1(t) - m_0^{\text{YOUNG}}(t) = 0 \\ m_2^{\text{YOUNG}}(t)/\tau_2 & \text{if } N + z_2(t) - z_1(t) - m_2^{\text{YOUNG}}(t) = 0 \end{cases} \\ \dot{m}_1^{\text{YOUNG}}(t) &= \dot{z}_1(t) - m_1^{\text{YOUNG}}(t)/\tau_1 \\ \dot{z}_2(t) &= m_1^{\text{YOUNG}}(t)/\tau_1 \\ \dot{m}_2^{\text{YOUNG}}(t) &= \dot{z}_2(t) - m_2^{\text{YOUNG}}(t)/\tau_2 \end{aligned} \quad (\text{ST-ECR})$$

In both system labels, the abbreviation “ST” stands for single-tier. The system of equations (ST-DCR) is easy to solve using the induction principle (\mathcal{I}). To find solutions of the system (ST-ECR), we can first look for the functions $(m_0^{\text{YOUNG}}, m_1^{\text{YOUNG}}, m_2^{\text{YOUNG}})$. Indeed, as a vector function h , they satisfy a coupled piecewise-linear differential system of the form $\dot{h}(t) = Ah(t) + b(t)$, where the matrix A can take two different values (both diagonalizable) depending upon which term attains the minimum in the equation on z_1 . Afterwards, we can integrate the obtained functions and get z_1 and z_2 .

Remark that in both cases, as long as not too many calls arrive, starting with N idle answering agents guarantees that at the start of the execution we first use the equation $z_1(t) = z_0(t - \tau_0)$ in (ST-DCR), and the equation $\dot{z}_1(t) = m_0^{\text{YOUNG}}(t)/\tau_0$ in (ST-ECR).

a) The input profile is a low-slope ramp. We first suppose that for all $t \geq 0$, $z_0(t) = \lambda t$, with λ not too high. This corresponds to arrivals of inbound calls with a constant rate. The resolution of both systems (ST-DCR) and (ST-ECR) provides the following counter evolution:



The solutions in the DCR model are simple:

$$z_1(t) = \lambda(t - \tau_0)\mathbb{1}_{[\tau_0, \infty)}(t) \quad \text{and} \quad z_2(t) = \lambda(t - \tau_0 - \tau_1)\mathbb{1}_{[\tau_0 + \tau_1, \infty)}(t).$$

This is very reasonable; it tells us that call pick ups and hang ups follow the pace imposed by the input, even though holding times of the places create a constant delay between z_0 , z_1 and z_2 . Since τ_0 is short, z_1 closely follows z_0 , contrary to z_2 which follows z_1 by an important delay due to the conversation time τ_1 .

For the ECR model, we obtain little more complex expressions:

$$z_1(t) = \lambda(t - \tau_0(1 - e^{-t/\tau_0})) \quad \text{and} \quad z_2(t) = \frac{\lambda}{\tau_1 - \tau_0} \left(\tau_1(t - \tau_1(1 - e^{-t/\tau_1})) - \tau_0(t - \tau_0(1 - e^{-t/\tau_0})) \right),$$

however it is easy to figure out that these functions, that are continuously differentiable, admit the same asymptotes than in the DCR case (and therefore with same interpretation).

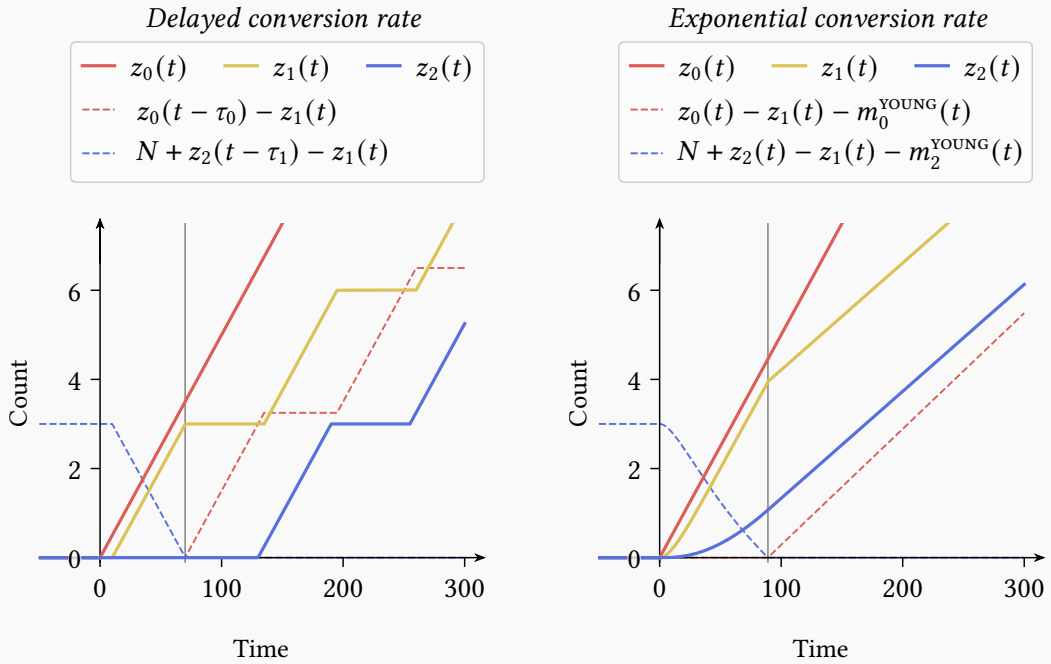
The major difference with the DCR situation is that the transition q_2 fires from the very start of the experience, which is not very realistic for this application.

In both cases, we have shown in dashed blue line by “how far” the system is from switching regime (in the ECR model, this is exactly the function m_2^{OLD}). When it hits zero, it means that no answering agents can be idle anymore, or in other words the internal part of the system is no longer ahead of the input. It is remarkable that in the two cases, the asymptotic value of this quantity is $N - \lambda(\tau_1 + \tau_2)$ (it is not too hard to obtain with DCR equations, but a little more involved with ECR ones). This means that the arrival rate λ^* defined by

$$\lambda^* := \frac{N}{\tau_1 + \tau_2}$$

plays the role of a *critical arrival rate*, and we are here in the situation where $\lambda \leq \lambda^*$. Note how λ^* can be interpreted as a mean cycle time, since we have N agents performing instructions requiring a total of $\tau_1 + \tau_2$ units of time.

b) The input profile is a high-slope ramp. If we still impose $z_0(t) = \lambda t$ but choose $\lambda \geq \lambda^*$, we get quite different behaviours of the counter functions.



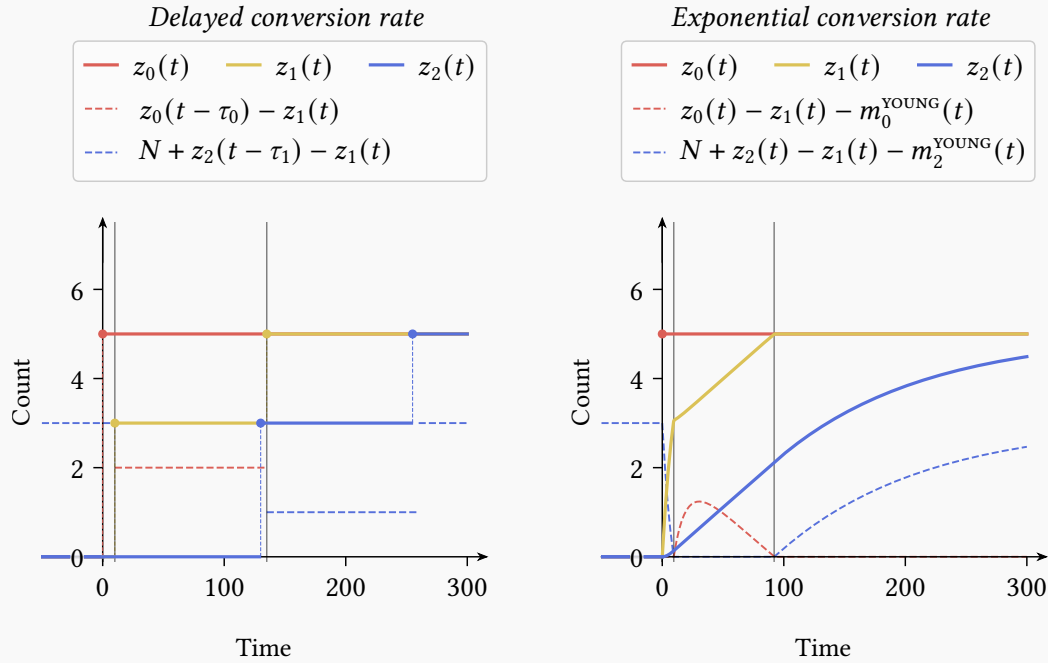
For the DCR case, we still have $z_1(t) = \lambda(t - \tau_0)$ shortly after τ_0 , but after taking N calls (which occurs at time $\tau_0 + N/\lambda$), all the agents are busy so no more calls can be picked up, and z_1 stays constant. The system is at this point in another regime, where the counter variables do not follow the input anymore. At time $\tau_0 + \tau_1 + \tau_2$, the agents having picked up the first calls (at time τ_0) become available again, hence the growth of z_1 resumes at pace λ . This scheme repeats indefinitely, alternating between *plateaus* of length $\tau_1 + \tau_2 - N/\lambda$ and ramps with slope λ and length N/λ . This “zigzag” pattern hence has pseudo-slope λ^* .

Using ECR equations, we also observe that the system enters in a new phase for some instant t_0 such that $m_2^{\text{OLD}}(t_0) = 0$. The change of slope of z_1 at t_0 is justified by the discontinuity of \dot{z}_1 in (St-ECR), but z_2 remains continuously differentiable. In this second regime, the

evolution of z_1 and z_2 is autonomous and occurs with ultimate slope λ^* ; in addition z_2 follows z_1 with an ultimate delay τ_1 .

An important difference between the two models is that in the DCR case (resp. ECR), z_2 begins its growth at $t = \tau_0 + \tau_1$ (resp. $t = 0$). This means that the comebacks of answering agents occurs sooner in ECR, which is why the phase switch happens later than the DCR model – and as a consequence the solutions of DCR lag behind those of ECR.

c) The input profile is a high step. We suppose that $z_0(t) = c\mathbb{1}_{[0,\infty)}(t)$ is a Heaviside-type function, representing a sudden bulk of calls arriving at $t = 0$ (we choose $c = 5$, so that $N < c < 2N$).



In the DCR model, none of these calls are visible to the idle agents before the instant τ_0 , so that we are technically in an input-driven regime. At $t = \tau_0$, the answering agents pick up as many calls as they can, so that z_1 suddenly jumps at N . These N calls are handed up τ_1 units of time later, and our agents become available again at time $\tau_0 + \tau_1 + \tau_2$. They are in sufficient staffing to pick up the remaining calls, so that the demand is caught up with and the system switches back in the first regime. The $c - N$ last calls are handed up at time $\tau_0 + 2(\tau_1 + \tau_2)$.

For the ECR setting, the fact that z_0 is discontinuous at time $t = 0$ also makes m_0^{YOUNG} discontinuous, and as a consequence z_1 non-differentiable. The counter z_1 quickly grows until all agents are saturated. In this second phase beginning with a new discontinuity of \dot{z}_1 , both counters have a nearly-linear growth with “slope” λ^* . When z_1 catches up with the input, it naturally stays flat, and z_2 is a simple exponential function afterwards. In particular, because the asymptotic value of m_1^{YOUNG} is null, the growth z_2 becomes slower and slower, so that the total catch-up of the input takes an infinite time.

1.3 CONFLICT-FREE TRANSFORMATION OF PRESELECTION ROUTINGS. — At this point, it is worth mentioning an operation on continuous-relaxed Petri nets that has no equivalent for discrete nets. Introduced by Gaujal and Giua in their work [GGo4b], it consists in replacing the conflict places ruled by proportional preselection routing by equivalent conflict-free patterns, thus referred to as a **conflict-free transformation**. If applied to an initial Petri net that do not feature any priority routing, this technique provides a transformed Petri net with no conflict place at all, or in other words, such that $|p^{\text{out}}| \leq 1$ for all places p in \mathcal{P} .

Specifically, given a continuous-relaxed Petri net $(\mathcal{P}, \mathcal{Q}, \mathcal{E})$ such that $\mathcal{P}_{\text{psel}} \neq \emptyset$, its *equivalent conflict-free* net is the Petri net $(\mathcal{P}', \mathcal{Q}, \mathcal{E}')$ such that

- Each place p in $\mathcal{P}_{\text{psel}}$ ruled by preselection in the original Petri net is duplicated as many times as its number of downstream transitions (into $\{p_q \mid q \in p^{\text{out}}\}$), so that in the transformed net

$$\mathcal{P}' := (\mathcal{P} \setminus \mathcal{P}_{\text{psel}}) \cup \bigcup_{p \in \mathcal{P}_{\text{psel}}} \{p_q \mid q \in p^{\text{out}}\},$$

- Each one of this new places now admits *a single* downstream transition, with unchanged weight

$$\forall p \in \mathcal{P}_{\text{psel}}, \quad \forall q \in p^{\text{out}}, \quad \alpha'_{qp} := \alpha_{qp},$$

- Each one of this new places admits *all* upstream transitions of their corresponding original place p as upstream transitions, with arc multipliers weighted by p 's preselection probabilities

$$\forall p \in \mathcal{P}_{\text{psel}}, \quad \forall q \in p^{\text{out}}, \quad \forall q' \in p^{\text{in}}, \quad \alpha'_{pqq'} := \pi_{qp} \alpha_{pq'}.$$

This construction is illustrated by the Figure 2.5. The Figure 2.5a depicts a pattern involving a place ruled by proportional preselection routing with two upstream transitions and two downstream transitions. As a result, its conflict-free equivalent still features four transitions but now two places instead of one. Observe on Figure 2.5b how each upstream transition now produces tokens in all the newly created places, with arc multipliers taking into account the former routing proportions.

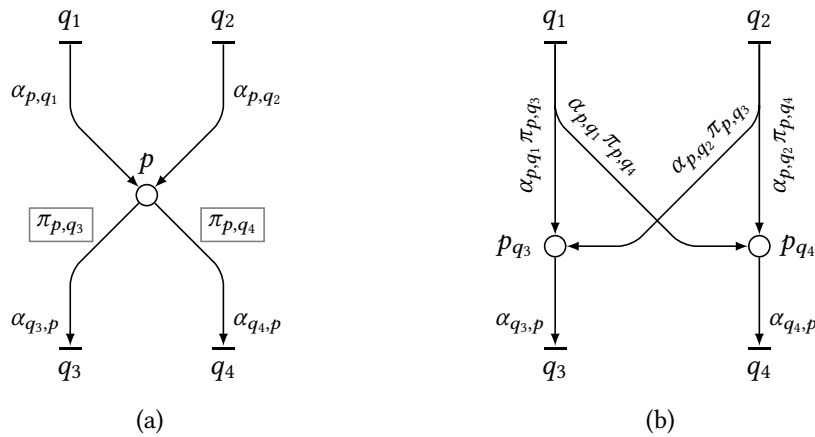


Figure 2.5: (a) A Petri net with a place ruled by proportional preselection routing. (b) Its conflict-free equivalent Petri net, all places have now a single downstream transition.

The main interest of the conflict-free transformation is naturally to avoid dealing with conflicts and confusion, but it also spares the need to distinguish dynamical equations for places in $\mathcal{P}_{\text{psel}}$

(by construction, $\mathcal{P}_{\text{psel}} = \emptyset$ in conflict-free transformed nets). Above all, we have the following result, once again for standard Petri nets (i.e., with no priority rules).

- **PROPOSITION 2.7** ([GG04b]). *The transition counter functions of a non-Zeno continuous-relaxed timed Petri net and of its equivalent conflict-free Petri net have the same dynamics, and as a consequence both nets have the same trajectories.*

Therefore, whenever it comes in handy, we shall use the conflict-free transformation in our statements and proofs.

2 Control aspects of monotonic discrete and fluid dynamics

We take in this section a more high-level view of the objects we are dealing with, whether they originate from discrete or continuous-relaxed Petri nets, and discuss how they fit in a control and operator approach setting. Since priority rules alter several structural properties of the dynamics, we only work in what follows with *monotonic* models, that is to say priority-free Petri nets. We shall come back to the study of models with priority in the Chapter 5.

- 2.1 ABSTRACT FORM OF THE SYSTEMS.** — For a (classical) Petri net with transitions set Q , let us write the counters that are not associated with inputs in vector form, i.e., $Z := (Z_q)_{q \in Q \setminus Q_{\text{input}}}$, and likewise, denote $Z_{\text{input}} := (Z_q)_{q \in Q_{\text{input}}}$ (in the context of continuous relaxed nets, we would write z and z_{input}). We have seen in Theorem 1.23 about discrete Petri nets dynamics that under the non-Zeno assumption (see Definition 1.18), the evolution of priority-free nets can be reduced to the master equation (M)

$$\left\{ \begin{array}{l} \forall t \geq 0 \quad Z(t) = F\left(\left(Z(t - \tau)\right)_{\tau \in \mathcal{T}}, \left(Z_{\text{input}}(t - \tau)\right)_{\tau \in \mathcal{T}_{\text{input}}}\right) \\ \forall t < 0 \quad Z(t) = \mathbf{0} \\ Z_{\text{input}} \text{ is prescribed externally over } \mathbb{R} \end{array} \right., \quad (\text{M})$$

where \mathcal{T} (resp. $\mathcal{T}_{\text{input}}$) is a finite subset of $\mathbb{R}_{>0}$ (resp. $\mathbb{R}_{\geq 0}$). The operator F encapsulates all the topological and conflict-solving aspects of the Petri net (incidence relations, arc multipliers, preselection mappings...) as long as initial markings, and apart from that, only features sum, minimum and integer part operations (in particular, it is nonlinear). Systems of the form (M) are **piecewise affine dynamical systems with delays**, and typically arise in discrete-event models. As it is emphasized on in the induction principle (I), the master equation (M) allows us to determine the present (i.e., $Z(t)$) using the past (the $(Z(t - \tau))_{\tau \in \mathcal{T}}$) and the inputs.

It can help understanding the above framework to compare it with the extensively studied and continuous counterpart setting of ordinary differential equations (ODEs). In this case, the master equation of the dynamics does not feature delays but derivatives of an unknown function x , and becomes $\dot{x}(t) = F(x(t), u(t))$, associated with a condition $x(t_0) = x_0$. In this context, the vector-valued function u plays the role of a control. By comparison, observe that Z_{input} can also be interpreted in (M) as a control chosen by an external agent, and the equation $Z|_{(-\infty, 0)} = \mathbf{0}$ is nothing but a (functional) initial condition. The operator F still encloses the dynamics governing the evolution of the state vector x .

In this light, and as already mentioned in Section 1.5 of Chapter 1, many questions that are customary in control theory arise. First, we may wonder whether the system is controllable, and

in particular if we are able to guarantee the follow-up of the inputs. Characterizing the trajectories of the system, determining if they reach an ultimate regime, and if yes at which speed are also questions of interest. All these aspects are hidden in the operator F .

Observe further that \mathcal{Q} is already partitioned into the subsets $\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$ and $\mathcal{Q}_{\text{input}}$, so that our system is naturally divided into an internal part and an external part. As usual in control, focusing first on the internal part (or in other words “disabling” the external inputs) is fruitful to understand the properties of F . In our setting, since all the coordinates of Z_{input} occur in (M) in minimum operations with components of Z , this “impulse response” (or “free-run regime”) can be obtained by letting $Z_q(t) = \infty$ for all q in $\mathcal{Q}_{\text{input}}$ and t in \mathbb{R} .

When the operator F is specifically built using the continuous-relaxed equations of Table 2.3 (and no more using discrete equations of Table 1.13), the integer parts disappear. This way, F is still nonlinear, but only involves sums and minimums. For the sum and minimum operations are both associative and because the sum distributes over the minimum, each component of F actually arises as a **minimum of finitely many affine functions** of the coordinates of $(Z(t - \tau))_{\tau \in \mathcal{T}}$ and $(Z_{\text{input}}(t - \tau))_{\tau \in \mathcal{T}_{\text{input}}}$ (with nonnegative “slope” coefficients). In other words, for all q in $\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$, there is an integer $\ell^{(q)} \geq 1$, there are real numbers $(a_i^{(q)})_{1 \leq i \leq \ell^{(q)}}$, nonnegative vectors $(b_i^{(q,\tau)})_{1 \leq i \leq \ell^{(q)}, \tau \in \mathcal{T}}$ and $(c_i^{(q,\tau)})_{1 \leq i \leq \ell^{(q)}, \tau \in \mathcal{T}_{\text{input}}}$ such that for all $t \geq 0$ we have

$$Z_q(t) = \min_{1 \leq i \leq \ell^{(q)}} \left(a_i^{(q)} + \sum_{\tau \in \mathcal{T}} \langle b_i^{(q,\tau)}, Z(t - \tau) \rangle + \sum_{\tau \in \mathcal{T}_{\text{input}}} \langle c_i^{(q,\tau)}, Z_{\text{input}}(t - \tau) \rangle \right),$$

where $\langle \cdot, \cdot \rangle$ denotes the usual dot-product. Note that in the Petri net models that we will study, the vectors $(b_i^{(q,\tau)})_{1 \leq i \leq \ell^{(q)}, \tau \in \mathcal{T}}$ and $(c_i^{(q,\tau)})_{1 \leq i \leq \ell^{(q)}, \tau \in \mathcal{T}_{\text{input}}}$ will typically be sparse – that is to say, with many zero coefficients.

Although the vector Z_{input} formally intervenes as a control, it will in our systems most often stand for the inbound demand, that is by essence not chosen but endured (see for instance Example 2.6). It is possible to transform the equation (M) to better account for this phenomenon, by simply blending in a new operator G the action of the operator F together with the effect of Z_{input} and its delays. Because Z_{input} is a function of the time, so is G , and the equation becomes

$$\begin{cases} \forall t \geq 0 & Z(t) = G\left((Z(t - \tau))_{\tau \in \mathcal{T}}, t\right) \\ \forall t < 0 & Z(t) = \mathbf{0} \end{cases}, \quad (\text{M}')$$

In the system (M'), the function Z arises as the solution of an autonomous but non-stationary dynamics, contrary to the previous non-autonomous but stationary form. The analysis of the trajectories of the system now requires to focus on the properties of the operator G . To that purpose, one can adopt an approach based on evolution operators, frequently introduced to deal with the solutions of dynamical systems.

2.2 THE OPERATOR APPROACH. – Recall from Theorem 1.23 and Theorem 2.5 that whether we work with discrete or continuous-relaxed models of (priority-free) timed Petri nets, the knowledge of inputs determines a unique trajectory of the system. More specifically, the induction principle (J) on which these two theorems are based shows that we can propagate a solution Z of the dynamics as soon as we know Z_{input} over \mathbb{R} and we know Z over $[-\tau_{\text{max}}, 0)$ – recall that τ_{max} is the greatest holding time of the net. So far, we have always fixed $Z|_{[-\tau_{\text{max}}, 0)} = \tilde{\mathbf{0}}$, since by definition transitions do not fire before the physical time $t = 0$. Observe from a more abstract perspective that any other choice of **functional initial condition** would lead to a fully determined

continuation as well. In particular, for $t_0 \geq 0$, we should be able to start from an initial condition drawn from the execution of a Petri net between the instants $-\tau_{\max} + t_0$ and t_0 , and recover the same trajectory after time t_0 .

To formalize this idea, we introduce the set $\mathcal{Z} := \mathcal{F}([-\tau_{\max}, 0), \mathbb{R}^{Q \setminus Q_{\text{input}}})$ of vector-valued functions over $[-\tau_{\max}, 0)$, and define the notion of evolution operator.

► **DEFINITION 2.8.** We call *evolution operator* the two-parameter family $(\mathcal{U}_s^t)_{0 \leq s \leq t}$ defined by

$$\mathcal{U}_s^t : \begin{cases} \mathcal{Z} & \longrightarrow \mathcal{Z} \\ Z^0 & \longmapsto \begin{cases} [-\tau_{\max}, 0) & \longrightarrow \mathbb{R}^{Q \setminus Q_{\text{input}}} \\ u & \longmapsto Z(u + t) \end{cases} \end{cases} \quad (2.7)$$

where Z is the function uniquely determined by the initial condition $Z(\cdot + s) = Z^0$ on $[-\tau_{\max}, 0)$, the input Z_{input} on $[-\tau_{\max}, \infty)$ and the equations of Table 1.13 or Table 2.3.

We shall explicitly write $\mathcal{U}^{\text{DISC}}$ and $\mathcal{U}^{\text{CONT}}$ to refer to the evolution operators respectively associated with dynamics of Table 1.13 and Table 2.3, but we use the general notation \mathcal{U} when our claims apply to both cases. For Z^0 in \mathcal{Z} , we denote by $\mathcal{U}_s^t[Z^0]$ the image of Z^0 by \mathcal{U}_s^t , and sometimes $\mathcal{U}_s^t[Z^0, Z_{\text{input}}]$ to recall the dependence on the input functions. We may refer to Z^0 as a *seed* function, and to \mathcal{Z} as the set of seeds functions, since they give rise to the trajectories.

The Definition 2.8 can be interpreted as follows. The action of \mathcal{U}_s^t is to “push” forward a function defined over $[-\tau_{\max} + s, s)$ by $t - s$ units of time according to the dynamical system and the inputs. From this perspective, we may think of \mathcal{U} as an operator of “initial condition transportation”, analogously to flow properties for ODEs.

If Z is the solution of the master equation (M) (or (M')), then, by choosing sufficiently small $\varepsilon > 0$ (actually any positive number in $(0, \tau_{\max})$ works), we have

$$\text{for all } t \geq 0, \quad Z(t) = \mathcal{U}_0^{t+\varepsilon}[\tilde{\mathbf{0}}](-\varepsilon), \quad (2.8)$$

since $Z(\cdot + 0)$ must be null on $[-\tau_{\max}, 0)$ (hence the choice $s = 0$ with respect to (2.7)). Observe that it is not possible to write $Z(t) = \mathcal{U}_0^t[\tilde{\mathbf{0}}](0)$ since elements of \mathcal{Z} cannot be evaluated in zero. This is somehow expected since the physical instant $t = 0$ is already part of the execution of the system, and thus we need to “push” the dynamics strictly more than t units of time from 0 to figure out the value of $Z(t)$.

The form (2.8) surely seems very burdensome to simply denote $Z(t)$, but the evolution operator \mathcal{U} is intended to act on much more general functions of \mathcal{Z} than $\tilde{\mathbf{0}}$. The following proposition formalizes the “bootstrapping” idea expressed in the beginning of the subsection.

► **PROPOSITION 2.9.** For nonnegative reals r, s and t such that $0 \leq r \leq s \leq t$, the evolution operator \mathcal{U} satisfies

$$\mathcal{U}_s^t \circ \mathcal{U}_r^s = \mathcal{U}_r^t.$$

Proof. Let $Z^0 \in \mathcal{Z}$. Let us denote by Z^1 the function defined over $[-\tau_{\max} + r, \infty)$ with values in $\mathbb{R}^{Q \setminus Q_{\text{input}}}$ such that $Z^1|_{[-\tau_{\max} + r, r)} = Z^0$ and Z^1 satisfies the discrete (or continuous) dynamics. We have for all $u \in [-\tau_{\max}, 0)$ that $\mathcal{U}_r^s[Z^0](u) = Z^1(s + u)$ and $\mathcal{U}_r^t[Z^0](u) = Z^1(t + u)$. Let us now denote by Z^2 the function defined over $[-\tau_{\max} + s, \infty)$ such that $Z^2|_{[-\tau_{\max} + s, s)} = \mathcal{U}_r^s[Z^0] = Z^1(s + \cdot)|_{[-\tau_{\max}, 0)}$, so that for all $u \in [-\tau_{\max}, 0)$, we have $\mathcal{U}_s^t[\mathcal{U}_r^s[Z^0]](u) = Z^2(t + u)$. Since Z^1 and Z^2 coincide over $[-\tau_{\max} + s, s)$, by uniqueness of the continuation under the dynamics, they coincide over $[-\tau_{\max} + s, \infty)$. In particular, because $t \geq s$, we have $Z^1(t + u) = Z^2(t + u)$ for all $u \in [-\tau_{\max}, 0)$. This amounts to saying that $\mathcal{U}_r^t[Z^0]$ and $\mathcal{U}_s^t[\mathcal{U}_r^s[Z^0]] = (\mathcal{U}_s^t \circ \mathcal{U}_r^s)[Z^0]$ are equal, as two functions of \mathcal{Z} . Since this holds for any Z^0 , the operators \mathcal{U}_r^t and $\mathcal{U}_s^t \circ \mathcal{U}_r^s$ are equal. \clubsuit

Example 2.10. We consider the same single-tier call-center than in the Figure 2.4 of Example 2.6, with continuous-relaxed dynamics (ST-DCR) and parameters $N = 1$, $\tau_0 = 10$ s, $\tau_1 = 120$ s and $\tau_2 = 5$ s. The arrivals of inbound calls in this case follow a smooth profile growing from 0 to 4 received calls between $t = 0$ s and $t = 400$ s.

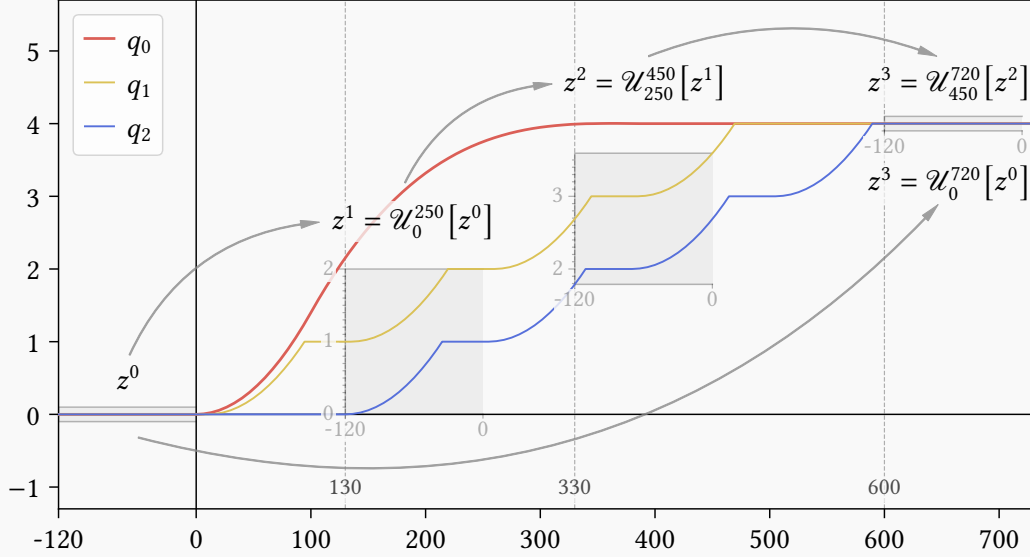


Figure 2.6: Counters functions evolution for the given input and fluid Petri net of Figure 2.4.

The initial condition is $z^0 = (z_1|_{[-120,0]}, z_2|_{[-120,0]}) = \tilde{\mathbf{0}}$. The solution of the dynamics would of course remain null if it was not for the increasing input z_{input} , that kicks the growth of z_1 (resp. z_2) at $t = \tau_0$ (resp. $t = \tau_0 + \tau_1$). Since there are not enough agents to pick up all the calls, we observe “zigzag” growth profiles for z_1 and z_2 , as in Example 2.6b.

The Figure 2.6 illustrates the composition property of the evolution operator \mathcal{U} . We denote by $z := (z_1, z_2)$ the vector function represented over $[-120, 730)$. Let the dynamics evolve for 250 s from $t = 0$ s and the initial null condition z^0 , and let us focus on the function $z^1 := z|_{[130,250)}(\cdot + 250) = \mathcal{U}_0^{250}[z^0]$. The “+250” translation guarantees that z^1 is also an element of \mathcal{Z} (which we represent by the grayed box in overlay). Therefore, we *could* compute $\mathcal{U}_0^t[z^1]$. However, this would amount to applying the dynamics from z^1 and use values of z_{input} over $[-120, t]$, which would not give back z . On the contrary, computing $\mathcal{U}_{250}^{450}[z^1]$ consists in propagating the dynamics from the initial condition $z^1(\cdot - 250) = z|_{[130,250)}$, starting from instant 130 and thus using values of z_{input} lying in $[130, 450)$. This gives back $z|_{[330,450)}$, from which we can define $z^2 := z|_{[330,450)}(\cdot + 450)$ so that z^2 is in \mathcal{Z} . Applying \mathcal{U}_{450}^{720} to z^2 , we end up with a function z^3 that is constant over $[-120, 0)$, since at $t = 600$ s the system has reached a final state with $z_1(t) = z_2(t) = 4$. The function z^3 is none other than $\mathcal{U}_0^{720}[z^0]$, since

$$\mathcal{U}_0^{720}[z^0] = \mathcal{U}_{450}^{720}[\mathcal{U}_{250}^{450}[\mathcal{U}_0^{250}[z^0]]] = \mathcal{U}_{450}^{720}[\mathcal{U}_{250}^{450}[z^1]] = \mathcal{U}_{450}^{720}[z^2] = z^3.$$

A special case occurs if there are no input transitions in (M), or if Z_{input} is constant. In this situation, the system is explicitly autonomous, with $Z(t) = F((Z(t - \tau))_{\tau \in \mathcal{T}})$, and the effect of the evolution operator \mathcal{U} is more conveniently represented by an **evolution semigroup**.

- **PROPOSITION 2.11.** *If there are no inputs or if the inputs counters are constant, then $\mathcal{U}_{s+r}^{t+r} = \mathcal{U}_s^t$ for all $r \geq 0$ and $0 \leq s \leq t$. Defining the operator $\mathcal{S}_t := \mathcal{U}_0^t$, we have that $(\mathcal{S}_t)_{t \geq 0}$ forms a one-parameter semigroup, i.e., $\mathcal{S}_0 = \text{id}_{\mathcal{X}}$ and for all $t, t' \geq 0$, we have $\mathcal{S}_{t+t'} = \mathcal{S}_t \circ \mathcal{S}_{t'}$.*

Proof. The invariance by translation property readily follows from the abstract form (M) and Definition 2.8. To prove the semigroup property, we first note that \mathcal{S}_0 is the identity operator over \mathcal{X} by definition, and that for $t, t' \geq 0$, we have $\mathcal{S}_t \circ \mathcal{S}_{t'} = \mathcal{U}_0^t \circ \mathcal{U}_0^{t'} = \mathcal{U}_{t'}^{t+t'} \circ \mathcal{U}_0^{t'} = \mathcal{U}_0^{t+t'} = \mathcal{S}_{t+t'}$ (the second equality has been obtained leveraging on the invariance by translation, and the third one using Proposition 2.9). \clubsuit

We now discuss a few structural and algebraical properties of the evolution operator \mathcal{U} . The set \mathcal{X} of seed functions is equipped with the usual partial ordering (\leq), so that for Z^1 and Z^2 in \mathcal{X} , the inequality $Z^1 \leq Z^2$ means that for all q in $\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$ and for all u in $[-\tau_{\text{max}}, 0)$, we have $Z_q^1(u) \leq Z_q^2(u)$ – and more generally all inequalities between vector-valued functions are understood this way. A fundamental property of the dynamics we study is that they preserve this partial ordering. Echoing the remark stated in Chapter 1 after the Figure 1.8, we point out that this would not hold with priority routings, since they break the monotonicity of the evolution.

- **PROPOSITION 2.12.** *The evolution operator \mathcal{U} is **order-preserving** (or **nondecreasing**) relatively to both the seed function and the input functions, i.e., for all Z^1 and Z^2 in \mathcal{X} , given an input $Z_{\text{input}} := (Z_q)_{q \in \mathcal{Q}_{\text{input}}}$, we have*

$$Z^1 \leq Z^2 \quad \Longrightarrow \quad \mathcal{U}_s^t[Z^1, Z_{\text{input}}] \leq \mathcal{U}_s^t[Z^2, Z_{\text{input}}],$$

and likewise, given a seed Z^0 in \mathcal{X} , we have

$$Z_{\text{input}}^1 \leq Z_{\text{input}}^2 \quad \Longrightarrow \quad \mathcal{U}_s^t[Z^0, Z_{\text{input}}^1] \leq \mathcal{U}_s^t[Z^0, Z_{\text{input}}^2].$$

Proof. Remark that up to translating the inputs, we can always reduce to the case where $s = 0$. We show the proposition building on the induction principle (I). For Z^1 and Z^2 in \mathcal{X} , let us still denote by an abuse of notation Z^1 and Z^2 the trajectories they uniquely determine over $[-\tau_{\text{max}}, \infty)$ under the same inputs. By definition, we have $Z^1(t) \leq Z^2(t)$ if $t \in [-\tau_{\text{max}}, 0)$, which is the initialization of (I). For the heredity, suppose that for some $k \in \mathbb{N}$, we have $Z^1(t) \leq Z^2(t)$ for all $t \in [-\tau_{\text{max}}, k\tau_{\text{min}})$. Let us take $t \in [k\tau_{\text{min}}, (k+1)\tau_{\text{min}})$, if $q \in \mathcal{Q}_{\text{sync}}$ and $p \in q^{\text{in}}$, it is clear that $\alpha_{qp}^{-1}(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} Z_{q'}^1(t - \tau_p)) \leq \alpha_{qp}^{-1}(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} Z_{q'}^2(t - \tau_p))$ since the multipliers are positive, and taking the integer part and the minimum preserve inequalities. If $q \in \mathcal{Q}_{\text{psel}}$, we use the same arguments and also the fact that for $p \in q^{\text{in}}$, the mapping Π_q^p is by definition nondecreasing. As a result, $Z^1(t) \leq Z^2(t)$ also holds on $[k\tau_{\text{min}}, (k+1)\tau_{\text{min}})$. For two inputs Z_{input}^1 and Z_{input}^2 , the heredity part is exactly the same and the initialization is trivial for the Z^1 and Z^2 have common seeds. \clubsuit

The multipliers and the proportions involved in the evolution of counters in continuous-relaxed nets raise the question of possible “explosion” (or decimation) of markings. To control the growth of the different counters, it is useful to introduce the notion of stoichiometric invariants.

- **DEFINITION 2.13.** *We say that a vector e of $\mathbb{R}^{\mathcal{Q}}$ is a **stoichiometric invariant** of a Petri net if*

$$\text{for all } q \text{ in } \mathcal{Q}_{\text{sync}}, \quad \text{for all } p \text{ in } q^{\text{in}}, \quad e_q = \alpha_{qp}^{-1} \sum_{q' \in p^{\text{in}}} \alpha_{pq'} e_{q'},$$

$$\text{for all } q \text{ in } \mathcal{Q}_{\text{psel}}, \quad \text{such that } q^{\text{in}} = \{p\}, \quad e_q = \pi_{qp} \alpha_{qp}^{-1} \sum_{q' \in p^{\text{in}}} \alpha_{pq'} e_{q'}.$$

The next proposition underlines another important structural property of the evolution operator in the continuous-relaxed setting, the sub-homogeneous character. In particular, it will ensure

that under the existence of a stoichiometric invariant, a single counter function cannot “blow up” because of a nudge affecting another coordinate of the trajectory.

- **PROPOSITION 2.14.** *If e is a positive stoichiometric invariant of the net, then the evolution operator $\mathcal{U}^{\text{CONT}}$ is **additively sub-homogeneous** with respect to e , that is to say, for all $x \geq 0$, for all $0 \leq s \leq t$ and for all z^0 in \mathcal{Z} :*

$$\mathcal{U}^{\text{CONT}}|_s^t[z^0 + x\tilde{e}] \leq \mathcal{U}^{\text{CONT}}|_s^t[z^0] + x\tilde{e} \quad \text{and} \quad \mathcal{U}^{\text{CONT}}|_s^t[z^0 - x\tilde{e}] \geq \mathcal{U}^{\text{CONT}}|_s^t[z^0] - x\tilde{e},$$

where \tilde{e} denotes the constant function of \mathcal{Z} such that $\tilde{e}(u) = e|_{Q \setminus Q_{\text{input}}}$ for all u in $[-\tau_{\text{max}}, 0)$.

Proof. We prove the result using induction principle (I), and like the proof of Proposition 2.12, we reduce to the case where $s = 0$. Let $z^0 \in \mathcal{Z}$ and $x \geq 0$, we denote by z (resp. z^x) the unique continuation of z^0 (resp. $z^0 + x\tilde{e}$) given some similar inputs and continuous-relaxed dynamics of Table 2.3. We know that $z^x(t) \leq z(t) + xe|_{Q \setminus Q_{\text{input}}}$ for all $t \in [-\tau_{\text{max}}, 0)$ (using the seed function, this inequality is the trivial relation $z^0(t) + xe|_{Q \setminus Q_{\text{input}}} \leq z^0(t) + xe|_{Q \setminus Q_{\text{input}}}$), and we want to show that this also holds for $t \geq 0$. Suppose that we have $k \in \mathbb{N}$ such that for all $t \in [-\tau_{\text{max}}, k\tau_{\text{min}})$, we indeed have $z^x(t) \leq z(t) + xe|_{Q \setminus Q_{\text{input}}}$. If $t \in [k\tau_{\text{min}}, (k+1)\tau_{\text{min}})$ and $q \in Q_{\text{sync}}$, we use the corresponding equation in Table 2.3, and we have

$$\begin{aligned} z_q^x(t) &= \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}^x(t - \tau_p) \right) \\ &\leq \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(m_p + \sum_{\substack{q' \in p^{\text{in}} \\ q' \notin Q_{\text{input}}}} \alpha_{pq'} (z_{q'}(t - \tau_p) + xe_{q'}) + \sum_{\substack{q' \in p^{\text{in}} \\ q' \in Q_{\text{input}}}} \alpha_{pq'} z_{q'}(t - \tau_p) \right) \\ &\leq \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} (z_{q'}(t - \tau_p) + xe_{q'}) \right) \\ &= \min_{p \in q^{\text{in}}} \left\{ \alpha_{qp}^{-1} \left(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}(t - \tau_p) \right) + x \alpha_{qp}^{-1} \sum_{q' \in p^{\text{in}}} \alpha_{pq'} e_{q'} \right\} \\ &= z_q(t) + xe_q \end{aligned}$$

where the first inequality has been obtained by using the induction hypothesis on internal transitions and the fact that the input profiles are the same on input transitions, and the second inequality is derived from Proposition 2.12, since increasing the inputs on the right-hand side preserves the inequality. The last line makes use of the defining property of stoichiometric invariant. The proof is similar for transitions in Q_{psel} . Finally, to show the property if $x \leq 0$, it suffices to apply our inequality to the seed $z^0 - x\tilde{e}$. \clubsuit

Remark that if there are no inputs, the same lines of proof show that the evolution operator $\mathcal{U}^{\text{CONT}}$ is additively homogeneous with respect to \tilde{e} , meaning that $\mathcal{U}^{\text{CONT}}|_s^t[z^0 + x\tilde{e}] = \mathcal{U}^{\text{CONT}}|_s^t[z^0] + x\tilde{e}$, for any z^0 in \mathcal{Z} , x in \mathbb{R} and $0 \leq s \leq t$. This is in general not true for $\mathcal{U}^{\text{DISC}}$ due to the additional integrality constraints (floor functions or preselection mappings).

Let us finish this section by elaborating further on the notion of stoichiometric invariants, that shall play a key role in the Section 3 to come and also in our Chapter 4. The two equations in Definition 2.13 express the conservation of the markings of the net if the transitions of the net are fired with weights $(e_q)_{q \in Q}$ (and taking into account the arc multipliers and routing proportions). In this sense, the notion of stoichiometric invariant slightly refines the one of T-invariant introduced in Definition 1.6, by imposing that in and out flows of tokens are balanced for all places but also respect the proportional preselection rules. The two definitions actually coincide in the setting of conflict-free nets, introduced in Section 1.3.

- **PROPOSITION 2.15.** *If e in \mathbb{R}^Q is a T-invariant of the pure and equivalent conflict-free net of some (priority-free) Petri net, then e is a stoichiometric invariant of this original Petri net.*

Proof. Let us denote by $(\mathcal{P}, \mathcal{Q}, \mathcal{E})$ the original priority-free Petri net, and by $(\mathcal{P}', \mathcal{Q}, \mathcal{E}')$ its conflict-free equivalent Petri net (by definition, they have common set of transitions). We denote by C and C' their weighted incidence matrices.

If e is a T-invariant of $(\mathcal{P}', \mathcal{Q}, \mathcal{E}')$, we have $C'e = 0$, or equivalently $C'^+e = C'^-e$. This vector identity provides $|\mathcal{P}'|$ scalar equalities accounting for the balance equations of all the places of the conflict-free net. The places in $\mathcal{P}' \setminus \mathcal{P}'_{\text{psel}}$ and their in/out incidences are unchanged with respect to the original net, and by definition, these places only have at most one downstream transition, so we can write

$$\forall p \in \mathcal{P} \setminus \mathcal{P}'_{\text{psel}}, \quad \sum_{q \in \mathcal{Q}} C_{pq}^+ e_q = \sum_{q \in \mathcal{Q}} C_{pq}^- e_q,$$

or equivalently using only nonzero multipliers,

$$\forall p \in \mathcal{P} \setminus \mathcal{P}'_{\text{psel}}, \quad \text{s.t. } p^{\text{out}} = \{q\}, \quad \sum_{q' \in p^{\text{in}}} \alpha_{pq'} e_{q'} = \alpha_{qp} e_q.$$

If $p \in \mathcal{P}'_{\text{psel}}$, coefficients of C and C' will differ on the corresponding rows, but if $p^{\text{out}} = \{q_1, \dots, q_k\}$, we have reusing the notation of Section 1.3 that $C_{pq_i}^- = \pi_{p,q_i} C_{p q_i q_i}^-$ and $C_{pq_i}^+ = C_{p q_i}^+$ for all $1 \leq i \leq k$. Therefore we obtain the identities

$$\forall p \in \mathcal{P}'_{\text{psel}}, \quad \forall q \in p^{\text{out}}, \quad \pi_{qp} \sum_{q' \in p^{\text{in}}} \alpha_{pq'} e_{q'} = \alpha_{qp} e_q.$$

We now observe that for all $q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$, there is some p in \mathcal{P} such that $q \in p^{\text{out}}$, hence the conditions for e to be a stoichiometric invariant are all satisfied. \clubsuit

On the model of Proposition 2.14, we shall pay a particular attention to stoichiometric invariants that are positive; or in other words whose support (set of nonzero components) is the whole set \mathcal{Q} . It was shown in both the discrete and the continuous contexts (see [CCS89] and [RTS99]) that live and bounded conflict-free Petri nets admit positive T-invariants that are unique up to a multiplicative constant. Conversely, Gaujal and Giua proved that in strongly connected and conflict-free (continuous-relaxed) Petri nets, there is at most one positive T-invariant up to scaling [GG04b, Proposition 12]. We study below a weaker topological condition to derive a similar result, also more suitable to take input transitions into account. To that purpose, we use a notion of accessibility between transitions that is the converse of the one featured in the proof of Theorem 1.23, i.e., we work in the directed graph with nodes set \mathcal{Q} and with directed arcs set $\{(q_1, q_2) \in \mathcal{Q}^2 \mid q_1^{\text{out}} \cap q_2^{\text{in}} \neq \emptyset\}$.

- **PROPOSITION 2.16.** *Suppose that a continuous-relaxed Petri net is conflict-free, and that there is a subset \mathcal{Q}_0 of \mathcal{Q} such that for all q in \mathcal{Q} , there is q_0 in \mathcal{Q}_0 that has access to q . Then, for all positive vector e_0 of $\mathbb{R}_{>0}^{\mathcal{Q}_0}$ there is at most one positive T-invariant e in $\mathbb{R}_{>0}^{\mathcal{Q}}$ such that $e|_{\mathcal{Q}_0} = e_0$*

Proof. Suppose that we have two positive T-invariants e and e' such that $e|_{\mathcal{Q}_0} = e'|_{\mathcal{Q}_0}$, but assume for contradiction that $e \neq e'$. Let us define $\beta := \min\{e_q/e'_q \mid q \in \mathcal{Q}\}$. Up to swapping the role of e and e' if $e' \leq e$, we have $\beta < 1$. By definition, the vector $\tilde{e} := e - \beta e'$ is still a T-invariant since $C\tilde{e} = Ce - \beta Ce' = 0 - 0$, and it is nonnegative. However it has a null component.

Because $\beta < 1$, a null component q of \tilde{e} must lie in $\mathcal{Q} \setminus \mathcal{Q}_0$. On one hand, q is accessible from some q_0 in \mathcal{Q}_0 , so q^{in} is nonempty and the places of q^{in} also admit upstream transitions. On the other hand, using the definition of the conflict-free net, we have denoting by p an upstream place of q that $\tilde{e}_q = \alpha_{qp}^{-1} \sum_{q' \in p^{\text{in}}} \alpha_{pq'} \tilde{e}_{q'}$. Because \tilde{e} is nonnegative and $\tilde{e}_q = 0$, we deduce that $\tilde{e}_{q'} = 0$ for all $q' \in p^{\text{in}}$. We can reapply this scheme (a finite number of times) until getting that $\tilde{e}_{q_0} = 0$ for some $q_0 \in \mathcal{Q}_0$, hence $e_{q_0} = \beta e'_{q_0} < e'_{q_0}$, and the contradiction. \clubsuit

We insist on the fact that the condition of Proposition 2.16 is about uniqueness and not existence of a positive T-invariant of the net, in particular the system $Ce = \mathbf{0}$ may be overdetermined and not admit any solution – even though it is possible to form a triangular subsystem. We retrieve the fact that strongly connected Petri nets have at most one positive T-invariant (up to a multiplicative constant) since in this case any singleton $\{q\}$ with q in Q fulfills the requirements of our claim. This is also true for single-input Petri nets where the input transition has access to all the “internal” transitions. More generally, we can use the Proposition 2.16 to conclude that there are at most $|Q_{\text{input}}|$ degrees of freedom for positive T-invariants of nets where any transition is accessible from an input.

3 Fluid Petri nets as large-scale limits of discrete models

Our main motivation for using continuous-relaxed dynamics in what precedes is to simplify the original discrete dynamics governing the evolution of timed Petri nets. By doing so, we necessarily introduce errors between a solution Z of the discrete equations and a solution z of their fluid counterparts. It is then quite a natural matter to know how good z is an approximation of Z .

Leveraging on the infinitely many divisible nature of fluid-tokens and their interpretation as a huge number of particles, we can draw the intuition that as more and more tokens circulate in a discrete Petri net (obtained by increasing initial markings and input transitions firings), its behaviour is better and better approximated by a “similar” continuous-relaxed Petri net with same markings. Indeed at large scales, integer-part effects or tokens preselection should play a reduced relative effect. This idea was mentioned in [CGQ95, Section V], although with no systematic treatment.

In this section, our purpose is to formalize the study of this phenomenon, and in particular to detail the notion of “similarity” between a classical Petri net and a continuous-relaxed Petri net. To take explicitly into account an increase of the initial resources and of the tokens produced by inputs, we shall make use of a positive integer **scaling parameter**, hereafter denoted by N . Such scaling techniques were used before in queueing networks to study heavy traffic limits and “state space collapse” properties – see [Bra98, Wil98, BDo1].

3.1 THE BOUNDED DEVIATIONS MODEL. – We first focus on a setting close from the deterministic Petri net models developed so far. To the purpose of studying the impact of the scaling parameter N which controls increases of initial markings and input transitions firings of a Petri net, we choose to introduce for both discrete and continuous-relaxed frameworks a sequence (or family) of **scaled Petri nets**, whose N -th term shall be referred to as the *scaled-by- N* or simply *scaled* Petri net. We impose in addition that these discrete and fluid families of Petri nets are such that

- All the scaled discrete and fluid Petri nets share the same topology, meaning that they have same places set \mathcal{P} , transitions set Q and weighted incidence relations \mathcal{E} , regardless of N . In particular, the arc multipliers are integers, to fit in the framework of discrete Petri nets.
- The scaled-by- N instances of discrete and fluid Petri nets have identical initial marking Nm , where m is a vector of $\mathbb{N}^{\mathcal{P}}$ independent of N . For a place p ruled by preselection, initial quantities of tokens dedicated to the firing of q in p^{out} also coincide for scaled discrete and fluid nets, so that the real quantity $\pi_{qp}^0 m_p$ of (2.6) is set to the integer value $m_{q,p}$ of (P4) by fixing the rational π_{qp}^0 .

- Both families of nets have identical and constant place-by-place holding times, that are independent of the scaling factor. We assume that these times ensure non-Zenoness of the Petri nets in the sense of Definition 1.18.
- If q is in $\mathcal{Q}_{\text{input}}$, the firing profile of q in the scaled-by- N fluid Petri net is $z_q^N(t) := Nz_q(t)$, where z_q is a reference nondecreasing, piecewise-continuous and càdlàg function of $\mathbb{R}_{\geq 0}$ to itself (hence being the firing profile of q for $N = 1$). We assume that the corresponding firing profile Z_q^N in the discrete scaled-by- N model has **bounded deviation** with respect to z_q^N , i.e., there is $\eta_q > 0$ independent of N such that for all $t \geq 0$, we have $|Z_q^N(t) - z_q^N(t)| \leq \eta_q$.
- All the conflicts in the fluid Petri nets are arbitrated by proportional preselection routings (see Section 1.1) with proportions independent of N . The conflicts in the discrete models are arbitrated by preselection mappings also independent of N , and we assume that these have **bounded deviations** with respect to the corresponding fluid proportions, meaning that for all p in $\mathcal{P}_{\text{psel}}$ and q in p^{out} , there is $\eta_{p,q} > 0$ (independent of N) such that $|\Pi_q^p(n) - \pi_{qp}n| \leq \eta_{p,q}$ for all n in \mathbb{N}^* . In particular we have $\lim_{n \rightarrow \infty} \Pi_q^p(n)/n = \pi_{qp}$.

For N in \mathbb{N}^* and given input profiles $(Z_q^N)_{q \in \mathcal{Q}_{\text{input}}}$, we shall denote by $Z^N := (Z_q^N)_{q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}}$ the trajectory of the transition counter functions of the scaled-by- N discrete Petri net, i.e., that obeys the discrete dynamical evolution equations and such that Z^N is null over $\mathbb{R}_{< 0}$. Similarly, we denote $z^N := (z_q^N)_{q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}}$ the trajectory of the scaled-by- N continuous-relaxed net. If $N = 1$, we write Z and z instead of Z^1 and z^1 .

Using the equations of Table 1.13 for the discrete nets and the equations of Table 2.3 for the fluid ones, we can write for all N in \mathbb{N}^* the dynamical equations verified by Z^N and z^N , gathered in the Table 2.7 opposite. Due to our strong hypotheses on the deviations of inputs and preselection mappings, we call these equations the **BOUNDED DEVIATIONS** setting. Our aim is now to *evaluate the difference* $Z^N - z^N$. Primarily, we would like to know in which extent our “low-level” bounded deviations assumptions reflect from a “high-level” and trajectory-focused perspective.

Notice that it is a direct consequence of Theorems 1.23 and 2.5 that under the knowledge of input profiles and some (functional) initial condition, the equations of Table 2.7 determine unique trajectories of the scaled-by- N nets for any N in \mathbb{N}^* . Therefore, on the model of Definition 2.8, these equations induce two sequences of evolution operators $(\mathcal{U}^{\text{DISC}, N})_{N \in \mathbb{N}^*}$ and $(\mathcal{U}^{\text{CONT}, N})_{N \in \mathbb{N}^*}$ enjoying in particular the order-preserving character, and for the continuous evolution operators, the additively sub-homogeneous character (see Propositions 2.12 and 2.14).

In addition, due to the fact that N operates a multiplicative scaling on input profiles and initial markings for the fluid scaled models, the evolution operator $\mathcal{U}^{\text{CONT}, N}$ also satisfies a **scaling-homogeneity** property, highlighted by the following result.

► **PROPOSITION 2.17.** *For all z^0 in \mathcal{X} , N in \mathbb{N}^* and $0 \leq s \leq t$, we have*

$$\mathcal{U}^{\text{CONT}, N} \Big|_s^t [Nz^0] = N \mathcal{U}^{\text{CONT}} \Big|_s^t [z^0].$$

Proof. Just like in the proof of Proposition 2.12, we reduce to the case $s = 0$ and apply induction principle (I). Denoting by z^N (resp. z) the trajectories uniquely determined by identical inputs, continuous-relaxed equations for scaled-by- N (resp. scaled-by-1, or unscaled) system in Table 2.7, and initial condition Nz^0 (resp. z^0) on $[-\tau_{\text{max}}, 0)$, it is clear that z^N and Nz coincide on $[-\tau_{\text{max}}, 0)$. Now suppose that they coincide on $[-\tau_{\text{max}}, t)$ for some $t \geq 0$, using equations of Table 2.7, the fact that all holding times are greater than τ_{min} and induction hypothesis, it is clear that the equations for z^N can always be factored by N , hence $z^N = Nz$ over $[t, t + \tau_{\text{min}})$. ♠

Another way to state this property, looking only at trajectories of the Petri nets initialized with $z^0 = \tilde{\mathbf{0}}$, is to write $(z^N)/N = z$. It is interesting to see this relationship as an invariance of the fluid

Type	Counter equation in the discrete model (BOUNDED DEVIATIONS setting)	
$q \in \mathcal{Q}_{\text{input}}$	$Z_q^N(t)$ is prescribed externally	with $ Z_q^N(t) - z_q^N(t) \leq \eta_q$
$q \in \mathcal{Q}_{\text{sync}}$	$Z_q^N(t) = \min_{p \in q^{\text{in}}} \left[\alpha_{qp}^{-1} \left(Nm_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} Z_{q'}^N(t - \tau_p) \right) \right]$	
$q \in \mathcal{Q}_{\text{psel}}$	$Z_q^N(t) = \left[\alpha_{qp}^{-1} \left(Nm_{p,q} + \Pi_q^p \left(\sum_{q' \in p^{\text{in}}} \alpha_{pq'} Z_{q'}^N(t - \tau_p) \right) \right) \right]$	with $ \Pi_q^p(n) - \pi_{qp}n \leq \eta_{pq}$
Type	Counter equation in the continuous model (BOUNDED DEVIATIONS setting)	
$q \in \mathcal{Q}_{\text{input}}$	$z_q^N(t) = Nz_q(t)$	with $z_q(t)$ prescribed externally
$q \in \mathcal{Q}_{\text{sync}}$	$z_q^N(t) = \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(Nm_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}^N(t - \tau_p) \right)$	
$q \in \mathcal{Q}_{\text{psel}}$	$z_q^N(t) = \alpha_{qp}^{-1} \left(Nm_{p,q} + \pi_{qp} \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}^N(t - \tau_p) \right)$	

Table 2.7: Dynamic equations followed by counter functions of scaled-by- N discrete and fluid Petri nets

dynamics under conjugation, in the sense that a premultiplication by N and a postmultiplication by $1/N$ do not affect the unscaled dynamics. Likewise, we may also think of $(Z^N)/N$ as a conjugate of Z by the scaling operation, although the relation $Z^N = NZ$ does not hold here due to the leeway of discrete inputs and preselection mappings – in a tubular neighborhood of their fluid counterparts – and also due to the side effects of integer-parts. In what follows, we shall then not focus on the deviation $Z^N - z^N$ but rather on the difference $(Z^N)/N - z$, which anyway appears as the right renormalized object to study (rescaling by $1/N$ after accelerating the arrivals by a factor of N).

We now state a technical lemma that plants the first seed for quantifying the deviation of the continuous-relaxed dynamics compared to the discrete one.

- **LEMMA 2.18.** *There is a positive constant η such that for all Z^0 in \mathcal{X} , all N in \mathbb{N} , all $s \geq 0$ and t in $[s, s + \tau_{\min})$, we have*

$$\mathcal{U}^{\text{DISC}, N} \Big|_s^t [Z^0] - \eta \tilde{\mathbf{1}} \leq \mathcal{U}^{\text{CONT}, N} \Big|_s^t [Z^0] \leq \mathcal{U}^{\text{DISC}, N} \Big|_s^t [Z^0] + \eta \tilde{\mathbf{1}} .$$

Proof. Let $s \geq 0$ and $t \in [s, s + \tau_{\min})$. We denote by Z^N (resp. z^N) the function defined over $[s - \tau_{\max}, \infty)$ and with values in $\mathbb{R}^{\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}}$ such that $Z^N(s + u) = Z^0(u)$ (resp. $z^N(s + u) = Z^0(u)$) for all $u \in [-\tau_{\max}, 0)$ and satisfying the scaled-by- N discrete (resp. continuous) dynamics of Table 2.7 with corresponding inputs. By definition, Z^N and z^N coincide on $[s - \tau_{\max}, s)$.

- ◇ We first unfold the proof when all the holding times of the net are positive, in particular, for all $p \in \mathcal{P}$, $\tau_p \geq \tau_{\min}$. We explain at the end of the proof how to adapt the reasoning in the presence of zero holding times.

Let us take $q \in \mathcal{Q}_{\text{sync}}$. Remark that if $p \in q^{\text{in}}$, because $t \in [s, s + \tau_{\min})$, we have $t - \tau_p \in [s - \tau_{\max}, s)$, so we can replace in the relations of Table 2.7 the delayed terms of the form $(Z_{q'}^N(t - \tau_p))_{q' \in p^{\text{in}}}$ by explicit

values of Z^0 or values of Z_{input}^N and z_{input}^N , which leads to

$$Z_q^N(t) = \min_{p \in q^{\text{in}}} \left[\alpha_{qp}^{-1} \left(Nm_p + \sum_{\substack{q' \in p^{\text{in}} \\ q' \notin Q_{\text{input}}}} \alpha_{pq'} Z_{q'}^0(t - \tau_p) + \sum_{q' \in Q_{\text{input}}} \alpha_{pq'} Z_{q'}^N(t - \tau_p) \right) \right],$$

$$z_q^N(t) = \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(Nm_p + \sum_{\substack{q' \in p^{\text{in}} \\ q' \notin Q_{\text{input}}}} \alpha_{pq'} Z_{q'}^0(t - \tau_p) + \sum_{q' \in Q_{\text{input}}} \alpha_{pq'} Nz_{q'}(t - \tau_p) \right).$$

Using on one hand the classical inequalities $x - 1 < \lfloor x \rfloor \leq x$ that are valid for all $x \in \mathbb{R}$, and on the other hand the inequalities $Nz_{q'}(u) - \eta_{q'} \leq Z_{q'}^N(u) \leq Nz_{q'}(u) + \eta_{q'}$ valid for all $q' \in Q_{\text{input}}$ and $u \in \mathbb{R}$, we obtain that

$$z_q^N(t) - M_q - 1 \leq Z_q^N(t) \leq z_q^N(t) + M_q,$$

where $M_q := \max_{p \in q^{\text{in}}} \left[\alpha_{qp}^{-1} \sum_{q' \in p^{\text{in}} \cap Q_{\text{input}}} \alpha_{pq'} \eta_{q'} \right]$.

◇ If we now take $q \in Q_{\text{psel}}$, with $p \in \mathcal{P}$ such that $q^{\text{in}} = \{p\}$, we have by similar arguments that

$$Z_q^N(t) = \min_{p \in q^{\text{in}}} \left[\alpha_{qp}^{-1} \left(Nm_{p,q} + \Pi_q^p \left(\sum_{\substack{q' \in p^{\text{in}} \\ q' \notin Q_{\text{input}}}} \alpha_{pq'} Z_{q'}^0(t - \tau_p) + \sum_{q' \in Q_{\text{input}}} \alpha_{pq'} Z_{q'}^N(t - \tau_p) \right) \right) \right],$$

$$z_q^N(t) = \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(Nm_{p,q} + \pi_{qp} \sum_{\substack{q' \in p^{\text{in}} \\ q' \notin Q_{\text{input}}}} \alpha_{pq'} Z_{q'}^0(t - \tau_p) + \pi_{qp} \sum_{q' \in Q_{\text{input}}} \alpha_{pq'} Nz_{q'}(t - \tau_p) \right),$$

and using in addition of previous inequalities the fact that $\pi_{qp}n - \eta_{pq} \leq \Pi_q^p(n) \leq \pi_{qp}n + \eta_{pq}$, we get

$$z_q^N(t) - M_q - \eta_{pq} - 1 \leq Z_q^N(t) \leq z_q^N(t) + M_q + \eta_{pq}, \quad (2.9)$$

with an identical definition of M_q (here reduced to one term in the minimum).

◇ To conclude the proof in the case where all the holding times are positive, we remark that letting

$$\eta := \max_{q \in Q \setminus Q_{\text{input}}} \{M_q\} + \max_{q \in Q_{\text{psel}}, q^{\text{in}} = \{p\}} \{\eta_{pq}\} + 1$$

gives the desired inequality $z^N(t) - \eta \mathbf{1} \leq Z^N(t) \leq z^N(t) + \eta \mathbf{1}$.

If some holding times are null (but the Petri net stays non-Zeno), we need to resort to substitutions schemes such as the one of Theorem 1.23 before replacing terms of the form $Z_{q'}(t - \tau_p)$ by values of Z^0 . In both models (discrete and continuous), this would lead to nested equations, possibly featuring nested preselection mappings, and in the discrete models, we would encounter nested integer parts. However, we still obtain the result claimed in the lemma by using the two deviation properties and the inequality on the integer part; the terms M_q would typically become more cumbersome (product of multipliers along maximal non-Zeno walks) and the term 1 in (2.9) could be replaced by a greater integer corresponding to the maximum number of nested floor functions. Nevertheless, the term η that we obtain still only depends upon the topology of the net and the preselection proportions. \mathcal{E}

The Lemma 2.18 is a local result, and we now want to extend it globally. Once more, we could formally use the induction principle (\mathcal{I}) to achieve this goal, but instead, we present with Theorem 2.19 a proof using the evolution operator machinery to emphasize its convenience – both techniques share the same spirit anyhow. Propagating the result of Lemma 2.18 beyond time-windows of length τ_{min} also requires to control the evolution of the error between two functions of \mathcal{X} under the application of the dynamic equations, hence the need for a stoichiometric invariant (see Definition 2.13) of the nets (that all share the same topology, hence the same invariants). Remark that this theorem answers questions raised in the end of [CGQ95] on the comparison between the evolution of discrete and continuous-relaxed nets.

- **THEOREM 2.19.** *Suppose that the Petri nets admit a positive stoichiometric invariant e , and denote by \tilde{e} the constant function of \mathcal{X} equal to e . Then, there are positive constants κ_0 and κ_1 such that for all Z^0 in \mathcal{X} , all N in \mathbb{N} and $0 \leq s \leq t$, we have*

$$-(\kappa_1(t-s) + \kappa_0)\tilde{e} \leq \mathcal{U}^{\text{CONT},N} \Big|_s^t [Z^0] - \mathcal{U}^{\text{DISC},N} \Big|_s^t [Z^0] \leq (\kappa_1(t-s) + \kappa_0)\tilde{e}.$$

In particular, denoting by Z^N the unique solution of the **scaled discrete** dynamics (by a factor N) and by z the unique solution of the **unscaled continuous** dynamics of the Petri net, we have

$$\text{for all } t \geq 0, \quad \frac{Z^N(t)}{N} \underset{N \rightarrow \infty}{=} z(t) + O\left(\frac{t}{N}\right). \quad (2.10)$$

Proof. Take $Z^0 \in \mathcal{X}$, $N \in \mathbb{N}$ and $0 \leq s \leq t$. Let us also choose a real number β in $(0, 1)$, intended to be very close to 1. We define the integer $L := \lceil (t-s)/\beta\tau_{\min} \rceil$ (the notation $\lceil \cdot \rceil$ stands for the upper integer part, or ceiling function), and for all integer ℓ such that $0 \leq \ell < L$, we denote $t_\ell := s + \ell\beta\tau_{\min}$, and finally $t_L := t$, so that $0 < t_{\ell+1} - t_\ell < \beta\tau_{\min}$ for all $0 \leq \ell < L$. Because e is a stoichiometric invariant, all its coordinates are positive (see Definition 2.13), we then introduce the positive number $\kappa_0 := \eta / \min_{q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}} \{e_q\}$ where η is the constant featured in Lemma 2.18.

- ◊ We are going to show by finite induction that $\mathcal{U}^{\text{CONT},N} \Big|_s^{t_\ell} [Z^0] - \mathcal{U}^{\text{DISC},N} \Big|_s^{t_\ell} [Z^0] \leq \ell\kappa_0\tilde{e}$ for all $1 \leq \ell \leq L$. The initialization for $\ell = 1$ is a slightly weaker version of the Lemma 2.18, since for all $q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$, we have $\eta \leq \kappa_0 e_q$ by definition of κ_0 . Suppose now that this inequality is satisfied for some $\ell < L$. We have:

$$\begin{aligned} \mathcal{U}^{\text{CONT},N} \Big|_s^{t_{\ell+1}} [Z^0] &= \mathcal{U}^{\text{CONT},N} \Big|_{t_\ell}^{t_{\ell+1}} \left[\mathcal{U}^{\text{CONT},N} \Big|_s^{t_\ell} [Z^0] \right] \\ &\leq \mathcal{U}^{\text{CONT},N} \Big|_{t_\ell}^{t_{\ell+1}} \left[\mathcal{U}^{\text{DISC},N} \Big|_s^{t_\ell} [Z^0] + \ell\kappa_0\tilde{e} \right] \\ &\leq \mathcal{U}^{\text{CONT},N} \Big|_{t_\ell}^{t_{\ell+1}} \left[\mathcal{U}^{\text{DISC},N} \Big|_s^{t_\ell} [Z^0] \right] + \ell\kappa_0\tilde{e} \\ &\leq \mathcal{U}^{\text{DISC},N} \Big|_{t_\ell}^{t_{\ell+1}} \left[\mathcal{U}^{\text{DISC},N} \Big|_s^{t_\ell} [Z^0] \right] + \ell\kappa_0\tilde{e} + \eta\tilde{1} \\ &\leq \mathcal{U}^{\text{DISC},N} \Big|_s^{t_{\ell+1}} [Z^0] + (\ell+1)\kappa_0\tilde{e}, \end{aligned}$$

where the successive (in)equalities respectively come from the composition property of the evolution operator (Proposition 2.9); the induction hypothesis combined with the nondecreasing character of the evolution operator (Proposition 2.12); the additively sub-homogeneous character of the evolution operator with respect to e (Proposition 2.14); the application of Lemma 2.18 to the function $\mathcal{U}^{\text{DISC},N} \Big|_s^{t_\ell} [Z^0]$ of \mathcal{X} with $s = t_\ell$ and $t = t_{\ell+1}$ (and we indeed have $t \in [s, s + \tau_{\min})$); and finally the composition property of the evolution operator again in addition to the inequality $\eta\tilde{1} \leq \kappa_0\tilde{e}$.

In particular our inductive property holds for $\ell = L$, and we obtain the exact same converse inequality by below using the other side of Lemma 2.18. Reusing the defining formula of L leads to

$$-\kappa_0 \left[\frac{t-s}{\beta\tau_{\min}} \right] \tilde{e} \leq \mathcal{U}^{\text{CONT},N} \Big|_s^t [Z^0] - \mathcal{U}^{\text{DISC},N} \Big|_s^t [Z^0] \leq \kappa_0 \left[\frac{t-s}{\beta\tau_{\min}} \right] \tilde{e}.$$

It only suffices to use the bound $\lceil x \rceil \leq x+1$ to obtain the result claimed in the theorem, with $\kappa_1 := \kappa_0/(\beta\tau_{\min})$. Both coefficients κ_0 and κ_1 depend only on the topology of the net. Observe that the previous inequalities are verified for all $0 < \beta < 1$, so by taking the limit when $\beta \rightarrow 1$, we could as well use $\kappa_1 := \kappa_0/\tau_{\min}$.

- ◊ For the last part of the theorem, we apply the shown inequality for $Z^0 = \tilde{\mathbf{0}}$, between instants $s = 0$ and $t + \varepsilon$ (for a small $\varepsilon > 0$). We also use the multiplicative homogeneity of the continuous evolution operator (Proposition 2.17) to obtain that $\mathcal{U}^{\text{CONT},N} \Big|_0^{t+\varepsilon} [\tilde{\mathbf{0}}] = N\mathcal{U}^{\text{CONT}} \Big|_0^{t+\varepsilon} [\tilde{\mathbf{0}}]$. Evaluating in $-\varepsilon$, we obtain

$$-(\kappa_1(t+\varepsilon) + \kappa_0)\tilde{e}(-\varepsilon) \leq N\mathcal{U}^{\text{CONT}} \Big|_0^{t+\varepsilon} [\tilde{\mathbf{0}}](-\varepsilon) - \mathcal{U}^{\text{DISC},N} \Big|_0^{t+\varepsilon} [\tilde{\mathbf{0}}](-\varepsilon) \leq (\kappa_1(t+\varepsilon) + \kappa_0)\tilde{e}(-\varepsilon).$$

Using the equation (2.8) and the two functions considered in the theorem, this is more conveniently rewritten as

$$-(\kappa_1(t+\varepsilon) + \kappa_0)e \leq Nz(t) - Z^N(t) \leq (\kappa_1(t+\varepsilon) + \kappa_0)e.$$

To finish, we divide all terms by N and observe that $Z^N(t)/N - z(t)$ is a bounded function of t/N . ♠

The double inequality of Theorem 2.19 is our most precise result in the sense that it is applicable to any starting point s , however, the equation (2.10) provides the answer to our question. Under bounded deviations assumptions, not only does the “rescaled” trajectory Z^N/N of the scaled-by- N discrete net converges towards the trajectory z of the unscaled fluid net, but we also know with which precision. In particular, even though it is not surprising to see the error between $Z^N(t)$ and $z^N(t)$ grow with t , asserting that this drift is at most linear shall be helpful to us when focusing on the throughput of Petri nets in the Chapter 4.

3.2 THE THREEFOLD STOCHASTIC MODEL. – In the same spirit as the BOUNDED DEVIATIONS setting, we now want to study how applying a scaling factor affects the execution of discrete Petri nets with stochastic behaviours. Recall that in the Section 3.3 of Chapter 1, we have explained that multiple sources of nondeterminism could be considered in discrete Petri nets: stochastic arrivals for inputs (typically Poisson), stochastic holding times for places, stochastic allocation for preselection routings (that is to say Bernoulli routings). All these effects tremendously increase the modeling power of Petri nets. However, even though they can be taken into account in counter functions, the resulting equations are hard to work with analytically.

Interestingly enough, nothing prevents in the previous section to consider that the counters $(Z_q^N)_{q \in Q, N \in \mathbb{N}^*}$ are subject to stochastic behaviours for inputs or preselections, as long as the bounded deviation assumptions hold (which is nevertheless not true for Poisson processes or Bernoulli laws). Still, Theorem 2.19 teaches us that $(Z^N)/N$ approaches the deterministic limit z as N tends to ∞ , suggesting that the continuous-relaxed approximation for Petri nets has a regularizing power.

To investigate this hypothesis, let us place ourselves in a fairly general setting with *all* the aforementioned stochastic effects, that we shall refer to as the THREEFOLD STOCHASTIC setting. Similarly to our approach in Section 3.1, we consider two families of scaled-by- N discrete and continuous Petri nets, such that

- As in the BOUNDED DEVIATIONS framework, all the considered Petri nets have identical topologies regardless of N . The arc multipliers are integer, the integer initial markings (and if applicable their allocation in preselection routings) coincide in the scaled-by-1 discrete and fluid instances, and get multiplied by N in scaled-by- N instances.
- For all q in Q_{input} , the firing profile \widehat{Z}_q^N of the input transition q in the scaled-by- N discrete Petri net is a realization of a nonhomogeneous Poisson process with intensity $N\lambda_q$, where λ_q is a measurable reference function from $\mathbb{R}_{\geq 0}$ to itself. We assume that the counterpart firing profile of q in the scaled-by- N fluid Petri net is the deterministic nondecreasing function $\Lambda_q: t \mapsto \int_0^t \lambda_q(s) ds$.
- We assume that all the conflicts in the discrete Petri nets are solved by Bernoulli routings (identical for all N), given for all places p in $\mathcal{P}_{\text{conflict}} = \mathcal{P}_{\text{psel}}$ by probability vectors $(\pi_{qp})_{q \in p^{\text{out}}}$ that are also independent of the scaling factor. These routings induce random preselection routing maps $(\widehat{\Pi}_q^p)_{p \in \mathcal{P}_{\text{psel}}, q \in p^{\text{out}}}$ as in Chapter 1. In the continuous-relaxed models, conflicts are solved by proportional preselection routings with same probabilities than the discrete models.
- Somehow most importantly, we assume that in the scaled-by- N discrete Petri nets, the holding times of the places are stochastic, given by realizations of random variables $(\widehat{\tau}_p^N)_{p \in \mathcal{P}, N \in \mathbb{N}^*}$ that all draw i.i.d. values from a same distribution measure μ_p (in particular for two different values of N , the generated holding times may differ but still have same law).

This last aspect prompts us to determine a continuous-relaxed equivalent of the stochastic holding times. We have seen with the equation (P2') in Chapter 1 that in this setting, the relation between \widehat{Y}_p and \widehat{X}_p (in an unscaled discrete Petri net) is much more complex than the mere $Y_p(t) = X_p(t - \tau_p)$ given by (P2), and we need to resolve to the $*$ operation defined in equation (1.6) to write $\widehat{Y}_p = \widehat{X}_p * \widehat{\tau}_p$. Remark that for nonhomogeneous Poisson processes and Bernoulli routings, the deterministic continuous-relaxed counterparts that we have proposed are given by expectations of these laws, so it is also natural to study the expected value of \widehat{Y}_p . When \widehat{X}_p is known and given by the nondecreasing function X_p from $\mathbb{R}_{\geq 0}$ to \mathbb{N} with jump instants $(T_k)_{k \in \mathbb{N}}$, it turns out that this computation is easy. It gives in particular exploiting (P2') and the fact that the infinite sums are actually finite for all $t \geq 0$:

$$\begin{aligned} \mathbb{E}(\widehat{Y}_p(t) - m_p \mid \widehat{X}_p = X_p) &= \sum_{k=1}^{\infty} \mathbb{E}(\mathbb{1}_{[T_k + \widehat{\tau}_k^{(p)}, \infty)}(t)) = \sum_{k=1}^{\infty} \mathbb{P}(t \geq T_k + \widehat{\tau}_k^{(p)}) = \sum_{k=1}^{\infty} \int_0^{t-T_k} d\mu_p(\tau) \\ &= \sum_{k=1}^{\infty} \int_0^{\infty} \mathbb{1}_{[0, t-T_k]}(\tau) d\mu_p(\tau) = \sum_{k=1}^{\infty} \int_0^{\infty} \mathbb{1}_{[T_k + \tau, \infty)}(t) d\mu_p(\tau) \\ &= \int_0^{\infty} \left[\sum_{k=1}^{\infty} \mathbb{1}_{[T_k + \tau, \infty)}(t) \right] d\mu_p(\tau) = \int_0^{\infty} (X_p(t - \tau) - m_p) d\mu_p(\tau). \end{aligned}$$

This results then suggests, at least as a heuristic, to assume that stochastic holding times for discrete models of Petri nets degenerate in **convolution operations** in continuous-relaxed Petri nets. This is consistent with the idea that distributed times bring an averaging or smoothing effect (observe in addition that the relationship $Y_p(t) = X_p(t - \tau_p)$ is retrieved by the above form if μ_p is a Dirac measure concentrated on the value τ_p). From a semantics point of view in continuous-relaxed Petri nets, this behaviour would mean that each time a fluid-token is produced in the place p , it is divided in infinitely many small fluid tokens, each one being associated with a different holding time according to the distribution μ_p . Remark that the above convolution form explains why in Example 1.25 of Chapter 1, we could see a smooth junction between two affine behaviour for counters X_1 and Z_2 .

Building on these ideas, we are finally able to propose two complete recursive equation sets for both discrete and fluid scaled and compatible models of Petri nets, that we summarize in the Table 2.8, specialized to an event ω .

For the sake of simplicity, we shall restrict in what follows our analysis to cases were the measures $(\mu_p)_{p \in \mathcal{P}}$ are **nonsingular**. In accordance with Lebesgue's decomposition theorem for measures on the real line (see for instance [Hal74, p. 182]), this means that for all p in \mathcal{P} , μ_p is the sum of an absolutely continuous measure μ_p^{ac} and a pure point measure μ_p^{pp} . In other words, there exists a measurable function f_{μ_p} from $\mathbb{R}_{\geq 0}$ to itself (that is the density of μ_p^{ac} with respect to Lebesgue's measure) and a countable set $\partial\mu_p \subset \mathbb{R}_{\geq 0}$ (the charged points of μ_p^{pp}) such that for all Borel set B of the real line we have

$$\mu_p(B) = \mu_p^{\text{ac}}(B) + \mu_p^{\text{pp}}(B) = \int_{s \in B} f_{\mu_p}(s) ds + \sum_{s \in \partial\mu_p} \mu_p^{\text{pp}}(\{s\}) \delta_s(B), \quad \text{with } \mu_p(\mathbb{R}_{\geq 0}) = 1. \quad (2.11)$$

Another simplification that we make is that we will assume that for all p in \mathcal{P} , the support of μ_p is included in $[\tau_{\min}, \infty)$, or equivalently $\mu_p([0, \tau_{\min})) = 0$. This can be seen as a non-Zeno condition much stronger than the one proposed in Definition 1.18. We shall later discuss how this hypothesis could be relaxed and made minimal.

Type	Counter equation in the discrete model (THREEFOLD STOCHASTIC setting)
$q \in \mathcal{Q}_{\text{input}}$	$Z_q^N(t; \omega)$ is a (cumulative) nonhomogeneous Poisson process with intensity $N\lambda_q$
$q \in \mathcal{Q}_{\text{sync}}$	$Z_q^N(t; \omega) = \min_{p \in q^{\text{in}}} \left[\alpha_{qp}^{-1} Y_p^N(t; \omega) \right]$
$q \in \mathcal{Q}_{\text{psel}}$	$Z_q^N(t; \omega) = \left[\alpha_{qp}^{-1} \left(Nm_{p,q} + \Pi_q^p(Y_p^N(t; \omega) - Nm_p; \omega) \right) \right]$
$p \in \mathcal{P}$	$X_p^N(t; \omega) = Nm_p + \sum_{q \in p^{\text{in}}} \alpha_{pq} Z_q^N(t; \omega)$ $Y_p^N(t; \omega) = (X_p^N * \tau_p^N)(t; \omega)$
Type	Counter equation in the continuous model (THREEFOLD STOCHASTIC setting)
$q \in \mathcal{Q}_{\text{input}}$	$z_q^N(t) = \int_0^t N\lambda_q(s) ds$
$q \in \mathcal{Q}_{\text{sync}}$	$z_q^N(t) = \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(Nm_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} \int_0^\infty z_{q'}^N(t - \tau) d\mu_p(\tau) \right)$
$q \in \mathcal{Q}_{\text{psel}}$	$z_q^N(t) = \alpha_{qp}^{-1} \left(Nm_{p,q} + \pi_{qp} \sum_{q' \in p^{\text{in}}} \alpha_{pq'} \int_0^\infty z_{q'}^N(t - \tau) d\mu_p(\tau) \right)$

Table 2.8: Dynamic equations followed by transitions counter functions

It also does not sound useless to detail a bit the sample space Ω involved here for the discrete scaled Petri nets in Table 2.8. For the unscaled instance ($N = 1$), we draw one sequence of positive interarrival times per input nonhomogeneous Poisson process to determine \widehat{Z}_q (q in $\mathcal{Q}_{\text{input}}$), we draw one sequence of holding times greater than τ_{\min} per place to determine the $(\widehat{\tau}_k^{(p)})_{p \in \mathcal{P}, k \in \mathbb{N}^*}$, and we draw $|p^{\text{out}}|$ sequences of $\{0, 1\}$ to determine the Bernoulli routing of each p in $\mathcal{P}_{\text{psel}}$, via the $(\widehat{R}_k^{(q)})_{q \in p^{\text{out}}, k \in \mathbb{N}^*}$ or equivalently the $(\widehat{\Pi}_q^p)_{q \in p^{\text{out}}}$ (see Section 3.3 of Chapter 1). As a result we have

$$\Omega_1 := ((\mathbb{R}_{\geq 0})^{\mathbb{N}^*})^{|\mathcal{Q}_{\text{input}}|} \times (([\tau_{\min}, \infty))^{\mathbb{N}^*})^{|\mathcal{P}|} \times (\{0, 1\}^{\mathbb{N}^*})^{|\mathcal{Q}_{\text{psel}}|},$$

and because we can draw the random variables of all scaled models at the same time, we shall have $\Omega := \Omega_1^{\mathbb{N}}$. All the random variables associated with a discrete Petri net model with a scaling factor N are assumed to be independent with respect to all the ones corresponding to other scaled models. For a given N in \mathbb{N}^* , all the $(\widehat{Z}_q^N)_{q \in \mathcal{Q}_{\text{input}}}$, $(\widehat{\tau}_{N,k}^{(p)})_{p \in \mathcal{P}, k \in \mathbb{N}^*}$, and $(\widehat{R}_{N,k}^{(q)})_{q \in \mathcal{Q}_{\text{psel}}^0, k \in \mathbb{N}^*}$ are assumed to be mutually independent, so long as the subset $\mathcal{Q}_{\text{psel}}^0$ of $\mathcal{Q}_{\text{psel}}$ does not feature two transitions with a common upstream place. Our sample space Ω is formally made into a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ by the Ionescu-Tulcea's theorem (see for instance [Nev65, Proposition 5.1.1]), using the Poisson

rates $(\lambda_q)_{q \in Q_{\text{input}}}$, the holding times distributions $(\mu_p)_{p \in \mathcal{P}}$ and the Bernoulli probability systems $(\pi_{qp})_{p \in \mathcal{P}_{\text{sel}}, q \in \mathcal{P}^{\text{out}}}$.

In spite of the increased complexity of the equations of Table 2.8 compared to the classical discrete and fluid models (systems Table 1.13 and Table 2.3), our assumption that the holding times are deterministically greater than τ_{\min} easily allows us to state the following well-posedness result on the trajectories $\widehat{Z}^N := (\widehat{Z}_q^N)_{q \in Q \setminus Q_{\text{input}}}$ and $z^N := (z_q^N)_{q \in Q \setminus Q_{\text{input}}}$ of discrete and fluid scaled-by- N Petri nets. As in the previous section, we write \widehat{Z} and z instead of \widehat{Z}^1 and z^1 when $N = 1$.

- **PROPOSITION 2.20.** *Under the complete knowledge of an event ω in Ω , the discrete scaled equations of Table 2.8 determine unique trajectories of all discrete scaled-by- N Petri nets with null initial conditions (that is to say, $Z^N(\cdot, \omega)|_{(-\infty, 0)} = \tilde{\mathbf{0}}$).*

Similarly, the fluid scaled equations of Table 2.8 determine unique trajectories of all fluid scaled-by- N Petri nets with null initial conditions (such that $z^N|_{(-\infty, 0)} = \tilde{\mathbf{0}}$). In addition, these fluid solutions verify the same *scaling-homogeneity* property of Proposition 2.17, and as a result for all N in \mathbb{N}^* , we have $z^N = Nz$.

..... *Proof.* The proof relies once again on the induction principle (I). Indeed, the knowledge of ω and $X_p(\cdot; \omega)$ on $[-\tau_{\max}, k\tau_{\min})$ allows one to compute $Y_p(\cdot; \omega)$ on $[-\tau_{\max}, (k+1)\tau_{\min})$, and then the $Z_q(\cdot; \omega)$ as well up to $(k+1)\tau_{\min}$. For the continuous-relaxed setting, we use the fact that $\mu_p([0, \tau_{\min})) = 0$ for all $p \in \mathcal{P}$, as a result the integrals over $[0, \infty)$ are actually taken over $[\tau_{\min}, \infty)$, which lets us apply (I). The scaling-homogeneity is transparent and works exactly as in the proof of Proposition 2.17. \clubsuit

We now roll up our sleeves and lay the grounds needed to prove an asymptotic result on the model of Theorem 2.19. We know that it is possible that some events of Ω lead to trajectories of discrete Petri nets that deviate a lot from an expected (or average) behaviour. However, these possible cases should be *negligible*, and the best convergence type we can hope for should be that rescaled dynamics of scaled-by- N discrete Petri nets converge *almost surely* towards a standard unscaled dynamics of the fluid model as N grows. Recall that a property is verified almost surely (abbreviated a.s.) when it is verified on a subset of Ω with measure 1. The fundamental tool to deal with such problems is the strong law of large numbers, that will indeed be one of the main ingredients of our proofs.

We state a first lemma that addresses the case of input profiles, whose behaviour is autonomous.

- **LEMMA 2.21.** *For all transitions q in Q_{input} , we have*

$$\text{for all } t \geq 0, \quad \frac{\widehat{Z}_q^N(t)}{N} \xrightarrow[N \rightarrow \infty]{\text{a.s.}} z_q(t) .$$

..... *Proof.* Let q in Q_{input} , by assumption \widehat{Z}_q^N is the cumulative arrival function of a nonhomogeneous Poisson process with (functional) rate $N\lambda_q$. As in the Section 3.3 of Chapter 1, we introduce the primitive Λ_q of λ_q such that

$$\Lambda_q(t) := \int_0^t \lambda_q(s) \, ds .$$

- ◆ We first show that for $t \geq 0$, $\widehat{Z}_q^N(t)$ has the same law than $\sum_{i=1}^N \widehat{G}_i(t)$, where the $(\widehat{G}_i)_{1 \leq i \leq N}$ are random variables given by the cumulative arrival functions of independent nonhomogeneous Poisson processes with rate λ_q . Recall that on one hand, we have by (1.11) that

$$\forall k \in \mathbb{N}, \quad \mathbb{P}(\widehat{Z}_q^N(t) = k) = \frac{(N\Lambda_q(t))^k}{k!} e^{-N\Lambda_q(t)}$$

On the other hand, we have for $k \in \mathbb{N}$ that

$$\begin{aligned} \mathbb{P}\left(\sum_{i=1}^N \widehat{G}_i(t) = k\right) &= \sum_{i_1+\dots+i_N=k} \mathbb{P}\left(\widehat{G}_1(t) = i_1 \text{ and } \dots \text{ and } \widehat{G}_N(t) = i_N\right) \\ &= \sum_{i_1+\dots+i_N=k} \prod_{j=1}^N \left(\frac{\Lambda_q(t)^{i_j}}{i_j!} e^{-\Lambda_q(t)}\right) \\ &= \frac{\Lambda_q(t)^k}{k!} e^{-N\Lambda_q(t)} \sum_{i_1+\dots+i_N=k} \frac{k!}{i_1! \dots i_N!} \\ &= \frac{\Lambda_q(t)^k}{k!} e^{-N\Lambda_q(t)} N^k, \end{aligned}$$

where we have successively used the decomposition of k as the sum of N nonnegative integers and mutual exclusion of all the possibilities; mutual independence of the $(\widehat{G}_i)_{1 \leq i \leq N}$ and the nonhomogeneous Poisson law; factorization by common terms; and finally Newton's multinomial law for $N^k = (1 + 1 + \dots + 1)^k$.

◇ We now know that $(1/N)\widehat{Z}_q^N(t)$ has the same law than $(1/N)\sum_{k=1}^N \widehat{G}_i(t)$, however, as a mean of N independent random variables with same law, the latter converges almost surely as $N \rightarrow \infty$ towards $\mathbb{E}(\widehat{G}_1(t))$, that is to say $\Lambda_q(t)$. This is also the (deterministic) value of $z_q(t)$, hence the result. \clubsuit

The objective of the next technical lemma is to delineate how the discrete “random time-shuffling” operation $*$ defined in (1.6) morphs into convolution integrals in the fluid context. The proof features in particular approximation of Lebesgue's integrals by piecewise constant functions, and dater-to-counter transformations, in addition to two instances of the strong law of large numbers.

► **LEMMA 2.22.** *Let $(\widehat{X}_N)_{N \in \mathbb{N}}$ be a sequence of nondecreasing random piecewise constant functions from \mathbb{R} to \mathbb{N} , such that for any N in \mathbb{N} we have $\widehat{X}_N|_{(-\infty, 0)} = \mathbf{\tilde{0}}$ (deterministically). Let $T > 0$ and x be a nondecreasing piecewise continuous function over \mathbb{R} such that for all $t < T$, we have $\lim_{N \rightarrow \infty} \widehat{X}_N(t)/N = x(t)$ almost surely.*

*In addition, we suppose that we have nonnegative random variables $(\widehat{\tau}_N)_{N \in \mathbb{N}}$ all with nonsingular probability measure μ whose support is contained in $[\tau_{\min}, \infty)$, and we denote $\widehat{Y}_N := \widehat{X}_N * \widehat{\tau}_N$ for all N in \mathbb{N} . Then we have that*

$$\text{for all } t < T + \tau_{\min}, \quad \frac{\widehat{Y}_N(t)}{N} \xrightarrow[N \rightarrow \infty]{\text{a.s.}} \int_0^\infty x(t-s) d\mu(s).$$

Proof. Let $t \in [0, T + \tau_{\min})$ (the proof is of course immediate if $t < 0$). First observe that the function x as well is null over $\mathbb{R}_{<0}$, and $\mu([0, \tau_{\min})) = 0$, hence we can write

$$\int_0^\infty x(t-s) d\mu(s) = \int_{s \in [\tau_{\min}, t]} x(t-s) d\mu(s).$$

Let $\varepsilon > 0$. Our goal is to show that there is a subset Ω_ε of Ω and an integer N_0 in \mathbb{N} such that

$$\mathbb{P}(\Omega_\varepsilon) = 1 \quad \text{and} \quad \forall \omega \in \Omega_\varepsilon, \quad \forall N \geq N_0, \quad \left| \frac{\widehat{Y}_N(t; \omega)}{N} - \int_{s \in [\tau_{\min}, t]} x(t-s) d\mu(s) \right| \leq \varepsilon$$

◇ Let us denote by ∂x the (finite) set of discontinuities of the càdlàg and piecewise-continuous function x over $[\tau_{\min}, t]$, i.e., $\partial x := \{s \in [\tau_{\min}, t] \mid x(s) > x(s^-)\}$, so that the function x^c defined by

$$x^c : s \mapsto x(s) - \sum_{r \in \partial x} (x(r) - x(r^-)) \mathbb{1}_{[r, \infty)}(s)$$

is continuous over $[\tau_{\min}, t]$. Using the Heine-Cantor theorem, we know that there is $\delta > 0$ such that for all $(s, s') \in [\tau_{\min}, t]^2$, $|s - s'| \leq \delta$ implies $|x^c(s) - x^c(s')| \leq \varepsilon/2$.

By assumption, the measure μ admits the form (2.11), so that denoting by $\partial\mu := \{s \in [\tau_{\min}, \infty) \mid \mu(\{s\}) > 0\}$ the countable set of charged points of μ , we have a continuous and nonnegative function f_μ , such that for any Borel set B of \mathbb{R} , we have $\mu(B) = \int_{s \in B} f_\mu(s) ds + \sum_{s \in \partial\mu} \mu(\{s\}) \delta_s(B)$.

Let $\tau_{\min} = s_0 < s_1 < \dots < s_{L-1} < s_L = t$ be a partition of $[\tau_{\min}, t]$ with step δ adapted to x and μ , i.e., for all $\ell \in \{0, 1, \dots, L-1\}$, verifying $0 < s_{\ell+1} - s_\ell \leq \delta$ and also $(\partial x \cup \partial\mu) \cap [\tau_{\min}, t] \subset S_\delta$, where $S_\delta := \{s_\ell \mid 0 \leq \ell \leq L\}$. For all integer ℓ such that $0 \leq \ell < L-1$, we denote $I_\ell := [s_\ell, s_{\ell+1})$, and $I_{L-1} := [s_{L-1}, s_L]$, so that $\bigsqcup_{\ell=0}^{L-1} I_\ell = [\tau_{\min}, t]$, i.e., the disjoint union of the $(I_\ell)_{0 \leq \ell < L}$ indeed gives $[\tau_{\min}, t]$.

We denote by μ^δ the discrete measure $\mu^\delta := \sum_{\ell=0}^{L-1} \mu(I_\ell) \delta_{s_\ell}$. In addition, when ν is a positive measure, we denote by $\langle \nu, x \rangle$ the scalar $\int_{[\tau_{\min}, t]} x(t-s) d\nu(s)$. This way, we have

$$\langle \mu, x \rangle = \int_{s \in [\tau_{\min}, t]} x(t-s) d\mu(s) \quad \text{and} \quad \langle \mu^\delta, x \rangle = \sum_{\ell=0}^{L-1} x(t-s_\ell) \mu(I_\ell).$$

Observe that using the decomposition properties of μ , $\langle \mu, x \rangle$ and $\langle \mu^\delta, x \rangle$ can respectively be rewritten as

$$\begin{aligned} \langle \mu, x \rangle &= \int_{s \in [\tau_{\min}, t]} x(t-s) f_\mu(s) ds + \sum_{s \in \partial\mu \cap [\tau_{\min}, t]} x(t-s) \mu(\{s\}) \\ \langle \mu^\delta, x \rangle &= \sum_{\ell=0}^{L-1} x(t-s_\ell) \left(\int_{s \in I_\ell} f_\mu(s) ds + \sum_{s \in \partial\mu \cap I_\ell} \mu(\{s\}) \right), \end{aligned}$$

and therefore, building on integral's additions and linearity properties, and because $\partial\mu \cap [\tau_{\min}, t] \subset S_\delta$, we have

$$\langle \mu, x \rangle - \langle \mu^\delta, x \rangle = \sum_{\ell=0}^{L-1} \int_{s \in I_\ell} (x(t-s) - x(t-s_\ell)) f_\mu(s) ds,$$

and this last term is also given using the continuous function x^c by

$$\sum_{\ell=0}^{L-1} \int_{s \in I_\ell} \left(x^c(t-s) - x^c(t-s_\ell) + \sum_{r \in \partial x} (x(r) - x(r^-)) (\mathbb{1}_{(-\infty, t-r]}(s) - \mathbb{1}_{(-\infty, t-r]}(s_\ell)) \right) f_\mu(s) ds. \quad (2.12)$$

However, since $\partial x \subset S_\delta$, for all $s \in [s_\ell, s_{\ell+1})$, the term $\mathbb{1}_{(-\infty, t-r]}(s) - \mathbb{1}_{(-\infty, t-r]}(s_\ell)$ is null (indeed we have $\partial x \cap [s_\ell, s_{\ell+1}) \subset \{s_\ell\}$). Only a possible discontinuity of x in $s_L = t$ could achieve a nonzero value of this term, but because f_μ is measurable and Lebesgue's measure is diffuse, it does not change the value of the integral. We can now use uniform continuity of x^c and because intervals $([s_\ell, s_{\ell+1}))_{0 \leq \ell < L-1}$ have length inferior to δ , we obtain

$$|\langle \mu, x \rangle - \langle \mu^\delta, x \rangle| \leq \sum_{\ell=0}^{L-1} \int_{s \in [s_\ell, s_{\ell+1})} |x^c(t-s) - x^c(t-s_\ell)| f_\mu(s) ds \leq \frac{\varepsilon}{2} \int_{\tau_{\min}}^t f_\mu(s) ds \leq \frac{\varepsilon}{2}.$$

Notice that we could also have defined $\mu^\delta := \sum_{\ell=0}^{L-1} \mu(I_\ell) \delta_{s_{\ell+1}}$, in which case all the terms " s_ℓ " in (2.12) would have been replaced by " $s_{\ell+1}$ ". The integrand in this modified equation still coincides almost everywhere with $x^c(t-s) - x^c(t-s_{\ell+1})$, so the bound $|\langle \mu, x \rangle - \langle \mu^\delta, x \rangle| \leq \varepsilon/2$ survives. Combining this last result with the first one, we can finally state that

$$-\frac{\varepsilon}{2} + \int_0^\infty x(t-s) d\mu(s) \leq \sum_{\ell=0}^{L-1} x(t-s_{\ell+1}) \mu(I_\ell) \quad \text{and} \quad \sum_{\ell=0}^{L-1} x(t-s_\ell) \mu(I_\ell) \leq \frac{\varepsilon}{2} + \int_0^\infty x(t-s) d\mu(s). \quad (2.13)$$

◇ Let $\omega \in \Omega$ (here $\Omega = \mathbb{R}_{\geq 0}^{N \times N^*} \times [\tau_{\min}, \infty)^{N \times N^*}$). For all $N \in \mathbb{N}$, we denote by $(T_{N,k}(\omega))_{k \in \mathbb{N}}$ the jump instants of $X_N(\cdot; \omega)$, possibly repeated, and indeed countable because $X_N(\cdot, \omega)$ is piecewise constant. We also denote by $(\tau_{N,k}(\omega))_{k \in \mathbb{N}}$ the realizations of $\widehat{\tau}_N$ for the event ω , so that we have

$$X_N(t; \omega) = \sum_{k=1}^{\infty} \mathbb{1}_{[T_{N,k}(\omega), \infty)}(t) \quad \text{and} \quad Y_N(t; \omega) = \sum_{k=1}^{\infty} \mathbb{1}_{[T_{N,k}(\omega) + \tau_{N,k}(\omega), \infty)}(t) \quad (2.14)$$

Remark that we have the following equality and inequality:

$$\begin{aligned} Y_N(t; \omega) &= \sum_{\ell=0}^{L-1} \sum_{k=1}^{\infty} \mathbb{1}_{[T_{N,k}(\omega) + \tau_{N,k}(\omega, \infty)}(t) \times \mathbb{1}_{I_\ell}(\tau_{N,k}(\omega)) \\ &\leq \sum_{\ell=0}^{L-1} \sum_{k=1}^{X_N(t-s_\ell; \omega)} \mathbb{1}_{I_\ell}(\tau_{N,k}(\omega)). \end{aligned} \quad (2.15)$$

The equality is a rewriting of the relation on $Y_N(t; \omega)$ in (2.14) using our partition of $[\tau_{\min}, t]$ and the fact that holding times $(\tau_{N,k}(\omega))_{k \in \mathbb{N}}$ greater than t do not contribute to $Y_N(t; \omega)$. The inequality (2.15) is obtained by remarking, for instance using the characterization of $X_N(t; \omega)$ given in (2.14), that

$$\begin{aligned} T_{N,k}(\omega) + \tau_{N,k}(\omega) \leq t \quad \text{and} \quad \tau_{N,k}(\omega) \geq s_\ell &\implies T_{N,k}(\omega) + s_\ell \leq t \\ &\implies k \leq X_N(t - s_\ell; \omega). \end{aligned}$$

and as a consequence $\mathbb{1}_{[T_{N,k}(\omega) + \tau_{N,k}(\omega, \infty)}(t) \times \mathbb{1}_{I_\ell}(\tau_{N,k}(\omega)) \leq \mathbb{1}_{[1, X_N(t-s_\ell; \omega)]}(k) \times \mathbb{1}_{I_\ell}(\tau_{N,k}(\omega))$. Likewise, it holds that

$$\begin{aligned} k \leq X_N(t - s_{\ell+1}; \omega) \quad \text{and} \quad \tau_{N,k}(\omega) \leq s_{\ell+1} &\implies T_{N,k}(\omega) \leq t - s_{\ell+1} \quad \text{and} \quad \tau_{N,k}(\omega) \leq s_{\ell+1} \\ &\implies T_{N,k}(\omega) + \tau_{N,k}(\omega) \leq t. \end{aligned}$$

As a result, we can derive a lower bound counterpart of inequality (2.15), providing after dividing all the terms by N the double inequality

$$\frac{1}{N} \sum_{\ell=0}^{L-1} \xi_{\ell+1}(t; N; \omega) \leq \frac{Y_N(t; \omega)}{N} \leq \frac{1}{N} \sum_{\ell=0}^{L-1} \xi_\ell(t; N; \omega), \quad (2.16)$$

where $\xi_\ell(t; N; \omega) := \sum_{k=1}^{X_N(t-s_\ell; \omega)} \mathbb{1}_{I_\ell}(\tau_{N,k}(\omega))$.

Recall that $\widehat{X}_N(s)/N$ converges almost surely towards x if $s < T$, and for all $\ell \in \mathbb{N}$ such that $0 \leq \ell \leq L$, we indeed have $t - s_\ell < T$ since $t < T + \tau_{\min}$ and $s_\ell \geq \tau_{\min}$. For such an ℓ , let us denote by Ω_ℓ^X a subset of Ω such that $\mathbb{P}(\Omega_\ell^X) = 1$ and for all $\omega \in \Omega_\ell^X$, $\lim_{N \rightarrow \infty} X_N(t - s_\ell; \omega)/N = x(t - s_\ell)$. If $x(t - s_\ell) > 0$, then we also have that $\lim_{N \rightarrow \infty} X_N(t - s_\ell; \omega) = \infty$. In particular, we have for N large enough that

$$\forall \omega \in \Omega_\ell^X, \quad \frac{\xi_\ell(t; N; \omega)}{N} = \frac{X_N(t - s_\ell; \omega)}{N} \times \frac{1}{X_N(t - s_\ell; \omega)} \sum_{k=1}^{X_N(t - s_\ell; \omega)} \mathbb{1}_{I_\ell}(\tau_{N,k}(\omega)), \quad (2.17)$$

in which we recognize an instance of the strong law of large numbers. Indeed, for all $M \in \mathbb{N}^*$, the quantity $w_{M,\ell}(\omega) := (1/M) \sum_{k=1}^M \mathbb{1}_{I_\ell}(\tau_{N,k}(\omega))$ gives the proportion of realizations of the $(\tau_{N,k})_{1 \leq k \leq M}$ that fall in the interval I_ℓ . Since these are i.i.d., this average converges almost surely towards $\mathbb{E}(\widehat{\tau} \in I_\ell) = \mu(I_\ell)$, so we have a subset $\Omega_{\ell,\varepsilon}^\tau$ of Ω with measure 1 such that for all $\omega \in \Omega_{\ell,\varepsilon}^\tau$, $\lim_{M \rightarrow \infty} w_{M,\ell}(\omega) = \mu(I_\ell)$. Remark that the second factor in (2.17) simply forms a subsequence of $(w_{M,\ell}(\omega))_{M \in \mathbb{N}^*}$, as a result we have

$$\forall \omega \in \Omega_{\ell,\varepsilon}^X \cap \Omega_{\ell,\varepsilon}^\tau, \quad \lim_{N \rightarrow \infty} \frac{\xi_\ell(t; N; \omega)}{N} = x(t - s_\ell) \mu(I_\ell).$$

We then denote $\Omega_{\ell,\varepsilon} := \Omega_{\ell,\varepsilon}^X \cap \Omega_{\ell,\varepsilon}^\tau$ and further, we denote $\Omega_\varepsilon := \bigcap_{\ell=0}^L \Omega_{\ell,\varepsilon}$. The previous arguments ensure that $\mathbb{P}(\Omega_\varepsilon) = 1$, and we obtain

$$\forall \omega \in \Omega_\varepsilon, \quad \sum_{\ell=0}^{L-1} x(t - s_{\ell+1}) \mu(I_\ell) \leq \liminf_{N \rightarrow \infty} \frac{Y_N(t; \omega)}{N} \leq \limsup_{N \rightarrow \infty} \frac{Y_N(t; \omega)}{N} \leq \sum_{\ell=0}^{L-1} x(t - s_\ell) \mu(I_\ell); \quad (2.18)$$

(indeed, we have not shown yet that $Y_N(t; \omega)/N$ has a limit). Observe that if for some ℓ in $[0, L] \cap \mathbb{N}$, we have $x(t - s_\ell) = 0$, the property $\lim_{N \rightarrow \infty} \xi_\ell(t; N; \omega)/N = x(t - s_\ell) \mu(I_\ell)$ still holds for all $\omega \in \Omega_{\ell,\varepsilon}^X$, thanks to the convergence of $X_N(t - s_\ell; \omega)/N$ towards $x(t - s_\ell) = 0$ and the boundedness of the sequence $(w_{M,\ell}(\omega))_{M \in \mathbb{N}^*}$.

◇ Using (2.18), we have the existence of some N_0 in \mathbb{N} such that

$$\forall \omega \in \Omega_\varepsilon, \quad \forall N \geq N_0, \quad \frac{Y_N(t; \omega)}{N} \geq \sum_{\ell=0}^{L-1} x(t - s_{\ell+1}) \mu(I_\ell) - \frac{\varepsilon}{2} \quad \text{and} \quad \frac{Y_N(t; \omega)}{N} \leq \sum_{\ell=0}^{L-1} x(t - s_\ell) \mu(I_\ell) + \frac{\varepsilon}{2}$$

We can thus finally reuse (2.13), to obtain

$$\forall \omega \in \Omega_\varepsilon, \quad \forall N \geq N_0, \quad -\varepsilon + \int_0^\infty x(t-s) d\mu(s) \leq \frac{Y_N(t; \omega)}{N} \leq \varepsilon + \int_0^\infty x(t-s) d\mu(s),$$

which achieves the proof, since $\mathbb{P}(\Omega_\varepsilon) = 1$. ♠

We are now ready to state the main result of this section, that it would not be unreasonable to call “Strong law of large numbers for Petri nets”. Alternatively, this theorem makes continuous-relaxed Petri nets arise as “mean-field approximations” of discrete stochastic Petri nets.

► **THEOREM 2.23.** *The rescaled trajectory of a scaled-by- N discrete Petri net in the THREEFOLD STOCHASTIC setting converges pointwise and almost surely towards the trajectory of the equivalent continuous-relaxed and fluid Petri net, i.e.,*

$$\text{for all } t \geq 0, \quad \frac{\widehat{Z}^N(t)}{N} \xrightarrow[N \rightarrow \infty]{\text{a.s.}} z(t).$$

Proof. We denote by $(\widehat{Z}^N)_{N \in \mathbb{N}^*}$ the random trajectories of scaled-by- N discrete Petri nets and by z the unique trajectory of the fluid unscaled Petri net. For any N in \mathbb{N}^* , the values of \widehat{Z}^N are deterministically zero over $(-\infty, 0)$, just like z , so the result is trivial for $t < 0$. This is the initial step to unravel the induction principle (I) another time.

Let us suppose that the claim of the theorem is verified for all $t \in (-\infty, k\tau_{\min})$, with $k \in \mathbb{N}$. We now choose $t \in [k\tau_{\min}, (k+1)\tau_{\min})$.

◇ Using equations of Table 2.8, our induction hypothesis and the Lemma 2.21, we are able to say that for all $s < k\tau_{\min}$ and for all $p \in \mathcal{P}$, we have

$$\frac{\widehat{X}_p^N(s)}{N} \xrightarrow[N \rightarrow \infty]{\text{a.s.}} m_p + \sum_{q \in \mathcal{P}^{\text{in}}} \alpha_{pq} z_q(s).$$

Applying the Lemma 2.22 to $\widehat{X}_N := \widehat{X}_p^N - Nm_p$, $\mu := \mu_p$, $T := k\tau_{\min}$, $x = \sum_{q \in \mathcal{P}^{\text{in}}} \alpha_{pq} z_q$, we obtain since $t < T + \tau_{\min}$ that

$$\forall p \in \mathcal{P}, \quad \frac{\widehat{Y}_p^N(t)}{N} \xrightarrow[N \rightarrow \infty]{\text{a.s.}} m_p + \int_0^\infty \sum_{q \in \mathcal{P}^{\text{in}}} \alpha_{pq} z_q(t-\tau) d\mu_p(\tau).$$

In what follows, we will denote by Ω_t a subset of Ω with probability measure 1 on which the above limit is indeed achieved.

◇ Let us choose $q \in \mathcal{Q}_{\text{sync}}$ and $\omega \in \Omega_t$. Using the corresponding equation in Table 2.8, we can say that

$$Z_q^N(t; \omega) = \min_{p \in \mathcal{Q}^{\text{in}}} \left[\alpha_{qp}^{-1} Y_p^N(t; \omega) \right] = \min_{p \in \mathcal{Q}^{\text{in}}} \left(\alpha_{qp}^{-1} Y_p^N(t; \omega) \right) + H_q(t; N; \omega),$$

where H_q is a function bounded by 1, accounting for the error made approximating the integer-part of a number by this number itself. It is thus clear that for all $\omega \in \Omega_t$, $\lim_{N \rightarrow \infty} Z_q^N(t; \omega)/N$ exists and

$$\begin{aligned} \forall \omega \in \Omega_t, \quad \lim_{N \rightarrow \infty} \frac{Z_q^N(t; \omega)}{N} &= \min_{p \in \mathcal{Q}^{\text{in}}} \alpha_{qp}^{-1} \lim_{N \rightarrow \infty} \frac{Y_p(t; \omega)}{N} \\ &= \min_{p \in \mathcal{Q}^{\text{in}}} \alpha_{qp}^{-1} \left(m_p + \sum_{q' \in \mathcal{P}^{\text{in}}} \alpha_{pq'} \int_0^\infty z_{q'}(t-\tau) d\mu_p(\tau) \right) = z_q(t). \end{aligned}$$

where we finished by identification using Table 2.8. Doing so for all $t \in [k\tau_{\min}, (k+1)\tau_{\min})$, this completes the induction for these transitions.

◇ If we now choose $q \in \mathcal{Q}_{\text{psel}}$, and we call p the unique upstream place of q , we write using Table 2.8 that

$$\forall \omega \in \Omega_t, \quad \frac{Z_q^N(t; \omega)}{N} = \alpha_{qp}^{-1} \left(m_{p,q} + \frac{\Pi_q^p(Y_p^N(t; \omega) - Nm_p; \omega)}{N} \right) + \frac{H_q(t; N; \omega)}{N},$$

where we hid again a bounded-by-1 error term in the function $H_q(t; N; \omega)$.

If $\sum_{q' \in p^{\text{in}}} \alpha_{pq'} \int_0^\infty z_{q'}(t - \tau) d\mu(\tau)$ is a positive number, we know that $\lim_{N \rightarrow \infty} Y_p^N(t) - Nm_p = \infty$. In this case, we have for N large enough that

$$\frac{\Pi_q^p(Y_p^N(t; \omega) - Nm_p; \omega)}{N} = \frac{Y_p^N(t; \omega) - Nm_p}{N} \times \frac{1}{Y_p^N(t; \omega) - Nm_p} \sum_{k=1}^{Y_p^N(t; \omega) - Nm_p} R_k^{(q)}(\omega), \quad (2.19)$$

where the random variables $(\widehat{R}_k^{(q)})_{k \in \mathbb{N}^*}$ were defined in Section 3.3 of Chapter 1. What matters is that the average $\frac{1}{M} \sum_{k=1}^M \widehat{R}_k^{(q)}$ converges almost surely by the strong law of large numbers towards $\mathbb{E}(\widehat{R}_1^{(q)}) = \pi_{qp}$. Hence we have a subset Ω_q of Ω such that $\mathbb{P}(\Omega_q) = 1$ and for all $\omega \in \Omega_q$, $\lim_{M \rightarrow \infty} \frac{1}{M} \sum_{k=1}^M R_k^{(q)}(\omega) = \pi_{qp}$. Remarking that (2.19) features a subsequence of this sequence, we have

$$\forall \omega \in \Omega_t \cap \Omega_q, \quad \lim_{N \rightarrow \infty} \frac{\Pi_q^p(Y_p^N(t; \omega) - Nm_p; \omega)}{N} = \pi_{qp} \sum_{q' \in p^{\text{in}}} \alpha_{pq'} \int_0^\infty z_{q'}(t - \tau) d\mu(\tau).$$

If on the contrary we have $\sum_{q' \in p^{\text{in}}} \alpha_{pq'} \int_0^\infty z_{q'}(t - \tau) d\mu(\tau) = 0$, the above still holds by boundedness of the averaging sum. It readily follows that on this intersection that still has probability measure 1, we have using Table 2.8 one last time that

$$\lim_{N \rightarrow \infty} \frac{Z_q^N(t; \omega)}{N} = \alpha_{qp}^{-1} \left(m_{p,q} + \pi_{qp} \sum_{q' \in p^{\text{in}}} \alpha_{pq'} \int_0^\infty z_{q'}(t - \tau) d\mu(\tau) \right) = z_q(t).$$

Again, this can be done for any $q \in \mathcal{Q}_{\text{psel}}$ and $t \in [k\tau_{\min}, (k+1)\tau_{\min})$, hence the induction scheme is completed. \clubsuit

The result of Theorem 2.23 is a strong argument to legitimate the study and the use of continuous-relaxed Petri nets (possibly involving non-constant holding times), as they constitute the best and deterministic approximation of discrete twin Petri nets subject to many stochastic effects, when conjugated by the application of a scaling-factor. By comparison to Theorem 2.19, observe here that we have only obtained the limit (almost surely) of $(\widehat{Z}^N)/N$ as $N \rightarrow \infty$, but no bound on its difference with z . Deriving such an error term amounts to showing a “large deviations”-type result, that sounds more challenging. Indeed, although we may control the deviation of “low-level” random behaviours of Poisson processes, Bernoulli laws or stochastic holding times, determining how the application of the dynamics alter these bounds is not trivial.

It is also worth-mentioning that we ended up uncovering a deterministic limit for $(\widehat{Z}^N)/N$ without resorting to the moments of \widehat{Z}^N , and notably $\mathbb{E}(\widehat{Z}^N)$, that is after all the “best” deterministic estimate of \widehat{Z}^N . In particular, we do not have the relationship $\mathbb{E}(\widehat{Z}^N) = z^N = Nz$, and we cannot show that the fluid dynamics is a good approximation of the discrete dynamics by simply applying expectation operators to the discrete equations of Table 2.8! This has especially to do with the fact that expectations do not mix too well with non-linear operators such as the minimum. Intuitively, $\mathbb{E}(\widehat{Z}) \neq z$ can be explained by the fact that the discrete system is much less “fluid” (in the sense that it works more slowly) than its continuous-relaxed counterpart, on account of integrality phenomena, so we should indeed expect $\mathbb{E}(\widehat{Z}) < z$. For the case of timed event graphs, positive mean ergodic results can be obtained on this aspect, see in particular [BL91].

Note that we have considered only nonsingular measures for the distribution of stochastic holding times. This is quite reasonable for many real-life usual laws, however we reckon that the more general case could also be addressed, for instance since pure point measures are dense in the space of probability measures over metrizable spaces [AB06, Theorem 15.10, p. 513], and only the proof of Lemma 2.22 would need changing. In our Section 2.8 of the Chapter 3 to come, we shall also discuss further how we could relax the assumption that the sojourn times distribution supports do not contain 0 in their adherence and instead are included in $[\tau_{\min}, \infty)$. Indeed, even if τ_{\min} can be chosen arbitrarily small and therefore this condition is quite acceptable for many applications, it is not much in line with our quest of taking into account null sojourn times.

3

SEMI-MARKOV DECISION PROCESSES UNDER WEAK NON-ZERO CONDITIONS

Contents

1	Basic terminology and recalls on Markov chains	87
1.1	Underlying random process and states classification	87
1.2	Long-run behaviour of irreducible Markov chains	88
1.3	Handling multiple subchains	89
2	The finite-horizon value of semi-Markov decision processes	92
2.1	Definition of the game	92
2.2	Strategies and policies	95
2.3	The different finite-horizon criteria	97
2.4	Existence of the value functions in presence of null sojourn times	100
2.5	Computing the value function	107
2.6	Backward of forward propagation?	115
2.7	The evolution semigroup approach	117
2.8	Extension to the case with non-constant sojourn times	120
3	Long-run properties of solutions of semi-Markov dynamics	121
3.1	Affine regimes and average-cost vector	121
3.2	Second order results in cases reducible to MDP	125
4	Stochastic Shortest Path configurations	126
4.1	Topological and cost-related restrictions	127
4.2	Convergence in the semi-Markov setting	128

In this chapter, we momentarily leave aside models of Petri nets to better come back to it in the next one. We focus here on the class of one-player stochastic games known as Markov decision processes (abb. MDP) in which a decision-maker seeks to minimize cost criteria under uncertainty.

More generally on the class of semi-Markov decision processes (abb. SMDP), in which the elapsing of the physical time is involved.

Following the seminal work of Bellman on MDPs [Bel57], semi-Markov decision processes (sometimes also referred to as Markov renewal programs) were introduced by Jewell [Jew63] and De Cani [DC64] in the early 1960's. These articles already featured the so-called finite-horizon dynamic programming equations of SMDPs, characterizing optimality of strategies implemented by the decision-maker to minimize his/her incurred overall cost when playing the SMDP for a given period of time fixed in advance. As customary in stochastic control problems, these dynamic programming equations are the gateway to both theoretical analysis of the random processes at stake and algorithmic methods to compute optimal strategies. Refinements revolving around these equations and extensions to milder cases include the work of Yushkevich [Yus82] or Huang and Guo [HG11]. We also mention the book of Puterman [Put14] that is a detailed introduction to MDPs and basic SMDP settings. The average-cost problem, in which one wants to incur the least possible cost per unit of time when the game is played indefinitely, has received a lot of attention in the MDP setting (see the survey [ABFG⁺93]) but also in the semi-Markov framework. Average-cost optimality equations derived from the finite-horizon criterion as well as related algorithmic aspects and analysis of convergence properties can be found in the work of Denardo and Fox [DF68] or in a series of articles of Schweitzer and Federgruen [SF77, SF78, SF79]. We also refer to the work of Feinberg [Fei94] and Jaśkiewicz [Jaś04] for a treatment of the problem in expressive contexts. In practice, SMDPs have been used in numerous applications, including preventive maintenance [CT05], control of queuing systems [YT91] or production scheduling [Sen89].

By contrast with the vast majority of the literature on SMDPs, we consider in this chapter models in which the player can take decisions that do not require him/her to wait before the next step of the game. This entails that several costs can be incurred over a zero time-period. We state minimal non-Zeno conditions under which these instantaneous moves can only occur finitely many times, so that classical cost criteria studied in the theory of SMDPs are still well-defined. In this sense, and although their proof require a more delicate treatment and rely on advanced operator-based techniques, the results of this chapter can mainly be seen as extensions of well-known properties of SMDPs under weak non-Zeno conditions and the readers already familiar with this literature should not be surprised by our statements. Another originality of our approach though is to use a formalism borrowed from control theory and based on an evolution

Type of cost	Notation	Exists and is finite under
Total cost in finite-horizon t (possibly with terminal and stopping costs)	$v^*(t)$ (3.5)	Assumption A
Total discounted cost in finite-horizon t (possibly with terminal and stopping costs)	$v_\gamma^*(t)$ (3.6)	Assumption B
Average-cost in infinite-horizon	g^* (3.25)	Assumption A
Total cost in infinite-horizon	u^* (3.28)	Assumptions A and C

Table 3.1: Cost criteria studied for semi-Markov decision processes in Chapter 3

semigroup, which allows us to highlight the essential properties of the finite-horizon dynamic programming equations and which facilitates a systematic treatment.

To help getting a clear picture of the cost criteria studied in this chapter from Section 2 onwards, we summarize them in the Table 3.1.

We devote Section 1 to preliminaries on Markov chains, since the analysis of MDPs and SMDPs builds on their transient and stationary properties. Readers already acquainted with these notions can easily skip this part. Our Section 2 introduces undiscounted and discounted finite-horizon value functions before exhibiting minimal conditions under which they are well-posed. It then focuses on the core of the SMDP machinery, that is, the dynamic programming equations and the existence of solutions under the same weak non-Zeno assumptions. We also introduce the corresponding evolution semigroup. The Section 3 deals with the asymptotic behaviour of the undiscounted value function and provides interpretation in terms of affine solutions to the dynamics induced by non-Zeno SMDPs. Finally, we extend in Section 4 to the semi-Markov case the results of Bertsekas and Tsitsiklis [BT91] on the stochastic shortest path problem, so far only stated for MDPs. This allows us to prove that solutions to SMDP dynamics converge under appropriate conditions.

1 Basic terminology and recalls on Markov chains

The aim of this section is to introduce the notation and recall the notions that we shall use on Markov chains, ranging from elementary definitions to little more advanced results on reducible chains. The propositions and theorems stated below are given without proofs; for a complete presentation on the topic, we refer the reader to the textbooks of Norris [Nor98], of Berman and Plemmons [BP79] and Meyer [Mey00].

- 1.1 UNDERLYING RANDOM PROCESS AND STATES CLASSIFICATION.** — We consider a random process $(\widehat{i}_n)_{n \in \mathbb{N}}$ on a finite set $S := \{1, 2, \dots, |S|\}$. This process is said to be a **Markov chain**, or to verify the *Markov property*, if for all n in \mathbb{N} and for all (i_0, i_1, \dots, i_n) in S^{n+1} , we have

$$\mathbb{P}(\widehat{i}_n = i_n \mid \widehat{i}_{n-1} = i_{n-1}, \dots, \widehat{i}_1 = i_1, \widehat{i}_0 = i_0) = \mathbb{P}(\widehat{i}_n = i_n \mid \widehat{i}_{n-1} = i_{n-1}), \quad (3.1)$$

in other words the value of the random process at any step depends only on its value at the previous step. The set S is called the **state space** of the Markov chain, and if for n in \mathbb{N} we have $\widehat{i}_n = j$, we say that the chain is in state j at **epoch** (or step) n . We shall in addition always suppose that the Markov chain is **homogeneous**, i.e., the quantity $\mathbb{P}(\widehat{i}_n = i_n \mid \widehat{i}_{n-1} = i_{n-1})$ above is actually independent of the epoch number n . This way, we can consider the $|S| \times |S|$ matrix P , defined by

$$\text{for all } i, j \text{ in } S, \quad P_{ij} := \mathbb{P}(\widehat{i}_n = j \mid \widehat{i}_{n-1} = i) \quad \text{regardless of } n \text{ in } \mathbb{N}.$$

The matrix P is stochastic (i.e., it has nonnegative entries and $P\mathbf{1} = \mathbf{1}$) and it is customarily referred to as the **transition matrix** of the Markov chain. For k in \mathbb{N} , the matrix P^k is still stochastic, and using the Markov property, one can show that for any n in \mathbb{N} , and i, j in S , we have $\mathbb{P}(\widehat{i}_{n+k} = j \mid \widehat{i}_n = i) = [P^k]_{ij}$, so that P^k is the k -step transition matrix of the Markov chain. In general, using the Bayes rule, we have for all n in \mathbb{N} and (i_0, i_1, \dots, i_n) in S^{n+1} that

$$\mathbb{P}(\widehat{i}_0 = i_0, \widehat{i}_1 = i_1, \dots, \widehat{i}_n = i_n) = \mathbb{P}(\widehat{i}_0 = i_0) \prod_{j=0}^{n-1} P_{i_j i_{j+1}}.$$

Given two states i and j in S , we say that j is accessible from i if there is some k in \mathbb{N} such that $[P^k]_{ij} > 0$. We say that the states i and j **communicate** if i is accessible from j and j is accessible from i . Communication between states defines an equivalence relation over S , whose equivalence classes are called the communication classes of S ; they form a partition of S by definition. In the case where S is made of a single communication class, we say that the Markov chain is **irreducible**.

Knowing if some state of S shall be visited or not during the execution of the chain depending on the initial state is a key question with many implications. To that purpose, it is useful to introduce for all i in S the *return time* \widehat{T}_i to i , which is a random variable giving the minimum number of steps needed to hit i :

$$\text{for all } i \text{ in } S, \quad \widehat{T}_i := \inf \{n \in \mathbb{N}^* \mid \widehat{i}_n = i\}.$$

We say that a state i in S is **recurrent** if $\mathbb{P}(\widehat{T}_i < \infty \mid \widehat{i}_0 = i) = 1$, otherwise it is called **transient**. If it is visited once by the chain, a recurrent state will actually be visited infinitely many times in average, while a transient state will in average be visited a finite number of times. If i in S is recurrent, it is sometimes useful to distinguish whether $\mathbb{E}(\widehat{T}_i \mid \widehat{i}_0 = i) < \infty$ or not. If this inequality holds, i is said to be positive recurrent, and null recurrent otherwise. However, in the case of Markov chains over finite state spaces, recurrent states are necessarily positive recurrent.

If the state i in S is recurrent, its period $d(i)$ is defined by

$$d(i) := \text{GCD} \{n \in \mathbb{N}^* \mid \mathbb{P}(\widehat{i}_n = i \mid \widehat{i}_0 = i) > 0\}$$

where GCD denotes the greatest common divisor of a collection of integers, so that any return to i must occur in multiples of $d(i)$ steps. If $d(i) > 1$, we say that i is a **periodic** state of the Markov chain, otherwise it is called **aperiodic**. In a communicating class, all states have the same period, therefore we say that a class (or an irreducible Markov chain) is periodic (resp. aperiodic) when its states are periodic (resp. aperiodic).

1.2 LONG-RUN BEHAVIOUR OF IRREDUCIBLE MARKOV CHAINS. — If two states of S do communicate, they are either both recurrent or both transient. In particular, since not all states of the finite set S can be transient, we know that all the states of an irreducible Markov chain are positive recurrent. It is then a question of interest in this case to know if a “limiting” (or steady-state) distribution exists, i.e., to know for i and j in S if the quantity $\mathbb{P}(\widehat{i}_n = j \mid \widehat{i}_0 = i)$ admits a limit when n tends to ∞ . Based on the previous observations, it amounts to knowing if the sequence $([P^n]_{ij})_{n \in \mathbb{N}}$ converges. In order to answer this question, it is useful to take advantage of the following instance of the fundamental Perron–Frobenius theorem. Recall that for a square matrix M , the *spectral radius* $\rho(M)$ of M is the greatest modulus of its eigenvalues (and we already know if P is a stochastic matrix that $\rho(P) \geq 1$ since $P\mathbf{1} = \mathbf{1}$).

➤ **THEOREM 3.1** (Perron–Frobenius for stochastic matrices). *If P in $\mathbb{R}_{>0}^{S \times S}$ is the transition matrix of an irreducible Markov chain, then we have $\rho(P) = 1$ and 1 is a simple eigenvalue of P (which means that it has algebraic and geometric multiplicities of one). In addition, there is a unique positive vector μ in $\mathbb{R}_{>0}^S$ such that $\mu^T P = \mu^T$ and $\|\mu\|_1 = 1$. All the other left-eigenvectors of P associated with the eigenvalue 1 are multiples of μ , and no left-eigenvector of P associated with an eigenvalue different than 1 is positive.*

The vector μ above is called the **invariant measure** of the (irreducible) Markov chain. Intuitively, if 1 is the only eigenvalue of P with modulus one, we expect (for instance using the Jordan form of P) that the sequence $(P^n)_{n \in \mathbb{N}}$ converges towards the outer product $\mathbf{1} \mu^T$, or in other words the

$|S| \times |S|$ matrix with all rows equal to μ^\top . However, the previous assumption is not always satisfied if the Markov chain is periodic. Indeed, one can show that if all the states of an irreducible Markov chain have period d in \mathbb{N}^* , then the complex numbers $(e^{2ik\pi/d})_{0 \leq k \leq d-1}$ of the unit circle are all eigenvalues of P (in this case we say in Perron–Frobenius theory that P is not primitive). We have the following proposition:

► **PROPOSITION 3.2.** *Let P be the transition matrix of an irreducible Markov chain, and μ its invariant measure.*

– *If the chain is aperiodic, we have*

$$\lim_{n \rightarrow \infty} P^n = \mathbf{1} \mu^\top .$$

– *If the chain is periodic with period d in $\mathbb{N} \setminus \{0, 1\}$, then for all i and j in S , there exists an integer $0 \leq r \leq d - 1$ such that*

$$\lim_{n \rightarrow \infty} [P^{nd+r}]_{ij} = d\mu_j .$$

Although one may not reach a steady-state distribution due to periodic behaviours, we can still wonder if the states of S have an ultimate “visit frequency”. The aperiodic case discussed in Proposition 3.2 indicates that for all i in S , μ_i gives the average fraction of time spent in state i . Computing this average no longer involves focusing on the sequence $(P^n)_{n \in \mathbb{N}}$, but on its associated Cesàro sequence $((1/n) \sum_{k=0}^{n-1} P^k)_{n \in \mathbb{N}}$. The following proposition confirms that doing so has a regularizing power, just like for sequences of real numbers.

► **PROPOSITION 3.3.** *Let P be the transition matrix of an irreducible Markov chain, and μ its invariant measure. Then, denoting by P^\star the Cesàro limit of the sequence $(P^n)_{n \in \mathbb{N}}$, we have*

$$P^\star := \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} P^k = \mathbf{1} \mu^\top .$$

When P is a stochastic matrix, the matrix P^\star defined above always exists because its eigenvalues with modulus one are necessarily semisimple (i.e., they have equal algebraic and geometric multiplicities), see for instance [Mey00]. The matrix P^\star is such that $PP^\star = P^\star P = P^\star$, and also $(P^\star)^2 = P^\star$, it is therefore sometimes called the **spectral projector** of the Markov chain. Finally, we may add that in the context of irreducible finite Markov chains, for all states i in S , we have $\mu_i = 1/\mathbb{E}(\widehat{T}_i \mid \widehat{i}_0 = i)$, so that the state i is visited in average once every $\mathbb{E}(\widehat{T}_i \mid \widehat{i}_0 = i)$ steps. This is consistent with the idea that the less the return time to i is, the more the visit frequency μ_i is.

1.3 HANDLING MULTIPLE SUBCHAINS. — Not all finite Markov chains are irreducible, or equivalently made of a single positive recurrent communicating class. Since two communicating states are either both recurrent or both transient as recalled in the beginning of Section 1.2, the partition of S into communicating classes mentioned in Section 1.1 allows us to build a slightly coarser partition of S into the form

$$S = T \uplus F_1 \uplus F_2 \uplus \cdots \uplus F_m \quad \text{for some } m \text{ in } \mathbb{N}^*, \quad (3.2)$$

where T is the set of transient states of the Markov chain (that may or may not communicate), and the $(F_k)_{1 \leq k \leq m}$ are communicating classes that by definition do not communicate with each other. The $(F_k)_{1 \leq k \leq m}$ are called the **final classes** (or sometimes the recurrent classes, or the ergodic classes) of the Markov chain, because once a state of F_k is visited for some $1 \leq k \leq m$, the

random process is *forced* to stay in F_k forever. Hence, the global Markov chain induces irreducible subchains over its final classes. Since the irreducible case is well understood, this motivates the following result, in which a convenient block-decomposition of the transition matrix P of the whole Markov chain is introduced.

- **THEOREM 3.4.** *Let P be the transition matrix of a Markov chain. Then, up to relabeling the states (or equivalently up to conjugation by a permutation matrix), P admits the following **standard form** adapted to the decomposition (3.2)*

$$P = \begin{pmatrix} Q_0 & Q_1 & Q_2 & \cdots & Q_m \\ 0 & P_1 & 0 & \cdots & 0 \\ 0 & 0 & P_2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & P_m \end{pmatrix}, \quad (3.3)$$

where (P_1, \dots, P_m) are the (stochastic and square) transition matrices associated with the final classes $(F_k)_{1 \leq k \leq m}$ of the Markov chain. The square matrix Q_0 is strictly substochastic, in particular we have $\rho(Q_0) < 1$.

The matrix Q_0 gives the probabilities that the random process initialized in a transient state stays in T in one step, while the (rectangular) matrices Q_1, Q_2, \dots, Q_m tell the probabilities for the random process initialized in a transient state to directly move in a final class. Remark that the eigenvalues of P are the eigenvalues of the matrices $Q_0, P_1, P_2, \dots, P_m$, consequently $\rho(P) = 1$ and from the semisimplicity result mentioned before, 1 is an eigenvalue of P with algebraic and geometric multiplicity m . It can easily be seen that if $\mu_1, \mu_2, \dots, \mu_m$ are the invariant measures of the matrices P_1, P_2, \dots, P_m , then for all $1 \leq k \leq m$, the vector

$$\tilde{\mu}_k := \left(\underbrace{\mathbf{0}}_{|T|}, \underbrace{\mathbf{0}}_{|F_1|}, \underbrace{\mathbf{0}}_{|F_2|}, \dots, \underbrace{\mu_k^\top}_{|F_k|}, \dots, \underbrace{\mathbf{0}}_{|F_m|} \right)^\top$$

is an invariant measure of P . Actually, the set of invariant measures of P is a convex polytope whose vertices are the $(\tilde{\mu}_k)_{1 \leq k \leq m}$.

We shall need to use the following lemma later in the chapter. Although it is not difficult to show, it is not exactly standard in the literature, thus its proof is given.

- **LEMMA 3.5.** *Let P be the transition matrix of a finite Markov chain. Let J be a subset of S and $P|_{J \times J}$ be the square submatrix of P restricted to states of J . Then, this submatrix has spectral radius 1 if and only if it contains a stochastic submatrix, or alternatively*

$$\rho(P|_{J \times J}) = 1 \iff \exists k \in \{1, 2, \dots, m\}, \quad F_k \subset J.$$

Proof. We take advantage of the decomposition (3.3). It is clear that if we have $F_k \subset J$ for some $1 \leq k \leq m$, we can find a right eigenvector associated with eigenvalue 1 for $P|_{J \times J}$. Conversely, by irreducibility, any strict submatrix of P_1, \dots, P_m is substochastic and then has spectral radius strictly less than 1, therefore $P|_{J \times J}$ cannot admit an eigenvector associated with 1 if J doesn't contain the whole support of some P_k for $1 \leq k \leq m$, i.e., a final class of P . \clubsuit

We now go back to the question of understanding the long-run behaviour of our random process in case the Markov chain is not irreducible. To that purpose, we still study the spectral projector introduced in Section 1.2 for irreducible chains. An essential difference with Proposition 3.3 is

The matrix H^* is frequently referred to as the **fundamental matrix** of the Markov chain. Note that it verifies $P^*H^* = H^*P^* = 0$, and $P^* + H^* = I + PH^*$. A key property that explains the importance of H^* and why it allows one to gain more insight on the long-run behaviour of the Markov chain beyond the mere knowledge of steady-state distributions is the following, for instance found in [Ber11, Proposition 5.1.1].

► **PROPOSITION 3.8.** *If P is the transition matrix of a Markov chain, then for all α in $(0, 1)$, the matrix $I - \alpha P$ is nonsingular, and*

$$(I - \alpha P)^{-1} \underset{\alpha \rightarrow 1}{=} \frac{1}{1 - \alpha} P^* + H^* + o(1 - \alpha).$$

2 The finite-horizon value of semi-Markov decision processes

We now discuss one of our main objects of interest in this thesis, the semi-Markov decision processes (SMDPs). Since these are extensions of the usual Markov decision processes (MDPs), we choose to directly state definitions and properties of SMDPs, that will be particularized to MDPs when relevant.

2.1 DEFINITION OF THE GAME. — The models of semi-Markov decision processes found in the literature can vary a little from one author to another, but they always revolve around the following ground objects:

- (1) A **state space** S , that can be either finite, countably infinite, compact, or even a Borel set of a topological space.
- (2) For all states i in S , a nonempty set A_i of **playable actions** from i , that can also range from finite to continuous collections. Contrary to a common practice in the literature, we choose without any loss of generality to impose that the $(A_i)_{i \in S}$ are disjoint, so that action-dependent objects can be non-ambiguously indexed by actions only, and not state-action pairs. We shall denote by A the disjoint union $\bigsqcup_{i \in S} A_i$ of all the playable actions.
- (3) For all actions a in A , a **transition probability distribution** p^a over S . In case the state space is finite, p^a is conveniently represented by a vector of $\mathbb{R}_{\geq 0}^S$.
- (4) For all actions a in A , a nonnegative **sojourn time** t^a , that may be given as the realization of a nonnegative random variable with a distribution ν_a depending only on a . If ν_a is a Dirac measure, we talk about deterministic sojourn time (t^a is constant).
- (5) For all actions a in A , a real-valued **cost** c^a . This cost may be a combination of a real-valued *impulse cost* c_1^a independent of t^a , and of a real-valued measurable *rate of cost* function c_2^a integrated between 0 and t^a , i.e., $c^a = c_1^a + \int_0^{t^a} c_2^a(s) ds$. If the sojourn time associated with an action is deterministic, so is the corresponding cost.

Occasionally, the following modeling aspects can be used, and we will progressively discuss their role in Section 2.3:

- (6) For all a in A , a nonnegative **discount factor** γ^a , taken into account to reflect the player's "preference for the present". To this purpose, setting $\gamma^a := \exp(-\alpha t^a)$ where α is a positive constant is a canonical choice, but we may also allow other values, even greater than 1.

- (7) For all i in S , a **terminal cost** $w(i, \cdot)$, that is paid by the player only once the game is finished in state i . We will explain in Section 2.3 why it makes sense that $w(i, \cdot)$ is a function.
- (8) For all i in S , a **stopping cost** $R(i, \cdot)$ (or leaving cost), that the player will pay if he/she decides to prematurely stop playing the game at some point from state i . Once again, $R(i, \cdot)$ is a function whose modeling interest will be explained in Section 2.3

Let us now describe how the game proceeds, letting aside momentarily the occasional objects (6)–(8). When arriving in a state i of S , the player must immediately choose an action a of A_i to play. By doing so, he/she immediately incurs the impulse cost c_1^a , and is held in state i during the (possibly nondeterministic) time t^a . Throughout this period, the player also pays a cost determined by the rate c_2^a . After the sojourn time t^a is elapsed, the player instantaneously goes to another state of S according to the distribution probability p^a (and beyond his/her control, so that we often talk about this steps as “nature moves”) and the whole process starts again. The time instants at which arrivals in states occur and therefore at which actions are picked are called the **decision epochs** of the process.

Several authors (for instance [Sch92, HG11]) consider a more general model in which the transition probability distributions attached to actions depend upon the sojourn time. To model this entangled behaviour between time and transition probabilities, **transition kernels** $(K_j^a)_{a \in A, j \in S}$ may be introduced, in such a way that the probability to go to i from j in S after t units of time is $K_j^a(t)$ when pulling the action a in A_i . In addition, we sometimes encounter models in which transition probability distributions $(p^a)_{a \in A}$ may be substochastic, i.e., with total measure less than 1, so that $1 - \int_{j \in S} dp^a(j)$ gives the probability that the game stops after playing a (sometimes called the death probability for a). This case can handily be covered by properly stochastic transition distributions up to adding a distinguished *cemetery* state that is absorbing.

We shall work in what follows with **finite** state space and **finite** playable actions sets. We shall also consider unless specified otherwise that sojourn times are **deterministic** and that costs are reduced to their impulse component. We also assume that transition probability distributions are stochastic vectors (for all a in A , $\sum_{j \in S} p_j^a = 1$) that are independent of the time. We will discuss in Section 2.8 how the properties we state extend to the case of stochastic sojourn times – but yet impulse costs and time-independent transition probabilities. The following principle epitomizes the core idea of SMDPs:

Semi-Markov decision processes are one-player games in which one evolves in a stochastic way through states by choosing actions at real-valued decision epochs. The chosen actions determine some costs that the player incur, as well as sojourn times holding him/her before continuing the game.

We talk about **Markov decision processes** when all the sojourn times are deterministically set to one (for all a in A , $t^a := 1$). This greatly simplifies the setting, since neither the transition probabilities nor the costs – reduced to their impulse part – associated with actions depend on the decision epochs, and the latter precisely occur at epochs 0, 1, 2, 3,... A major difference with far-reaching consequences between MDP and SMDPs is that the semi-Markov decision processes run with respect to the physical time, contrary to Markov decision processes which operate relatively to a logical time. This important aspect was already discussed in the Section 3.1 of Chapter 1.

Example 3.9. We illustrate below in Figure 3.2 an SMDP with state space $\{1, 2, 3\}$ and playable action sets $A_1 := \{a_{11}, a_{12}\}$, $A_2 := \{a_{21}, a_{22}\}$ and $A_3 := \{a_{31}, a_{32}\}$. The states are depicted by circles, and actions by squares attached to their parent states. Each action is labeled with a pair (c, t) indicating the immediate cost and the sojourn time incurred by the player as a result of choosing this action (in this example, we assume that there are no discount factors). Thin arrows joining actions to states represent the stochastic moves that are possible after pulling an action, the corresponding probabilities being given along these arrows if not equal to one.

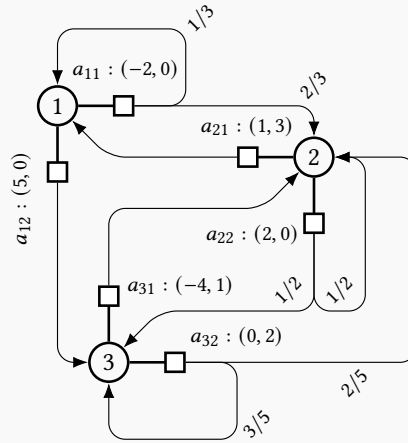


Figure 3.2: A semi-Markov decision process with three states and six actions

In this SMDP, observe that there are three “costly” actions a_{12} , a_{21} and a_{22} , while the two actions a_{11} and a_{31} actually give the player rewards (negative costs). Playing the action a_{32} gives no cost or reward. Playing one of the three actions a_{11} , a_{22} and a_{32} exposes to stochastic moves, while the evolution after pulling one of the three actions a_{12} , a_{21} and a_{31} is known in advance. Finally remark that three actions (a_{11} , a_{12} and a_{22}) cause no positive sojourn time, making the player instantaneously move to the next state when they are selected, contrary to actions a_{21} , a_{31} and a_{32} which respectively impose to hold for 2, 1 and 3 units of time.

Any execution of the game in n epochs is characterized by the knowledge of the $n + 1$ first visited states and the n first chosen actions. Indeed the knowledge of an action also determines the corresponding transition probabilities, sojourn time, incurred cost and potential discount factor. The chosen actions must at each epoch belong to the playable actions set of the parent state that is also visited at this epoch; consequently it is useful to define the set X of admissible state–actions pairs by

$$X := \{(i, a) \in S \times A \mid a \in A_i\}.$$

For n in \mathbb{N}^* , the set $H_n := X^n \times S$ is called the set of **histories of length n** of the game, and any element $(i_0, a_0, i_1, a_1, \dots, i_{n-1}, a_{n-1}, i_n)$ of H_n is called a history of length n . Remark that if we were working with non-deterministic sojourn times, we would have defined $H_n := (X \times [0, \infty))^n \times S$, for in this case the knowledge of the k -th chosen action is not sufficient to determine the k -th sojourn time and this information needs adding specifically. We shall also denote by H_∞ the set of histories with infinite length, isomorphic to $X^\mathbb{N}$.

2.2 STRATEGIES AND POLICIES. — The objective of a decision-maker playing the game described by an SMDP will be to minimize his/her overall incurred cost, either over a finite or infinite time period. To this purpose, some playing schemes naturally turn out to be more interesting than others. Indeed, the central question is the following; knowing that from each state we have possibly multiple choices of actions that in turn yield costs we incur, sojourn time during which we cannot play, and stochastic moves to other states, what actions should we pick to minimize our cost? This question actually hinges on somewhat subtle considerations, and we choose to discuss them progressively in this subsection in order to draw the attention of non-expert readers to the key aspects of the problem.

The simplest way to play is to apply a **policy** (in [Put14], the term **decision rule** is used instead), which is a mapping σ from S to A such that for all states i in S , $\sigma(i)$ is in A_i . If the decision-maker adopts the policy σ , it means that whenever he/she arrives in state i , he/she will always choose to pull the action $\sigma(i)$. We denote by \mathfrak{S} the set of policies, and in the setting of finite state and actions spaces, \mathfrak{S} is isomorphic to $\prod_{i \in S} A_i$. In particular we see that the number of policies is exponential with respect to the size of S or the $(A_i)_{i \in S}$. A policy σ induces a homogeneous Markov chain over S , since anytime the state i in S is visited, the probabilities $(p_j^{\sigma(i)})_{j \in S}$ to evolve towards state j as a result of choosing action $\sigma(i)$ are independent of the epoch. We shall denote by $P^\sigma := (p_j^{\sigma(i)})_{i,j \in S}$ the $|S| \times |S|$ transition matrix associated with the policy σ , and since the costs, sojourn times and discount factors incurred by applying a policy are only state-dependent as well, we shall later use the notation $c^\sigma := (c^{\sigma(i)})_{i \in S}$, $t^\sigma := (t^{\sigma(i)})_{i \in S}$ and $\gamma^\sigma := (\gamma^{\sigma(i)})_{i \in S}$.

However, much more general playing schemes than policies can be made up. First, policies are history-independent, in the sense that none of the previously visited states and chosen actions are taken into account in the selection of the action to play at a current epoch, and we may rightfully reckon that these extra pieces of information could yield a better action pick. Second, it is not trivial that a single choice of action a_k at each epoch k is optimal; and in the “Do not put all your eggs in one basket” spirit, the player might want to spread his/her choice over several actions $a_{k,1}, a_{k,2}, \dots, a_{k,m}$ with different weights. These two issues are captured by the standard notion of **strategies** (thoroughly discussed in [Put14, Section 2.1]).

Formally, a **randomized and history-dependent strategy** f of the game is a sequence $(f_n)_{n \in \mathbb{N}}$ where for all n in \mathbb{N} , f_n is a map from H_n to the set of probability measures over A , so that if $h_n = (i_0, a_0, \dots, i_n)$, then $f_n(h_n)$ is a probability measure over A_{i_n} . Using the strategy f , the player will choose the action a in A_{i_n} at epoch n with probability $f_n(h_n)(a)$, often rewritten $f_n(a | h_n)$. This way, all the past information is used to make present choices. Observe that for A finite, the set Δ of admissible probability measures over A is here isomorphic to a product of simplices, with $\Delta \cong \prod_{i \in S} \{x \in \mathbb{R}_{\geq 0}^A : \|x\|_1 = 1\}$. We shall denote by \mathbb{F} the set of all randomized and history-dependent strategies.

The strategy f of \mathbb{F} is called **pure** (or deterministic) if for all n in \mathbb{N} , $f_n(h_n)$ is a Dirac measure over A_{i_n} , in other words f_n can be thought of as a map from H_n to A . If for all n in \mathbb{N} , the probability measure $f_n(h_n)$ does not depend on the whole past history h_n but only on the current state i_n , we say that the strategy f is **Markovian** (as $f_n(\cdot | h_n) = f_n(\cdot | i_n)$ echoes the Markov property (3.1)). A Markovian strategy $(f_n)_{n \in \mathbb{N}}$ can be identified with a sequence of mappings from S to the set Δ of admissible probability measures over A (such that $f_n(i)$ is concentrated on A_i for all i in S). We say that a Markovian strategy is **stationary** if it does not depend upon the epoch n , meaning that if the player hits the state i in S at two different epochs, he/she will pick actions to play from A_i with same probabilities at these two steps. Markovian strategies therefore write (f_0, f_0, f_0, \dots) , where f_0 can be identified with a mapping from S to the set of probability measures over A . Pure Markovian strategies can be identified to a sequence of policies, and pure stationary (hence Markovian)

strategies can be identified to a single policy. We respectively denote by \mathbb{F}^{HD} , \mathbb{F}^{MR} , \mathbb{F}^{MD} , \mathbb{F}^{SR} , \mathbb{F}^{SD} the sets of: history-dependent deterministic strategies, Markovian randomized strategies, Markovian deterministic strategies, stationary randomized strategies and stationary deterministic strategies, with the classical inclusions $\mathbb{F}^{\text{SD}} \subset \mathbb{F}^{\text{SR}} \subset \mathbb{F}^{\text{MR}} \subset \mathbb{F}$, $\mathbb{F}^{\text{SD}} \subset \mathbb{F}^{\text{MD}} \subset \mathbb{F}^{\text{MR}} \subset \mathbb{F}$, and $\mathbb{F}^{\text{SD}} \subset \mathbb{F}^{\text{MD}} \subset \mathbb{F}^{\text{HD}} \subset \mathbb{F}$.

In the coming sections, we will wonder if some strategies are “optimal” in the sense that they minimize the incurred cost of the player. In particular, we will be interested in the **sufficiency** of some subsets of \mathbb{F} , for example, can the very specific pure stationary strategies can do as well as the very general randomized and history-dependent strategies? In what follows, we shall encounter cases where we need to use history-dependent strategies, and we point out that allowing randomized strategies is for instance essential in constrained Markov decision problems, not considered here (see [Alt99]). A question related to the sufficiency problem is to wonder, what are the minimal information to keep from complete past histories to still obtain optimal strategies? This issue is of particular interest when the game has to be played during a known finite period of time, called the **planning horizon**, i.e., when the game stops at the first epoch for which the accumulation of sojourn times incurred as a result of picking actions exceeds this horizon (or equivalently, the game may continue but costs collected beyond this point are not taken into account). The notion of policy (or of pure stationary strategy) is by design blind to the concept of planning horizon, because identical decisions are always taken from identical states. On the contrary, choices of actions in a more general strategy may vary if we are close from the end of the game, for at this stage we may prefer riskier and more quickly rewarding actions.

It is classical in the field of MDPs or in stochastic control problems in general to show that optimal strategies for overall costs incurred in finite-horizon depend on the remaining number of decisions to take. In the SMDP setting, replacing the logical time by the physical time, one expects a similar outcome based on the **remaining time to play**. Miller shows such a result in his work [Mil68], see also the additional references provided by Mamer in [Mam86]. We may then be tempted to authorize strategies to explicitly depend on the planning horizon, seen as a contextual parameter. Yet, for a fixed planning horizon $t \geq 0$, a randomized and history-dependent strategy that has no explicit knowledge of t can emulate a strategy which does. To clarify this question in the next sections, we shall adopt a formalism with **augmented histories**, on the model of [HG11]. For n in $\mathbb{N} \cup \{\infty\}$, instead of considering a generic element of the set H_n to be of the form $h_n = (i_0, a_0, i_1, a_1, \dots, i_{n-1}, a_{n-1}, i_n)$, we shall rather write $h_n = (i_0, s_0, a_0, i_1, s_1, a_1, \dots, i_{n-1}, s_{n-1}, a_{n-1}, i_n)$, where for all $1 \leq k \leq n$, $s_k := \sum_{\ell=0}^{k-1} t^{a_\ell}$ is the total physical time elapsed at the beginning of epoch k for this history (and by convention $s_0 = 0$). The $(s_k)_{1 \leq k \leq n}$ are clearly redundant information (and they would also be in presence of non-deterministic sojourn time since in this case the $(t^{a_k})_{1 \leq k \leq n}$ are included in the histories), but we choose to keep them for the important role that they will play; adding the $(s_k)_{1 \leq k \leq n}$ in the histories indeed allow the strategies of \mathbb{F} to explicitly use them.

A strategy f in \mathbb{F} and an initial state i in S give rise to a random process $(\widehat{i}_n, \widehat{a}_n)_{n \in \mathbb{N}}$ over X of visited states and chosen actions (with $\widehat{i}_0 = i$). The random process $(\widehat{i}_n)_{n \in \mathbb{N}}$ is not a Markov chain in general. Equivalently, thinking of the set H_∞ of all infinite histories of the SMDP as a sample space, the $(\widehat{i}_n)_{n \in \mathbb{N}}$ and the $(\widehat{a}_n)_{n \in \mathbb{N}}$ can be seen as random variables over H_∞ , defined if h in H_∞ writes $h = (i_0, a_0, i_1, a_1, \dots)$ by $\widehat{i}_n(h) := i_n$ and $\widehat{a}_n(h) := a_n$ for all n in \mathbb{N} . Similarly, we define for all n in \mathbb{N} and h in H_∞ the random variables $\widehat{c}_n(h) := c^{a_n}$, $\widehat{t}_n(h) := t^{a_n}$, $\widehat{s}_n(h) := \sum_{\ell=0}^{n-1} t^{a_\ell}$ and $\widehat{y}_n(h) := \gamma^{a_n}$.

Denoting \mathcal{A} the σ -algebra associated with H_∞ (in our case the power set of $X^{\mathbb{N}}$), we know by the Ionescu-Tulcea’s theorem (see again [Nev65, Proposition 5.1.1] or [BS78, Proposition 7.8])

that a strategy $f = (f_n)_{n \in \mathbb{N}}$ in \mathbb{F} and an initial state i in S induce a unique probability measure \mathbb{P}_i^f on \mathcal{A} such that $\mathbb{P}_i^f(\widehat{i}_0 = i) = 1$ and for all n in \mathbb{N} , h_n in H_n , j in S and a in A , we have

$$\mathbb{P}_i^f(\widehat{a}_n = a \mid \widehat{h}_n = h_n) = f_n(a \mid h_n) \quad \text{and} \quad \mathbb{P}_i^f(\widehat{i}_{n+1} = j \mid \widehat{h}_n = h_n, \widehat{a}_n = a) = p_j^a,$$

where for $h = (i_0, a_0, i_1, a_1, \dots)$ in H_∞ , \widehat{h}_n refers to the random variable defined by $\widehat{h}_n(h) := (i_0, a_0, i_1, a_1, \dots, i_{n-1}, a_{n-1}, i_n)$. We will denote by \mathbb{E}_i^f the expectation operator associated with \mathbb{P}_i^f . We refer to [Jaśo4] for similar constructions if the sojourn times are non-deterministic.

2.3 THE DIFFERENT FINITE-HORIZON CRITERIA. — We now want to formalize the idea previously touched upon that different strategies applied by the player result in different incurred costs, and we want to find out if some strategies can achieve a minimal cost. Recall that even if we apply a fixed strategy, we are subject to stochastic moves, which is the reason why we shall only talk about costs of strategies through expectation operators. The total cost of a strategy also depends on how long the game is played, i.e., the planning horizon, and also on the starting state. These elements motivate the introduction of the **value function** v^* of an SMDP, which is a function from $S \times \mathbb{R}$ to \mathbb{R} such that for i in S and t in \mathbb{R} , $v^*(i, t)$ denotes the minimum expected cost (over all strategies) incurred by the player up to time t , by starting in state i at the instant 0:

$$v^*(i, t) := \inf_{f \in \mathbb{F}} \mathbb{E}_i^f \left(\sum_{k=0}^{\widehat{N}_t} \widehat{c}_k \right), \quad (3.5)$$

where \widehat{N}_t is the random variable from H_∞ to \mathbb{N} such that for all h in H_∞

$$\widehat{N}_t(h) := \sup \{ n \in \mathbb{N} \mid \widehat{s}_n(h) = \sum_{k=0}^{n-1} \widehat{t}_k(h) \leq t \},$$

since the cost \widehat{c}_n is incurred at time \widehat{s}_n , therefore we count only costs collected up to epoch \widehat{N}_t included. In the MDP case where all sojourn times equal one, we retrieve for any integer-valued epoch t in \mathbb{N} that $\widehat{N}_t = t$. By convention on empty sums, remark that the equation (3.5) implies $v^*(i, t) = 0$ for all i in S and $t < 0$, so we shall mainly focus on cases where $t \geq 0$.

The value function v^* can be thought of as the most elementary and basic cost criterion that a decision-maker would like to minimize. Alternatively, as announced in Section 2.1, the model can be enriched by several aspects leading to alternative cost criteria. To begin with, recall that we can take discount factors in consideration. Collecting the discount factor $\gamma \geq 0$ at epoch n in \mathbb{N} means that all the future costs $(c_k)_{k \geq n}$ will be multiplied by γ to the eyes of the player. Since in general we take $\gamma \leq 1$, it often amounts to have a preference for some gain in the present than the same gain in the future, and we customarily say that discounted models lead to short-sighted choices (important costs in the future finally do not matter much if shrunk by discount factors). Note that the effect is cumulative, in the sense that the real cost c_n incurred at epoch n will in fact weigh $\gamma^{a_0} \gamma^{a_1} \dots \gamma^{a_{n-1}} c_n$ for the player. On the model of (3.5), we can then introduce the **discounted value function** v_γ^* , such that for all i in S and t in \mathbb{R}

$$v_\gamma^*(i, t) := \inf_{f \in \mathbb{F}} \mathbb{E}_i^f \left(\sum_{k=0}^{\widehat{N}_t} \left(\prod_{\ell=0}^{k-1} \widehat{\gamma}_\ell \right) \widehat{c}_k \right). \quad (3.6)$$

The value function v^* , that we may also call “undiscounted value function” can be seen as a particular case of the discounted value function v_γ^* . For the sake of generality, we will state results

on both undiscounted and discounted value functions until the end of Section 2.6. From Section 2.7 onwards, we will only focus on undiscounted settings.

Another modification of the cost criterion occurs when the player has to pay a terminal cost when the game ends, i.e., when the planning horizon t is exceeded. Indeed, remark that a last action is chosen at epoch \widehat{N}_t (equivalently at physical time $\widehat{s}_{\widehat{N}_t} \leq t$), and therefore a last action cost $\widehat{c}_{\widehat{N}_t}$ is incurred, which makes the player arrive in a last state $\widehat{i}_{\widehat{N}_t+1}$ at physical time $\widehat{s}_{\widehat{N}_t+1} > t$. At this point, no more action can be pulled since the horizon has been crossed, but we can still take into account a last payment $w(\widehat{i}_{\widehat{N}_t+1})$ that represents the cost of terminating the game in state $\widehat{i}_{\widehat{N}_t+1}$. A low (resp. high) value of $w(i)$ for i in S would encourage (resp. discourage) the player to finish his/her ride in i and consequently modify the endgame strategy. It is in fact also reasonable to allow this last payment to depend upon the overshooting $\widehat{s}_{\widehat{N}_t+1} - t$ of the horizon, in order to penalize or reward the fact of exceeding t by far. This leads to the definition of the **value function v_w^* with terminal cost w** for all i in S and t in \mathbb{R} :

$$v_w^*(i, t) := \inf_{f \in \mathbb{F}} \mathbb{E}_i^f \left(\sum_{k=0}^{\widehat{N}_t} \widehat{c}_k + w(\widehat{i}_{\widehat{N}_t+1}, t - \widehat{s}_{\widehat{N}_t+1}) \right) \quad (3.7)$$

where w is a function of $[-t_{\max}, 0) \rightarrow \mathbb{R}^S$ and $t_{\max} := \max_{a \in A} \{t^a\}$ is the greatest sojourn time of the SMDP (remark indeed that $-t_{\max} \leq t - \widehat{s}_{\widehat{N}_t+1} < 0$). Up to slightly changing the definition of \widehat{N}_t , we can convene that when $t < 0$, we have $\widehat{N}_t := -1$ so that v_w^* and w coincide on $[-t_{\max}, 0)$. Criteria of the form (3.7) were already considered by Yushkevich [Yus82].

A final modeling aspect frequently found in stochastic control problems is the possibility to stop the game prematurely (see for instance [Zuc78] in the semi-Markov case), meaning, at some given epoch M in \mathbb{N} , to decide not to pull any action in A_{i_M} , and consequently to stop incurring costs of actions, but in exchange paying the stopping cost (or leaving cost) $R(\widehat{i}_M, t - \widehat{s}_M)$. Here again, it makes sense that this stopping payment depends on the remaining time of the game $t - \widehat{s}_M$, so that it accounts not only for an abandonment penalty in state i but also for a penalty of undershooting with respect to the horizon, and we can then discourage the player to abandon the game too early. The associated total cost criterion is the **value function v_R^* with stopping costs R** , defined for all i in S and t in \mathbb{R} (and denoting by \wedge the minimum operand between two numbers) by

$$v_R^*(i, t) := \inf_{\substack{f \in \mathbb{F} \\ M \in \mathbb{N}}} \mathbb{E}_i^f \left(\sum_{k=0}^{\widehat{N}_t \wedge (M-1)} \widehat{c}_k + R(\widehat{i}_M, t - \widehat{s}_M) \right), \quad (3.8)$$

where R is a function from $[0, \infty)$ to $(\mathbb{R} \cup \{\infty\})^S$. Having $R(i, t) = \infty$ for some state i in S formally means that is the player cannot stop the game in state i if it remains t units of time to play.

As frequently done in the MDP and sometimes in the SMDP literature, the discount factors, the terminal costs and the stopping costs can merrily be combined together. We shall not focus specifically on every one of the eight possible combinations, but we reckon it is worth writing the formal definition of the value function $v_{w,R}^*$ with terminal *and* stopping costs. By definition, the player can stop the game only if the latter is not already terminated, so that we have for all i in S and t in \mathbb{R}

$$v_{w,R}^*(i, t) := \inf_{\substack{f \in \mathbb{F} \\ M \in \mathbb{N}}} \mathbb{E}_i^f \left(\sum_{k=0}^{\widehat{N}_t \wedge (M-1)} \widehat{c}_k + w(\widehat{i}_{\widehat{N}_t+1}, t - \widehat{s}_{\widehat{N}_t+1}) \mathbb{1}_{M > \widehat{N}_t} + R(\widehat{i}_M, t - \widehat{s}_M) \mathbb{1}_{M \leq \widehat{N}_t} \right). \quad (3.9)$$

The Example 3.10 below will in particular illustrate this definition and explain why it is well-founded relatively to our model. We simply mention the most general cost criterion allowed by our framework, the value function $v_{\gamma, w, R}^*$ with discount factors, terminal and stopping costs defined by

$$v_{\gamma, w, R}^*(i, t) := \inf_{\substack{f \in \mathbb{F} \\ M \in \mathbb{N}}} \mathbb{E}_i^f \left(\sum_{k=0}^{\widehat{N}_t \wedge (M-1)} \gamma^{[k-1]} \widehat{c}_k + \gamma^{[\widehat{N}_t]} w(\widehat{i}_{\widehat{N}_t+1}, t - \widehat{s}_{\widehat{N}_t+1}) \mathbb{1}_{M > \widehat{N}_t} + \gamma^{[M]} R(\widehat{i}_M, t - \widehat{s}_M) \mathbb{1}_{M \leq \widehat{N}_t} \right)$$

with for short $\gamma^{[k]} := \gamma^{a_0} \gamma^{a_1} \dots \gamma^{a_k}$. The elementary value function v^* is recovered from $v_{\gamma, w, R}^*$ by letting $\gamma^a = 1$ for all a in A , $w = 0$ and $R = \infty$.

Example 3.10. We illustrate in Figure 3.3 three trials resulting from playing the game described by the SMDP of Figure 3.2. These histories can for instance be obtained by applying the randomized and stationary strategy that selects from each state one of its two playable actions with a fifty-fifty probability. We suppose that the game has to be played in finite horizon $t = 8$ units of time (typically seconds) and always starts in state 3.

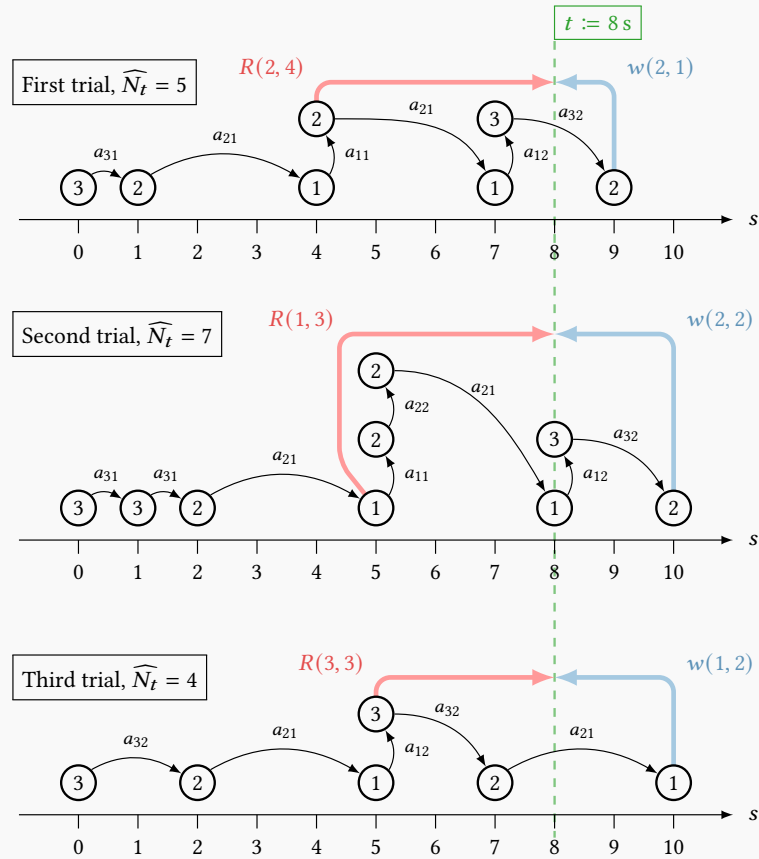


Figure 3.3: Three trials resulting from playing the SMDP of Figure 3.2 in finite-horizon

In the first trial, the player begins by pulling the action $\widehat{a}_0 = a_{31}$ from state $\widehat{i}_0 = 3$, hence incurs a cost of $\widehat{c}_0 = -4$, waits $\widehat{t}_0 = 1$ second and goes to state $\widehat{i}_1 = 2$. He/she activates the

action $\widehat{a}_1 = a_{21}$, pays $\widehat{c}_1 = 1$, waits $\widehat{t}_1 = 3$ and arrives in state $\widehat{i}_2 = 1$ at time $\widehat{s}_2 = \widehat{t}_0 + \widehat{t}_1 = 4$. The game goes on, with sometimes moves achieved in null sojourn times, so that the physical time does not progress (here we have $\widehat{s}_3 = \widehat{s}_2$ because $\widehat{t}_3 = 0$). When the player is in state $\widehat{i}_5 = 3$ at epoch 5, he/she plays the action $\widehat{a}_5 = a_{32}$, incurs a cost $\widehat{c}_5 = 0$ and arrives in state $\widehat{i}_6 = 2$ at the physical time $\widehat{s}_6 = 9$. Because the sixth epoch is the first to strictly exceed the planning horizon, we have $\widehat{N}_t = 5$.

If the terminal cost w is taken into account, the player has to pay the terminal cost $w(2, 1)$, since the game stopped in state 2 with an overshoot of 1 unit of time. The total cost of this trial is $\sum_{k=0}^5 \widehat{c}_k + w(2, 1) = 3 + w(2, 1)$. Alternatively, the player may have preferred to stop playing at epoch $M = 3$ from state $\widehat{i}_3 = 2$ at physical time $\widehat{s}_3 = 4$, i.e., keeping his/her gains $\sum_{k=0}^2 \widehat{c}_k$, not choosing actions any further, but paying the cost $R(2, 4)$ for stopping $t - \widehat{s}_3 = 4$ units of time before the planned end of the game in state 2, which represents a total cost of $-1 + R(2, 4)$ for the trial.

Two other trials of the game are presented. We let the reader check that they are respectively such that $\widehat{N}_t = 7$ and $\widehat{N}_t = 4$. These three sample trials would all intervene (weighed by suitable probabilities) in the computation of $v_{w,R}^*(3, 8)$, the minimum expected cost of playing the game in horizon $t = 8$ and starting from state 3 knowing that there are terminal costs w and stopping costs R . Observe also that for the second trial, there are three visited states at the physical time $s = 5$, owing to two moves in null time.

2.4 EXISTENCE OF THE VALUE FUNCTIONS IN PRESENCE OF NULL SOJOURN TIMES. —

In the SMDP literature, most authors either suppose that the sojourn times of the game are deterministic and all positive, or if they are stochastic that we always have a positive probability ε to draw sojourn times greater than a positive value δ . An originality of our setting is that we have chosen in Section 2.1 to work with deterministic and **nonnegative** sojourn times $(t^a)_{a \in A}$, therefore allowing “always null” times. This automatically questions the well-posedness of the value functions introduced previously, even for the basic definition of v^* in (3.5). Indeed if too many zero sojourn times are encountered for some history, we may accumulate infinitely many costs in a null time period (based on Figure 3.3, imagine states piling up indefinitely for some physical instant). In addition these costs can be either nonnegative or nonpositive, so it may lead to an infinite sum that fails to have a limit. We shall refer to this undesirable phenomenon as the **Zeno behaviour** for SMDPs. Note that it is totally specific to the semi-Markov case, with no equivalent issue in the MDP framework. If we work with v_γ^* in the discounted setting, this effect may in some cases be counterbalanced by the geometric shrinking brought by discount factors when they are smaller than one. On the other hand, discount factors greater than one worsen the problem since the weight of costs incurred in zero time increases. Observe that as far as the well-posedness of the criteria is concerned, the terminal or stopping costs do not cause any particular difficulties, therefore we focus only on v^* and v_γ^* in this section.

To prevent Zeno behaviours, we shall restrict ourselves to the SMDPs verifying adequate non-Zeno conditions. The Assumption A that follows will be used to tackle the undiscounted problem. Its statement only uses the concept of policies introduced in the beginning of Section 2.2 and basic tools from the field of Markov chains (see Section 1.3, and recall that policies induce Markov chains over the state space S).

- **ASSUMPTION A.** For all policies σ in \mathfrak{S} , for all final classes F of the Markov chain induced by σ , there is at least one state i in F with $t^{\sigma(i)} > 0$, i.e. with positive sojourn time attached to the chosen action $\sigma(i)$ from i .

This assumption, milder than imposing a positive sojourn time for all actions of the game, is the same as the one that Schweitzer and Federgruen considered in their study [SF78] of an average-cost criterion for SMDPs. However, to our knowledge, it was never used to show the well-posedness of the finite-horizon value (3.5). The core idea behind Assumption A is that when we apply a policy σ , we end up only visiting the states of some final class of the Markov chain induced by σ , and we want to rule out the possibility that this irreducible final chain is travelled through in zero time. What happens before reaching a final class is actually not relevant for the well-posedness of the value function, since from Proposition 3.7, this phase is expected to last only for a finite number of steps.

Until the end of this subsection, we are going to show that under Assumption A, we have $\sup_{f \in \mathbb{F}} \mathbb{E}_i^f(\widehat{N}_t) < \infty$, so that the quantity $v^*(i, t)$ in (3.5) is finite for all i in S and t in \mathbb{R} (recall that we work in a model where there is a finite number of actions with real costs, hence it suffices to control $\mathbb{E}_i^f(\widehat{N}_t)$ to bound $\mathbb{E}_i^f(\sum_{k=0}^{\widehat{N}_t} \widehat{c}_k)$). The more general case of v_y^* will be addressed in a second phase. The main challenge is to control the **Zeno actions**, that is to say actions in the set $A_Z := \{a \in A \mid t^a = 0\}$ which do not make the physical time progress. Alternatively, the state space S of the SMDP can be partitioned into two sets S_Z and S_{NZ} , respectively the **possibly Zeno states**, so that $S_Z := \{i \in S \mid \exists a \in A_i, t^a = 0\}$ and the **necessarily non-Zeno states** $S_{NZ} := \{i \in S \mid \forall a \in A_i, t^a > 0\}$. Intuitively, for Zeno behaviours to be prevented, it suffices that states of S_{NZ} are “positive recurrent” (in a sense similar to Section 1.1, meaning that they have finite average return times).

Instead of tackling the problem with probabilistic tools only, we choose to develop an approach that stays close from the initial SMDP, with a couple of intermediary results that although sometimes technical do highlight the difficulties of handling null sojourn times and in this light shall be reused later. Our proof begins by considering an auxiliary undiscounted Markov Decision Process (in which neither the time nor the discount factors are taken into account) with larger state space and equipped with a **layered** structure. We formally introduce it in the following paragraphs, and illustrate our construction in Figure 3.4, by representing the layered MDP built from the SMDP of Figure 3.2.

Let K be a positive integer, and define for all k in $\{1, 2, \dots, K\}$ the states of the **k -th layer**:

$$S^{(k)} := \left\{ i_1^{(k)}, i_2^{(k)}, \dots, i_{|S|}^{(k)} \right\}.$$

We may think of all these layers as *exact copies* of the initial state space S of our SMDP. Consequently, we have natural projections $(\pi_k: S^{(k)} \rightarrow S)_{1 \leq k \leq K}$ mapping a state of any layer to the state of S from which it is derived (we shall drop the index k and only denote these projections by π when nonambiguous). We also introduce a **cemetery** state denoted by \perp , so that the complete state space \widetilde{S} of our new MDP writes $\widetilde{S} := \left(\bigcup_{k=1}^K S^{(k)} \right) \cup \{\perp\}$. The decomposition of S between possibly Zeno states and necessarily non-Zeno states carries over to the layers $(S^{(k)})_{1 \leq k \leq K}$ and we shall later use the partitions $S^{(k)} = S_Z^{(k)} \cup S_{NZ}^{(k)}$ for all k in $\{1, 2, \dots, K\}$.

We impose that from the cemetery state, only one action a_\perp can be pulled (and then $\widetilde{A}_\perp = \{a_\perp\}$). From any other state $i^{(k)}$ within the k -th layer, the player can pick exactly as many actions as in $A_{\pi(i^{(k)})}$:

$$\widetilde{A}_{i^{(k)}} := \left\{ a_1^{(k)}, a_2^{(k)}, \dots, a_{|A_{\pi(i^{(k)})}|}^{(k)} \right\}.$$

Similarly to the states, the actions of the original SMDP are here copied K times, so that each layer has the *same* state-action structure than the SMDP. Once again there are natural projections

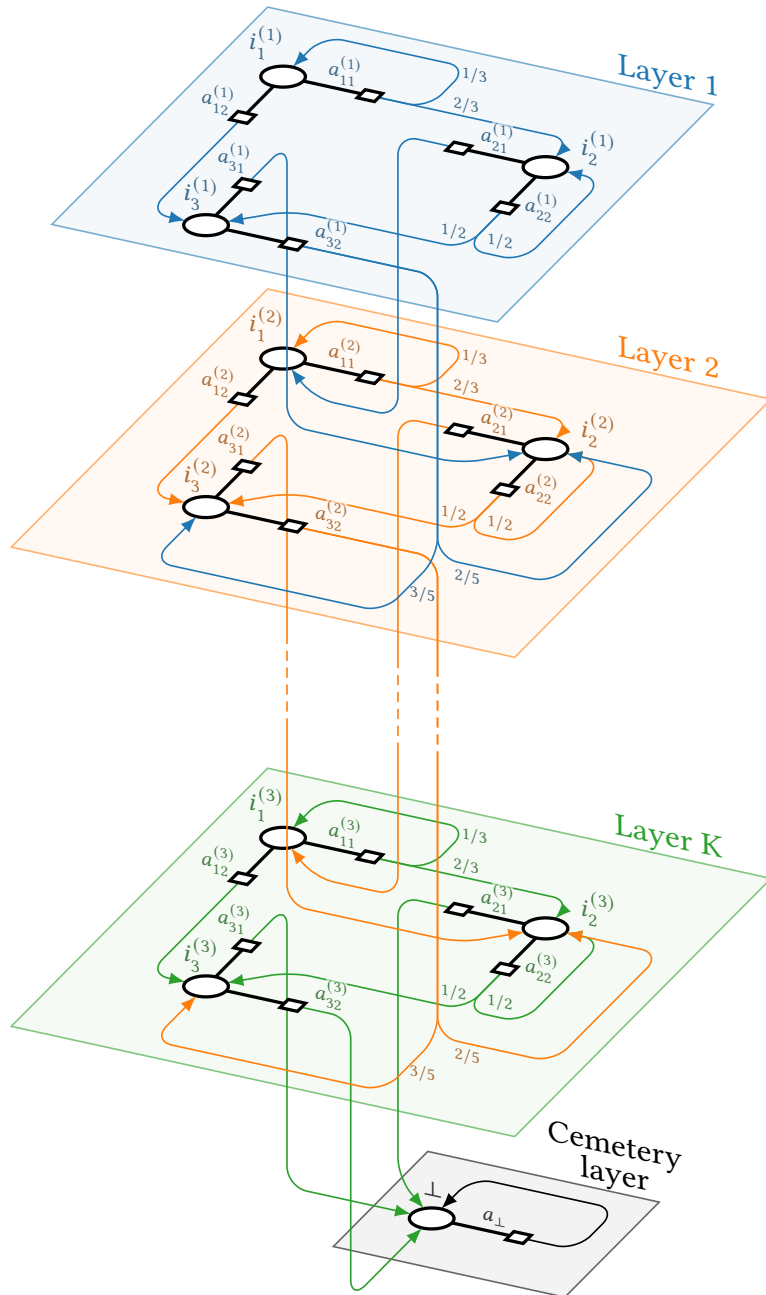


Figure 3.4: The layered MDP corresponding to the SMDP of Figure 3.2

from these newly created sets of playable actions towards those of the original SMDP, and for the sake of simplicity we still denote these maps by π .

For all $i^{(k)}$ in $S^{(k)}$, the transition probabilities $(p_j^a)_{j \in \tilde{S}}$ governing the evolution towards other states of \tilde{S} after playing action a in $\tilde{A}_{i^{(k)}}$ are defined by

$$\text{if } t^{\pi(a)} > 0, \text{ then } p_j^a := \begin{cases} p_{\pi(j)}^{\pi(a)} & \text{if } k < K \text{ and } j \in S^{(k+1)} \\ 1 & \text{if } k = K \text{ and } j = \perp \\ 0 & \text{otherwise} \end{cases}$$

$$\text{if } t^{\pi(a)} = 0, \text{ then } p_j^a := \begin{cases} p_{\pi(j)}^{\pi(a)} & \text{if } j \in S^{(k)} \\ 0 & \text{otherwise} \end{cases}$$

and finally $p_{\perp}^{a_{\perp}} = 1$ (the cemetery state is absorbing).

In other words, all actions in the new game derived from Zeno actions in the original SMDP make the player stay in the current layer, while picking some new action derived from a non-Zeno initial action forces to dive into the next layer (and to join the cemetery state if already in the K -th layer). We attach unit rewards $r^a := 1$ to all the actions a in $\tilde{A} \setminus \{a_{\perp}\}$ of this new MDP except in the cemetery state \perp where we impose that no more rewards can be earned, hence $r^{a_{\perp}} := 0$.

Observe that the admissible trajectories associated with this new layered game correspond to histories of the initial SMDP, until potentially hitting \perp after which a single action can be chosen, so the strategies of the original game precisely induce all the strategies of the new one. Since any f in \mathbb{F} and initial state i in \tilde{S} define a random process $(i, \widehat{a}_0, \widehat{r}_0, \widehat{i}_1, \widehat{a}_1, \widehat{r}_1, \dots)$, we can introduce the vector $N^{[K]}$ of $\mathbb{R}_{\geq 0}^{\tilde{S}}$ where

$$\text{for all } i \text{ in } \tilde{S}, \quad N^{[K]}(i) := \sup_{f \in \mathbb{F}} \mathbb{E}_i^f \left(\sum_{k=0}^{\infty} \widehat{r}_k \right). \quad (3.10)$$

Because $\widehat{r}_k = 1$ if and only if we are not in the cemetery layer, $N^{[K]}(i)$ can be interpreted as the maximum (over all strategies) expected number of moves to play before reaching \perp by starting in state i of \tilde{S} . If there is a strategy for which \perp is never attained, it implies $N^{[K]}(i) = \infty$. The following lemma explains how $N^{[K]}$ is tied to our problem of controlling the quantity $\mathbb{E}_i^f(\widehat{N}_t)$ for the original SMDP. We denote by t_{\min} the smallest positive sojourn time of this SMDP, that is, $t_{\min} := \min\{t^a \mid a \in A, t^a > 0\}$.

► **LEMMA 3.11.** *Let $t \geq t_{\min}$ and denote $K := \lceil t/t_{\min} \rceil$. We have for all i in S that*

$$N^{[1]}(i^{(1)}) \leq \sup_{f \in \mathbb{F}} \mathbb{E}_i^f(\widehat{N}_t) \leq N^{[K]}(i^{(1)}),$$

where the second inequality actually holds for all $t \geq 0$.

..... *Proof.* From what precedes, we understand that for all i in S and k in $\{1, 2, \dots, K\}$, the quantity $N^{[K]}(i^{(k)})$ gives in the original SMDP the maximum expected number of moves starting from i and before performing $K + 1 - k$ non-Zeno moves.

..... In particular for any admissible trajectory of the process, the number \widehat{N}_t of moves performed in this original SMDP before exceeding a total physical elapsed time t must be less than the number of moves before pulling K non-Zeno actions (costing a total time of at least $Kt_{\min} \geq t$ units of time), hence the second inequality. On the other hand if $t \geq t_{\min}$, it requires to pull at least one non-Zeno action before possibly exceeding a total physical elapsed time of t , hence the first inequality. ♠

Let us denote by $\tilde{\mathfrak{S}} := \prod_{i \in \tilde{S}} \tilde{A}_i$ the set of policies in our new layered MDP. Similarly to the original SMDP, for all policies $\tilde{\sigma}$, the matrix with entries $(p_j^{\tilde{\sigma}(i)})_{i,j \in \tilde{S}}$ is denoted by $P^{\tilde{\sigma}}$. It is interesting to remark that up to reordering the states of \tilde{S} , any $P^{\tilde{\sigma}}$ has the form

$$\begin{pmatrix} P_{Z,Z}^{(1)} & P_{Z,NZ}^{(1)} & Q_{Z,Z}^{(1)} & Q_{Z,NZ}^{(1)} & \cdots & 0 & 0 & 0 & S_Z^{(1)} \\ 0 & 0 & P_{NZ,Z}^{(1)} & P_{NZ,NZ}^{(1)} & \cdots & 0 & 0 & 0 & S_{NZ}^{(1)} \\ 0 & 0 & P_{Z,Z}^{(2)} & P_{Z,NZ}^{(2)} & \cdots & 0 & 0 & 0 & S_Z^{(2)} \\ 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & S_{NZ}^{(2)} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & P_{Z,Z}^{(K)} & P_{Z,NZ}^{(K)} & Q_{Z,\perp}^{(K)} & S_Z^{(K)} \\ 0 & 0 & 0 & 0 & \cdots & 0 & 0 & \mathbf{1} & S_{NZ}^{(K)} \\ 0 & 0 & 0 & 0 & \cdots & 0 & 0 & \mathbf{1} & \perp \end{pmatrix}$$

$$\begin{matrix} S_Z^{(1)} & S_{NZ}^{(1)} & S_Z^{(2)} & S_{NZ}^{(2)} & \cdots & S_Z^{(K)} & S_{NZ}^{(K)} & \perp \end{matrix}$$

where all the blocks depend on $\tilde{\sigma}$. Indeed, we already justified that playing actions from necessarily non-Zeno states forces to join the next layer, while playing actions from possibly Zeno states can either result in staying on the current layer (if a Zeno action is chosen) or going to the next one (otherwise).

To show finiteness of the $(N^{[K]}(i^{(1)}))_{i \in S}$, let us first introduce the self-map B of $\mathbb{R}^{\tilde{S} \setminus \{\perp\}}$ by

$$\text{for all } x \text{ in } \mathbb{R}^{\tilde{S} \setminus \{\perp\}}, \quad B(x) := \mathbf{1} + \max_{\tilde{\sigma} \in \tilde{\mathfrak{S}}} \left\{ P^{\tilde{\sigma}} \Big|_{(\tilde{S} \setminus \{\perp\}) \times (\tilde{S} \setminus \{\perp\})} x \right\}, \quad (3.11)$$

where the maximum is taken componentwise. We have the following:

► **LEMMA 3.12.** *The vector $N^{[K]}$ with its \perp -coordinate dropped is a fixed-point of B :*

$$N^{[K]} \Big|_{\tilde{S} \setminus \{\perp\}} = B(N^{[K]} \Big|_{\tilde{S} \setminus \{\perp\}}).$$

Proof. We recognize in (3.10) an expected total reward as defined in [Put14, Section 7.1.1]. All the rewards are nonnegative so [Put14, Theorem 7.1.3] applies and we know that $N^{[K]}$ verifies the optimality equation

$$N^{[K]} = \max_{\tilde{\sigma} \in \tilde{\mathfrak{S}}} \{ r^{\tilde{\sigma}} + P^{\tilde{\sigma}} N^{[K]} \}, \quad (3.12)$$

Remark that the vector $r^{\tilde{\sigma}}$ does not actually depend on $\tilde{\sigma} \in \tilde{\mathfrak{S}}$ (it has only unit components except on the last coordinate where it is null). By the definition (3.10), we know that $N^{[K]}(\perp) = 0$. This last fact gives $(P^{\tilde{\sigma}} N^{[K]}) \Big|_{\tilde{S} \setminus \{\perp\}} = P^{\tilde{\sigma}} \Big|_{(\tilde{S} \setminus \{\perp\}) \times (\tilde{S} \setminus \{\perp\})} N^{[K]} \Big|_{\tilde{S} \setminus \{\perp\}}$, hence the result of the lemma using (3.12). \clubsuit

The Lemma 3.12 does not show that $N^{[K]}$ has finite entries, however it prompts us to focus on the properties of the operator B instead. The latter belongs to the wider class of piecewise-affine maps, and the following proposition, resulting from operator-based and nonlinear Perron–Frobenius

techniques of [AGN11, AGQS19] provides a sufficient condition for such maps to be contractions in weighted sup-norm, and therefore to admit unique fixed-points. Recall that for u in $\mathbb{R}_{>0}^n$, the application $\|\cdot\|_u: x \mapsto \sup_{1 \leq i \leq n} |x_i/u_i|$ is a norm over \mathbb{R}^n , referred to as the weighted sup-norm associated with u . Recall also that for $0 \leq \lambda < 1$ that a self-map F over \mathbb{R}^n is a λ -contraction relatively to the norm $\|\cdot\|$ if for all x and y in \mathbb{R}^n , we have $\|F(x) - F(y)\| \leq \lambda\|x - y\|$. Weighted sup-norm contractions play a central role in dynamic programming, see for instance the review of Bertsekas [Ber12c].

- **PROPOSITION 3.13.** *Let n in \mathbb{N}^* and let \mathcal{E} be a finite set. Let $(C^{(k)})_{k \in \mathcal{E}}$ be matrices of $\mathbb{R}_{\geq 0}^{n \times n}$ and $(d^{(k)})_{k \in \mathcal{E}}$ be vectors of \mathbb{R}^n . We introduce the following two piecewise-affine self-maps of \mathbb{R}^n :*

$$\text{for all } x \text{ in } \mathbb{R}^n, \quad C^{\max}(x) := \max_{k \in \mathcal{E}} \{C^{(k)}x + d^{(k)}\} \quad \text{and} \quad C^{\min}(x) := \min_{k \in \mathcal{E}} \{C^{(k)}x + d^{(k)}\},$$

where the maximum and the minimum operate componentwise. For (k_1, \dots, k_n) in \mathcal{E}^n , we denote by $C^{(k_1 \dots k_n)}$ the $(n \times n)$ -matrix whose i -th row is given by the i -th row of $C^{(k_i)}$.

If all matrices $(C^{(k_1 \dots k_n)})_{(k_1, \dots, k_n) \in \mathcal{E}^n}$ have a spectral radius strictly less than one, then there exists λ in $[0, 1)$, and u in $\mathbb{R}_{>0}^n$ such that C^{\max} and C^{\min} are λ -contracting in weight sup-norm $\|\cdot\|_u$, and thus both C^{\max} and C^{\min} admit a (unique) fixed-point in \mathbb{R}^n .

Proof. We suppose that for all (k_1, \dots, k_n) in \mathcal{E}^n , we have $\rho(C^{(k_1 \dots k_n)}) < 1$ and we want to show that C^{\max} and C^{\min} admit a fixed-point.

We first introduce the homogeneous part \check{C}^{\max} (resp. \check{C}^{\min}) of operator C^{\max} (resp. C^{\min}) defined by $\check{C}^{\max}(x) := \max_{k \in \mathcal{E}} \{C^{(k)}x\}$ (resp. $\check{C}^{\min}(x) := \min_{k \in \mathcal{E}} \{C^{(k)}x\}$) for x in \mathbb{R}^n . We want to apply [AGQS19, Theorem 1], which claims that for $\lambda \geq 0$ and $u \in \mathbb{R}_{>0}^n$, the order-preserving operator C^{\max} (resp. C^{\min}) of \mathbb{R}^n is λ -contracting in weighted sup-norm $\|\cdot\|_u$ if and only if $\check{C}^{\max}(u) \leq \lambda u$ (resp. $\check{C}^{\min}(u) \leq \lambda u$). Actually, since for all x in $\mathbb{R}_{\geq 0}^n$, we have $\mathbf{0} \leq \check{C}^{\min}(x) \leq \check{C}^{\max}(x)$, we shall just focus on \check{C}^{\max} in what follows.

Although it is not linear (as a supremum of linear maps), we may define a notion of spectral radius for \check{C}^{\max} by

$$\rho(\check{C}^{\max}) := \sup \{ \lambda \geq 0 \mid \exists x \in \mathbb{R}_{\geq 0}^n \setminus \{0\}, \check{C}^{\max}(x) = \lambda x \}$$

Since all the maps $x \mapsto C^{(k)}x$ are order-preserving on $\mathbb{R}_{\geq 0}^n$, we obtain according to [AGN11, Proposition 8.1] that

$$\rho(\check{C}^{\max}) = \max_{(k_1, \dots, k_n) \in \mathcal{E}^n} \rho(C^{(k_1 \dots k_n)}).$$

(indeed, to satisfy the selection property needed in [AGN11, Proposition 8.1], we must here take into account all the matrices with rows shuffled).

Therefore, our initial assumption entails that $\rho(\check{C}^{\max}) < 1$. However, from the generalized Collatz–Wielandt theorem (see [AGN11, Theorem 1.1]), it also holds that

$$\rho(\check{C}^{\max}) = \inf \{ \lambda > 0 \mid \exists x \in \mathbb{R}_{>0}^n, \check{C}^{\max}(x) \leq \lambda x \},$$

so by taking a specific λ in $[\rho(\check{C}^{\max}), 1)$, we know that there exists a positive vector u such that $\mathbf{0} \leq \check{C}^{\max}(u) \leq \lambda u$. From our previous bound on \check{C}^{\min} and by applying [AGQS19, Theorem 1], we conclude that the (no longer homogeneous) operators C^{\max} and C^{\min} are λ -contracting in the weighted sup-norm relative to u with $0 \leq \lambda < 1$. It immediately follows that both of them admit a (unique) fixed-point in \mathbb{R}^n by usual iterating schemes on contraction mappings. \clubsuit

Building on Proposition 3.13, we can now connect the properties of B together with the transition matrices of our layered MDP and of our initial SMDP, to finally go back to the Assumption A. It stands out here that what matters to avoid Zeno behaviours is to have contraction properties of the policies' transition matrices on the possibly Zeno states, while no restrictions are needed on the necessarily non-Zeno states.

► **PROPOSITION 3.14.** *The following are equivalent:*

- (1) *the operator B admits a (unique) finite positive fixed-point,*
- (2) *for all $\tilde{\sigma}$ in $\tilde{\mathfrak{S}}$, $\rho(P^{\tilde{\sigma}}|_{(\tilde{S}\setminus\{\perp\})\times(\tilde{S}\setminus\{\perp\})}) < 1$,*
- (3) *for all σ in \mathfrak{S} , $\rho(P^\sigma|_{S_Z\times S_Z}) < 1$,*
- (4) *the conditions of Assumption A are fulfilled.*

Proof. The implication (1) \implies (2) is an application of Proposition 3.13 to the particular case of the operator B defined in Lemma 3.12, observing that working with policies causes the collection $(P^{\tilde{\sigma}})_{\tilde{\sigma}\in\tilde{\mathfrak{S}}}$ to already contain all the matrices that can be constructed by the row-selection procedure. The unique fixed-point x of B must first lie in $\mathbb{R}_{\geq 0}^n$ since it can be obtained as the limit of a sequence of nonnegative vectors of type $(B^n(\mathbf{0}))_{n\in\mathbb{N}}$. But more precisely, x must be in $\mathbb{R}_{> 0}^n$, since there exists a policy $\tilde{\sigma} \in \tilde{\mathfrak{S}}$ such that $x = \mathbf{1} + P^{\tilde{\sigma}}|_{(\tilde{S}\setminus\{\perp\})\times(\tilde{S}\setminus\{\perp\})}x$.

- ◊ We prove (2) \implies (1) by supposing that B admits a finite positive fixed-point x and assuming by contradiction that there is a particular $\tilde{\sigma}$ in $\tilde{\mathfrak{S}}$ such that $\rho(P^{\tilde{\sigma}}|_{(\tilde{S}\setminus\{\perp\})\times(\tilde{S}\setminus\{\perp\})}) = 1$. By Lemma 3.5, the latter matrix contains a stochastic square submatrix, which then admits an invariant measure (i.e., a nonnegative left eigenvector for the eigenvalue 1). Up to completing this invariant measure with zeroes, we can construct a nonnegative left eigenvector $\mu \in \mathbb{R}_{\geq 0}^{\tilde{S}\setminus\{\perp\}}$ for the value 1 for $P^{\tilde{\sigma}}|_{(\tilde{S}\setminus\{\perp\})\times(\tilde{S}\setminus\{\perp\})}$. On the other hand, the fixed-point property guarantees that

$$x \geq P^{\tilde{\sigma}}|_{(\tilde{S}\setminus\{\perp\})\times(\tilde{S}\setminus\{\perp\})}x + \mathbf{1}.$$

Taking the dot-product with the nonnegative vector μ on both sides leads to $\mu^T x \geq \mu^T x + \mu^T \mathbf{1}$, and thus $\mathbf{0} \geq \mathbf{1}$ hence the contradiction.

- ◊ To obtain (2) \iff (3), we remark for example by computing the characteristic polynomial of the big triangular matrix written before the Lemma 3.12 that for $\tilde{\sigma} \in \tilde{\mathfrak{S}}$, we have by denoting $(\sigma_1, \dots, \sigma_K) \in \mathfrak{S}^K$ the restrictions of $\tilde{\sigma}$ to layers $(S^{(k)})_{1 \leq k \leq K}$ that

$$\rho(P^{\tilde{\sigma}}|_{(\tilde{S}\setminus\{\perp\})\times(\tilde{S}\setminus\{\perp\})}) = \max_{1 \leq k \leq K} \rho(P_{Z,Z}^{(\sigma_k)}) = \max_{1 \leq k \leq K} \rho(P^{\sigma_k}|_{S_Z \times S_Z}),$$

and as a result

$$\max_{\tilde{\sigma} \in \tilde{\mathfrak{S}}} \rho(P^{\tilde{\sigma}}|_{(\tilde{S}\setminus\{\perp\})\times(\tilde{S}\setminus\{\perp\})}) = \max_{\sigma \in \mathfrak{S}} \rho(P^\sigma|_{S_Z \times S_Z}).$$

- ◊ The equivalence (3) \iff (4) follows from Lemma 3.5 (with $J = S_Z$) since by negating its claim for all policies we have

$$\forall \sigma \in \mathfrak{S}, \rho(P^\sigma|_{S_Z \times S_Z}) < 1 \iff \forall \sigma \in \mathfrak{S}, \forall F \in \mathcal{F}(\sigma), F \not\subseteq S_Z,$$

the last assertion corresponding exactly to any final class of any policy having at least one positive sojourn time, that is to say, Assumption A. \mathfrak{E}

The previous results allow us to conclude on the well-posedness of v^* in (3.5). It also covers the case of the undiscounted value functions with terminal or stopping costs defined with (3.7), (3.8) and (3.9). Interestingly, the following theorem shows that Assumption A cannot be weakened.

► **THEOREM 3.15.** *The undiscounted value function v^* is well-defined if and only if Assumption A holds.*

Proof. Suppose that Assumption A holds. Then, by Proposition 3.14, the operator B defined by (3.11) admits a unique positive fixed-point, which is finite. However, Lemma 3.12 tells us that this fixed-point coincides with $N^{[K]}|_{\tilde{S}\setminus\{\perp\}}$. It follows that $N^{[K]}$ is finite for all K in \mathbb{N}^* . The Lemma 3.11 therefore guarantees for all i in S and $t \geq 0$ that $\sup_{f \in \mathbb{F}} \mathbb{E}_i^f(\widehat{N}_t) < \infty$.

- ◊ Conversely, if Assumption A is not verified, we reuse Proposition 3.14 to conclude that $N^{[1]}$ must have non-finite value on some coordinate i in S . Lemma 3.11 entails that for the same i and some $t \geq t_{\min}$ we have $\sup_{f \in \mathbb{F}} \mathbb{E}_i^f(\widehat{N}_t) = \infty$. \mathfrak{E}

When working with the discounted value function v_Y^* , the quantity that should stay finite is no longer $\sup_{f \in \mathbb{F}} \mathbb{E}_i^f(\widehat{N}_t)$ but $\sup_{f \in \mathbb{F}} \mathbb{E}_i^f(\sum_{k=0}^{\widehat{N}_t} \prod_{\ell=0}^{k-1} \widehat{\gamma}_k)$. If some of the discount factors $(\gamma^a)_{a \in A}$ are bigger than one, this is not implied by Assumption A. Building on the complete line of proof of Theorem 3.15, we can thus propose a stronger assumption for the general discounted case. For a policy σ in \mathfrak{S} , we will use the notation $\Gamma^\sigma := \text{diag}((\gamma^{\sigma(i)})_{i \in S})$.

- **ASSUMPTION B.** For all policies σ in \mathfrak{S} , we have $\rho((\Gamma^\sigma P^\sigma)|_{S_Z \times S_Z}) < 1$.

The Assumption B cannot be interpreted simply in terms of sojourn times in the same way as Assumption A, because it deals with the mixed effect of these sojourn times *and* the discount factors. However, it stills constrains the choice of the parameters *only* on possibly Zero states, that is to say S_Z . We have a result equivalent to Theorem 3.15 in the discounted case:

- **THEOREM 3.16.** The discounted value function v_Y^* is well-defined if and only if Assumption B holds.

Proof. It suffices to adapt the whole reasoning carried out for v^* . We still work with an auxiliary layered MDP but this time with discount factors derived from the original SMDP. The criterion (3.10) has to be replaced by

$$\text{for all } i \in \widetilde{S}, \quad N^{[K]}(i) := \sup_{f \in \mathbb{F}} \mathbb{E}_i^f \left(\sum_{k=0}^{\infty} \left(\prod_{\ell=0}^{k-1} \widehat{\gamma}_k \right) \widehat{r}_k \right),$$

in which case (3.12) becomes $N^{[K]} = \max_{\widetilde{\sigma} \in \widetilde{\mathfrak{S}}} \{r^{\widetilde{\sigma}} + (\Gamma^{\widetilde{\sigma}} P^{\widetilde{\sigma}}) N^{[K]}\}$. The contraction properties of the $((\Gamma^\sigma P^\sigma)|_{S_Z \times S_Z})_{\sigma \in \mathfrak{S}}$ are then used in lieu of those of the $(P^\sigma|_{S_Z \times S_Z})_{\sigma \in \mathfrak{S}}$ before. \clubsuit

2.5 COMPUTING THE VALUE FUNCTION. — The definitions of type (3.5) and (3.6) are of little use in practice. On the contrary, and as customarily done in dynamic programming, it is fruitful to exploit a set of recursive **optimality equations** that the value functions satisfy. Although a great part of the SMDP literature deals with infinite-horizon problems, such equations for the finite-horizon case were for instance stated in [Jew63, DC64, Sch80, Yus82] for varied models. These authors do not establish the optimality equations in the particular presence of deterministic zero sojourn times, and it will be the object of Corollary 3.19 to show that the undiscounted value function v^* satisfies the equations (DP):

$$\forall t \geq 0, \quad \forall i \in S, \quad v(i, t) = \min_{a \in A_i} \left\{ c^a + \sum_{j \in S} p_j^a v(j, t - t^a) \right\}, \quad (\text{DP})$$

and that the discounted value function v_Y^* satisfies the equations (DP- γ):

$$\forall t \geq 0, \quad \forall i \in S, \quad v_Y(i, t) = \min_{a \in A_i} \left\{ c^a + \gamma^a \sum_{j \in S} p_j^a v_Y(j, t - t^a) \right\}. \quad (\text{DP-}\gamma)$$

The dynamic programming equations (DP) tell us that the minimum expected cost $v^*(\cdot, t)$ that we can incur in horizon t is related to the same quantity at prior horizons $(v^*(\cdot, t - t^a))_{a \in A}$ in a “one-step-of-the-game” fashion. Indeed, for i in S , $v^*(i, t)$ is given by choosing among all the actions of A_i an action a that minimizes the immediate cost c^a on top of a “terminal” cost that is the expectation $\sum_{j \in S} p_j^a v^*(j, t - t^a)$ of the cost subsequent to the choice of a . The same goes for v_Y^* where the terminal part gets discounted by γ^a .

In the (undiscounted) Markov decision process setting ($t^a = 1$ for all a in A), the optimality equations are conveniently written under vector form and using policies as:

$$\forall t \in \mathbb{N}, \quad v(t) = \min_{\sigma \in \mathfrak{S}} \{c^\sigma + P^\sigma v(t - 1)\}, \quad (\text{DP-MDP})$$

where $v(t) := (v(i, t))_{i \in S}$. On the contrary, the distinct sojourn times considered in the semi-Markov case lead to some asynchronicity phenomena, which prevents us from writing such a compact formula.

Optimality equations can be applied recursively in order to “go back in time” as long as nonnegative physical instants are involved (note indeed that these equations are only valid for $t \geq 0$). It is instructive to study how the nature of the sojourn times $(t^a)_{a \in A}$ affects the computation effort needed to obtain a solution v of (DP). For example, the equation (DP-MDP) shows for MDPs that for all t in \mathbb{N} , $v(t)$ is entirely determined by $v(t-1)$. The vector $v(t)$ can thus be computed from $v(-1)$ by applying $t+1$ times the so-called **Bellman operator** T , which is the self-map of \mathbb{R}^S defined by

$$\forall x \in \mathbb{R}^S, \quad T(x) := \min_{\sigma \in \mathfrak{E}} \{c^\sigma + P^\sigma x\}. \quad (3.13)$$

According to the discussion following the introduction of the criterion v_w^* in Section 2.3, remark that the vector $v(-1)$ coincides with a terminal cost vector, although it is here used as the initial condition for applying the dynamic programming method and recover $v(t) = T^{t+1}(v(-1))$ (as usual with this type of methods, see for instance [Ber12b]). This computational technique is known as the **value iteration algorithm**.

In the case where the sojourn times take positive integer values, the decision epochs still lie in \mathbb{N} , but *stricto sensu*, we leave the MDP setting owing to the non-uniformity of the $(t^a)_{a \in A}$. However, we can easily reduce to it by considering an augmented vector of the form $\tilde{v}(t) := (v(t), v(t-1), v(t-2), \dots, v(t-t_{\max}+1))$, where t_{\max} denotes the maximal sojourn time. This way, $\tilde{v}(t)$ is determined by $\tilde{v}(t-1)$, and we can apply a MDP-type scheme in larger dimension; in particular we shall eventually need to know the terminal costs $v(-1), v(-2), \dots, v(-t_{\max})$, corresponding to overshooting a planning horizon for respectively 1, 2, ..., t_{\max} units of time. Going a step further, if the sojourn times took positive rational values and D denoted the least common multiple of denominators of the $(t^a)_{a \in A}$, we could still apply the MDP iterative technique to a large vector $\tilde{v}(t) = (v(t-k/D))_{0 \leq k < Dt_{\max}}$ by steps of D^{-1} for all t in $D^{-1}\mathbb{N}$, and any value of $v(t)$ could be determined by the “sole” knowledge of the $(v(-k/D))_{1 \leq k \leq Dt_{\max}}$.

When sojourn times take arbitrary nonnegative real values though, we find ourselves in a “proper” semi-Markov framework where we cannot resolve to finite dimension settings of the previous forms. Indeed, computing v over an interval of $\mathbb{R}_{\geq 0}$ using (DP) iteratively will in general require the knowledge of v over the whole interval $[-t_{\max}, 0)$ as an initial condition of the method; even computing $v(t)$ only for all admissible decision epochs t requires to know v over $[-t_{\max}, 0) \cap \sum_{a \in A} t^a \mathbb{Z}$, that is in general a dense set, and quite equivalent to $[-t_{\max}, 0)$ itself if we impose elementary regularity conditions on v . This motivates the introduction of an infinite dimensional space that is more natural to work with, namely the space \mathcal{V} of bounded $|S|$ -dimensional vector functions over $[-t_{\max}, 0)$. For the sake of simplicity we will also suppose that these functions are piecewise-continuous. These functions shall act as initial conditions for computing a solution of (DP) (or (DP- γ)) iteratively, but they are in fact terminal cost functions as discussed in Section 2.3.

The following proposition shows that, given such an initial condition in \mathcal{V} , there is a unique “trajectory” satisfying the dynamic programming equations (either undiscounted or discounted). The challenge again comes from the presence of null sojourn times in our model. Indeed, the equations (DP) and (DP- γ) that are usually explicit turn out to be implicit if for some i in S we have a in A_i such that $t^a = 0$, meaning that components of $v(t)$ appear on both sides of some

equation. We tackle this issue by building on contraction properties (in weighted sup-norm) similar to those used to demonstrate Theorem 3.15.

Recall from Chapter 1 that a function of a real variable is càdlàg if it is right-continuous and has left-limits everywhere.

- **PROPOSITION 3.17.** *Let v^0 be an initial condition in \mathcal{V} . Under Assumption A, there is a unique function v defined on $[-t_{\max}, \infty)$ with values in \mathbb{R}^S satisfying the dynamic programming equation (DP) and $v(t) = v^0(t)$ for all t in $[-t_{\max}, 0)$. In addition, if v^0 is càdlàg (resp. piecewise-constant, piecewise-affine), then v is càdlàg (resp. piecewise-constant, piecewise-affine). The same is true replacing (DP) by (DP- γ) and Assumption A by Assumption B.*

Proof. We show the result for the discounted value function v_γ^* under Assumption B, since it also proves the result for v^* by setting the discount factors to one.

Our line proof is it to show correctness of the following inductive scheme: “let $t_0 \geq 0$, if v is known on $[-t_{\max}, t_0)$, then v is uniquely determined on $[t_0, t_0 + t_{\min})$ ”, since it suffices to apply this reasoning step by step for $t_0 \in t_{\min}\mathbb{N}$ to prove the proposition.

- ◇ Suppose then that $t_0 \geq 0$ is fixed and v is known on $[-t_{\max}, t_0)$. Let $t \in [t_0, t_0 + t_{\min})$ be a fixed instant as well. We reuse the partition of the state space $S = S_Z \uplus S_{NZ}$ introduced in the Section 2.4. We shall also denote $A_{i,Z} := \{a \in A_i \mid t^a = 0\} \neq \emptyset$ the set of Zeno actions playable from i if i is in S_Z .

If $i \in S_{NZ}$, $v(i, t)$ is explicitly determined by the knowledge of v on $[-t_{\max}, t_0)$ and the (DP- γ) equation

$$v(i, t) = \min_{a \in A_i} \left\{ c^a + \gamma^a \sum_{j \in S} p_j^a v(j, t - t^a) \right\}, \quad (3.14)$$

for by definition, all the $(t^a)_{a \in A_i}$ are positive and greater than t_{\min} , consequently the $(t - t^a)_{a \in A_i}$ lie in $[-t_{\max} + t_0, t_0)$. However, the restriction $v_Z(t) := (v(i, t))_{i \in S_Z}$ of v to states of S_Z is solution of the implicit equation

$$v_Z(t) = \varphi_t(v_Z(t)) := \min(C_{Z,t}(v_Z(t)), b_t),$$

where we have

$$\text{for all } i \in S_Z, \quad [b_t]_i := \min_{a \in A_i \setminus A_{i,Z}} \left\{ c^a + \gamma^a \sum_{j \in S} p_j^a v(j, t - t^a) \right\},$$

$$\text{for all } a \in \bigcup_{i \in S_Z} A_{i,Z}, \quad c_t^a := c^a + \gamma^a \sum_{j \in S_{NZ}} p_j^a v(j, t), \quad \text{and}$$

$$\text{for all } i \in S_Z, \quad [C_{Z,t}(x)]_i := \min_{a \in A_{i,Z}} \left\{ c_t^a + \gamma^a \sum_{j \in S_Z} p_j^a x_j \right\}.$$

Observe that b_t actually depends upon t only via positive time-delays, and so do the (c_t^a) because we can reinject expressions previously computed for states of S_{NZ} . Therefore, they are explicitly determined by the knowledge of v on $[-t_{\max}, t_0)$ as well. We keep the index “ t ” nevertheless to emphasize that we shall determine v_Z pointwise on $[t_0, t_0 + t_{\min})$.

Observe that $C_{Z,t}$ can be written in the compact form

$$C_{Z,t}(x) = \min_{\sigma \in \mathfrak{S}_Z} \left\{ c_t^\sigma|_{S_Z} + (\Gamma^\sigma P^\sigma)|_{S_Z \times S_Z} x \right\},$$

where $\mathfrak{S}_Z := \prod_{i \in S_Z} A_{i,Z}$ and by an abuse of notation, for σ in \mathfrak{S}_Z , $c_t^\sigma|_{S_Z}$ (resp. $(\Gamma^\sigma P^\sigma)|_{S_Z \times S_Z}$) denotes the vector $(c_t^{\sigma(i)})_{i \in S_Z}$ (resp. the matrix with entries $(\gamma^{\sigma(i)} p_j^{\sigma(i)})_{i,j \in S_Z}$) (indeed, the use of the restriction bars is actually not correct but we keep them to help visualizing the dimensions of the objects at stake).

In particular, we recognize that $C_{Z,t}$ has the form C^{\min} studied by Proposition 3.13. Because Assumption B is verified, we obtain following a result of the type of Proposition 3.14 the existence of $\lambda \in [0, 1)$ and $u \in \mathbb{R}_{>0}^{S_Z}$ such that $C_{Z,t}$ is λ -contracting in weighted sup-norm $\|\cdot\|_u$. This property extends to φ_t , (since taking the minimum with the constant application b_t does not modify the contraction rate of $C_{Z,t}$). As a result, $v_Z(t)$ arises as the unique fixed-point of φ_t , not necessarily positive but still finite. This shows that $v(t)$ is uniquely determined for all t in $[t_0, t_0 + t_{\min})$, hence achieving our inductive scheme.

- ◇ Now that we have proved the uniqueness of $v(t)$ and in particular the one of $v_Z(t)$, we can present for the latter a more convenient form. Indeed, we know that $v_Z(t)$ satisfies the following implicit relation:

$$v_Z(t) = \min \left(\min_{\sigma \in \mathfrak{S}_Z} \left\{ c_t^\sigma |_{S_Z} + (\Gamma^\sigma P^\sigma) |_{S_Z \times S_Z} v_Z(t) \right\}, b_t \right). \quad (3.15)$$

Equation (3.15) holds if and only if for all σ in \mathfrak{S}_Z , we have $v_Z(t) \leq c_t^\sigma |_{S_Z} + (\Gamma^\sigma P^\sigma) |_{S_Z \times S_Z} v_Z(t)$ and also $v_Z(t) \leq b_t$, with an equality achieved for each coordinate. Since by Assumption **B**, we have $\rho((\Gamma^\sigma P^\sigma) |_{S_Z \times S_Z}) < 1$ for all σ of \mathfrak{S}_Z , the first set of inequalities can be replaced by $v_Z(t) \leq (I - (\Gamma^\sigma P^\sigma) |_{S_Z \times S_Z})^{-1} c_t^\sigma |_{S_Z}$ for all σ in \mathfrak{S}_Z (we can use a series expansion to show that this inverse matrix has only nonnegative entries and thus is order-preserving). We therefore obtain the following explicit expression:

$$v_Z(t) = \min \left(\min_{\sigma \in \mathfrak{S}_Z} \left\{ (I - (\Gamma^\sigma P^\sigma) |_{S_Z \times S_Z})^{-1} c_t^\sigma |_{S_Z} \right\}, b_t \right). \quad (3.16)$$

If v is the function determined by $v^0 \in \mathcal{V}$ and (DP- γ), we easily show that if v^0 is càdlàg (resp. piecewise-affine, piecewise-constant), then is so is v , by supposing that v is indeed càdlàg (resp. piecewise-affine, piecewise-constant) on $[-t_{\max}, t_0]$ for $t_0 \geq 0$ and extending this property on $[t_0, t_0 + t_{\min}]$ using equations (3.14) and (3.16) which preserve the càdlàg character (resp. piecewise-affine, piecewise-constant) and the fact that $t \mapsto (c_t^a)_{a \in A_Z}$ and $t \mapsto b_t$ are càdlàg (resp. piecewise-affine, piecewise-constant) mappings of $[t_0, t_0 + t_{\min}]$ as well. \clubsuit

The proof of Proposition 3.17 shows that we can successfully determine a solution of (DP) or (DP- γ) over $[0, \infty)$ from an initial condition of \mathcal{V} by applying an induction scheme very close from the induction principle (**I**) of Chapter 1, except that there is this time an implicit part to handle. We choose once again to highlight this fact in order to reproduce it later.

- **INDUCTION PRINCIPLE (\mathcal{I}_Z)**. *In a semi-Markov decision process verifying Assumption **A** (resp. Assumption **B**), denoting by $t_{\min} := \min_{a \in A} \{ t^a \mid t^a > 0 \}$ and $t_{\max} := \max_{a \in A} \{ t^a \mid t^a > 0 \}$, the properties of a solution of the dynamic programming equations (DP) (resp. (DP- γ)) over $[0, \infty)$ can be studied by induction from an initial condition defined over $[-t_{\max}, 0)$, since for all k in \mathbb{N} , and after possible inversion of an implicit “Zeno-part” of the system, the restriction of a solution over $[kt_{\min}, (k+1)t_{\min}]$ is explicitly given by its restriction over $[kt_{\min} - t_{\max}, kt_{\min}]$.*

Building on Proposition 3.17, we are now going to establish that v_w^* defined by (3.7) as the minimum expected cost with terminal cost w over all strategies in finite-horizon satisfies (DP). As a particular case, we will obtain that v^* (value function with no terminal cost) also satisfies (DP). Let w be a terminal cost function in \mathcal{V} ; we have justified in Section 2.3 that $v_w^* |_{[-t_{\max}, 0)} = w$. We are actually going to show that v_w^* coincides over $[-t_{\max}, \infty)$ with \bar{w} , the solution of (DP) uniquely determined by w . The fact that \bar{w} satisfies the dynamic programming equations gives us for all i in S and $t \geq 0$ a choice $\sigma(i, t)$ of action in A_i that realizes the minimum in (DP), i.e., such that

$$\forall t \geq 0, \quad \forall i \in S, \quad \bar{w}(i, t) = c^{\sigma(i, t)} + \sum_{j \in S} p_j^{\sigma(i, t)} \bar{w}(j, t - t^{\sigma(i, t)}). \quad (3.17)$$

For a fixed value of t , $(\sigma(i, t))_{i \in S}$ can be identified with a policy of \mathfrak{S} . In addition, σ can be chosen to be a piecewise-constant and càdlàg application from $S \times \mathbb{R}_{\geq 0}$ to \mathfrak{S} , since \bar{w} is itself piecewise-continuous càdlàg according to Proposition 3.17 and \mathfrak{S} is a discrete set.

If the planning horizon $t \geq 0$ is known, then σ allows us to define the particular strategy $f^* = (f_n^*)_{n \in \mathbb{N}}$ of \mathbb{F} , by letting for all epochs n in \mathbb{N} and all histories h_n of length n of the form $h_n = (i_0, s_0, a_0, i_1, s_1, a_1, \dots, i_{n-1}, s_{n-1}, a_{n-1}, i_n)$ in H_n (see Section 2.2 for a detailed discussion on the form of the histories):

$$f_n^*(h_n) := \sigma(i_n, \max(t - s_n, 0)).$$

The strategy f^* is pure but it is not strictly speaking Markovian. It does depend on the history only through the current state *and* the current total elapsed time since 0, or equivalently the time remaining to play before hitting the planning horizon t . As a consequence f^* is not a stationary strategy; observe although that it still has *some form* of stationarity. Indeed, the choice of actions to play at different epochs varies but does not depend on the epoch number itself, therefore the *same* horizon-dependent rule is applied at all epochs, and it is “Whatever the epoch, pick the action given by σ evaluated in the current state and the time remaining to play”. Huang and Guo call these strategies “**horizon-relevant stationary**” (see [HG11]). We can now state the following expected result in the undiscounted setting.

- **THEOREM 3.18.** *Under Assumption A, the function v_w^* , that is by (3.7) the undiscounted value function in finite-horizon with terminal cost w , is the unique function coinciding with w on $[-t_{\max}, 0)$ and that verifies the dynamic programming equations (DP). Moreover, it is achieved by the strategy f^* .*

Proof. When f is a strategy of \mathbb{F} , we define:

$$v^f(i, t) := \mathbb{E}_i^f \left(\sum_{k=0}^{\widehat{N}_t} \widehat{c}_k + w(\widehat{i}_{\widehat{N}_t+1}, t - \widehat{s}_{\widehat{N}_t+1}) \right) = \mathbb{E}_i^f \left(\sum_{k=0}^{\infty} \widehat{c}_k \mathbb{1}_{\widehat{s}_k \leq t} + w(\widehat{i}_{\widehat{N}_t+1}, t - \widehat{s}_{\widehat{N}_t+1}) \right).$$

Our goal is to prove that for all i in S and $t \geq 0$, we have $\bar{w}(i, t) = \inf_{f \in \mathbb{F}} v^f(i, t)$.

- ◊ We are first going to prove that v^{f^*} verifies (DP). To insist on the key fact that f^* has the knowledge of the planning horizon, we write f_t^* instead of f^* (only in this part of the proof). Let i in S and $t \geq 0$. We have

$$\begin{aligned} v^{f^*}(i, t) &= \mathbb{E}_i^{f_t^*} (\widehat{c}_0 \mathbb{1}_{0 \leq t}) + \mathbb{E}_i^{f_t^*} \left(\sum_{k=1}^{\infty} \widehat{c}_k \mathbb{1}_{\widehat{s}_k \leq t} + \sum_{k=0}^{\infty} w(\widehat{i}_{k+1}, t - \widehat{s}_{k+1}) \mathbb{1}_{\widehat{s}_k \leq t < \widehat{s}_{k+1}} \right) \\ &= c^{\sigma(i, t)} + \mathbb{E}_i^{f_t^*} \left(\mathbb{E}_i^{f_t^*} \left(\sum_{k=1}^{\infty} \widehat{c}_k \mathbb{1}_{\widehat{s}_k \leq t} + \sum_{k=0}^{\infty} w(\widehat{i}_{k+1}, t - \widehat{s}_{k+1}) \mathbb{1}_{\widehat{s}_k \leq t < \widehat{s}_{k+1}} \mid \widehat{s}_0, \widehat{i}_0, \widehat{a}_0, \widehat{s}_1, \widehat{i}_1 \right) \right) \\ &= c^{\sigma(i, t)} + \sum_{j \in S} p_j^{\sigma(i, t)} \mathbb{E}_i^{f_t^*} \left(\sum_{k=1}^{\infty} \widehat{c}_k \mathbb{1}_{\widehat{s}_k \leq t} + \sum_{k=0}^{\infty} w(\widehat{i}_{k+1}, t - \widehat{s}_{k+1}) \mathbb{1}_{\widehat{s}_k \leq t < \widehat{s}_{k+1}} \mid \widehat{a}_0 = \sigma(i, t), \widehat{i}_1 = j \right) \end{aligned}$$

where the first equality is obtained separating the first term from the sum and using the characterization of \widehat{N}_t ; the first bit of the second equality follows from the fact that \widehat{c}_0 is deterministically equal to $c^{\sigma(i, t)}$ under the application of the strategy f_t^* starting from state i , the other transformation used to obtain the second equality is the application of the law of total expectation; and finally the third equality is entailed by the fact that $\widehat{s}_0, \widehat{i}_0, \widehat{a}_0$ and \widehat{s}_1 are deterministic under the choice of $\widehat{i}_0 = i$ and under the strategy f_t^* . As a first case, we are going to suppose that $t - t^{\sigma(i, t)} \geq 0$, so that the condition $t < \widehat{s}_1$ is false, and we can get rid of the first term in the sum with the terminal cost.

Consider an admissible random process $(\widehat{i}_0, \widehat{s}_0, \widehat{a}_0, \widehat{i}_1, \widehat{s}_1, \widehat{a}_1, \dots)$ (called “process A”) generated by the initial state $i \in S$ and the strategy f_t^* and such that $\widehat{i}_1 = j$ for $j \in S$ fixed. Then, the random process $(\widehat{i}'_0, \widehat{s}'_0, \widehat{a}'_0, \widehat{i}'_1, \widehat{s}'_1, \widehat{a}'_1, \dots) := (\widehat{i}_1, \widehat{s}_1 - t^{\sigma(i, t)}, \widehat{a}_1, \widehat{i}_2, \widehat{s}_2 - t^{\sigma(i, t)}, \widehat{a}_2, \dots)$ (that we may call “process B”) has the same law as the random process generated by initial state j and the strategy $f_{t-t^{\sigma(i, t)}}^*$. Indeed the k -th chosen action in process B is given by $\sigma(\widehat{i}'_k, t - t^{\sigma(i, t)} - \widehat{s}'_k)$, coinciding with $\sigma(\widehat{i}_{k+1}, t - \widehat{s}_{k+1})$, the $(k+1)$ -th chosen action in process A. Hence

$$\begin{aligned} v^{f^*}(i, t) &= c^{\sigma(i, t)} + \sum_{j \in S} p_j^{\sigma(i, t)} \mathbb{E}_j^{f_{t-t^{\sigma(i, t)}}^*} \left(\sum_{k=1}^{\infty} \widehat{c}'_{k-1} \mathbb{1}_{t^{\sigma(i, t)} + \widehat{s}'_{k-1} \leq t} + \sum_{k=1}^{\infty} w(\widehat{i}'_k, t - t^{\sigma(i, t)} - \widehat{s}'_k) \mathbb{1}_{\widehat{s}'_{k-1} \leq t - t^{\sigma(i, t)} < \widehat{s}'_k} \right) \\ &= c^{\sigma(i, t)} + \sum_{j \in S} p_j^{\sigma(i, t)} \mathbb{E}_j^{f_{t-t^{\sigma(i, t)}}^*} \left(\sum_{k=0}^{\infty} \widehat{c}'_k \mathbb{1}_{\widehat{s}'_k \leq t - t^{\sigma(i, t)}} + \sum_{k=0}^{\infty} w(\widehat{i}'_{k+1}, t - t^{\sigma(i, t)} - \widehat{s}'_{k+1}) \mathbb{1}_{\widehat{s}'_k \leq t - t^{\sigma(i, t)} < \widehat{s}'_{k+1}} \right) \\ &= c^{\sigma(i, t)} + \sum_{j \in S} p_j^{\sigma(i, t)} v^{f^*}(j, t - t^{\sigma(i, t)}) \end{aligned}$$

This identity allows us to prove using induction principle (\mathcal{I}_Z) that v^{f^*} and \bar{w} that we already know to coincide on $[-t_{\max}, 0)$ actually coincide on $[-t_{\max}, \infty)$. In particular, v^{f^*} , satisfies equations (DP). We would have obtained the same equality in case $t - t^{\sigma(i,t)} < 0$ since in this case we would have directly written an identity featuring the terminal cost w .

- ◇ We now want to conclude that $\bar{w} = v^{f^*} \leq v^* = \inf_{f \in \mathbb{F}} v^f$. Using computations similar to the previous ones, we can demonstrate that if $f = (f_0, f_1, f_2, \dots)$ is a randomized and history-dependent strategy, we have for all i in S and $t \geq 0$

$$v^f(i, t) = \sum_{a \in A_i} f_0(a | i) \left(c^a + \sum_{j \in S} p_j^a v^{f_a}(j, t - t^a) \right), \quad (3.18)$$

where for i in S and a in A_i , f'_a denotes the randomized and history-dependent strategy $(f'_{a,0}, f'_{a,1}, \dots)$ such that for all k in \mathbb{N} and h_k in H_k , $f'_{a,k}(h_k) := f_{k+1}((i, a, h_k))$. This equation is hard to work with, but becomes much simpler if f is a Markovian strategy, i.e. when $f_k(h_k) = f_k(i_k)$, because in this case all the $(f'_a)_{a \in A}$ coincide with the strategy $f^{(1)} := (f_1, f_2, \dots)$. Hopefully, we can take advantage of the fact that for any randomized history-dependent strategy f in \mathbb{F} , there is a randomized Markovian strategy f_{MR} in \mathbb{F}^{MR} so that the random processes generated by f and f_{MR} have the same law (see [Put14, Theorem 5.5.1]), and as a result $\inf_{f \in \mathbb{F}} v^f = \inf_{f \in \mathbb{F}^{\text{MR}}} v^f$.

We now apply induction principle (\mathcal{I}_Z) to show that $\inf_{f \in \mathbb{F}^{\text{MR}}} v^f(i, t) \geq v^{f^*}(i, t)$ for all i in S and $t \geq 0$. We do that by supposing that this inequality holds for all t in $[-t_{\max}, t_0)$ and proving that it carries over if $t \in [t_0, t_0 + t_{\min})$. Let $f \in \mathbb{F}^{\text{MR}}$, according to our previous remark, (3.18) takes the form

$$v^f(i, t) = \sum_{a \in A_i} f_0(a | i) \left(c^a + \sum_{j \in S} p_j^a v^{f^{(1)}}(j, t - t^a) \right). \quad (3.19)$$

Since (3.19) features a sum of $|A_i|$ quantities weighted by probabilities $(f_0(a | i))_{a \in A_i}$, we can bound below the whole sum by its minimum term, and we get

$$v^f(i, t) \geq \min_{a \in A_i} \left(c^a + \sum_{j \in S} p_j^a v^{f^{(1)}}(j, t - t^a) \right). \quad (3.20)$$

Similarly to the proof of Proposition 3.17, we distinguish the “easy and explicit case” where $i \in S_{NZ}$ and the “more intricate and implicit case” where $i \in S_Z$. Indeed, if $i \in S_{NZ}$, we can apply our induction hypothesis to $f^{(1)}$ (that is a Markovian history-dependent strategy) since we have $t - t^a \leq t_0$ for all $a \in A_i$, and heredity step of (\mathcal{I}_Z) is accomplished since v^{f^*} verifies (DP). Let us focus on the restriction $v_Z^f(t)$ of v^f to states of S_Z . Still like the proof of Proposition 3.17, we can introduce the self-map φ_i^* on \mathbb{R}^S defined by

$$[\varphi_i^*]_i(x) := \min \left(\min_{a \in A_i \setminus A_{i,Z}} \left\{ c^a + \sum_{j \in S} p_j^a v^{f^*}(j, t - t^a) \right\}, \min_{a \in A_{i,Z}} \left\{ c^a + \sum_{j \in S_{NZ}} p_j^a v^{f^*}(j, t - t^a) + \sum_{j \in S_Z} p_j^a x_j \right\} \right),$$

such that we have thanks to our induction hypothesis $v_Z^f(t) \geq \varphi_i^*(v_Z^{f^{(1)}}(t))$. Denoting by $f^{(n)}$ the strategy $(f_n, f_{n+1}, f_{n+2}, \dots)$ if n is in \mathbb{N} , we obtain by immediate induction for φ_i^* is order-preserving that for all n

$$v_Z^f(t) \geq (\varphi_i^*)^n(v_Z^{f^{(n)}}(t)).$$

Note the difference with the fixed-point equation $v_Z^{f^*}(t) = \varphi_i^*(v_Z^{f^*}(t))$ satisfied by $v_Z^{f^*}(t)$. Now denote for short $x_n := v_Z^{f^{(n)}}(t)$, $x^* := v_Z^{f^*}(t)$ and $\varphi := \varphi_i^*$, we claim that we still have $\lim_{n \rightarrow \infty} \varphi^n(x_n) = x^*$. Indeed as seen in Proposition 3.17, we have $u \in \mathbb{R}_{>0}$ such that φ is λ -contracting in weighted sup-norm $\|\cdot\|_u$ and admits x^* as a fixed-point. Also observe that following Lemma 3.11, we have an upper bound of the form $\|x_n\|_u \leq M$ (we could for instance take $M := \max(u) \max_{i \in S_Z} \{N^{\lceil t/t_{\min} \rceil}(i^{(1)})\} \max_{a \in A} \{c^a\}$). So we have for n in \mathbb{N} that $\varepsilon_n := \|\varphi^n(x_n) - x^*\|_u \leq \lambda^n \|x_n - x^*\|_u \leq 2M\lambda^n$. Since $\lim_{n \rightarrow \infty} \varepsilon_n = 0$, we deduce that for all f in \mathbb{F}^{MR} , $v_Z^f(t) \geq \lim_{n \rightarrow \infty} (\varphi_i^*)^n(v_Z^{f^{(n)}}(t)) = v_Z^{f^*}(t)$, which shows correctness of our inductive scheme (\mathcal{I}_Z) also on the Zeno part.

- ◇ We can finally conclude by the inequalities $v^{f^*}(i, t) \leq \inf_{f \in \mathbb{F}^{\text{MR}}} v^f(i, t) = \inf_{f \in \mathbb{F}} v^f(i, t) = v^*(i, t) \leq v^{f^*}(i, t)$ that the value function v_w^* is given by the unique solution of the equations (DP) with initial condition w . \clubsuit

We state below a corollary concerning the value function v^* and give without demonstration the similar result for the discounted value function v_γ^* , since its proof is syntactically equivalent to the one of Theorem 3.18.

- **COROLLARY 3.19.** *Under Assumption A, the undiscounted value function in finite-horizon v^* defined by (3.5) is the unique function null on $[-t_{\max}, 0)$ that verifies the dynamic programming equations (DP), and it is achieved by a strategy depending at each epoch only on the current state and the time remaining to play.*

⋮ *Proof.* As mentioned in Section 2.3, the value function v^* is recovered from v_w^* by taking $w = \tilde{0}$. ♠

- **THEOREM 3.20.** *Under Assumption B, the discounted value function in finite-horizon v_γ^* defined by (3.6) is the unique function null on $[-t_{\max}, 0)$ that verifies the dynamic programming equations (DP- γ), and it is achieved by a strategy depending at each epoch only on the current state and the time remaining to play.*

The Theorems 3.18 and 3.20 enlighten in particular the discussion in Section 2.2 on the sufficiency of some subsets of strategies. In the context of finite-horizon problems, it is sound to work only with pure (or deterministic) strategies, however considering only (Markovian) stationary strategies is not sufficient to achieve a minimal expected cost. It would be a little misleading to say that we cannot avoid working with pure history-dependent strategies of \mathbb{F}^{HD} , since it is actually sufficient to consider strategies depending at each epoch on the current state and the time remaining to play, which is a set much smaller than \mathbb{F}^{HD} . The following example illustrates the strategy f^* used in practice.

Example 3.21. We depict in the Figure 3.5 a simple SMDP with states $\{1, 2, 3\}$ and actions $\{a_1, a_2, b_2, a_3, b_3\}$. Like in Figure 3.2, actions are labeled with pairs of a cost (for instance in dollars) and a sojourn time (typically in seconds).

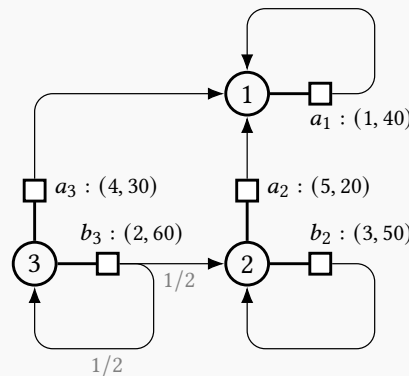


Figure 3.5: A semi-Markov decision process with four policies

Since $|A_1| = 1$ and $|A_2| = |A_3| = 2$, there is a total of four policies, that we can enumerate in the form

$$\mathfrak{S} = \{(a_1, a_2, a_3), (a_1, b_2, a_3), (a_1, a_2, b_3), (a_1, b_2, b_3)\}.$$

Observe that the player has no leverage on the minimization of his/her total incurred cost when starting in state 1, since he/she is bound to play action a_1 indefinitely and stay in state 1.

The dynamic programming equations (DP) verified by v^* write

$$\begin{aligned} v^*(1, t) &= 1 + v^*(1, t - 40) \\ v^*(2, t) &= \min(5 + v^*(1, t - 20), 3 + v^*(2, t - 50)) \\ v^*(3, t) &= \min\left(4 + v^*(1, t - 30), 2 + \frac{1}{2}v^*(2, t - 60) + \frac{1}{2}v^*(3, t - 60)\right) \end{aligned} \quad (3.21)$$

The Figure 3.6 represents the three components of v^* for planning horizons ranging from 0 to 300 seconds. Recall that $v^*(t)$ is null if $t < 0$, therefore, $v^*(0)$ is given by the minimum costs that we can obtain by playing the game in zero time. Since all actions of the SMDP are time-consuming, this amounts to doing a single move, and we naturally choose actions with the smallest costs, i.e., we play (a_1, b_2, b_3) to obtain a total cost $v^*(0) = (1, 3, 2)$. These three actions form the most interesting policy to play for $t \leq 100$.

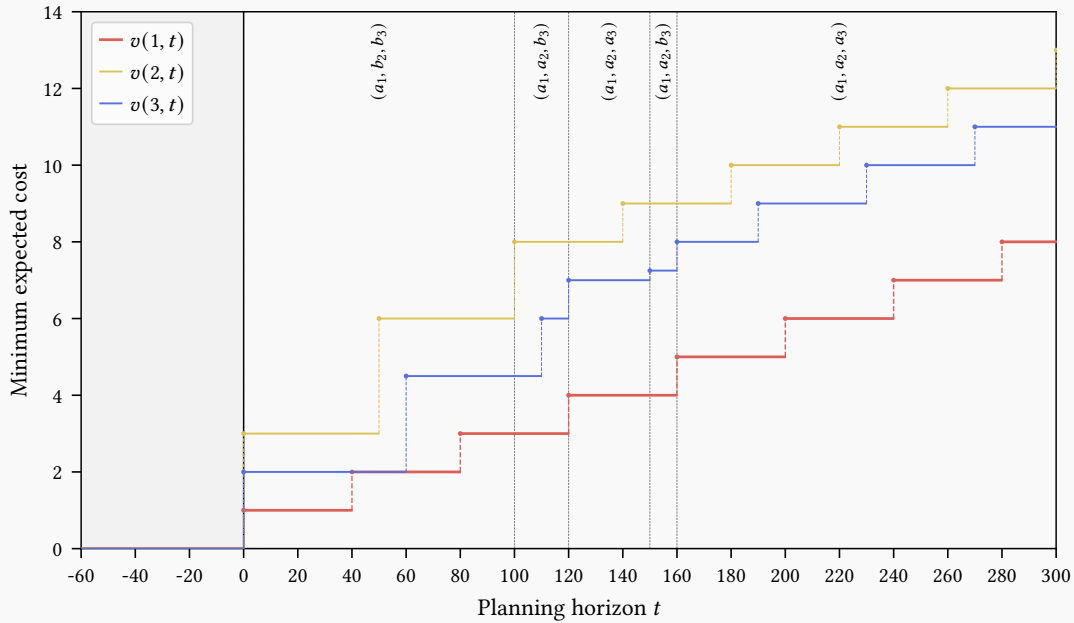


Figure 3.6: Finite-horizon value function of the SMDP of Figure 3.5.

Indeed, if the planning horizon is in $[100, 120)$ the player ought to play a_2 by starting in state 2, hence incurring a more important cost to begin with, but smaller costs afterwards (first 5\$, but after 20 seconds, 1\$ every 40 seconds, which is better than 3\$ every 50 seconds). Mind that it does not mean that the player should *always* pick a_2 from state 2 at *any moment* of the game if the latter lasts 100 seconds, but that the player should pick a_2 from state 2 in *every* game whenever it remains 100 seconds to play. This way, for any value $s \leq t$ inferior to the horizon t , there is always an optimal policy $\sigma^*(t - s)$ to play when it remains $t - s$ seconds to play, as it was introduced in (3.17).

The vertical bars separate the physical instants at which σ^* switches to a different policy. In this example, we see that the player should always apply policy (a_1, a_2, a_3) if it remains more than 160 seconds to play and the policy (a_1, b_2, b_3) if it remains less than 100 seconds to play, which well illustrates the difference between short-term and long-term decisions. Inbetween, the action a_2 from state 2 is always chosen, but the choice from state 3 varies.

We also retrieve on the Figure 3.6 that the value function is piecewise-constant and càdlàg, as stated in Proposition 3.17 since the initial condition $\tilde{\mathbf{0}}$ is constant over $[-t_{\max}, 0)$.

To finish this section, we come back to the last modeling aspect left aside in this section so far, the stopping costs. The corresponding value function of this problem still verifies dynamic programming equations, but that take into account these stopping options at each epoch. This form is well-known in the context of MDPs, see for instance [Ber11, Section 4.4], or [TVR99].

► **THEOREM 3.22.** *Under Assumption A, the undiscounted value function v_R^* in finite-horizon with stopping cost R defined by (3.8) verifies the dynamic programming equation*

$$\forall t \geq 0, \quad \forall i \in S, \quad v(i, t) = \min \left(\min_{a \in A_i} \left\{ c^a + \sum_{j \in S} p_j^a v(j, t - t^a) \right\}, R(i, t) \right). \quad (\text{DP-STOP})$$

Proof. To preclude an indigestion of too many similar proofs, we only give the ingredients of demonstration that differ from the proof of Theorem 3.18. We still leverage on the existence of a unique solution to (DP-STOP) given an initial condition in \mathcal{V} (and the knowledge of the function R), and on the existence of an application $\sigma: S \times \mathbb{R} \rightarrow A$ that minimizes the non-stopping part of (DP-STOP). This induces a strategy f^* of \mathbb{F} . Applying this strategy starting from i , we still have the choice on the stopping time M . If we take it to be null, the game stops directly and we get $R(i, t)$, otherwise we get $c^{\sigma(i,t)} + \sum_{j \in S} p_j^{\sigma(i,t)} v^{f^*}(j, t - t^{\sigma(i,t)})$, hence the apparition of the minimum of the two quantities, and we can show that f^* coincides with the solution of (DP-STOP). The second part of the proof is analogous to the one of Theorem 3.18. \clubsuit

As mentioned in Section 2.3, observe that taking $R = \infty$ (which *de facto* forbids the player to stop) allows us to recover the dynamic programming equations of the setting with no stopping costs. Such as the end of Section 2.3, we do not state further results on all the possible combinations of models involving discount factors, terminal and stopping costs, since Theorems 3.18, 3.20 and 3.22 suffice to figure out by syntactic analogy the dynamic programming equations of all these settings.

2.6 BACKWARD OF FORWARD PROPAGATION? — This subsection is optional, but we reckon it is worth commenting further on the form of the dynamic programming equations (DP) that we study, and more particularly on the fact that we express values of v^* in the present based on values in the past. Some readers that are more familiar with problems using dynamic programming algorithms might have expected to find an equation in which the present values are expressed in terms of values in the future.

To avoid reintroducing the notation of the basic and general problem of stochastic control (presented as such for instance in [Ber11, Section 1.3]), let us discuss the particular case of Markov Decision Processes in horizon t (with $t \in \mathbb{N}^*$) and with terminal costs. A decision-maker wants to minimize his/her total cost resulting from visiting states $\hat{i}_0, \hat{i}_1, \dots, \hat{i}_{t-1}, \hat{i}_t$, and pulling actions $(\hat{a}_k)_{0 \leq k < t}$ (with \hat{a}_k lying in $A_{\hat{i}_k}$ for all $0 \leq k < t$), hence determining costs $(\hat{c}_k)_{0 \leq k < t}$, knowing that finishing in state \hat{i}_t at epoch t yields a **terminal cost** $r(\hat{i}_t)$. We introduce the optimal cost-to-go $J_k^t(i)$ when we are in state i of S at epoch $0 \leq k \leq t$ by $J_k^t(i) := \inf_{f \in \mathbb{F}} \mathbb{E}_i^f(r(\hat{i}_t) + \sum_{\ell=k}^{t-1} \hat{c}_\ell)$ (in particular the terminal costs are such that $J_t^t(i) = r(i)$).

What we are really interested in is to determine the vector $J_0^t := (J_0^t(i))_{i \in S}$ of the minimum costs that we can expect from each starting state when there are t decision epochs ahead of us. The dynamic programming principle ensures that a strategy that is optimal from 0 to t is also optimal when truncated from k to t , with $0 \leq k \leq t$. Applied step-by-step, it gives us that $J_k^t = T(J_{k+1}^t)$ for all $0 \leq k \leq t$, where T is the Bellman operator such that for all x in \mathbb{R}^S , $T(x) = \min_{\sigma \in \mathcal{G}} \{c^\sigma + P^\sigma x\}$, as defined before by (3.13). Since the value $J_t^t = r$ is known for the

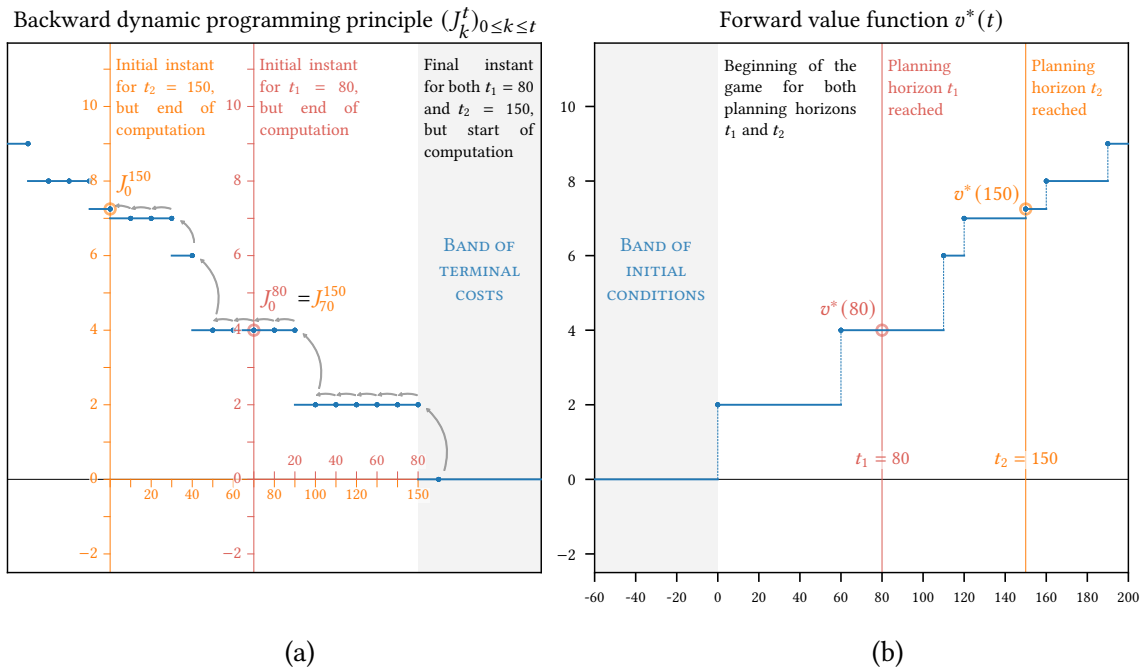


Figure 3.7: Backward and forward propagation are two mirrored methods

last epoch, it suffices to apply t times the operator T to **go backward in time** and determine J_0^t , i.e., the total cost-to-go for the initial epoch.

In this setting, the horizon t is fixed in advance and does not change, so “going backward in time” actually refers to the intermediary time-steps $0 \leq k \leq t$, and not to the horizon itself. To determine the total cost-to-go function for two different horizons t_1 and t_2 , we must start our algorithm twice from the same value $J_{t_1}^{t_1} = J_{t_2}^{t_2} = r$, and go backwards from t_1 steps in one case and from t_2 steps in the other case. Actually, if $t_1 \leq t_2$, it is not necessary to redo the t_1 first steps to compute $J_0^{t_2}$ and according to the optimality principle, it suffices to perform $t_2 - t_1$ steps from $J_0^{t_1} = J_{t_2-t_1}^{t_2}$.

On the other hand, when we want to focus on the influence of the horizon t , we can let $v^*: t \mapsto J_0^t$ to be the value function that maps any finite-horizon t to the minimum expected cost in t decision epochs. In that case, $v^*(t) = J_0^t = T(J_1^t) = T(J_0^{t-1}) = T(v^*(t-1))$, so that v^* is computed using a **forward time scheme**, now genuinely relatively to the horizons and not to intermediary steps. Compared to the dynamic programming algorithm, the value function somehow “sets back the time in the correct physical direction”. In particular, what was in the latter a terminal cost of the practical underlying model can now be interpreted as an initial condition of the computation scheme.

The Figure 3.7 illustrates these two sides of the same coin, based on the minimum expected cost of state 2 in Figure 3.5 earlier, when studying two different horizons $t_1 = 80$ s and $t_2 = 150$ s. Figure 3.7a shows the iterative and backward dynamic programming procedure (by steps of 10 seconds), starting with the terminal cost from t_1 or t_2 and going back to 0, i.e., computation ends when hitting the initial instant of the game. The Figure 3.7b shows that it suffices to mirror this method to retrieve the value function v^* . Note a subtlety here, although the general

dynamic programming algorithm described in this informal subsection is played in horizon t , the problem (DP-MDP) in Section 2.5 is actually played in horizon $t + 1$ (with $v(t)$ recovered by applying $t + 1$ times the operator T from $v(-1)$), explaining why there is an initial step from the null terminal cost to the costs-to-go $J_{80}^{80} = J_{150}^{150} = 2$ on Figure 3.7a. We chose to keep this hiatus to insist on the fact that our propagation schemes working with SMDPs *do not* propagate a scalar value, but propagate a function of \mathcal{V} ; recall indeed from the discussion in Section 2.5 that a scalar initial condition at $t = -1$ in the simplified setting of MDP corresponds to a functional initial condition given over $[-t_{\max}, 0)$ in the context of SMDPs.

2.7 THE EVOLUTION SEMIGROUP APPROACH. — We want in this section to understand and characterize the fundamental properties of the solutions of the dynamic programming equations (DP) in the undiscounted semi-Markov setting and in the absence of stopping costs. Until the end of the chapter, we assume that Assumption A is satisfied. The Proposition 3.17 is already a key result, as it proves that it is well-founded to talk about solutions of (DP), and that these are entirely determined by their restriction over $[-t_{\max}, 0)$. In the same logic as the one presented in the Section 2.2 of Chapter 2, we adopt here a more high-level look on the dynamics (DP) by studying its action over the set of \mathcal{V} of initial conditions. To that purpose, we define the following **evolution semigroup**.

► **DEFINITION 3.23.** For $t \geq 0$, we define the operator \mathcal{S}_t by

$$\mathcal{S}_t: \begin{cases} \mathcal{V} & \longrightarrow \mathcal{V} \\ v^0 & \longmapsto \begin{cases} [-t_{\max}, 0) & \longrightarrow \mathbb{R}^S \\ s & \longmapsto v(t+s) \end{cases} \end{cases} \quad (3.22)$$

where v is the function uniquely determined by the initial condition v^0 and the equations (DP), in accordance with Proposition 3.17.

We first establish that $(\mathcal{S}_t)_{t \geq 0}$ is a semigroup:

► **PROPOSITION 3.24.** The family $(\mathcal{S}_t)_{t \geq 0}$ forms a one-parameter semigroup, meaning that $\mathcal{S}_0 = \text{id}_{\mathcal{V}}$ and for all $t, t' \geq 0$, we have $\mathcal{S}_{t+t'} = \mathcal{S}_t \circ \mathcal{S}_{t'}$.

⋮ *Proof.* The fact that \mathcal{S}_0 is the identity mapping of \mathcal{V} is direct. Then, take nonnegative numbers t and t' , a test function v^0 in \mathcal{V} and s in $[-t_{\max}, 0)$. The value $(\mathcal{S}_t \circ \mathcal{S}_{t'})[v^0](s)$ is equal to $v'(t+s)$, where v' is the solution of the dynamics equations determined by the initial condition $\mathcal{S}_{t'}[v^0]$, but the latter is given by $r \mapsto v(t'+r)$ for r in $[-t_{\max}, 0)$ where v is the solution of the dynamics equations determined by v^0 . Since v' and $r \mapsto v(t'+r)$ coincide on $[-t_{\max}, 0)$, they are also equal on $[-t_{\max}, \infty)$ and we obtain $v'(t+s) = v(t'+t+s)$. As a result, $\mathcal{S}_t \circ \mathcal{S}_{t'} = \mathcal{S}_{t+t'}$. ◻

Recall from Theorem 3.18 and Corollary 3.19 that the value functions v_w^* (with terminal cost w) and v^* (with no terminal cost) for SMDPs are respectively the unique solution of (DP) that coincide with w over $[-t_{\max}, 0)$ and the solution of (DP) that is null over $[-t_{\max}, 0)$. As a result, the minimum expected cost with terminal payments $v_w^*(t)$ in finite-horizon t can be expressed as $v_w^*(t) = \mathcal{S}_{t+\varepsilon}[w](-\varepsilon)$ for all ε sufficiently small (any value ε in $(0, t_{\max})$ convenes), while its counterpart with no terminal cost v^* is such that $v^*(t) = \mathcal{S}_{t+\varepsilon}[\tilde{\mathbf{0}}](-\varepsilon)$, where $\tilde{\mathbf{0}}$ denotes as usual the null function of \mathcal{V} . Recall that functions of \mathcal{V} cannot be evaluated at the instant 0, hence the use of ε . The operator \mathcal{S}_t for SMDPs with $t \geq 0$ plays the exact same role as the t -fold iterate T^t of the Bellman operator for MDPs (with in that case t in \mathbb{N}), where T is defined by (3.13). The same way that T needs to be applied $t + 1$ times to recover $v^*(t)$ based on the initial condition $v^*(-1)$ in the context of MDPs (see Section 2.5), it is not surprising to see that the dynamics (DP)

need to be unraveled during a period strictly greater than t to find out what the value $v^*(t)$ is. Due to these similarities, we still refer to the application of the semigroup $(\mathcal{S}_t)_{t \geq 0}$ to an initial condition of \mathcal{V} as the **value iteration algorithm** for semi-Markov decision processes.

As explained in Section 2.2 of Chapter 2, the good way to think as the operator $(\mathcal{S}_t)_{t \geq 0}$ is to picture it as an “initial condition transportation” operator, an idea that we detail in the following example.

Example 3.25. The Figure 3.8 gives a solution v of the system (3.21) built from the SMDP of Figure 3.5 in the Example 3.21, but this solution is not the value function v^* since it is not initialized with the null function over $[-\tau_{\max}, 0)$ (recall here that $\tau_{\max} = 60$ s). Instead, we choose as initial condition a function v^0 in \mathcal{V} such that $v^0(1, s) = s/40$ for all s in $[-t_{\max}, 0)$, which will ensure by (3.21) that we will have $v(1, t) = t/40$ for all $t \geq 0$. The components $v^0(2, \cdot)$ and $v^0(3, \cdot)$ are set to two shifted sine functions.

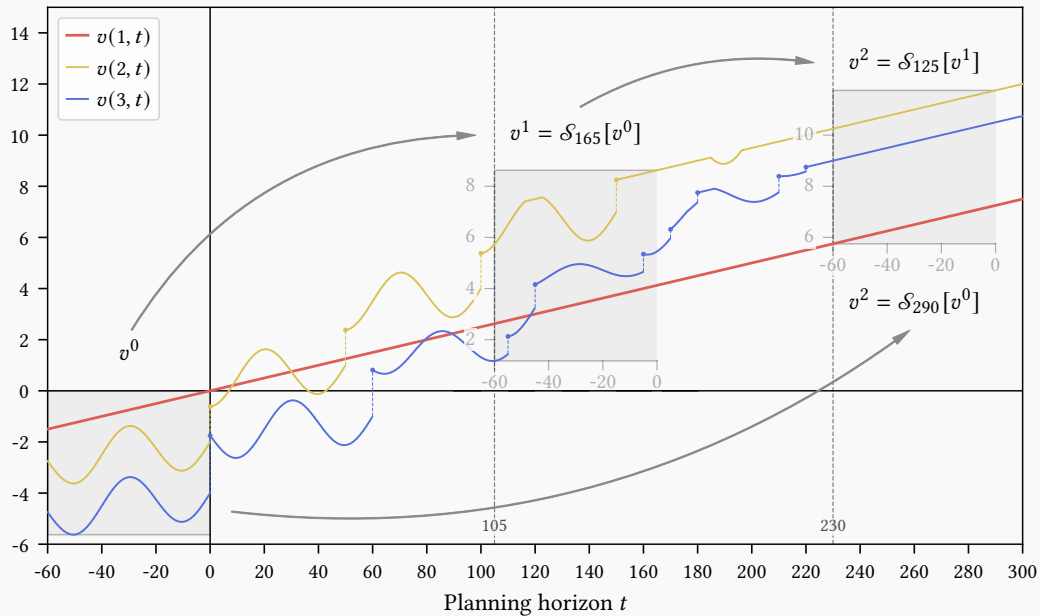


Figure 3.8: Dynamics (3.21) applied to a non-zero initial condition.

Applying the dynamics (DP) yields a solution v that seems for short horizons pretty erratic, with several discontinuities and a behaviour hard to predict. However, for long horizons ($t \geq 220$), the solution v stabilizes to an affine behaviour.

For $t \geq 0$, the effect of \mathcal{S}_t can be visualised by drawing a local frame between $t - t_{\max}$ (included) and t (excluded). The values taken there by v are (up to translation to put back the function over $[-t_{\max}, 0)$) precisely those of $\mathcal{S}_t[v^0]$. We illustrate on the figure the composition property of the semigroup. It is equivalent starting from v^0 to unravel the dynamics first for 165 seconds, hence obtaining v^1 , and then for another 125 seconds from v^1 , than to start from v^0 and apply the dynamics for $165+125=290$ seconds.

We list below some important properties of the operators $(\mathcal{S}_t)_{t \geq 0}$. We equip the space \mathcal{V} with the usual pointwise partial ordering \leq , and with the sup-norm $\|v^0\|_{\infty} := \sup_{s \in [-t_{\max}, 0)} \|v^0(s)\|_{\infty}$. We denote as before by $\tilde{1}$ the constant function of \mathcal{V} with all coordinates equal to 1.

► **PROPOSITION 3.26.** Let $t \geq 0$, v^0 and v'^0 be two initial conditions in \mathcal{V} . The operator \mathcal{S}_t is

(1) *additively homogeneous:*

$$\forall \alpha \in \mathbb{R}, \quad \mathcal{S}_t[v^0 + \alpha \tilde{\mathbf{1}}] = \mathcal{S}_t[v^0] + \alpha \tilde{\mathbf{1}},$$

(2) *monotone (or order-preserving):*

$$v^0 \leq v'^0 \implies \mathcal{S}_t[v^0] \leq \mathcal{S}_t[v'^0],$$

(3) *nonexpansive:*

$$\|\mathcal{S}_t[v'^0] - \mathcal{S}_t[v^0]\|_\infty \leq \|v'^0 - v^0\|_\infty,$$

(4) *continuous for both the uniform and pointwise convergence topologies.*

Proof. To show that the evolution semigroup is additively-homogeneous and order-preserving, we reuse the induction principle (\mathcal{J}_Z) and some of the notation defined in the proof of Proposition 3.17, being understood that all the discount factors are here taken equal to one.

- ◊ We first prove the additively-homogeneous character. Suppose v^0 and v'^0 are in \mathcal{V} , denote v and v' the functions they uniquely determine on $[-t_{\max}, \infty)$. In the particular case where $v'^0 = v^0 + \alpha \tilde{\mathbf{1}}$ with α in \mathbb{R} , if we suppose that $v'(t) = v(t) + \alpha \mathbf{1}$ for all t in $[-t_{\max}, t_0)$, then we easily see that this property still holds on $[t_0, t_0 + t_{\min})$ for non-Zeno states using explicit expression (3.14). On Zeno states, we make use of the explicit expression (3.16). Following from the non-Zeno part, it is clear that $b'_t = b_t + \alpha \mathbf{1}$, and for all σ in \mathfrak{S} , we also have $c_t^{\sigma'} = c_t^\sigma + \alpha P^\sigma|_{S_{NZ} \times S_Z} \mathbf{1}$. However for $P^\sigma|_{S_{NZ} \times S_Z} \mathbf{1} + P^\sigma|_{S_Z \times S_Z} \mathbf{1} = \mathbf{1}$ (fundamental property of stochastic matrices), we also have $c_t^{\sigma'} = c_t^\sigma + \alpha(I - P^\sigma|_{S_Z \times S_Z})\mathbf{1}$, and it follows readily from (3.16) that $v'_Z(t) = v_Z(t) + \alpha \mathbf{1}$, which completes the proof of the additively homogeneous character.
- ◊ For the monotonicity, take $v'^0 \geq v^0$ and let us suppose that $v' \geq v$ on $[-t_{\max}, t_0)$. Using explicit expression (3.14), we still obtain that this carries along on $[t_0, t_0 + t_{\min})$ for non-Zeno states. On states of S_Z , observe that for a fixed t in $[t_0, t_0 + t_{\min})$ we have $b'_t \geq b_t$ and $c_t^{a'} \geq c_t^a$ for all a , hence explicit expression (3.16) provides $v'_Z(t) \geq v_Z(t)$ by order-preserving character of the inverse matrices at stake.
- ◊ The nonexpansiveness itself is a consequence of both the additively homogeneous and order-preserving characters, as it was shown by Crandall and Tartar in [CT80].
- ◊ It is clear that \mathcal{S}_t is continuous for the uniform topology for all $t \geq 0$ from the nonexpansiveness, for the latter amounts to the 1-Lipschitz property. The continuity with respect to the pointwise convergence is a little more subtle. We build on the fact (for instance following from the proof of Proposition 3.17 by induction once again) that for all $t \geq 0$ and $s \in [-t_{\max}, 0)$, the number $\mathcal{S}_t[v^0](s)$ continuously depend on finitely many values taken by $v^0 \in \mathcal{V}$. Indeed, we can write

$$\mathcal{S}_t[v^0](s) = \psi_{t+s} \left(\left\{ v^0(\tau) \mid \tau \in \left(t+s - \sum_{a \in A} t^a \mathbb{N} \right) \cap [-t_{\max}, 0) \right\} \right),$$

the mapping ψ_{t+s} is continuous because it is obtained by taking minimums and linear combinations (including inversions of Cramer systems, see equations (3.14) and (3.16)), and the number of values of v^0 used is finite as the intersection of a discrete countable set (whose points are all isolated) and a bounded interval. We then use the sequential characterization of the continuity to deduce that if $(v_{(k)}^0)_{k \in \mathbb{N}}$ converge pointwise towards v_∞^0 as functions of \mathcal{V} , then for all $t \geq 0$ and for all s in $[-t_{\max}, 0)$, the sequence $(\mathcal{S}_t[v_{(k)}^0](s))_{k \in \mathbb{N}}$ converges towards $\mathcal{S}_t[v_\infty^0](s)$. \clubsuit

In Section 3 and Section 4, we will build on these properties to derive rich convergence results on the solutions of (DP). Remark that we have implicitly used the fact that the dynamics (DP) is autonomous to define the semigroup $(\mathcal{S}_t)_{t \geq 0}$. Had we considered a problem with stopping cost R , we would have been compelled to use the dynamics (DP-STOP) instead, that is non-autonomous since it relies on the knowledge of R . Hence, we would have had to resort to the more general formalism of the evolution operator discussed in the Section 2.2 of Chapter 2.

► **REMARK 3.27.** It shall turn convenient in Chapter 4 to consider an SMDP induced by a subset of actions. This may also require to restrict ourselves to an adequate subset of states. To that purpose, we say that a pair $(S', A' := \biguplus_{i \in S'} A'_i)$ with $S' \subset S$ and $\emptyset \neq A'_i \subset A_i$ for all i in S is a **consistent subset of states and actions** if for all $a \in A'$, we have $\sum_{j \in S'} p_j^a = 1$.

If (S', A') is such a pair, the equations (DP) restricted to actions of A' and to states of S' induce an evolution semigroup that we may denote by $(\mathcal{S}_t|_{S', A'})_{t \geq 0}$ on the set of “initial condition” functions with suitable dimension, as set out in Proposition 3.17 and Definition 3.23. Any restricted semigroup of this form also enjoys the properties listed in Proposition 3.26.

2.8 EXTENSION TO THE CASE WITH NON-CONSTANT SOJOURN TIMES. — To finish this section on the finite-horizon value functions, we discuss more informally how the preceding results are affected in the case where the sojourn times of the SMDP are random variables and follow action-dependent probability distributions $(\nu_a)_{a \in A}$. We consider here only the undiscounted case. We still want the value functions to be well-defined, while allowing some sojourn times to be deterministically null (i.e., $\nu_a = \delta_0$ for some actions a in A). As noted in [HG11], the mildest condition needed for the undiscounted value-function (3.5) to be well-posed at all horizons is that $\widehat{s}_\infty := \sum_{k=0}^{\infty} \widehat{t}_k$ is almost surely infinite for all strategies and starting states. This way, for all $t \geq 0$, there exists k in \mathbb{N} such that $\widehat{N}_t = k$.

Let us write down several SMDP “well-posedness” conditions:

- (i) there exists $t_{\min} > 0$ such that for all actions a in A , $\nu_a([t_{\min}, \infty)) = \mathbb{P}(\widehat{t}^a > t_{\min}) = 1$
- (ii) for all actions a in A , $\int_{\mathbb{R}_{>0}} s \, d\nu_a(s) = \mathbb{E}(\widehat{t}^a) > 0$
- (iii) for all policies σ , for all final classes F in $\mathcal{F}(\sigma)$, there exists i in F such that $\nu_{\sigma(i)}(\{0\}) < 1$
- (iv) for all states i in S and strategies f in \mathbb{F} , $\mathbb{P}_i^f(\widehat{s}_\infty = \infty) = 1$.

The condition (i) is here the strongest. It is clear that it prevents any non-Zeno behaviour, but it does not allow one to model null sojourn times. The condition (ii) is more general, and is the one very frequently made in the SMDP literature (the requirement that the average sojourn time $\mathbb{E}(\widehat{t}^a)$ is positive is equivalent to the existence of positive scalars δ and ε such that $\mathbb{P}(\widehat{t}^a > \delta) > \varepsilon$), see for instance [Jew63, Fei94, Yus82, Jašo4]. Although (ii) allows the sojourn times to sometimes take null values, almost surely zero sojourn times are also forbidden. On the model of Assumption A, we thus propose the condition (iii) above, which demands that the hypothesis (ii) is verified only at least for one state per final class of any policy (the condition $\nu_a(\{0\}) < 1$ being equivalent to $\mathbb{E}(\widehat{t}^a) > 0$). We have (i) \implies (ii) \implies (iii) \iff (iv), as a result (iii) appears to be a good setting to work with SMDPs with finite state and action sets (the equivalence (iii) \iff (iv) is elementary for pure stationary strategies associated with policies, while general lines of proof are close to those of Theorem 3.15).

With stochastic sojourn times and under any of the above conditions, we can establish that the value function in finite horizon (with or without terminal cost) satisfies the following dynamic programming equations:

$$\forall t \geq 0, \quad \forall i \in S, \quad v(i, t) = \min_{a \in A_i} \left\{ c^a + \sum_{j \in S} p_j^a \int_0^t v(j, t-s) \, d\nu_a(s) \right\}. \quad (\text{DP}')$$

The equation (DP') is the stochastic-times counterpart of (DP), where the generic delayed term $v(j, t - t^a)$ has been replaced by a convolution integral.

It remains to study the well-posedness of this new system, i.e., showing existence and uniqueness of solutions, the same way we have done in this section for the system (DP). Under the previous condition (i), the convolution integral in (DP') is actually taken over $[t_{\min}, t]$, hence it is still possible to resolve to an induction principle and determine $v(t)$ explicitly in terms of $v|_{[-t_{\max}, t-t_{\min}]}$. But as we have said before, it does not include the case of null sojourn times. If 0 is in the support of some measures $(\nu_a)_{a \in A}$, then the system (DP') features an implicit part again.

When $|A_i| = 1$ for all states i in S , (DP') takes the form of a system of so-called renewal equations, studied by Bellman and Cooke in [BC63, Chapters 7 and 8]. They prove if the $(\nu_a)_{a \in A}$ are absolutely continuous that there exists a unique solution of this system, and precise further that it also holds if the $(\nu_a)_{a \in A}$ admit a pure point part (that is, these distributions are nonsingular). More generally, this type of equations belongs to the class of system of Volterra equations of the second kind, in this case linear, for which the continuity of the kernels is often required to prove the existence of solutions (we refer on this topic to the book of Brunner [Bru17]). Showing that the condition (iii) suffices or not for (DP') to have unique trajectories seems nontrivial and is left for future work, and so is the more general case where $|A_i| \geq 1$ for some states i in S . This indeed provides nonlinear renewal equations, for which fixed-point results used in the linear case should be adapted, or reduction to the linear case by means of time-dependent and piecewise-constant choices of policies could be implemented.

3 Long-run properties of solutions of semi-Markov dynamics

The (undiscounted) value function v^* defined in (3.5) actually gives the solutions of a continuum of optimization problems, indeed, for all finite horizons $t \geq 0$, $v^*(t)$ yields the minimum expected cost when an SMDP has to be played during t units of time. In this section, we are interested in the behaviour of $v^*(t)$ when t tends to ∞ , and more generally to the behaviour of $v(t)$ for large values of t , when v is a solution of (DP) (not necessarily v^* , which means that nonzero terminal costs are taken into account).

In most cases, we expect that $\lim_{t \rightarrow \infty} v^*(t) = \infty$, for instance if all the action costs in the SMDP are positive and hence we are doomed to lose more and more money. Conversely if all the costs are negative (so that they are actually rewards), we can have $\lim_{t \rightarrow \infty} v^*(t) = -\infty$ and get infinitely rich. Alternatively, taking into account discount factors $(\gamma^a)_{a \in A}$ that are strictly smaller than one in order to increase (resp. reduce) the importance of costs incurred in short horizons (resp. far horizons) is a way to deal with a criterion that remains finite when t approaches infinity (indeed, we have in this case $-\infty < \liminf_{t \rightarrow \infty} v_y^*(t) \leq \limsup_{t \rightarrow \infty} v_y^*(t) < \infty$).

3.1 AFFINE REGIMES AND AVERAGE-COST VECTOR. — In the absence of discount factors, the most natural quantity to study is the **ultimate rate of cost** that the player incurs per unit of time depending of his/her starting state, i.e., the vector $\liminf_{t \rightarrow \infty} v^*(t)/t$. More generally, since any solution of (DP) can be expressed using an initial condition v^0 in \mathcal{V} and the dynamics semigroup $(\mathcal{S}_t)_{t \geq 0}$, we want to characterize the rate of cost $\liminf_{t \rightarrow \infty} \mathcal{S}_t[v^0](s)/t$, where s is in $[-t_{\max}, 0)$. The following proposition teaches us a good property of this limit.

➤ **PROPOSITION 3.28.** *Let v^0 be an initial condition of \mathcal{V} and s in $[-t_{\max}, 0)$. Then, the vector*

$$\liminf_{t \rightarrow \infty} \frac{\mathcal{S}_t[v^0](s)}{t}$$

is independent of both v^0 and s .

Proof. Let v^0, v'^0 be in \mathcal{V} and s, s' be in $[-t_{\max}, 0)$. Without loss of generality, let us suppose that $s' \geq s$. First remark that for all $t \geq 0$, we have thanks to the Definition 3.23 and the Proposition 3.24 that

$$\mathcal{S}_t[v'^0](s') - \mathcal{S}_t[v^0](s) = \mathcal{S}_t[\mathcal{S}_{s'-s}[v'^0]](s) - \mathcal{S}_t[v^0](s).$$

The result claimed by the proposition is a key consequence of the nonexpansiveness of the semigroup $(\mathcal{S}_t)_{t \geq 0}$, shown in Proposition 3.26–(3). Indeed, we have for all $t \geq 0$ that

$$\|\mathcal{S}_t[\mathcal{S}_{s'-s}[v'^0]] - \mathcal{S}_t[v^0]\|_{\infty} \leq \|\mathcal{S}_{s'-s}[v'^0] - v^0\|_{\infty},$$

hence, we obtain for all $t > 0$ (componentwise)

$$-\frac{1}{t} \times \|\mathcal{S}_{s'-s}[v'^0] - v^0\|_{\infty} \mathbf{1} \leq \frac{\mathcal{S}_t[v'^0](s')}{t} - \frac{\mathcal{S}_t[v^0](s)}{t} \leq \frac{1}{t} \times \|\mathcal{S}_{s'-s}[v'^0] - v^0\|_{\infty} \mathbf{1},$$

which in turn yields by taking the \liminf when $t \rightarrow \infty$

$$\liminf_{t \rightarrow \infty} \frac{\mathcal{S}_t[v'^0](s')}{t} = \liminf_{t \rightarrow \infty} \frac{\mathcal{S}_t[v^0](s)}{t}.$$

♣

The fact that the ultimate rate of cost does not depend upon the initial condition v^0 in \mathcal{V} has a natural explanation when we think about v^0 as the terminal cost of the problem (3.7). Indeed, for the game is played over a horizon that gets bigger and bigger, the relative importance of the terminal cost fades away. Because this ultimate rate of cost depends only the parameters of the game (sojourn times, costs of actions, transition probabilities) or equivalently on the dynamics (DP) governed by the semigroup $(\mathcal{S}_t)_{t \geq 0}$, it is licit to denote it by $\chi(\mathcal{S})$, or in other words

$$\chi(\mathcal{S}) := \liminf_{t \rightarrow \infty} \frac{v^*(t)}{t} = \liminf_{t \rightarrow \infty} \frac{\mathcal{S}_t[v^0](s)}{t} \quad \text{regardless of } v^0 \text{ in } \mathcal{V}, s \text{ in } [-t_{\max}, 0). \quad (3.23)$$

It is enlightening to look at the special case of Markov decision processes, in which all the sojourn times are equal to one, since in this case we can write $v^*(t) = T^t(v^*(0))$ for all t in \mathbb{N} , using the Bellman operator T defined by (3.13). The quantity $\chi(\mathcal{S})$, that we may as well denote by $\chi(T)$ in this special case, arises as the **escape rate** of the self-map T of \mathbb{R}^S (still unique by nonexpansiveness). The following theorem of Kohlberg guarantees the existence and the finiteness of $\lim_{t \rightarrow \infty} T^t(v^*(0))/t$.

- **THEOREM 3.29** (Kohlberg, [Koh80]). *Let n in \mathbb{N} and suppose $F: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is piecewise-affine and nonexpansive (in any norm). Then, there exists two vectors g and h in \mathbb{R}^n such that $s \mapsto sg + h$ (with s in $\mathbb{R}_{\geq 0}$) is an **invariant half-line** of F , i.e., for all $s \geq 0$ we have $F(sg + h) = (s + 1)g + h$.*

For T is piecewise-affine and nonexpansive in sup-norm, we know that $\chi(T)$ coincides with the vector g of Theorem 3.29. Coming back to the general case of semi-Markov decision processes, we may want to look for an equivalent result on $(\mathcal{S}_t)_{t \geq 0}$. Given two vectors g and h in \mathbb{R}^S , we shall say that the affine function $v^{\text{aff}}: t \mapsto gt + h$ of \mathcal{V} is an **affine regime of the SMDP** (or an invariant half-line) if the affine function $t \mapsto gt + h$ of $[-t_{\max}, \infty) \rightarrow \mathbb{R}^S$ (still denoted by v^{aff}) is solution of (DP). For affine regimes, applying the dynamics of the SMDP amounts to a translation in time, because for all $t \geq 0$ and s in $[-t_{\max}, 0)$, we have $\mathcal{S}_t[v^{\text{aff}}](s) = v^{\text{aff}}(t + s)$. The following result is somehow standard in the SMDP literature but not customarily presented in this form, therefore it is given with a proof (see [DY79] for close statements in the MDP case).

► **PROPOSITION 3.30.** Let g and h be two vectors of \mathbb{R}^S . The following are equivalent:

- (i) there exists a nonnegative number t_0 such that $t \mapsto g(t + t_0) + h$ is an affine regime of the SMDP,
- (ii) g and h are solution of the two optimality equations (OE1)–(OE2)

$$g(i) = \min_{a \in A_i} \left\{ \sum_{j \in S} p_j^a g(j) \right\} \quad (\text{OE1})$$

$$h(i) = \min_{a \in A_i^*} \left\{ c^a - t^a g(i) + \sum_{j \in S} p_j^a h(j) \right\} \quad (\text{OE2})$$

where for all i in S , A_i^* is the subset of A_i where the minimum is achieved in (OE1).

Proof. We proceed by double implication.

◊ Suppose that (i) holds, by definition the function $t \mapsto g(t + t_0) + h$ satisfies (DP) and we have

$$\forall t \geq 0, \forall i \in S, \quad g(i)(t + t_0) + h(i) = \min_{a \in A_i} \left\{ c^a + \sum_{j \in S} p_j^a g(j)t + \sum_{j \in S} p_j^a (g(j)(t_0 - t^a) + h(j)) \right\}.$$

The function in the right-hand side is ultimately affine, so we have for any i in S and t large enough that

$$\begin{aligned} g(i)(t + t_0) + h(i) &= \min_{a \in A_i} \left\{ c^a + \sum_{j \in S} p_j^a g(j)t + \sum_{j \in S} p_j^a (g(j)(t_0 - t^a) + h(j)) \right\} \\ &= \min_{a \in A_i} \left\{ \sum_{j \in S} p_j^a g(j) \right\} t + \min_{a \in A_i^*} \left\{ c^a + \sum_{j \in S} p_j^a (g(j)(t_0 - t^a) + h(j)) \right\} \\ &= \min_{a \in A_i} \left\{ \sum_{j \in S} p_j^a g(j) \right\} t + \min_{a \in A_i^*} \left\{ c^a + g(i)(t_0 - t^a) + \sum_{j \in S} p_j^a h(j) \right\}, \end{aligned}$$

where we have used for the last inequality the fact that $a \in A_i^*$ implies $\sum_{j \in S} p_j^a g(j) = g(i)$. At the end of the day, we just need to identify the slope and the intercept of two ultimately affine functions, and we obtain that g and h satisfy (OE1) and (OE2).

◊ Conversely, assume that (ii) is true, i.e., g and h verify (OE1) and (OE2), and we want to prove that $t \mapsto g(t + t_0) + h$ is solution of (DP) for some t_0 . Equivalently, we need to show that

$$\forall t \geq 0, \forall i \in S, \quad \min_{a \in A_i} \left\{ \left(\sum_{j \in S} p_j^a g(j) - g(i) \right) (t + t_0) + \left(c^a - g(i)t^a + \sum_{j \in S} p_j^a h(j) - h(i) \right) \right\} = 0. \quad (3.24)$$

Let $i \in S$ and $a \in A_i$. If a achieves the minimum in (OE2), so that it also achieves the minimum in (OE1), then the two terms in the right-hand side of (3.24) vanish and the equality holds. If a achieves the minimum in (OE1) but not in (OE2), then the coefficient of $(t + t_0)$ in (3.24) still vanishes, and by (OE2), $c^a - g(i)t^a + \sum_{j \in S} p_j^a h(j) - h(i) \geq 0$. Finally, if a does not achieve the minimum in (OE1), this entails that the gap $\delta := \sum_{j \in S} p_j^a g(j) - g(i)$ is positive. Hence, since $t \geq 0$, the expression in (3.24) can be bounded below by $C + t_0\delta$ for some real constant C , and as a consequence, for t_0 large enough, this expression takes a nonnegative value, which entails that (3.24) holds. \clubsuit

The equations (OE1)–(OE2) are the **average-cost optimality equations** that were extensively studied by Schweitzer and Federgruen in their work [SF78], see also the work of Schäl [Sch92]. These equations play a fundamental role in identifying the minimal solutions of the **average-cost criterion** g^* defined for all i in S by

$$g^*(i) := \inf_{f \in \mathbb{F}} \liminf_{t \rightarrow \infty} \frac{1}{t} \mathbb{E}_i^f \left(\sum_{k=0}^{\widehat{N}_t} \widehat{c}_k \right), \quad (3.25)$$

in which we recognize a definition very close from the one of the undiscounted value function (3.5) (note that different notions of average-costs can be defined, see for instance [Ros70], [Fei94], or [JX04] for a recent review of the literature on the topic).

When studying average-cost criteria, it is important to take into account the underlying chain structure of the SMDP. Recall from Section 2.2 that every policy σ in \mathfrak{S} induces a Markov chain on S , with associated transition matrix that we have already denoted by P^σ . For σ in \mathfrak{S} , let us denote by $\mathcal{F}(\sigma)$ the set of final classes of the Markov chain induced by the policy σ , defined in Section 1.3. If each policy admits a unique final class (i.e., $|\mathcal{F}(\sigma)| = 1$ for all σ in \mathfrak{S}), we say that the configuration of the SMDP is **unichain**. This means that if we apply the stationary strategy corresponding to any policy, we shall eventually end up in a single irreducible class (possibly after having visited some transient states) and in this case, the ultimate rate of cost incurred by the player only depends on the chosen policy and not on the initial state. Otherwise, we say that the SMDP is **multichain**. This means that the final class of the game in which the player eventually ends up when applying the stationary strategy associated with some policy can depend on the choice of the initial state, via the access probabilities $(\phi_{F,i})_{F \in \mathcal{F}(\sigma), i \in S}$ defined in Proposition 3.6.

In what follows, we reuse from Section 1.3 the notation $(\mu_F^\sigma)_{F \in \mathcal{F}(\sigma)}$ to denote the (zero-padded) invariant measures of the final classes of the Markov chain associated with policy σ , and the notation $(\phi_{F,i}^\sigma)_{F \in \mathcal{F}(\sigma), i \in S}$ to denote the access probabilities of these final classes starting from any state of S . For σ in \mathfrak{S} , we also respectively denote by $(P^\sigma)^\star$ and $(H^\sigma)^\star$ the spectral projector and the fundamental matrix associated with the Markov chain induced by σ (see Proposition 3.3 and Proposition 3.7 in Section 1). Finally, we use the symbol \odot to denote the Hadamard (componentwise) product of two vectors.

► **THEOREM 3.31** (Schweitzer and Federgruen, [SF78]). *Suppose that Assumption A holds. Then,*

- (1) *there is a pure stationary strategy in \mathbb{F}^{SD} achieving minimal average-cost in (3.25),*
- (2) *any solution (g, h) of (OE1)–(OE2) verifies $g = g^*$, and we have*

$$g^*(i) = \min_{\sigma \in \mathfrak{S}} \sum_{F \in \mathcal{F}(\sigma)} \phi_{F,i}^\sigma \frac{\langle \mu_F^\sigma, c^\sigma \rangle}{\langle \mu_F^\sigma, t^\sigma \rangle}, \quad (3.26)$$

- (3) *for all scalars α, β in \mathbb{R} , the vectors $(g^*, h^* + \alpha \mathbf{1} + \beta g^*)$ are solutions of (OE1)–(OE2), where*

$$h^* := \min_{\sigma \in \mathfrak{S}^*} ((H^\sigma)^\star + (P^\sigma)^\star)(c^\sigma - g^* \odot t^\sigma),$$

$$\text{and } \mathfrak{S}^* := \{\sigma \in \mathfrak{S} \mid P^\sigma g^* = g^*\}.$$

Remark how the Assumption A that our Theorem 3.15 proved to be minimal for the finite-horizon value function to be well-defined is here also the right condition for the ratios in (3.26) to exist. In the special case where the SMDP is unichain, we verify for instance thanks to the matrix forms featured in Proposition 3.6 that g^* is uniform and arises as a ratio of costs by sojourn times, weighed by the invariant measure of the sole final class of an optimal policy, i.e., the long-run fraction of time spent in each state of this final class. The formula (3.26) expresses that this phenomenon partially survives in the multichain case, but different final classes can lead to different ultimate rates of cost, hence the weighted sum depending upon the access probabilities. Remark that in the unichain case, since g^* is uniform, the equation (OE1) is actually trivial. It is also worth-mentioning that a solution of (OE1)–(OE2) can be found constructively, as Denardo and Fox showed in [DF68] for instance by proving termination in a finite number of steps of

Howard's policy iteration algorithm (see [How60]). The item (3) of the theorem gives us a glimpse on the degrees of freedom on h , while g must be unique. The vector h can typically have up to $|\mathcal{F}(\sigma)|$ degrees of freedom if σ is in \mathfrak{S}^* ; we refer to [SF78] for advanced and subtle discussion on these aspects.

The links with our questions on the long-run behaviour of solutions of (DP) is established by the following corollary of Theorem 3.31.

► **COROLLARY 3.32.** *Suppose that Assumption A holds. Then,*

- (1) *there are affine regimes of the SMDP, with rate (or slope vector) $\chi(\mathcal{S}) = g^*$,*
- (2) *for all initial conditions v^0 in \mathcal{V} and s in $[-t_{\max}, 0)$, we have*

$$\mathcal{S}_t[v^0](s) \underset{t \rightarrow \infty}{=} \chi(\mathcal{S})t + O(1). \quad (3.27)$$

Proof. The first item is a consequence of the existence of solutions of (OE1)–(OE2) by Theorem 3.31 and the equivalence of (3.30). Let us take an affine regime $v^{\text{aff}}: t \mapsto g^*t + h$, by definition of $\chi(\mathcal{S})$, we have $\chi(\mathcal{S}) = \liminf_{t \rightarrow \infty} \mathcal{S}_t[v^{\text{aff}}](s)/t = \liminf_{t \rightarrow \infty} v^{\text{aff}}(t+s)/t = \liminf_{t \rightarrow \infty} (g^*(t+s) + h)/t = g^*$, where we have also used the defining property of affine regimes. Finally, we have already shown in the proof of Proposition 3.28 that there is a constant C such that $-C1 \leq \mathcal{S}_t[v^0](s) - \mathcal{S}_t[v^{\text{aff}}](s) \leq C1$, hence the result of the second item because $\mathcal{S}_t[v^{\text{aff}}](s) = \chi(\mathcal{S})(t+s) + h$. ♠

Observe that the equality $\chi(\mathcal{S}) = g^*$ can be thought of as a “limit inversion”-type result, comparing (3.23) and (3.25). We invite the reader to take another look at Examples 3.21 and 3.25 in the light of Corollary 3.32. Indeed, Example 3.25 was already featuring an affine regime of the SMDP of Figure 3.5, although it emerged only from the instant $t = 220$ s in Figure 3.8. In Example 3.21, no proper affine regime was reached but the value function still admitted a pseudo-affine growth from the instant $t = 160$ s in Figure 3.5, with similar slope than Example 3.25. We shall thoroughly explain in the next chapter how to determine this slope based on (3.26).

► **REMARK 3.33.** When we work in SMDP models where the sojourn times are non-deterministic, it is remarkable that the optimality equations (OE1) and (OE2) still take the same form, replacing for all a in A the constant time t^a by its expectation $\int_0^\infty s \, d\nu_a(s)$ (see again [Sch92]). This means that the average sojourn times associated with the actions are representative of the expected escape rate of the dynamics. In particular, this allows us to extend Corollary 3.32 in the context of stochastic sojourn times, provided the evolution semigroup is well-posed as discussed in Section 2.8.

3.2 SECOND ORDER RESULTS IN CASES REDUCIBLE TO MDP. — Obtaining finer convergence results than (3.27) is in general difficult. In the Section 4 to come, we shall focus on a particular case of interest that remains in the semi-Markov setting. On the other hand, we present in this subsection a couple of results that can be applied to the MDP framework, and more generally to situations where the sojourn times are multiple of a same positive real number, see the discussion on the reductions in Section 2.5. Hence, we focus again on operators that share common properties with the Bellman operator T , and as in the previous subsection we write $\chi(T)$ instead of $\chi(\mathcal{S})$ for the escape rate of the MDP-type dynamics.

On the model of Proposition 3.2 presented in Section 1, the evolution of the value function often involves periodic phenomena. First, the next theorem addresses the case in which $\chi(T) = 0$. We denote by $\text{Sym}(n)$ the symmetric group (set of permutations) on n letters.

- **THEOREM 3.34** ([Mar89, LN12], [LS05] and [AG03]). *Suppose that $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is nonexpansive in a polyhedral norm and that the orbits $\{T^k(x) \mid k \in \mathbb{N}\}$ of T are bounded for all x in \mathbb{R}^n . Then, for every x in \mathbb{R}^n , there exists an integer c bounded only as a function of the norm such that $T^{kc}(x)$ converges as $k \rightarrow \infty$. Moreover, if T is order-preserving and nonexpansive in weighted sup-norm, then $c \leq \binom{n}{\lfloor n/2 \rfloor}$. If in addition T is concave, then c is the order of an element of $\text{Sym}(n)$.*

The first part of the theorem was proved in [Mar89] and in several other works, see the discussion in [LN12]. The bound of c in the order-preserving and sup-norm nonexpansive case is established in [LS05]. The bound in the concave case is established in [AG03]. The following result deals with a special case of dynamics where $\chi(T)$ is a scalar multiple of $\mathbf{1}$, which typically occurs in the theory of unichain Markov decision processes as explained before. Recall that for x in \mathbb{R}^n , the quantity $\|x\|_H := \max_{1 \leq i \leq n} \{x_i\} - \min_{1 \leq i \leq n} \{x_i\}$ is called the Hilbert seminorm of x .

- **THEOREM 3.35.** *Suppose that $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is order-preserving and additively homogeneous. Suppose in addition that for all x in \mathbb{R}^n , the sequence of Hilbert's seminorms $(\|T^k(x)\|_H)_{k \in \mathbb{N}}$ is bounded. Then, for all x in \mathbb{R}^n , there exists an integer c that can be bounded as in Theorem 3.34 such that for all $0 \leq r \leq c - 1$, $T^{kc+r}(x) - \chi(T)(kc + r)$ converges as $k \rightarrow \infty$.*

Proof. It follows from [GG04a] that T has an additive eigenvector, meaning that there exists $u \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}$ such that $T(u) = u + \lambda \mathbf{1}$. Then, the map $G := T - \lambda \mathbf{1}$ has a fixed-point, and it is still order-preserving and sup-norm nonexpansive. It follows from Theorem 3.34 that for every $x \in \mathbb{R}^n$, there exists an integer c such that $G^{kc}(x)$ converges as $k \rightarrow \infty$.

Since G is continuous, $G^{kc+r}(x) = T^{kc+r}(x) - (kc + r)\lambda \mathbf{1} = T^{kc+r}(x) - (kc + r)\chi(T)$ also converges as $k \rightarrow \infty$. The bounds on c follow from the ones of Theorem 3.34. \clubsuit

The next theorem is stated in [SF77] for operators of multichain Markov decision processes. We provide a more abstract (equivalent) statement.

- **THEOREM 3.36** (Schweitzer and Federgruen, [SF77]). *Suppose that $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is concave, order-preserving, additively homogeneous, and piecewise linear. Then, for all x in \mathbb{R}^n , there exists an integer c that is the order of an element of $\text{Sym}(n)$, such that for all $0 \leq r \leq c - 1$, the sequence $T^{kc+r}(x) - (kc + r)\chi(T)$ converges as $k \rightarrow \infty$.*

By comparison with Theorem 3.35, the map T is required in addition to be *concave* and *polyhedral*. The concavity assumption leads to a refined explicit formula on the period c (given by a combinatorial invariant of a certain critical graph depending only on optimal stationary randomized policies, see [AG03]). The polyhedrality assumption allows one to avoid the restriction to maps whose orbits are bounded in Hilbert's seminorm.

4 Stochastic Shortest Path configurations

In general and as illustrated by Corollary 3.32–(2), taking the limit (or the limit inferior for the sake of definition) of the value function in (3.5) when t tends to ∞ leads to a non-finite result, which is why one often studies the discounted cost or the average-cost criteria, as they remain finite over an infinite horizon. However, in their article [BT91], Bertsekas and Tsitsiklis have identified a subclass of Markov decision processes and a set of assumptions under which the criterion

$$u^*(i) := \liminf_{t \rightarrow \infty} v^*(i, t) \tag{3.28}$$

stays finite for all i in S .

4.1 TOPOLOGICAL AND COST-RELATED RESTRICTIONS. — Following the terminology introduced by Bertsekas and Tsitsiklis, we say that an MDP is in **Stochastic Shortest Path configuration** (abb. SSP configuration) if there is a distinguished state of S denoted by 0 such that any playable action from state 0 has null cost and forces to stay in 0 (for all a in A_0 , $c^a = 0$ and $p_0^a = 1$). Although we are not aware of any study of the stochastic shortest path problem in the semi-Markov setting, the notion of SSP configuration is essentially topological and carries over to SMDPs. To have compatibility with Assumption **A**, we may only impose that $t^a > 0$ for all a in A_0 .

It can be seen that as soon as a decision-maker reaches state 0 while playing the game described by an SMDP in SSP configuration, he/she will stay in state 0 indefinitely. Due to the null costs, his/her total accumulated cost no longer evolves from this moment onwards. Hence the game virtually stops and it makes sense to introduce (3.28). The “stochastic shortest path” denomination comes from the fact that if we are bound to reach 0 from state i in S , the limit (3.28) of the value function corresponds to the minimal cost of going from i to 0 , knowing that we are subject to some non-deterministic moves after picking actions. If these transition probabilities were zero or one, we would retrieve the well-known shortest path problem in a graph.

The mere SSP configuration in terms of topology is not sufficient to obtain the finiteness of u^* . To ensure that the SSP problem is well-posed and that ultimate reachability of state 0 is guaranteed, the notion of proper policy is often introduced.

➤ **DEFINITION 3.37.** *We say that a policy σ in \mathfrak{S} is **proper** if for all i in S , $\lim_{n \rightarrow \infty} [(P^\sigma)^n]_{i0} = 1$. A non-proper policy is called **improper**.*

An equivalent formulation of the fact that the policy σ in \mathfrak{S} is proper is that $\mathcal{F}(\sigma) = \{\{0\}\}$, meaning that, the singleton $\{0\}$ is the unique final class of the Markov chain induced by σ . This way, we almost surely end up “trapped” in state 0 when we apply the stationary strategy associated with σ , regardless of the initial state. It remains to guarantee that proper policies are more interesting to play than improper ones. Indeed, a player has no interest in preferring a proper policy if an improper one gives him/her negative costs (that is, positive rewards) in the long term. Alternatively, one can see based on Corollary 3.32–(2) that a restriction is in order to ensure $\chi(\mathcal{S}) = \mathbf{0}$, that is a necessary condition for (3.28) to be defined. These ideas justify the second item of the Assumption **C** enounced by Bertsekas and Tsitsiklis in [BT91], where the cost of a policy is understood as the expected cost of the stationary strategy associated with this policy, see Section 2.2.

➤ **ASSUMPTION C.** *The SMDP is in SSP configuration, and is such that*

(1) *there exists a proper policy,*

(2) *every improper policy yields a cost of $+\infty$ for at least one initial state.*

One of the main results on SSP configurations, shown by Bertsekas and Tsitsiklis for the MDP case, is the following.

➤ **THEOREM 3.38** (Bertsekas and Tsitsiklis, [BT91]). *Suppose that Assumption **C** holds for a Markov decision process in SSP configuration. Then, the equations*

$$\begin{aligned} \forall i \in S \setminus \{0\}, \quad u(i) &= \min_{a \in A_i} \left\{ c^a + \sum_{j \in S} p_j^a u(j) \right\} \\ \text{and} \quad u(0) &= 0 \end{aligned} \tag{3.29}$$

admit a unique solution in \mathbb{R}^S , which coincides with the value u^ defined in (3.28).*

In other words, Theorem 3.38 claims that the criterion u^* of (3.28) is finite, and that it arises as the unique fixed-point u of the Bellman operator T for the MDP problem defined in (3.13) such that $u(0) = 0$. Note actually that the SSP value (3.28) in [BT91] is only defined in terms of pure Markovian strategies (that is to say elements of \mathbb{F}^{MD}), but Bertsekas recalls in his book [Ber11] the adequacy of Markovian strategies and in the joint work with Yu [BY13] justifies that randomized strategies do not improve the criterion defined over pure ones. As a consequence, Theorem 3.38 does apply to the definition (3.28), where an infimum over \mathbb{F} is hidden.

4.2 CONVERGENCE IN THE SEMI-MARKOV SETTING. — The other key result shown in [BT91] is that in the MDP setting, we have for all x in \mathbb{R}^S that $\lim_{t \rightarrow \infty} T^t(x) = u^*$ if $x_0 = 0$, in other words, the total cost u^* arises as the limit of the value iteration procedure applied to any starting vector (when the latter has a null coordinate associated with state 0).

We want to determine if an equivalent result holds in the framework of semi-Markov decision processes, where sojourn times are involved and knowing that we still authorize some sojourn times to be null. In this subsection, we suppose that we study an SMDP in SSP configuration and we denote by $(\mathcal{S}_t)_{t \geq 0}$ its corresponding evolution semigroup. Because the state 0 is cost-free, the set of initial conditions on which this semigroup naturally acts is the subset \mathcal{V}_0 of \mathcal{V} defined by $\mathcal{V}_0 := \{v \in \mathcal{V} \mid v_0(s) = 0 \text{ for all } s \in [-t_{\max}, 0)\}$.

We will now follow a line of proof that is close to the one of [BT91, Proposition 2] although resorting to the semigroup formalism, which avoids some cumbersome steps. In what follows, we denote by \tilde{u}^* the constant function of \mathcal{V}_0 that is equal to the vector u^* determined by Theorem 3.38.

► **LEMMA 3.39.** *Suppose that Assumptions A and C hold. If v^0 in \mathcal{V}_0 is a fixed-point of \mathcal{S}_t for all $t \geq 0$, then $v^0 = \tilde{u}^*$. Conversely, \tilde{u}^* is a fixed-point of \mathcal{S}_t for all $t \geq 0$.*

Proof. We begin by proving the first assertion. Suppose we have such a v^0 , so that

$$\forall t \geq 0, \quad \forall s \in [-t_{\max}, 0), \quad \mathcal{S}_t[v^0](s) = v^0(s).$$

By Definition 3.23, letting $t = -s$ in the above yields $v^0(s) = v(0)$ for all $s \in [-t_{\max}, 0)$ where v is the function uniquely determined by v^0 and (DP). As a result v^0 is constant. Reusing the notation $(A_{i,Z})_{i \in S}$ from the proof of Proposition 3.17, the same (DP) relationship tells us what this constant value is, since $v(0)$ verifies that for all $i \in S$,

$$\begin{aligned} v(i, 0) &= \min_{a \in A_{i,Z}} \left\{ c^a + \sum_{j \in S} p_j^a v(j, 0) \right\} \wedge \min_{a \notin A_{i,Z}} \left\{ c^a + \sum_{j \in S} p_j^a v^0(j, -t^a) \right\} \\ &= \min_{a \in A_i} \left\{ c^a + \sum_{j \in S} p_j^a v(j, 0) \right\}, \end{aligned}$$

where we have reused the fact that $v^0(s) = v(0)$ for s in $[-t_{\max}, 0)$. Thus, not only do we have that v is null on state 0, but we also have $v(0) = T(v(0))$. Since the Bellman operator T admits a unique fixed-point with null component on state 0 according to Theorem 3.38, we have $v(0) = u^*$, hence $v^0 = \tilde{u}^*$.

◊ Conversely, we can apply the induction principle (\mathcal{I}_Z) to show that $\mathcal{S}_t[\tilde{u}^*] = \tilde{u}^*$ for all $t \geq 0$. Let us call v the function uniquely determined by \tilde{u}^* and (DP), we suppose that there is $t_0 \geq 0$ such that $v(s) = u^*$ for all s in $[-t_{\max}, t_0)$. If $t \in [t_0, t_0 + t_{\min})$, we have using (3.14) that $v(i, t) = u^*(i)$ for all non-Zeno state i in S_{NZ} . By replacing these information in the equations (3.16), given by (DP) for the possibly Zeno states, we see that $v_Z(t)$ is solution of the dynamics if and only if $u := (u^*|_{NZ}, v_Z(t))$ satisfies $u = T(u)$, hence $v_Z(t) = u^*|_Z$, so the induction is shown. \clubsuit

The following lemma relies on a large extent on the fact that the evolution semigroup is order-preserving, and is the only place where we need to use its continuity with respect to the pointwise topology.

- **LEMMA 3.40.** *Suppose that Assumptions A and C hold. If there is v^0 in \mathcal{V}_0 such that for all $t \geq 0$, we have $\tilde{u}^* \leq \mathcal{S}_t[v^0] \leq v^0$, then $\mathcal{S}_t[v^0]$ converges pointwise towards \tilde{u}^* . The result also holds if conversely, we have $v^0 \leq \mathcal{S}_t[v^0] \leq \tilde{u}^*$.*

Proof. Suppose that we have for all $t \geq 0$ the pointwise inequality $\tilde{u}^* \leq \mathcal{S}_t[v^0] \leq v^0$ in \mathcal{V} . By monotonicity of the semigroup (Proposition 3.26–(2)) and by the fixed-point property of \tilde{u}^* (Lemma 3.39), we have for all $t' \geq 0$ that $\tilde{u}^* \leq \mathcal{S}_{t+t'}[v^0] \leq \mathcal{S}_{t'}[v^0]$. It follows that for all $s \in [-t_{\max}, 0)$, $t \mapsto \mathcal{S}_t[v^0](s)$ is a non-increasing and bounded from below function. Therefore, $\mathcal{S}_t[v^0]$ converges pointwise towards a function $w \in \mathcal{V}_0$. Taking the limit when $t' \rightarrow \infty$ in $\mathcal{S}_{t+t'}[v^0] = \mathcal{S}_t[\mathcal{S}_{t'}[v^0]] \leq \mathcal{S}_{t'}[v^0]$ yields by pointwise continuity of the semigroup (Proposition 3.26–(4)) $\mathcal{S}_t[w] \leq w$ for all $t \geq 0$. On the other hand, for w is a non-increasing limit we have for all $t, t' \geq 0$ that $w \leq \mathcal{S}_{t+t'}[v^0] = \mathcal{S}_t[\mathcal{S}_{t'}[v^0]]$. Letting $t' \rightarrow \infty$ entails $w \leq \mathcal{S}_t[w]$ for all $t \geq 0$. As a result, w is a fixed-point of \mathcal{S}_t in \mathcal{V}_0 for all $t \geq 0$, so we must have $w = \tilde{u}^*$ by Lemma 3.39. \clubsuit

We now state that the value iteration algorithm also converges for SMDPs in SSP configurations.

- **THEOREM 3.41.** *Suppose that Assumptions A and C hold. Then, for all initial conditions v^0 in \mathcal{V}_0 and for all s in $[-t_{\max}, 0)$, we have*

$$\lim_{t \rightarrow \infty} \mathcal{S}_t[v^0](s) = u^* .$$

Proof. Take v^0 in \mathcal{V}_0 . We let $\delta := \|v^0 - \tilde{u}^*\|_\infty$ and denote by $\Delta \in \mathbb{R}^S$ the vector such that $\Delta_0 = 0$ and $\Delta_i = \delta$ if $i \in S \setminus \{0\}$. We denote by $\tilde{\Delta}$ the constant function $s \mapsto \Delta$ of \mathcal{V}_0 . By definition we have $-\tilde{\Delta} \leq v^0 - \tilde{u}^* \leq \tilde{\Delta}$ (finer than a double inequality involving $\delta \mathbf{1}$ since v^0 and \tilde{u}^* are both null on coordinate 0).

We introduce the two functions \underline{v} and \bar{v} in \mathcal{V}_0 defined by $\underline{v} := \tilde{u}^* - \tilde{\Delta}$ and $\bar{v} := \tilde{u}^* + \tilde{\Delta}$. It is clear that we have

$$\underline{v} \leq v^0 \leq \bar{v} \quad \text{and} \quad \underline{v} \leq \tilde{u}^* \leq \bar{v} .$$

The monotonicity of the semigroup for all $t \geq 0$ (Proposition 3.26–(2)) and the fixed-point property of \tilde{u}^* (Lemma 3.39) lead to

$$\mathcal{S}_t[\underline{v}] \leq \mathcal{S}_t[v^0] \leq \mathcal{S}_t[\bar{v}] \quad \text{and} \quad \mathcal{S}_t[\underline{v}] \leq \tilde{u}^* \leq \mathcal{S}_t[\bar{v}] .$$

By non-expansiveness of the evolution semigroup (Proposition 3.26–(3)), we have

$$\|\mathcal{S}_t[\tilde{u}^*] - \mathcal{S}_t[\underline{v}]\|_\infty \leq \|\tilde{u}^* - \underline{v}\|_\infty$$

which we can rewrite as follows using the monotonicity of the semigroup (we recall that \mathcal{V}_0 is stable under the action of \mathcal{S}_t for all $t \geq 0$), the fact that $\tilde{u}^* - \underline{v} = \tilde{\Delta}$ and $\mathcal{S}_t[\tilde{u}^*] = \tilde{u}^* \geq \mathcal{S}_t[\underline{v}]$:

$$\tilde{\mathbf{0}} \leq \mathcal{S}_t[\tilde{u}^*] - \mathcal{S}_t[\underline{v}] \leq \tilde{\Delta} .$$

By the fixed-point property once again, the right inequality above provides $\underline{v} = \tilde{u}^* - \tilde{\Delta} \leq \mathcal{S}_t[\underline{v}]$.

Conversely, we also have by similar arguments that

$$\|\mathcal{S}_t[\tilde{u}^*] - \mathcal{S}_t[\bar{v}]\|_\infty \leq \|\tilde{u}^* - \bar{v}\|_\infty$$

which leads to the double inequality

$$-\tilde{\Delta} \leq \mathcal{S}_t[\tilde{u}^*] - \mathcal{S}_t[\bar{v}] \leq \tilde{\mathbf{0}}$$

where we retain the left inequality to obtain $\mathcal{S}_t[\bar{v}] \leq \tilde{u}^* + \tilde{\Delta} = \bar{v}$.

We have satisfactorily shown that for all $t \geq 0$, we have

$$\underline{v} \leq \mathcal{S}_t[\underline{v}] \leq \tilde{u}^* \leq \mathcal{S}_t[\bar{v}] \leq \bar{v} ,$$

we can then invoke Lemma 3.40 for both \underline{v} and \bar{v} , and we obtain that the limits of $\mathcal{S}_t[\underline{v}]$ and $\mathcal{S}_t[\bar{v}]$ when t tends to ∞ do exist and are equal to \tilde{u}^* . Using the bounds stated on $\mathcal{S}_t[v^0]$, we conclude that $\lim_{t \rightarrow \infty} \mathcal{S}_t[v^0] = \tilde{u}^*$. \clubsuit

Remark how the Theorem 3.41 states that the minimum ultimate total expected cost of an SMDP in SSP configuration is the same than in a corresponding MDP with identical topology and costs but with unit sojourn times. Although counter-intuitive in the first place, this result is a consequence of the fact that the criterion (3.28) is a time-indifferent quantity, as an infinite-horizon limit. This means that if a player wants to minimize his/her total cost in an SMDP in SSP configuration, he/she should play exactly the same optimal strategy than if all the sojourn times were equal to one, for he/she actually does not care on the horizon. Naturally, another (horizon-dependent) strategy would be used if the game had to be played in finite-horizon, but this is not the case here.

4

THE CORRESPONDENCE BETWEEN MONOTONIC FLUID PETRI NETS AND SEMI-MARKOV DECISION PROCESSES

Contents

1	A running example: an emergency call center for medical needs	136
1.1	Description of the SAMU model	136
1.2	A simplified model	138
2	The finite-horizon correspondence theorems	140
2.1	Statement of the results	140
2.2	Illustration on the example EMS-A	144
3	Studying the throughput of fluid Petri nets through SMDP tools	148
3.1	Existence and universality	148
3.2	The optimality equations	150
3.3	The throughput complex	153
3.4	Linear programming formulation	158
4	Deviation to a congestion-free regime	159
4.1	Reduction to a stochastic shortest path problem	160
4.2	Transience time needed to catch up the input	165
5	Convergence time of the value iteration algorithm for semi-Markov SSP problems	168
5.1	The different convergence phases	168
5.2	Hierarchical SSP configurations	173
5.3	Convergence by above	174
5.4	Convergence by below	176
6	Application to a staffing problem of the SAMU	185

We arrive at the cornerstone chapter of the thesis. At this point, the reader who has had a careful reading of the previous chapters may already have guessed what is on our mind.

The dynamical equations of the fluid models of Petri nets that we have discussed in Chapter 2 (with no priority rules) indeed bear much resemblance with the dynamic programming equations of the semi-Markov decision processes introduced in Chapter 3. The first goal of this chapter is to dig into this direction and study what these two families of models have in common, which takes the form of four **correspondence theorems**. The first stone of this approach was laid by Cohen, Gaubert and Quadrat in the articles [CGQ98, CGQ95]. The second objective in this chapter is to reap the benefits of these correspondences, stating what the known properties on SMDPs can bring to the realm of Petri nets. In particular, the correspondence theorems enable us to derive precise results on the most insightful measure of performance of the long-run behaviour of Petri nets, their throughput. By comparison with [CGQ95] who only studied the unichain case, we derive more general characterizations of the throughput vector, valid in the multichain case and. We refine further their asymptotic expansion of the Petri net trajectories, also derived by Gaujal and Giua in [GG04b] under milder assumptions. Most importantly, our more general correspondence theorems with SMDPs allow us to accurately deal with inputs of Petri nets.

More remarkable than the first-order estimate that is the throughput, we are also able to obtain results on the transience regimes of controlled Petri nets and compute upper bounds on the time needed for these nets to reach an input-driven regime. These quantitative techniques and other second-order results take the analysis of monotonic fluid Petri nets a stage further. Before, the question of convergence to a periodic or stationary regime has received much attention for the subclass of timed event graphs (that we talked about in Section 1 of Chapter 1) – see for instance [BCOQ92, Theorem 3.109] and also [HOvdW05, Chapters 8 and 9] with an application to the Dutch railway network. For such models, the duration of the transient behaviour, also sometimes called the “coupling time”, has been extensively studied, in particular by techniques of max-plus spectral theory [BG01, SS12, MNS14, MNS21]. The same problem has arisen in the setting of deterministic dynamic programming [HA99], and in the analysis of distributed algorithms [CBFN13].

Here, we extend these results passing from timed event graphs to fluid timed Petri nets with proportional preselection routings, or equivalently, passing from deterministic (semi-)Markov decision processes to stochastic (semi-)Markov decision processes. By comparison with the “deterministic” case, our proofs require new tools. Theorem 4.13 for instance exploits techniques of non-linear Perron–Frobenius theory [AGN11, AGQS19] in order to characterize the property of convergence in finite time of the value iteration algorithm for SMDPs. To our knowledge, no characterization of the finite time convergence was known, even in the setting of MDPs. Note that estimating the speed of convergence to the stationary regime for Markov decision processes is indeed a difficult and classical issue. General asymptotic convergence results, like the ones of [SF77, SF78, SF79], show that a convergence does occur with an ultimately geometric rate. However, they lead to bounds and speed estimates that are nonconstructive. An explicit bound of the time needed to enter in the geometric convergence regime was given in [Bon07], for stochastic shortest path configurations, supposing that all costs are positive. In contrast, we consider here the property of *finite time* convergence, leading to different bounds.

We endeavour to illustrate our theoretical results on a simple model of an idealized emergency call center presented in Section 1, originating from a real-life application and partnership with the medical emergency “SAMU” services of Paris area. This case-study shall be referred to extensively throughout the chapter and we hope that these interludes will give the reader an in-depth understanding of it. We would like to insist on how extraordinary fruitful this apparent

toy model has been in the development of our correspondence results, and again thank the SAMU for their help and support. The Section 2 is dedicated to the establishment of formal correspondence theorems between timed Petri net and SMDP models. The most salient elements of these correspondences are summarized in the Table 4.1. We then analyze the throughput vector of Petri nets by means of our correspondence tools in Section 3. In order to focus on transient regimes of Petri nets, we show in Section 4 that it suffices to study the convergence of the value iteration for SMDPs in stochastic shortest path configurations. Our original convergence results for both the general and hierarchical cases are presented in Section 5 (with explicit bounds given for the latter). We finish this chapter in Section 6 with an illustration of the practical staffing recommendations we have delivered to the four SAMU of AP-HP (SAMU 75, 92, 93 and 94) in [ABG20b], building on our analytical formulas.

Timed Petri nets	Semi-Markov decision processes
Transitions	States
Places	Actions
Holding times	Sojourn times
Initial markings	Costs
Physical time	Planning horizon
Counter functions	Finite-horizon value function
Synchronization	Multiple playable actions
Meeting places	Probabilistic moves
Throughput	Average-cost
Bottleneck places	Optimal policies
Congestion phases	Cells of the average-cost complex
Input catch-up times	Convergence time of value iteration

Table 4.1: Correspondence between Petri nets and semi-Markov decision processes

1 A running example: an emergency call center for medical needs

So far, the only place in the thesis where we have talked about models of call centers was the Example 2.6 in Chapter 2, where a “single-tier” organization was introduced, represented by means of a continuous-relaxed Petri net, and its dynamics was simulated in three simple cases.

We present in this section models of Petri nets that are a little more complex and represent real-life medical emergency call centers. They are drawn from a partnership with the Emergency Medical Services (EMS) of Paris and its inner suburbs, also called SAMU.¹ We refer the reader to the Section 2 of the [Introduction](#) where we give a more detailed view of the different French emergency services. In this section, we only detail the call handling procedures that are useful to understand the case of Parisian EMS.

1.1 DESCRIPTION OF THE SAMU MODEL. — In case of medical distress or need for urgent medical advice, French people can call the nation-wide phone number 15, and the calls are then dispatched to regional call centers located in hospital facilities. The calls are first answered by a pool of operators referred to as **medical regulation assistants** (MRA). These agents are not physicians, but are trained to categorize the requests and direct the patients to the correct interlocutor. Indeed, after taking note of essential personal information, medical regulation assistants transfer the calls to one of the following two types of physicians, depending on the estimated severity of the case:

- (i) an **emergency doctor**, able to dispatch Mobile Intensive Care Units or first-responding ambulances and to swiftly send the patient to the most appropriate hospital unit;
- (ii) a **general practitioner** (GP), who can dispatch ambulances and provide medical advice.

Alternatively, a MRA may also handle the call without transferring it, if a conversation with a physician is not needed (report from a medical partner, dial error, etc.).

There is a major difference of treatment between cases (i) and (ii). When dealing with a life-threatening or severe situation (case (i)), the MRA must wait for an emergency physician to be available before transferring the call, in order to report the details of the request. In this way, the patient is constantly kept on line with an interlocutor – this mechanism may be referred to as “escorted transfer”. In case (ii), patients are left on hold in a phone waiting room, and general practitioners pick up the calls in the order of arrival once they are available.

We depict in Figure 4.2 a Petri net that represents this organization. We describe its behaviour, insisting in particular on the “nature” of the tokens, since as underlined in Section 1.2 of Chapter 1, tokens in different places of a Petri net can model objects of different types. We use the color blue (resp. purple, resp. red) to highlight the circuits involving the MRAs (resp. the GPs, resp. the emergency physicians). For the sake of readability, patient exits are not shown.

The uppermost transition of the net is an input transition, that controls the arrivals of inbound calls, represented by the counter function z_0 (to keep notation light, we do not give labels anymore to transitions and refer to them directly by their counters). The arrived calls are stocked in a first place before being picked up by medical regulation assistants. The total number of MRAs is N_A , and they are initially in a pool waiting for calls to arrive, modeled by means of a place with zero holding time. When a call token and a MRA token are both available, they are consumed to produce a “patient+MRA” token in a place with zero holding time and ruled by preselection with three downstream transitions, so that this token is immediately directed towards one of the three main lanes of the system, according to long-run proportions π_A , π_G and π_P .

¹ In French *Service d’Aide Médicale Urgente*.

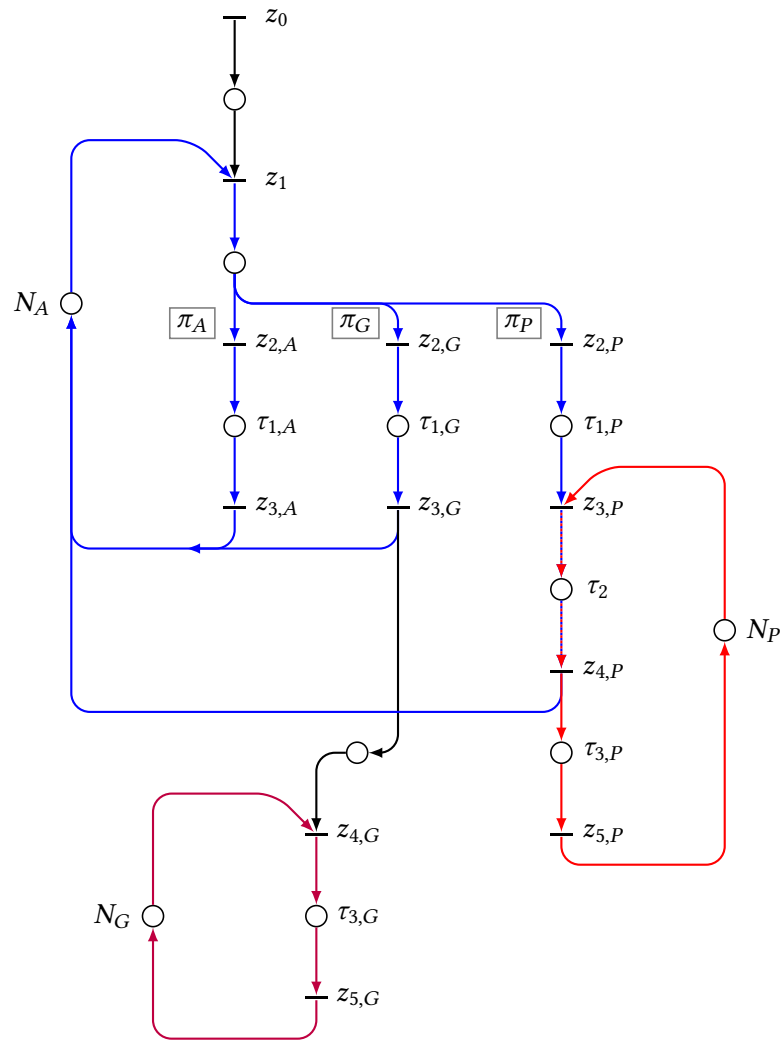


Figure 4.2: A medical emergency call center with two types of physicians

The calls that are not transferred to emergency physicians or GPs are counted by $z_{2,A}$. After a short time $\tau_{1,A}$ spent in a conversation place, these calls are hung up by MRAs who go back to their pool (returns counted by $z_{3,A}$), i.e., the “patient+MRA” token is consumed to produce an MRA token again. The counter $z_{2,G}$ gives the number of calls picked up by MRAs and directed to the lane of general practitioners. MRAs need a time $\tau_{1,G}$ to collect basic information on patients and send them in the waiting room (with count $z_{3,G}$), this last operation consumes a “patient+MRA” token and produces one MRA token in their pool place, and one patient token in the waiting room. The latter is by essence a place with no holding time and as soon as GPs (N_G doctors in total), are available from their base pool, they can pick up calls from the waiting room, thus incrementing the count $z_{4,G}$ (this creates a “patient+GP” token by consuming a GP token and a patient token). After a consultation with duration $\tau_{3,G}$, patients leave the system and GPs re-enter their pool (i.e., transition $z_{5,G}$ counts the number of times where a “patient+GP” token is consumed to produce a GP token).

Finally, calls allocated to emergency physicians, in amount $z_{2,P}$, take an instruction time $\tau_{1,P}$ to MRAs, and after this time is elapsed (which means that the “patient+MRA” token is mature), an additional wait may be experienced before the availability of an emergency physician

in the pool with staffing N_P . At this point, a “patient+MRA” token and an emergency doctor token are consumed, which increments $z_{3,P}$ and produces a “patient+MRA+physician” token in the place with holding time τ_2 . This duration represents the time needed by the MRA to give the emergency physician a very short briefing on the call. By firing the transition $z_{4,P}$, a “patient+MRA+physician” token is split into an MRA token on one hand, who can return in the MRA pool, and a “patient+physician” token on the other hand, which enters a consultation place with holding time $\tau_{3,P}$. At the end of this conversation, the patient is released from the system, and the physician becomes ready again to pick up new calls.

The holding times that are considered here are constant. It surely is not very realistic, but for several reasons that we shall discuss in Section 2 and in Remark 4.6, it is a fruitful first approximation. As introduced in Section 3.3 of Chapter 1, we can alternatively suppose that the handling times of the calls are stochastic, drawn from place-dependent distributions.

In addition, and similarly to the single-tier call center of Example 2.6, we assume that the callers have an **infinite patience** in the places with zero holding time, meaning that they do not hang up on their own. This can be seen as the main blind spot of our model. In our Part II where call centers are studied using numerical simulation and the real patience of callers is taken into account, we shall see that analytical results derived from infinite-patience models remain very good approximations of realistic-patience models in many cases. We will thoroughly discuss there (see in particular the Section 5 of Chapter 7 and the Section 3 of Chapter 8) the impact of the patience in practice.

1.2 A SIMPLIFIED MODEL. — A few simplifications can be brought to the Petri net of Figure 4.2 without quite altering its dynamics. To begin with, the place ruled by preselection (downstream of transition z_1) that dispatches the picked up calls in the three lanes is convenient to account for the different handling times by the MRA of the three types of calls,² but on the other hand produces more parameters to work with. In addition, this place is somehow an artefact since in practice, the flow of picked up calls is split after the conversation with the MRA. An alternative design is then to have a single place of handling by the MRA with time τ_1 , still ruled by preselection to distribute the calls in the lanes. To ensure consistency between the two models, we would typically take $\tau_1 := \pi_A \tau_{1,A} + \pi_G \tau_{1,G} + \pi_P \tau_{1,P}$ when working with constant holding times, or construct an average distribution if working with stochastic holding times.

Besides, we will want as a first step to focus on the coupling between the answering operator and the emergency physician; indeed the synchronization step at transition $z_{3,P}$ can be seen to be the most critical link of the system, the only one involving two different agents. On the other hand, we remark that the circuit of general practitioners is quite independent of the rest of the system, and from the MRA point of view, finishing the instruction of a call handed over to a GP or not transferred to a physician at all is equivalent; they just go back to their pool. As a result, we will consider in what follows a simplified model where these two lanes are “merged” and where GPs are ignored. In this call center, only two types of inbound calls can occur, the ones which require the MRA to wait for an emergency physician, and the ones which do not. Again, to keep consistency with the previous model, these two types of calls would arise in proportions π_P and $\pi_A + \pi_G$.

² Based on the data analysis of emergency call logs from four different SAMU call centers, it turns out that it takes less time to handle calls not transferred to physicians ($\tau_{1,A}$), and that it is longer for an MRA to complete the instruction of a case to be passed to an emergency physician ($\tau_{1,P}$) than to a general practitioner ($\tau_{1,G}$). Typical values are $\tau_{1,A} \approx 60$ s, $\tau_{1,G} \approx 110$ s and $\tau_{1,P} \approx 130$ s.

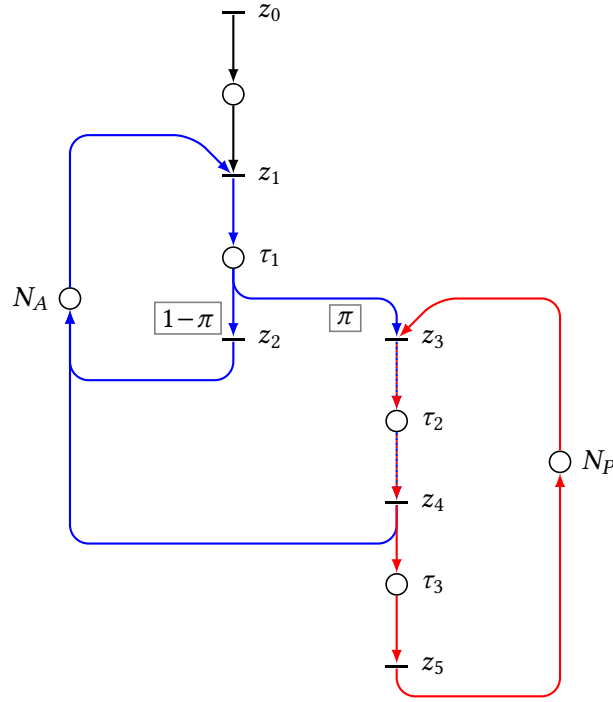


Figure 4.3: The medical emergency call center EMS-A, simplified version of the one in Figure 4.2

This simplified model is depicted in Figure 4.3. We still denote by N_A (resp. N_P) the total number of MRAs (resp. emergency doctors). The fraction of calls requiring a transfer to emergency physicians (formerly π_P) is denoted by π . Medical regulation assistants need a characteristic holding time τ_1 to handle inbound calls, a time τ_2 to transfer cases to emergency doctors, and the latter need a time τ_3 to conduct a phone consultation with patients. Remark that the transition z_3 is downstream of a place ruled by preselection, and that it admits another upstream place, this pattern is thus to be understood in the sense of the reduction of Figure 1.6.

Although the real call center is subject to integrality constraints (there are no such things as half a call or half an agent), we choose to study its continuous-relaxed dynamics. As justified in Chapter 2, this is a sound approximation whose validity increases with the size of the model (number of tokens navigating through it). Applying the fluid equations of Table 2.3 gives the following dynamical model governing the Petri net of Figure 4.3. We use the symbol \wedge to denote the minimum operator (i.e., $x \wedge y = \min(x, y)$).

$$\begin{cases} z_1(t) = z_0(t) \wedge (N_A + z_2(t) + z_4(t)) \\ z_2(t) = (1 - \pi)z_1(t - \tau_1) \\ z_3(t) = \pi z_1(t - \tau_1) \wedge (N_P + z_5(t)) \\ z_4(t) = z_3(t - \tau_2) \\ z_5(t) = z_4(t - \tau_3) \end{cases} \quad (\text{EMS-A})$$

In the rest of the thesis, we will refer to this system of equations as (EMS-A). Note that this model is priority-free. An alternative organization of a SAMU call center, implementing priority rules, will be presented in Section 3 of Chapter 5, and referred to as (EMS-B) by comparison.

2 The finite-horizon correspondence theorems

We are now going to establish a formal correspondence between monotonic fluid Petri nets and SMDPs. In the whole chapter, the Petri net models are supposed to be **priority-free** (hence associated with standard monotonic dynamics as highlighted in Section 2 of Chapter 2), and in what follows we use the terms fluid or continuous-relaxed indifferently. To get a full understanding of the results presented here, it is important to have in mind the objects introduced in Chapters 2 and 3. Moreover, we encourage the reader to read all the proofs of the current section.

2.1 STATEMENT OF THE RESULTS. — Our first correspondence theorem compares in particular the recursive equations of Table 2.3 in Chapter 2 verified by the counters of a fluid Petri net (a priori *controlled*, i.e., that admits input transitions) and the dynamic programming equations (**DP-STOP**) featured in Chapter 3 for SMDPs with stopping costs.

► **CORRESPONDENCE THEOREM 1.** *The counter function dynamics of the transitions in a continuous-relaxed and non-Zeno timed Petri net is equivalent to the discounted finite-horizon dynamic programming equations of a semi-Markov decision process with stopping costs.*

Proof. Because the Petri net is priority-free, we have $\mathcal{P}_{\text{prio}} = \emptyset$. According to the conflict-free transformation introduced in Section 1.3 of Chapter 2, we can suppose up to considering an equivalent conflict-free net that $\mathcal{P}_{\text{psel}} = \emptyset$ as well without changing the dynamics, as shown by Proposition 2.7. We are thus left with a timed Petri net such that $Q = Q_{\text{input}} \cup Q_{\text{sync}}$ and $|p^{\text{out}}| \leq 1$ for all places p in \mathcal{P} . The Table 2.3 of Chapter 2 gives us the dynamics equations followed by transitions of Q_{sync} :

$$\forall t \geq 0, \quad \forall q \in Q_{\text{sync}}, \quad z_q(t) = \min_{p \in q^{\text{in}}} \left\{ \alpha_{qp}^{-1} \left(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}(t - \tau_p) \right) \right\}. \quad (4.1)$$

Recall from Section 1.5 of Chapter 1 that without loss of generality, each input transition admits a unique downstream place, that in turn admits a single downstream transition of Q_{sync} . Conversely, if we denote by $\mathcal{P}_{\text{input}} := \bigcup_{q \in Q_{\text{input}}} q^{\text{out}}$, we can assume that for all $q \in Q$ we have $|q^{\text{in}} \cap \mathcal{P}_{\text{input}}| \leq 1$, that is to say any internal transition is constrained by at most one input. Moreover, up to changing the input profiles, we can suppose that multipliers involving places of $\mathcal{P}_{\text{input}}$ are unitary and that the associated initial markings and holding times are null. The previous equation can consequently be refined in the form

$$\forall t \geq 0, \quad \forall q \in Q_{\text{sync}}, \quad z_q(t) = \min_{\substack{p \in q^{\text{in}} \\ p \notin \mathcal{P}_{\text{input}}}} \left\{ \alpha_{qp}^{-1} \left(m_p + \sum_{\substack{q' \in p^{\text{in}} \\ q' \in Q_{\text{sync}}}} \alpha_{pq'} z_{q'}(t - \tau_p) \right) \right\} \wedge \min_{\substack{q' \in Q_{\text{input}} \\ q^{\text{in}} \cap q'^{\text{out}} \neq \emptyset}} \{ z_{q'}(t) \}, \quad (4.2)$$

where \wedge denotes the minimum as before, and the last minimum in the right-hand side is taken either on the empty set (hence it equals ∞) or on a singleton.

Introducing for all $p \in \mathcal{P} \setminus \mathcal{P}_{\text{input}}$ with unique downstream transition q the nonnegative quantities $b^p := \alpha_{qp}^{-1} m_p$, $\kappa^p := \alpha_{qp}^{-1} \sum_{q' \in p^{\text{in}}} \alpha_{pq'}$ and for all $q' \in p^{\text{in}}$ the quantity $\beta_{q'}^p := \kappa_p^{-1} \alpha_{qp}^{-1} \alpha_{pq'}$, we obtain that $(\beta_{q'}^p)_{q' \in Q}$ is a probability vector, and equation (4.1) can be rewritten as

$$\forall t \geq 0, \quad \forall q \in Q_{\text{sync}}, \quad z_q(t) = \min_{\substack{p \in q^{\text{in}} \\ p \notin \mathcal{P}_{\text{input}}}} \left\{ b^p + \kappa^p \sum_{q' \in Q} \beta_{q'}^p z_{q'}(t - \tau_p) \right\} \wedge \min_{\substack{q' \in Q_{\text{input}} \\ q^{\text{in}} \cap q'^{\text{out}} \neq \emptyset}} \{ z_{q'}(t) \}. \quad (4.3)$$

This last equation needs to be compared with the generic dynamic programming equation of a discounted SMDP with stopping costs (see equations (**DP- γ**) and (**DP-STOP**)):

$$\forall t \geq 0, \quad \forall i \in S, \quad v(i, t) = \min_{a \in A_i} \left\{ c^a + \gamma^a \sum_{j \in S} p_j^a v(j, t - t^a) \right\} \wedge R(i, t)$$

in which we see the following exact one-to-one correspondence; the counter functions of the Petri net follow the dynamic programming equations of an SMDP with $|Q \setminus Q_{\text{input}}|$ states, from the state corresponding to

$q \in \mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$ there are $|q^{\text{in}} \cap (\mathcal{P} \setminus \mathcal{P}_{\text{input}})|$ playable actions, each one of these actions correspond to an upstream place p of q , yielding a cost b^p , a sojourn time τ_p , a discount factor κ^p and transition probabilities $(\beta_{q'}^p)_{q' \in \mathcal{Q}}$, while there is also a possible stopping cost $z_{q'}$, where $q' \in \mathcal{Q}_{\text{input}}$ is the only input transition directly upstream of q (if such an input transition does not exist, it disappears of the equation). \clubsuit

The proof of Correspondence Theorem 1 and in particular the equation (4.3) shows that the transitions (resp. the places) of a Petri net are in formal correspondence with the states (resp. the actions) of an SMDP. More specifically, the *upstream* places of a transition in a Petri net correspond to the playable actions *from* a state in an SMDP, so that this is actually a **backward correspondence** (recall that tokens of a Petri net are produced in downstream places of transitions, and fire downstream transitions of places). The Figure 4.4 shows how the most elementary pattern from transition q_1 to transition q_2 with one intermediary place in a Petri net corresponds to a pattern going from state 2 to state 1 via one playable action in A_2 in the SMDP setting.



Figure 4.4: The Petri net / SMDP correspondence goes backward

It also readily follows from the proof of Correspondence Theorem 1 that the input profiles of the input transitions of the Petri net correspond to stopping costs in the SMDP setting. Before discussing any further the common features of monotonic fluid Petri nets and SMDPs, let us recall from Chapter 2 that although this is essential to our application, taking into account these externally-prescribed input profiles for Petri nets (or these stopping costs for SMDPs) increases the cumbersomeness of our formalism. We are however going to focus on the particular case where the input profiles (or the stopping costs) are affine functions. In the context of Petri nets it amounts to constant-rate arrivals of fluid-tokens, and it is a sane approximation of discrete (homogeneous) Poisson arrivals (see Section 3.2 of Chapter 2). Actually, studying the affine case is also much informative on the general case since any admissible input profile (nondecreasing function from $\mathbb{R}_{\geq 0}$ to $\mathbb{R}_{\geq 0}$) can ultimately be bounded above or below by an affine function (depending on if it is superlinear or sublinear). The following proposition shows in the context of SMDPs that *exogenous* stopping costs that are affine in the planning-horizon can be **emulated** by *endogenous* extra states.

- **PROPOSITION 4.1.** *The dynamic programming equations of an SMDP with state space S , actions A and affine stopping costs coincide with the dynamic programming equations of an SMDP with augmented state space $S \cup S_{\perp}$ and augmented action space $A \cup A_{\perp}$, with no stopping costs but with suitable terminal costs and such that all the states of S_{\perp} are absorbing.*

Proof. As recalled at the end of the proof of Correspondence Theorem 1, the generic dynamic programming equation of a discounted SMDP with stopping cost R writes

$$\forall t \geq 0, \quad \forall i \in S, \quad v(i, t) = \min_{a \in A_i} \left\{ c^a + \gamma^a \sum_{j \in S} p_j^a v(j, t - t^a) \right\} \wedge R(i, t). \quad (4.4)$$

Let us suppose that the stopping costs are affine functions, which means that for all $i \in S$, we either have two scalars λ_i and η_i such that $R(i, t) = \lambda_i t + \eta_i$ for all $t \geq 0$, or $R(i, t) = \infty$ for all $t \geq 0$. We denote $S_{<\infty} := \{i \in S \mid R(i, \cdot) < \infty\}$ the set of states where the stopping costs are effective.

For all i in $S_{<\infty}$, we introduce a new absorbing state i_{\perp} , and we denote $S_{\perp} := \{i_{\perp} \mid i \in S_{<\infty}\}$. For all i in $S_{<\infty}$, we introduce a unique playable action $a_{i_{\perp}}$ (so that $A_{i_{\perp}}$ is a singleton), with cost $\lambda_i t_{\max}$, sojourn time t_{\max} and unit discount factor (recall that $t_{\max} := \max_{a \in A} \{t^a\}$). To ensure that these states are absorbing,

we have in addition $p_j^{a_j} = 1$ for all $j \in S_\perp$. Finally, for all i in $S_{<\infty}$, we augment the set of playable actions A_i by allowing an additional action with zero cost and zero sojourn time making the player deterministically join the state i_\perp .

The dynamic programming equations of this modified SMDP are the following:

$$\forall t \geq 0, \quad \begin{cases} \forall i_\perp \in S_\perp, & v(i_\perp, t) = \lambda_i t_{\max} + v(i_\perp, t - t_{\max}) \\ \forall i \in S_{<\infty}, & v(i, t) = \min_{a \in A_i} \left\{ c^a + \gamma^a \sum_{j \in S} p_j^a v(j, t - t^a) \right\} \wedge v(i_\perp, t) \\ \forall i \in S \setminus S_{<\infty}, & v(i, t) = \min_{a \in A_i} \left\{ c^a + \gamma^a \sum_{j \in S} p_j^a v(j, t - t^a) \right\} \end{cases} \quad (4.5)$$

Now, we add to the model specifications the terminal cost $w: [-t_{\max}, 0) \rightarrow \mathbb{R}^{S \cup S_\perp}$, such that for all i in $S_{<\infty}$, we have $w(i_\perp, t) = \lambda_i t + \eta_i$ for all $t \in [-t_{\max}, 0)$. We can reuse the first equation of (4.5) to conclude that $v(i_\perp, t) = \lambda_i t + \eta_i = R(i, t)$ for all $t \in [-t_{\max}, \infty)$, and it guarantees that a solution of (4.5) restricted to states of S coincides with the solution of the initial dynamic programming equations (4.4). \clubsuit

We can combine Correspondence Theorem 1 and Proposition 4.1 to deduce that the dynamics of a continuous-relaxed Petri net (that is non-Zeno and priority-free) whose inputs have affine profiles are equivalent to the dynamic programming equations of a discounted SMDP with no stopping costs. In a sense, the Proposition 4.1 extends the Proposition 2.11, because it tells us that we can describe the dynamics of timed Petri net with affine inputs by means of an evolution semigroup and not necessarily using the more general evolution operator, as if there were no inputs (these two formalisms having been introduced in the Section 2.2 of Chapter 2).

The following result, stated directly under the assumption that the input profiles are affine, articulates a very suitable setting in which the correspondence between Petri nets and semi-Markov decision processes takes its full dimension. It hinges on the existence of a stoichiometric invariant of the continuous-relaxed Petri net (see Definition 2.13), or equivalently if the latter is conflict-free, on the existence of a T-invariant (see Definition 1.6 and Proposition 2.15).

► **CORRESPONDENCE THEOREM 2.** *The counter function dynamics in a continuous-relaxed timed Petri net with affine inputs and that admits a positive stoichiometric invariant is equivalent to the finite-horizon dynamic programming equations of an undiscounted semi-Markov decision process.*

Proof. Let e be a positive stoichiometric invariant of the net (that we again suppose conflict-free for convenience), i.e., according to Definition 2.13, a vector in $\mathbb{R}_{>0}^Q$, such that

$$\forall q \in Q_{\text{sync}}, \quad \forall p \in q^{\text{in}}, \quad e_q = \alpha_{qp}^{-1} \sum_{q' \in p^{\text{in}}} \alpha_{pq'} e_{q'}.$$

We have recalled in the proof of Correspondence Theorem 1 that without loss of generality each transition in $Q \setminus Q_{\text{input}}$ admits at most one upstream place in $\mathcal{P}_{\text{input}}$ and the latter admits a single downstream transition, with unit multipliers. Hence, again without loss of generality, for all $q_{\text{input}} \in Q_{\text{input}}$ and $q \in Q \setminus Q_{\text{input}}$, it holds that $q^{\text{in}} \cap q_{\text{input}}^{\text{out}} \neq \emptyset \implies e_{q_{\text{input}}} = e_q$. For $p \in \mathcal{P} \setminus \mathcal{P}_{\text{input}}$ such that $p^{\text{out}} = \{q\}$ and for all $q' \in Q$, we introduce the quantity $\xi_{q'}^p := \alpha_{qp}^{-1} \alpha_{pq'} e_{q'} / e_q$. It is direct that $\xi^p := (\xi_{q'}^p)_{q' \in Q \setminus Q_{\text{input}}}$ is a probability vector.

Now, we start back from the equation (4.2) of the proof of Correspondence Theorem 1, and we are able to write that for all $t \geq 0$,

$$\forall q \in Q_{\text{sync}}, \quad z_q(t) = \min_{\substack{p \in q^{\text{in}} \\ p \notin \mathcal{P}_{\text{input}}}} \left\{ \alpha_{qp}^{-1} m_p + \alpha_{pq}^{-1} \sum_{q' \in Q} \alpha_{pq'} e_{q'} (z_{q'}(t - \tau_p) / e_{q'}) \right\} \wedge \min_{\substack{q' \in Q_{\text{input}} \\ q^{\text{in}} \cap q'^{\text{out}} \neq \emptyset}} \{ z_{q'}(t) \}.$$

and as a result, for all $t \geq 0$,

$$\forall q \in Q_{\text{sync}}, \quad z_q(t)/e_q = \min_{\substack{p \in q^{\text{in}} \\ p \notin P_{\text{input}}}} \left\{ \alpha_{qp}^{-1} m_p / e_q + \sum_{q' \in Q} \xi_{q'}^p(z_{q'}(t - \tau_p) / e_{q'}) \right\} \wedge \min_{\substack{q' \in Q_{\text{input}} \\ q^{\text{in}} \cap q'^{\text{out}} \neq \emptyset}} \{z_{q'}(t) / e_{q'}\}. \quad (4.6)$$

We are now in the situation where the counters $(z_q/e_q)_{q \in Q_{\text{sync}}}$ follow a dynamics that is syntactically equivalent to the one of an undiscounted SMDP, still with stopping costs. Because the $(z_q)_{q \in Q_{\text{input}}}$ are affine, it suffices to build on Proposition 4.1 to obtain that the evolution of the counters $(z_q/e_q)_{q \in Q}$ can also be represented by dynamic programming equations of an undiscounted SMDP with no stopping costs in which input transitions correspond to absorbing states. \clubsuit

The proof of Correspondence Theorem 2 points out the fact that when e is a positive stoichiometric invariant of the continuous-relaxed Petri net, the appropriate objects to look at are the **renormalized counters** $(z_q/e_q)_{q \in Q}$. This is consistent with the fact that in the context of Petri nets, the vector e plays the same role as the vector $\mathbf{1}$ for SMDPs. Recall indeed from the remark subsequent to Proposition 2.14 that the Petri net fluid dynamics are additively homogeneous with respect to e , while those of SMDPs are additively homogeneous with respect to $\mathbf{1}$ according to Proposition 3.26–(1).

The most important comment to make is once again that transitions of the Petri net correspond to states of the SMDP, that places of the Petri net correspond to actions of the SMDP, and that the correspondence goes **backward**. In the light of (4.6), we are able to precise further that the actions of the SMDP have a sojourn time equal to the holding times of their corresponding places in the Petri net, and have immediate costs equal to the initial markings of their corresponding places up to renormalization by the stoichiometric invariant and places' post-multipliers. The transition probabilities after playing an action of the SMDP are determined by the arc multipliers of the Petri net, “conjugated” by the stoichiometric invariant.

In addition, observe how elementary incidence patterns in Petri nets discussed in Section 1.3 of Chapter 1 correspond to meaningful counterpart patterns in SMDPs. A synchronization transition in Petri nets, that have multiple upstream places, will correspond in the world of SMDPs to a state from which multiple actions can be played. Conversely, a “join” place that admits multiple upstream transitions corresponds to an action leading to multiple state, that is to say proper probabilistic moves, while places with unique upstream transitions correspond to actions with deterministic evolution. Finally, the conflict (or fork) patterns for places in Petri nets, equivalently represented by concurrency patterns for transitions using the conflict-free transformation, correspond to several actions (playable from different states) that lead to a same state. The correspondence elements discussed so far allow us to fill the first part of the Table 4.1 featured in the chapter introduction.

The Correspondence Theorems 1 and 2 explain quite naturally why we had to face for both Petri nets and semi-Markov decision processes identical difficulties in the treatment of Zeno phenomena, when null holding times or sojourn times are involved. However, the Assumption A uncovered in Chapter 3 for SMDPs using the notion of policies and final classes is more general than the Definition 1.18 of non-Zeno Petri nets, where a more naive condition on total holding times over circuits was required. Building on the Correspondence Theorem 2, the Assumption A appears to be the correct (and minimal) setting to work with non-Zeno Petri nets admitting stoichiometric invariants as well. Likewise, it is not surprising that we resorted to the two very close induction principles (\mathcal{I}) for Petri nets and (\mathcal{I}_Z) for SMDPs, the latter being a bit more general since it uses weaker non-Zeno assumptions.

Observe further that the syntactic correspondences exposed by Correspondence Theorems 1 and 2 also carry over cases where holding times of Petri nets or sojourn times of SMDPs are

non-deterministic. Indeed, compare the generic “mean-field” dynamical equation for stochastic Petri nets found out in Section 3.2 of Chapter 2

$$z_q(t) = \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} \int_0^\infty z_{q'}(t - \tau) d\mu_p(\tau) \right),$$

being understood in this case that $z_q(t) = 0$ for all q in $\mathcal{Q}_{\text{sync}}$ and $t < 0$, versus the generic dynamic programming equation much studied in the SMDP literature and recalled in the Section 2.8 of Chapter 3

$$v(i, t) = \min_{a \in A_i} \left\{ c^a + \sum_{j \in S} p_j^a \int_0^t v(j, t - s) dv_a(s) \right\}.$$

It is clear in this setting that our two correspondence theorems still apply, with identical definition of costs, transition probabilities, and discount factor. We see in addition that for a place of a Petri net with holding time distribution μ , the sojourn time distribution ν of the corresponding action in the twin SMDP is such that $\mu = \nu$ almost everywhere. Although we mainly focus in what follows on Petri nets with constant holding times, it is worth remembering that some of the coming properties can also be stated for the more general class of stochastic models.

In the rest of the chapter, we will present and discuss many implications of Correspondence Theorem 2. Because we have justified that monotonic fluid Petri nets and SMDPs are essentially the same, we will use in parallel both formalisms, switching from one to another whenever it is more convenient. These acrobatics are eased up by the fact that their respective notation have almost no collision (the letter p that designates a place in Petri nets but transition probabilities in SMDPs shall always have a clear meaning from the context).

2.2 ILLUSTRATION ON THE EXAMPLE EMS-A. — Let us now detail how Correspondence Theorem 2 applies to our running example (EMS-A). To begin with, observe that the positive vector $e := (1, 1, 1 - \pi, \pi, \pi, \pi)^\top$ is a stoichiometric invariant of the Petri net represented in Figure 4.3 (recall that π gives the proportion of inbound calls transferred to the emergency physician), since it is in the right-nullspace of the incidence matrix C of the net below and that the ratios e_2/e_1 and e_3/e_1 agree with the proportions of the preselection routing:

$$Ce = \begin{pmatrix} 1 & -1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 1 & 0 \\ 0 & 1 & -1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & -1 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 1 - \pi \\ \pi \\ \pi \\ \pi \end{pmatrix} = \mathbf{0}.$$

The Figure 4.5 depicts the equivalent conflict-free Petri net of the net of Figure 4.3 (in accordance with the conflict-free transformation introduced in Section 1.3 of Chapter 2, the place upstream of the proportional preselection routing has been duplicated). The same figure also shows the corresponding SMDP, whose actions are labeled with cost and sojourn time pairs in the same manner as Chapter 3. As set out by the correspondence theorems, each transition of the Petri net (still directly labeled by their counters z_0, z_1, z_2, z_3, z_4 and z_5) corresponds to a state of the SMDP, so that the state space of the latter is $\{0, 1, 2, 3, 4, 5\}$, and we can verify that each place of the Petri net (there are seven in total) corresponds to an action of the SMDP. Observe once again how the direction of the arcs are flipped, and how the transition z_0 that is a source in the Petri

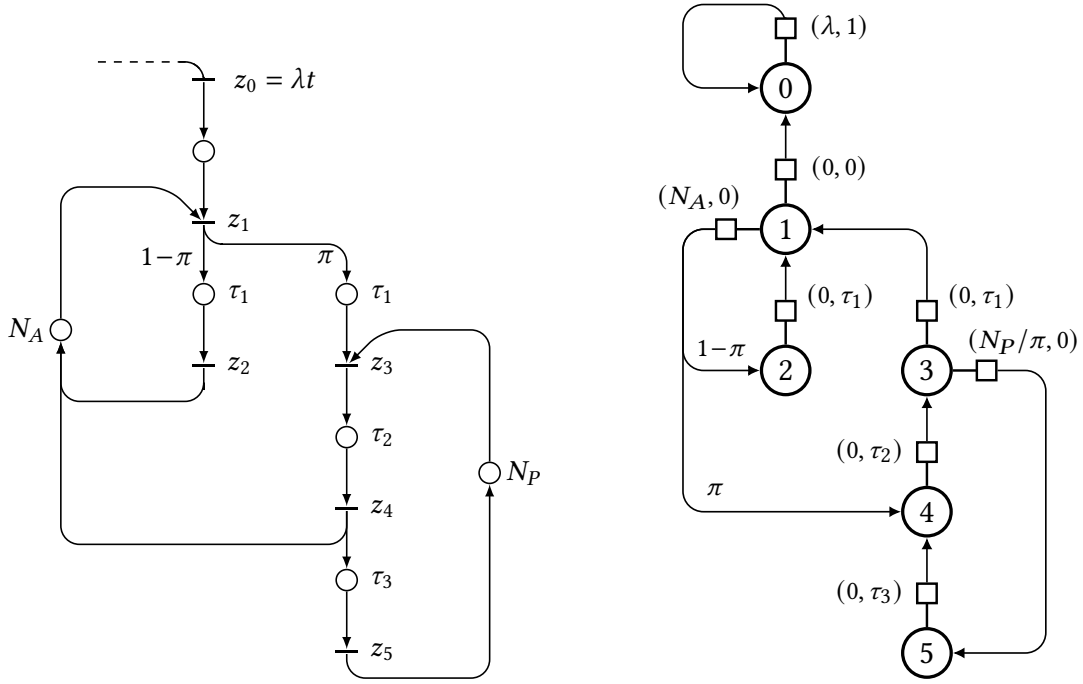


Figure 4.5: The (conflict-free) Petri net (EMS-A) (left) and the corresponding undiscounted SMDP (right).

net formalism becomes a sink in the SMDP one. Remark also as an effect of the renormalization by e how the probabilities π and $1 - \pi$ get moved, and how the initial marking N_P (number of emergency physicians) is increased by a factor $1/\pi$. We let the reader appreciate in the light of these figures the previous discussion on the correspondence between incidence patterns of Petri nets and SMDPs. Finally, observe how the input transition z_0 with a prescribed affine profile $t \mapsto \lambda t$ in the Petri net is replaced by an absorbing state with a cost-rate λ , consistently with Proposition 4.1 and its proof.

Applying the formulas of Chapter 3 and still using the symbol \wedge to denote the minimum, the dynamic programming equations of the SMDP of Figure 4.5 are the following

$$\begin{cases} v(1, t) = v(0, t) \wedge (N_A + (1 - \pi)v(2, t) + \pi v(4, t)) \\ v(2, t) = v(1, t - \tau_1) \\ v(3, t) = v(1, t - \tau_1) \wedge (N_P/\pi + v(5, t)) \\ v(4, t) = v(3, t - \tau_2) \\ v(5, t) = v(4, t - \tau_3) \end{cases} \quad (4.7)$$

that is exactly the same system as (EMS-A) after renormalization by e .

In the same spirit as Examples 1.22, 1.25, 2.6 and 2.10 for Petri nets and Examples 3.21 and 3.25 for SMDPs, we represent in the Figure 4.6 the trajectory of the Petri net (i.e., the counter functions) and of its twin SMDP (i.e., the value function) of our running example. In what follows, we concisely comment the behaviour of the solutions, highlighting phenomena that we will explain in details in Section 3.

We always suppose that the characteristic handling time τ_1 of an inbound call by the MRA is $\tau_1 := 150$ s, the typical transfer time between the MRA and the emergency doctor is $\tau_2 := 20$ s and the duration of a consultation between the doctor and the patient is $\tau_3 := 200$ s. In addition, we suppose that the calls arrive with a constant rate λ and there is no initial stock, so that $v(0, t) = \lambda t$ for all $t \geq 0$. We choose $\lambda = 0.025 \text{ s}^{-1}$, corresponding to one call every 40 seconds, or 90 calls per hour, and we assume that three tenth of them require a consultation with the physician, or in other words $\pi := 0.3$. However, we consider four different scenarios for the call center staffing, with the pair (N_A, N_P) (numbers of MRAs and emergency physicians) taking values in $\{(4, 2), (3, 2), (4, 1), (3.5, 1.5)\}$. The trajectories are given in the first 20 minutes of the execution. We mark all the break of slopes of the trajectories by a tick on the time axis, while we symbolize by a dotted vertical line a change of minimizing term in the systems (EMS-A) or (4.7).

In the first case where $N_A = 4$ and $N_P = 2$, one can see that the trajectories are exactly affine. Naturally, the trajectory of the SMDP is a rescaled version of the one of the Petri net, and since all profiles are parallel for the SMDP, we directly read that we have $z_0(t) = z_1(t) = \lambda t$, $z_2(t) = (1 - \pi)\lambda(t - \tau_1)$, $z_3(t) = \pi\lambda(t - \tau_1)$, $z_4(t) = \pi\lambda(t - \tau_1 - \tau_2)$ and $z_5(t) = \pi\lambda(t - \tau_1 - \tau_2 - \tau_3)$. This means, exactly such as the first experiment of Example 2.6 for a single-tier call center, that all the inbound calls are picked up and handled as soon as possible at each step, since the delays between the counters of transitions are the incompressible holding times of their upstream places.

The second scenario involves $N_A = 3$ MRAs and $N_P = 2$ emergency physicians. It can be seen that all the inbound calls are picked up by MRAs until 120 s, after which z_1 stays constant for a bit, because all the MRAs are busy on the phone at this instant. In particular the minimum in the first equation of (EMS-A) or (4.7) is no longer determined by the arrivals of calls. From this moment forward, not all calls are immediately handled and the difference $z_0(t) - z_1(t)$ gives the amount of pending calls. The trajectory of the Petri net shows an oscillating profile on top of an affine growth, similar to the second experiment of Example 2.6, with slopes proportional to e since it appears on the SMDP trajectory that all states of $\{1, 2, 3, 4, 5\}$ have the same growth-rate, smaller than λ . Observe that ultimately, the break of slopes occur only at regular instants, multiple of 10 seconds.

The third simulation supposes that $N_A = 4$ and $N_P = 1$. On the rescaled trajectory for the SMDP, one can see that the growth of $v(3, \cdot)$ stops at 283.33 s, indicating that there are not enough emergency physicians to treat all the calls handed over by the MRAs. This also affects the MRAs who cannot pick up all calls from the instant 316.67 s, moment at which the counters z_0 and z_1 consequently start to differ. It is worth mentioning that the counters z_3 , z_4 and z_5 admit ultimate profiles with distinct and long plateaus, while z_1 and z_2 are much smoother.

Finally, the fourth and last situation is intermediary to the previous ones, with $N_A = 3.5$ MRAs and $N_P = 1.5$ physicians; these numbers are of course not realistic but yet permitted in the continuous-relaxed framework. The particularity of this example is its high number of switches in the dynamics, recall indeed that each dashed vertical line marks a change in one minimizing terms of (EMS-A) or (4.7). It is here very difficult to predict which part of the dynamics will be limiting for horizons between 300 s and 1200 s! In addition, the trajectories suffer many break of slopes (precisely 189 in this figure) over short times (up to 0.028 s) and at non-trivial instants.

So why do small changes in the staffing impact that much the solutions of the call center dynamics? Can these behaviours be predicted and quantitatively analyzed? Those are questions that we shall now answer.

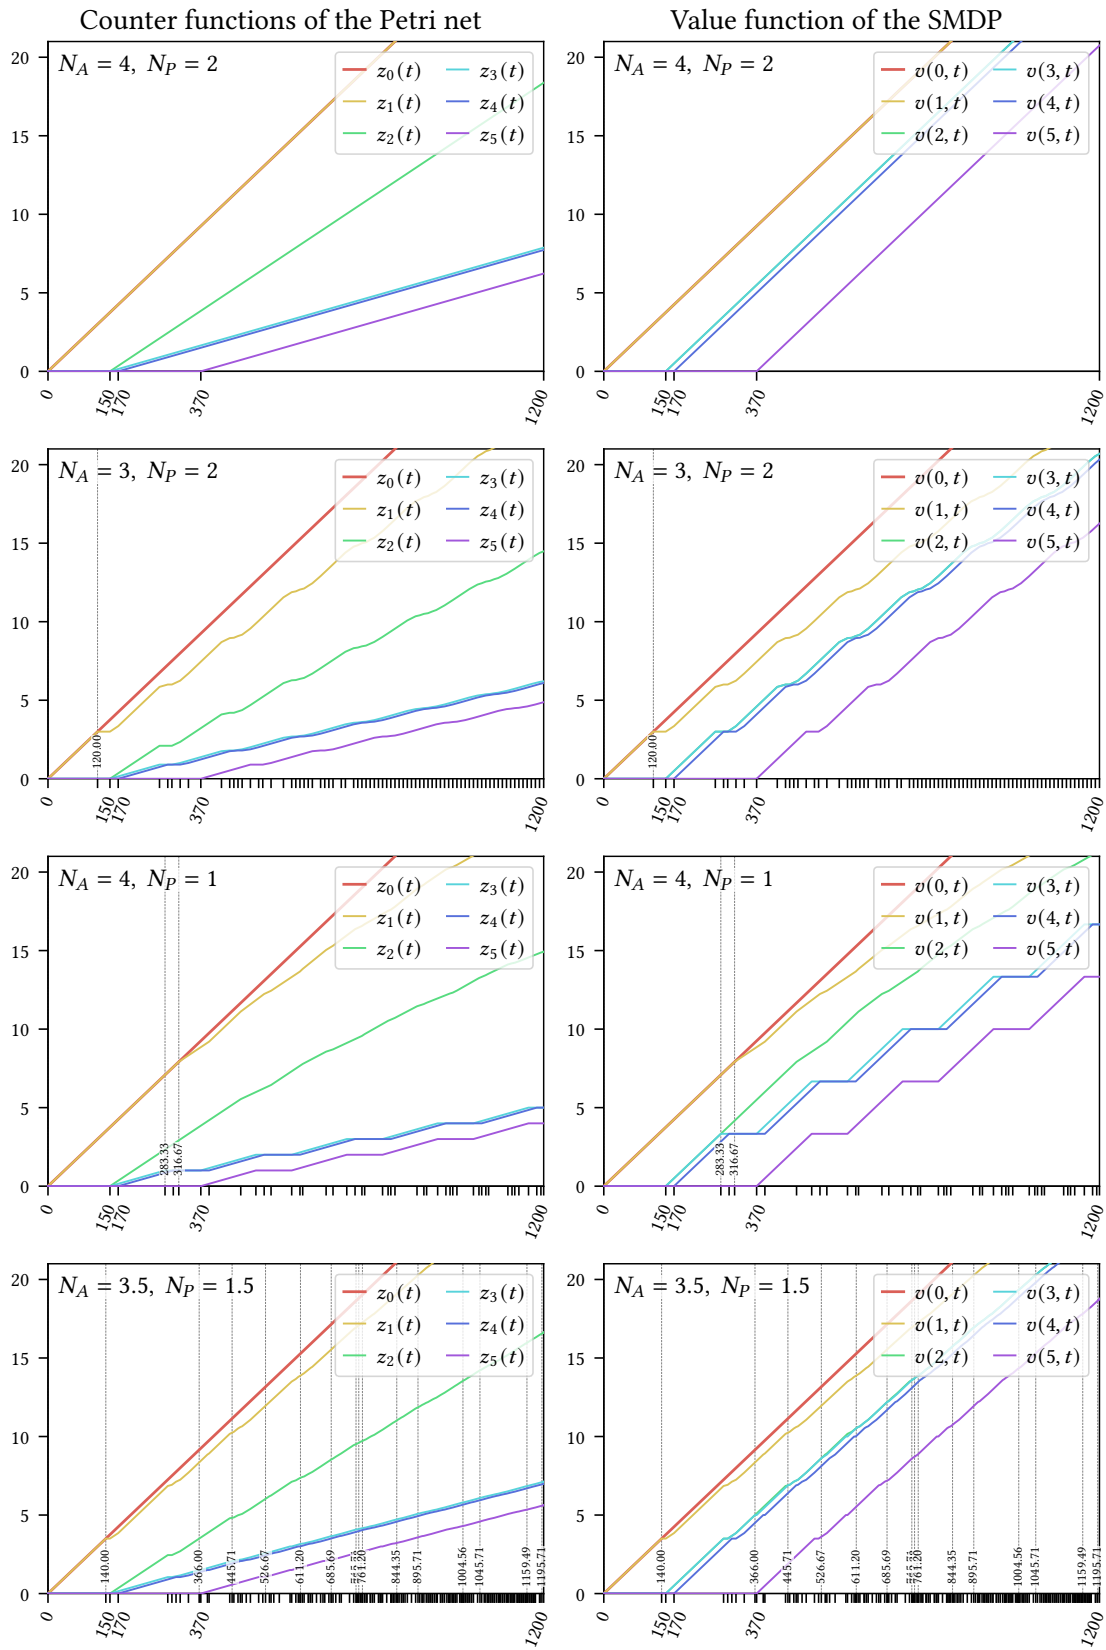


Figure 4.6: Simulation of counters and value functions for the Petri net and the SMDP of Figure 4.5

3 Studying the throughput of fluid Petri nets through SMDP tools

In this section, our goal is to study and characterize the **throughput** of continuous-relaxed and priority-free timed Petri nets, defined to be the ultimate number of firings per unit of time of the transitions of the net, or more formally as the vector ρ of $(\mathbb{R}_{\geq 0} \cup \{\infty\})^Q$ such that

$$\rho := \lim_{t \rightarrow \infty} \frac{z(t)}{t},$$

where $z := (z_q)_{q \in Q}$ and when the limit exists. The Figure 4.6 (specifically its left column) illustrates four situations in which the trajectory of the Petri net is pseudo-affine, and therefore where ρ seems to be finite; we are going to show that this is actually true in general. However, the throughput differs from one case to another, and we want to precisely explain why.

The throughput of Petri nets, as the first-order estimate of the trajectories of the system and one of the performance measures with the most practical interest, has already received much attention in the literature. The originality of our approach is nevertheless to analyze the throughput of timed Petri nets in the light of the following essential key result, which states a correspondence between the throughput of a Petri net and the average-cost vector of an SMDP. Recall that \odot denotes the Hadamard (componentwise) product (conversely, we shall soon use the symbol \oslash to denote the Hadamard quotient).

- **CORRESPONDENCE THEOREM 3.** *The throughput ρ of a continuous-relaxed non-Zeno timed Petri net with affine inputs and that admits a positive stoichiometric invariant e verifies*

$$\rho = \chi(\mathcal{S})_e \odot e, \quad (4.8)$$

where $\chi(\mathcal{S})_e$ is the ultimate rate-of-cost vector (3.23) of the undiscounted semi-Markov decision process with no stopping costs that corresponds to the Petri net according to Correspondence Theorem 2.

⋮ *Proof.* We reuse the result of Correspondence Theorem 2, telling us that $z = v \odot e$, where v is the value function of the SMDP. We apply Corollary 3.32–(2) to this value function, to obtain that $\lim_{t \rightarrow \infty} v(t)/t = \chi(\mathcal{S})$, and therefore $\rho = \chi(\mathcal{S}) \odot e$. ♠

We have purposively added the subscript e to insist on the fact that the undiscounted SMDP built by Correspondence Theorem 2 has costs and transition probabilities that depend upon e . However, the whole right-hand side of (4.8) is homogeneous in e , in the sense that scaling e by a positive scalar (that still yields a positive stoichiometric invariant) does not change the result. In fact, we have mentioned in the Section 2.2 of Chapter 2 that the set of positive stoichiometric invariants can have multiple degrees of freedom (which trivially happens for Petri nets with multiple connected components, but can also arise in connected nets). Nonetheless, while $\chi(\mathcal{S})_e$ and e change, their product remains the same.

Note also that (4.8) is not inconsistent with the fact that ρ is nonnegative, because Correspondence Theorem 2 showed that the costs of an SMDP corresponding to a Petri net are rescaled markings, thus they are nonnegative and so is $\chi(\mathcal{S})_e$.

- 3.1 EXISTENCE AND UNIVERSALITY.** — The first consequence of Correspondence Theorem 3 is the existence and finiteness of the throughput vector of Petri nets, stated in the following corollary. Recall from Theorem 1.23 and the discussion in Section 2.2 of Chapter 2 that the solutions of the dynamical equations of continuous-relaxed, non-Zeno and priority-free Petri net are uniquely

determined by *seed functions* defined over $[-\tau_{\max}, 0)$, that are elements of \mathcal{Z} , and we refer to these complete solutions over $[-\tau_{\max}, \infty)$ as trajectories of a Petri net.

- **COROLLARY 4.2.** *In a continuous-relaxed non-Zeno Petri net that admits a stoichiometric invariant,*
- (i) *there exists an affine stationary regime, i.e., two vectors (ρ, u) in $(\mathbb{R}_{\geq 0})^Q \times \mathbb{R}^Q$ such that, initializing the dynamics with the seed function $z(t) = \rho t + u$ for t in $[-\tau_{\max}, 0)$, we end up with the affine trajectory $z(t) = \rho t + u$ for all $t \geq 0$;*
 - (ii) *the vector ρ in (i) is universal, i.e., the solution z of the dynamics initialized with any seed function of \mathcal{Z} satisfies*

$$z(t) \underset{t \rightarrow \infty}{=} \rho t + O(1).$$

⋮ *Proof.* Both items are consequences of Corollary 3.32 and naturally Correspondence Theorem 3. ♠

Observe the implications in practice of Corollary 4.2–(ii), for example on our emergency call center model (EMS-A). The insensitivity of the throughput to the initial condition means that whatever the size of the queues of calls when the execution of the call center starts, the latter will always reach an ultimate handling speed exclusively determined by its internal parameters: arrivals rate, conversation times, proportions of the different types of calls, and number of agents.

Compared to [CGQ95], the existence of the throughput in Corollary 4.2 is proved in a multi-chain setting, that is to say, the corresponding semi-Markov decision process does not necessarily have a strongly connected underlying structure, and therefore ρ may not be uniform (i.e., a multiple of 1). In addition, the second item of the corollary can be seen as a second-order result since it tells us that the deviation $z(t) - \rho t$ of the trajectory to its ultimate growth-rate is bounded, which is more informative than the equality $\lim_{t \rightarrow \infty} z(t)/t = \rho$ featured in [CGQ95] and also studied in [GG04b].

- **REMARK 4.3.** Let us briefly come back to the original “discrete” models of Petri nets introduced in Chapter 1. Recall indeed that in Chapter 2, we have shown with Theorem 2.19 that the trajectory Z^N of a scaled-by- N Petri net differed at most from the trajectory z^N of the counterpart continuous-relaxed net by an affine function $t \mapsto \kappa_1 t + \kappa_0$, provided both inputs and preselection mappings verified so-called bounded deviations conditions. In addition, we had $z^N = Nz$. As a consequence, we know that there exists a constant $\kappa \geq 0$ such that $|Z^N(t)/Nt - z(t)/t| \leq \kappa/N$. This does not prove the existence of the throughput $\lim_{t \rightarrow \infty} Z(t)/t$ for discrete Petri nets (see the discussion in [GG04b, Section 5] on this aspect), however we can define $\rho_N^{\text{disc}} := \liminf_{t \rightarrow \infty} Z^N(t)/Nt$ to be a minimum *rescaled* throughput of the scaled-by- N discrete net. Our bound guarantees that $|\rho_N^{\text{disc}} - \rho| \leq \kappa/N$, in particular even for the scaled-by-1 instance we know that the real growth rate of the discrete Petri net cannot be too far away from the throughput of the equivalent fluid Petri net. Further, we retrieve the fact that the continuous-relaxed approximation is more and more precise as N grows, with $\lim_{N \rightarrow \infty} \rho_N^{\text{disc}} = \rho$. This is illustrated by the diagram below.

$$\begin{array}{ccc}
 \frac{Z^N(t)}{Nt} & \xrightarrow{\lim_{N \rightarrow \infty}} & \frac{z(t)}{t} \\
 \liminf_{t \rightarrow \infty} \downarrow & & \downarrow \lim_{t \rightarrow \infty} \\
 \rho_N^{\text{disc}} & \xrightarrow{\lim_{N \rightarrow \infty}} & \rho \\
 \in [\rho - \kappa/N, \rho + \kappa/N] & &
 \end{array}$$

3.2 THE OPTIMALITY EQUATIONS. — We are now interested in methods to compute the throughput vector. The following result makes it arise as the solution of a set of $2 \times |\mathcal{Q}|$ optimality equations.

► **COROLLARY 4.4.** *In a continuous-relaxed non-Zeno and conflict-free timed Petri net with affine inputs ($t \mapsto \lambda_q t + \eta_q$) $_{q \in \mathcal{Q}_{\text{input}}}$ and that admits a stoichiometric invariant, the throughput ρ is the unique vector of $(\mathbb{R}_{\geq 0})^{\mathcal{Q}}$ such that $\rho_q = \lambda_q$ for all q in $\mathcal{Q}_{\text{input}}$ and such that there exists a vector u in $\mathbb{R}^{\mathcal{Q}}$ verifying for all q in $\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$:*

$$\rho_q = \min_{p \in q^{\text{in}}} \left\{ \alpha_{qp}^{-1} \sum_{q' \in \mathcal{Q}} \alpha_{pq'} \rho_{q'} \right\} \quad (\text{L1})$$

$$u_q = \min_{p \in q_*^{\text{in}}} \left\{ \alpha_{qp}^{-1} m_p - \rho_q \tau_p + \alpha_{qp}^{-1} \sum_{q' \in \mathcal{Q}} \alpha_{pq'} u_{q'} \right\} \quad (\text{L2})$$

where q_*^{in} is the subset of q^{in} for which the minimum is achieved in (L1).

Proof. Let us denote by e a positive stoichiometric invariant. The Correspondence Theorem 3 and the Theorem 3.31 ensure that ρ is the throughput vector of the Petri net if and only if there exists a vector u such that $\rho \oslash e$ and $u \oslash e$ are solutions of the equations (OE1)–(OE2) of SMDPs, i.e.,

$$\begin{aligned} \rho_q / e_q &= \min_{p \in q^{\text{in}}} \left\{ \alpha_{qp}^{-1} \sum_{q' \in \mathcal{Q}} \alpha_{pq'} e_{q'} e_q^{-1} (\rho_{q'} / e_{q'}) \right\} \\ u_q / e_q &= \min_{p \in q_*^{\text{in}}} \left\{ \alpha_{qp}^{-1} m_p / e_q - \tau_p \rho_q / e_q + \alpha_{qp}^{-1} \sum_{q' \in \mathcal{Q}} \alpha_{pq'} e_{q'} e_q^{-1} (u_{q'} / e_{q'}) \right\}. \end{aligned}$$

It suffices to multiply these equations by e_q on both sides to obtain the equations (L1)–(L2) for all $q \in \mathcal{Q}_{\text{input}}$. Further, emulating the input transitions by absorbing states as done in Proposition 4.1 provides for all $q \in \mathcal{Q}_{\text{input}}$ the extra equation $(\rho_q, u_q) = (\rho_q, \lambda_q \tau_{\text{max}} - \rho_q \tau_{\text{max}} + u_q)$, which indeed implies $\rho_q = \lambda_q$. The value of u_q is not constrained for $q \in \mathcal{Q}_{\text{input}}$. \mathcal{E}

In [ABGz1], we have shown without using the notion of stoichiometric invariant that the existence of a solution (ρ, u) of (L1)–(L2) is equivalent to the existence of an ultimately affine solution of the form $t \mapsto \rho t + u$ of the dynamics of the Petri net – with a proof syntactically identical to the one of Proposition 3.30. The Corollary 4.4 actually tells us that the existence of a stoichiometric invariant is a sufficient condition for (L1)–(L2) to have a solution.

The Corollary 4.4 deals with conflict-free Petri nets, which is handy to treat the transitions ruled by proportional preselection routing or mere synchronization under the same formalism (the routing proportions become integrated in the arc multipliers). In practice, we may rather want to keep track of each parameter of the original Petri net, before it has been transformed in a conflict-free net. To avoid distinguishing the cases of transitions in $\mathcal{Q}_{\text{sync}}$ and $\mathcal{Q}_{\text{psel}}$, it is licit to set $\pi_{qp} := 1$ for all q in $\mathcal{Q}_{\text{sync}}$ and p in q^{out} . Further, the two optimality equations (L1) and (L2) can be compacted into a single equation by using the lexicographic order on real tuples; for two tuples (x, y) and (x', y') in \mathbb{R}^2 , we shall write $(x, y) \leq (x', y')$ if and only if $x \leq x'$, or $x = x'$ and $y \leq y'$. Denoting by \min^{LEX} the minimum operation relatively to this lexicographic ordering, we obtain the throughput optimality equation, valid for all transitions q in $\mathcal{Q} \setminus \mathcal{Q}_{\text{input}}$:

$$(\rho_q, u_q) = \min_{p \in q^{\text{in}}}^{\text{LEX}} \left\{ \left(\pi_{qp} \alpha_{qp}^{-1} \sum_{q' \in \mathcal{Q}} \alpha_{pq'} \rho_{q'}, \pi_{qp} \alpha_{qp}^{-1} m_p - \rho_q \tau_p + \pi_{qp} \alpha_{qp}^{-1} \sum_{q' \in \mathcal{Q}} \alpha_{pq'} u_{q'} \right) \right\}. \quad (\text{L})$$

Remark how the lexicographic minimum exactly encodes the fact that for all q , the term u_q in (L2) is only determined by upstream places of q realizing the minimum in (L1).

Gathering all the equations of type (L) together provides a lexicographic system that it is possible to solve by an enumerative approach, consisting in choosing in each equation which tuple on the right-hand side achieves the minimum. These equalities lead to a value of the throughput, while the non-minimizing terms yield inequalities that are validity conditions under which the chosen tuple candidates are indeed minimizing terms.

*
* *

Let us illustrate this optimality system on our running example (EMS-A) in which all the arc multipliers are equal to one in the original model (or alternatively only two arcs have nontrivial weight in the conflict-free model). When working on real tuples like below, we also use the operator \wedge to stand for the lexicographic minimum. Applying the equation (L) to our idealized call center provides

$$\begin{cases} (\rho_1, u_1) = (\lambda, u_0) \wedge (\rho_2 + \rho_4, N_A + u_2 + u_4) \\ (\rho_2, u_2) = ((1 - \pi)\rho_1, (1 - \pi)u_1 - \rho_2\tau_1) \\ (\rho_3, u_3) = (\pi\rho_1, \pi u_1 - \rho_3\tau_1) \wedge (\rho_5, N_P + u_5) \\ (\rho_4, u_4) = (\rho_3, u_3 - \rho_4\tau_2) \\ (\rho_5, u_5) = (\rho_4, u_4 - \rho_5\tau_3) \end{cases}$$

Consistently with Corollary 4.4, the rate λ of inbound calls arrivals appears in the equation derived from z_1 since the latter is the only transition constrained by the input. The simple equations relative to z_2, z_4 and z_5 enable us to eliminate six variables by substitution, and to focus on the reduced system

$$\begin{cases} (\rho_1, u_1) = (\lambda, u_0) \wedge ((1 - \pi)\rho_1 + \rho_3, N_A + (1 - \pi)(u_1 - \rho_1\tau_1) + u_3 - \rho_3\tau_2) \\ (\rho_3, u_3) = (\pi\rho_1, \pi u_1 - \rho_3\tau_1) \wedge (\rho_3, N_P + u_3 - \rho_3(\tau_2 + \tau_3)) \end{cases} \quad (4.9)$$

in which only the four variables ρ_1, ρ_3, u_1 and u_3 are to be determined. The variable ρ_1 corresponds to the throughput of z_1 , i.e., the number of emergency calls picked up by the medical regulation assistants per unit of time. Similarly, ρ_3 gives the pick up rate of calls needing a consultation with an emergency doctor. These two indicators arguably are the most important of the organization!

Let us detail the complete resolution of the system (4.9) with the enumerative approach, by first writing the four inequalities (on real tuples) associated with the previous lexicographic system (we label them based on the transition index and cost-like variables that appear on the right-hand side).

$$(\rho_1, u_1) \leq (\lambda, u_0) \quad (E1.\lambda)$$

$$(\rho_1, u_1) \leq ((1 - \pi)\rho_1 + \rho_3, N_A + (1 - \pi)(u_1 - \rho_1\tau_1) + u_3 - \rho_3\tau_2) \quad (E1.A)$$

$$(\rho_3, u_3) \leq (\pi\rho_1, \pi u_1 - \rho_3\tau_1) \quad (E3.0)$$

$$(\rho_3, u_3) \leq (\rho_3, N_P + u_3 - \rho_3(\tau_2 + \tau_3)) \quad (E3.P)$$

Since either one of the inequalities (E1. λ) or (E1.A) is an equality, and likewise either one of the inequalities (E3.0) or (E3.P) is an equality, we have four cases in total to study.

Minimum reached by (E1.λ) and (E3.0). – We derive from the ρ -parts of (E1.λ) and (E3.0) that $\rho_1 = \lambda$ and $\rho_3 = \pi\lambda$. However, looking at the ρ -part of (E1.A), we see that $(1 - \pi)\rho_1 + \rho_3 = \lambda$ as well, so the minimum between (E1.λ) and (E1.A) is in fact arbitrated by their u -parts. The same phenomenon occurs for (E3.0) and (E3.P) that have identical ρ -values (respectively $\pi\rho_1$ and ρ_3 , here both equal to $\pi\lambda$).

Since we assumed that the u -parts of (E1.λ) and (E3.0) achieve the minimum, the u -part of (E3.0) gives the equality $u_3 = \pi u_1 - \rho_3 \tau_1 = \pi u_1 - \pi\lambda\tau_1$, that we can re-inject in the inequality $u_1 \leq N_A + (1 - \pi)(u_1 - \lambda\tau_1) + u_3 - \pi\lambda\tau_2$ brought by the u -part of (E1.A) to conclude that $N_A \geq \lambda(\tau_1 + \pi\tau_2)$. Finally, the u -part of (E3.P) gives the inequality $u_3 \leq N_P + u_3 - \pi\lambda(\tau_2 + \tau_3)$ or equivalently $N_P \geq \pi\lambda(\tau_2 + \tau_3)$.

Minimum reached by (E1.A) and (E3.0). – Using the ρ -part of (E3.0), we know that $\rho_3 = \pi\rho_1$. Therefore, the ρ -part of (E1.A) just amounts to the trivial relation $\rho_1 = \rho_1$. However, we can use the u -part of the same (E1.A) to obtain the equality $u_1 = N_A + (1 - \pi)u_1 - (1 - \pi)\rho_1\tau_1 + u_3 - \pi\rho_1\tau_2$. On the other hand, we have $u_3 = \pi u_1 - \pi\rho_1\tau_1$ according to the u -part of (E3.0), thus by combination we obtain $\rho_1 = \rho_A$ (and still $\rho_3 = \pi\rho_A$) where ρ_A is the shorthand notation

$$\rho_A := \frac{N_A}{\tau_1 + \pi\tau_2}.$$

The ρ -part of (E1.λ) gives the inequality $\rho_1 \leq \lambda$, that we can write as $N_A \leq \lambda(\tau_1 + \pi\tau_2)$ or $\rho_A \leq \lambda$. As in the previous case, the u -part of (E3.P) still leads to $N_P \geq \rho_3(\tau_2 + \tau_3)$, and since $\rho_3 = \pi\rho_A$, we have $\pi\rho_A \leq \rho_P$, where ρ_P is the shorthand notation

$$\rho_P := \frac{N_P}{\tau_2 + \tau_3}.$$

Minimum reached by (E1.λ) and (E3.P). – The ρ -part of (E1.λ) ensures that $\rho_1 = \lambda$, while the u -part of (E3.P) provides $u_3 = N_P + u_3 - \rho_3(\tau_2 + \tau_3)$ so $\rho_3 = \rho_P$. On one hand, remark that we have the inequality $\rho_3 \leq \pi\rho_1 \leq \pi\lambda$ by the ρ -part of (E3.0), but on the other hand we know from the ρ -part of (E1.A) that $\lambda \leq (1 - \pi)\lambda + \rho_3$, so $\rho_3 \geq \pi\lambda$. As a result we have $\rho_3 = \pi\lambda$ and in particular $N_P = \pi\lambda(\tau_2 + \tau_3)$.

Remark that we have $u_3 \leq \pi u_1 - \rho_3\tau_1$ according to the u -part of (E3.0), hence the inequality $u_1 \leq N_A + (1 - \pi)(u_1 - \lambda\tau_1) + u_3 - \rho_3\tau_2$ becomes $u_1 \leq N_A + (1 - \pi)u_1 - (1 - \pi)\lambda\tau_1 + \pi u_1 - \rho_3\tau_1 - \rho_3\tau_2$. We can reuse the fact that $\rho_3 = \pi\lambda$ to obtain $N_A \geq \lambda(\tau_1 + \pi\tau_2)$, or equivalently $\rho_A \geq \lambda$.

Minimum reached by (E1.A) and (E3.P). – We use the u -parts of both (E1.A) and (E3.P) that respectively yield the equalities $u_3 = N_P + u_3 - \rho_3(\tau_2 + \tau_3)$ and $\rho_1 = (1 - \pi)\rho_1 + \rho_3$, so we have $\rho_3 = \rho_P$ and $\rho_1 = \rho_P/\pi$.

The ρ -part of (E1.λ) gives $\rho_1 \leq \lambda$ so it follows that $\rho_3 \leq \pi\lambda$, and in particular $N_P \leq \pi\lambda(\tau_2 + \tau_3)$ (alternatively we may have written $\rho_P \leq \pi\lambda$). We have the equality $u_1 = N_A + (1 - \pi)(u_1 - \rho_1\tau_1) + u_3 - \rho_3\tau_2$ given by the u -part of (E1.A), in which we can inject the inequality $u_3 \leq \pi u_1 - \rho_3\tau_1$ coming from the u -part of (E3.0). We obtain $u_1 \leq N_A + u_1 - \rho_1\tau_1 + \pi\rho_1\tau_1 - \rho_3\tau_1 - \rho_3\tau_2$ which means that $N_A \geq \rho_1(\tau_1 + \pi\tau_2)$ and therefore that $\rho_P \leq \pi\rho_A$.

We postpone the interpretation of these throughput values to the next subsection, where an alternative method to obtain them is introduced. However, we let the reader observe that in this resolution, we did find solutions only for the vector ρ , and not for the vector u . Trying to do so is in general pointless, and this is due to the fact that optimality equations (OE1)–(OE2) for SMDPs can have multiple independent h -solutions, as discussed in the remark subsequent to the Theorem 3.31. Observe also that the value u_0 did not play any role in our resolution.

3.3 THE THROUGHPUT COMPLEX. — Recall once again that Correspondence Theorems 1 and 2 establish that the transitions of Petri nets and their upstream places correspond to states of semi-Markov decision processes and their playable actions. We have seen in Chapter 3 that the notion of policies plays a key role in the average-cost problem, therefore it prompts us to also define **policies for Petri nets**. Building on the correspondences, we keep the same notation than the ones for SMDPs, saying that a map σ from \mathcal{Q} to $\prod_{q \in \mathcal{Q}} q^{\text{in}}$ is a policy of the Petri net (equivalently, σ is a map of $\mathcal{Q} \rightarrow \mathcal{P}$ such that the place $\sigma(q)$ lies in q^{in} for all q in \mathcal{Q}) and denoting by \mathfrak{S} the set of all policies. Defined this way, a policy of the Petri net is nothing more than a choice of upstream places for all transitions. Remark that such a choice is non-trivial only for transitions q in \mathcal{Q} such that $|q^{\text{in}}| > 1$ (which in priority-free nets is only allowed for synchronization transitions, elements of $\mathcal{Q}_{\text{sync}}$).

To each policy σ of \mathfrak{S} of the Petri net is associated the stochastic matrix P^σ of the policy σ in the corresponding SMDP. It enables us in particular to talk about the final classes of a Petri net (as subsets of \mathcal{Q}) under a certain policy, defined in Section 1.3 of Chapter 3. We shall indifferently denote by $\mathcal{F}(\sigma)$ the final classes of a Petri net or those of its twin SMDP under the policy σ . In the Petri net context, the invariant measures of σ (as vectors of $\mathbb{R}^{\mathcal{Q}}$) are still denoted by $(\mu_F^\sigma)_{F \in \mathcal{F}(\sigma)}$, and the absorption probabilities by $(\phi_{F,q}^\sigma)_{F \in \mathcal{F}(\sigma), q \in \mathcal{Q}}$. We can further define for the Petri net the vectors $m^\sigma := (m^{\sigma(q)})_{q \in \mathcal{Q}}$ and $\tau^\sigma := (\tau^{\sigma(q)})_{q \in \mathcal{Q}}$, that are the initial markings and holding times vector seen under the policy σ . In fact, consistently with the construction of Proposition 4.1, if q lies in $\mathcal{Q}_{\text{input}}$, then $m^{\sigma(q)}$ (resp. $\tau^{\sigma(q)}$) refers to the virtual marking $\lambda_q \tau_{\text{max}}$ (resp. virtual holding time τ_{max}).

► **COROLLARY 4.5.** *In a continuous-relaxed non-Zeno timed Petri net with affine inputs and that admits a positive stoichiometric invariant e , the throughput ρ is given by*

$$\forall q \in \mathcal{Q}, \quad \rho_q = e_q \min_{\sigma \in \mathfrak{S}} \sum_{F \in \mathcal{F}(\sigma)} \phi_{F,q}^\sigma \frac{\langle \mu_F^\sigma, D^\sigma m^\sigma \rangle}{\langle \mu_F^\sigma, \tau^\sigma \rangle}, \quad (4.10)$$

where the minimum is taken componentwise and for all σ in \mathfrak{S} , D^σ is the diagonal matrix defined by $D^\sigma := \text{diag}((e_q^{-1} \alpha_{q\sigma(q)}^{-1} \pi_{q\sigma(q)})_{q \in \mathcal{Q}})$.

⋮ *Proof.* This is a consequence of Correspondence Theorem 3 and Theorem 3.31–(2), adapting the notation to the Petri net formalism. The matrices $(D^\sigma)_{\sigma \in \mathfrak{S}}$ are used to transform the initial marking in the form of (4.6) (recall that $\pi_{qp} := 1$ whenever $q \notin \mathcal{Q}_{\text{psel}}$). ♠

Contrary to the optimality system (L) of which the throughput ρ is the solution, the expression (4.10) is a closed form that gives ρ as a function of all the parameters of the Petri net. It features explicitly the input rates, the initial markings and the holding times, and featuring “implicitly” the arc multipliers or the routing proportions via the different Markov chain analysis quantities (observe once again that scaling the stoichiometric invariant e of the formula still does not change the result, since the matrices $(D^\sigma)_{\sigma \in \mathfrak{S}}$ contravariantly depends upon e). The formula (4.10) holds in the much general multichain setting, hence it refines the throughput expression of “unchain” Petri nets that was stated in [CGQ95].

We know by Theorem 3.31–(1) that there exists a policy σ^* of \mathfrak{S} that realizes the throughput in (4.10) (and the SMDP average-cost in (3.26)). This minimizing policy has a clear interpretation, for all q in \mathcal{Q} , $\sigma^*(q)$ is the **bottleneck upstream place** of q , that will ultimately not contain enough fluid-tokens and effectively constrain the throughput of q .

The expression (4.10) also shows that the throughput ρ_q of the transition q , as a minimum of linear forms, is a **nondecreasing, concave, and piecewise-affine** function of the initial marking

vector m of $(\mathbb{R}_{\geq 0})^{\mathcal{P}}$. As is customary in polyhedral geometry, we associate with this map a polyhedral complex (i.e., a collection \mathcal{L} of polyhedra such that for all L in \mathcal{L} , any face F of L is also in \mathcal{L} and for L_1 and L_2 in \mathcal{L} , the polyhedron $L_1 \cap L_2$ is a face of both L_1 and L_2 , see [DLRS10]). If Σ is a set of policies, we define the polyhedral cell C_Σ to be the set of initial markings m such that the argument of the minimum in (4.10) is Σ (note that the cell C_Σ may be empty for some choices of Σ). The space $(\mathbb{R}_{\geq 0})^{\mathcal{P}}$ is covered by the cells of maximal dimension, that can be interpreted as **congestion phases**, or equivalently to a choice of bottleneck places for each q in \mathcal{Q} (non-trivial as soon as such that $|q^{\text{in}}| > 1$).

- **REMARK 4.6.** As stressed in the Remark 3.33, even though the formula (4.10) and the optimality equations (L1)–(L2) are stated here for holding times that are constant, they are in virtue of the more general correspondence stated in Section 2 also valid in the much more expressive setting of stochastic holding times, replacing the constant values by the **average** holding times of the places. This “free lunch” result is very advantageous and legitimizes our choice to focus mainly on frameworks with constant times to study throughput estimates, since final expressions are the same.

*
* *

We come back to our example (EMS-A) again to illustrate Corollary 4.5. Instead of using the formula (4.10) to directly compute the throughput of the Petri net, we are going to compute the average-cost $\chi(\mathcal{S})_e$ of the SMDP of Figure 4.5 using (3.26) and then Correspondence Theorem 3 with the relation $\rho = \chi(\mathcal{S})_e \odot e$. Because there are two states of the SMDP with two playable actions (states 1 and 3), there are four policies. They naturally correspond to the four cases studied in Section 3.2, since the method employed there was actually to enumerate all the choices of bottleneck upstream places. We label these four policies in accordance with the four equations featured in Section 3.2, that is, $\mathfrak{S} = \{(\lambda, 0), (A, 0), (\lambda, P), (A, P)\}$, corresponding to selecting the pairs of action that respectively make the inequalities (E1.λ)–(E3.0), (E1.A)–(E3.0), (E1.λ)–(E3.P) and (E1.A)–(E3.P) become equalities.

We detail in Tables 4.8 and 4.9 all the quantities involved in the computation of $\chi(\mathcal{S})_e$ for the stoichiometric invariant $e = (1, 1, 1 - \pi, \pi, \pi, \pi)^\top$, that is to say for every one of these four policies: the transition matrix, the spectral projector, the cost vector, the sojourn time vector, the final classes with their invariant measures and absorption probabilities, and finally the average-cost of this specific policy. We also reuse the shorthand notation ρ_A and ρ_P introduced in Section 3.2.

Observe that under the policy $(\lambda, 0)$, there is a unique final class, and it is reduced to the singleton $\{0\}$, therefore with trivial invariant measure and absorption vector. For all the other policies, $\{0\}$ remains a final class but there is always another one, that is either $\{1, 2, 3, 4\}$ or $\{3, 4, 5\}$. This shows that (EMS-A) features a multichain structure. Taking the minimum of the average-cost of each one of the four policies allows us to recover the global throughput vector ρ of the Petri net:

$$\rho = \begin{pmatrix} \rho_0 \\ \rho_1 \\ \rho_2 \\ \rho_3 \\ \rho_4 \\ \rho_5 \end{pmatrix} = \chi(\mathcal{S})_e \odot e = \left(\begin{pmatrix} \lambda \\ \lambda \\ \lambda \\ \lambda \\ \lambda \\ \lambda \end{pmatrix} \wedge \begin{pmatrix} \lambda \\ \rho_A \\ \rho_A \\ \rho_A \\ \rho_A \\ \rho_A \end{pmatrix} \wedge \begin{pmatrix} \lambda \\ \lambda \\ \lambda \\ \rho_P/\pi \\ \rho_P/\pi \\ \rho_P/\pi \end{pmatrix} \wedge \begin{pmatrix} \lambda \\ \rho_P/\pi \\ \rho_P/\pi \\ \rho_P/\pi \\ \rho_P/\pi \\ \rho_P/\pi \end{pmatrix} \right) \odot \begin{pmatrix} 1 \\ 1 \\ 1 - \pi \\ \pi \\ \pi \\ \pi \end{pmatrix} = \begin{pmatrix} \lambda \\ \rho^* \\ (1 - \pi)\rho^* \\ \pi\rho^* \\ \pi\rho^* \\ \pi\rho^* \end{pmatrix}, \quad (4.11)$$

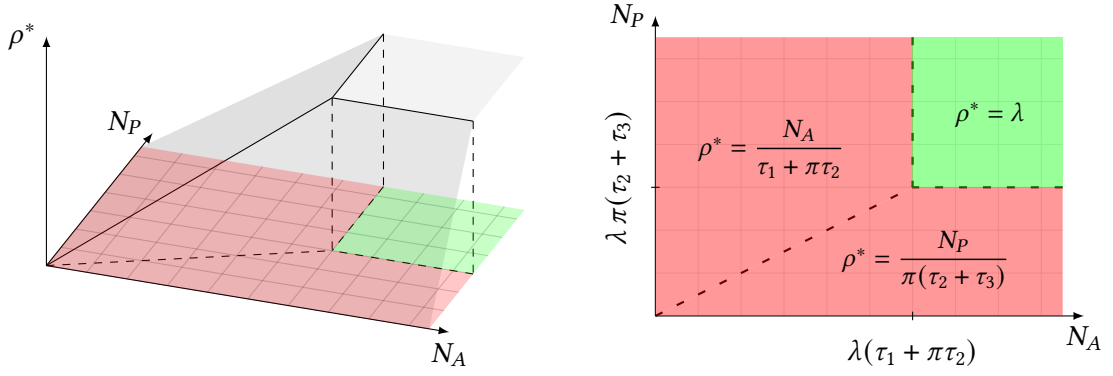


Figure 4.7: The throughput complex (or congestion diagram) of the (EMS-A) system (right), seen as the linearity regions of the map $(N_A, N_P) \mapsto \rho^*$ (left)

where we denote for short

$$\rho^* := \lambda \wedge \rho_A \wedge (\rho_P/\pi) = \lambda \wedge \frac{N_A}{\tau_1 + \pi\tau_2} \wedge \frac{N_P}{\pi(\tau_2 + \tau_3)}.$$

We retrieve the exact same result as in Section 3.2, indeed we let the reader check using the validity conditions that we always had $\rho_1 = \rho^*$ and $\rho_3 = \pi\rho^*$ (and the above values of ρ_2 , ρ_4 and ρ_5 would have easily followed).

We depict in Figure 4.7 the nondecreasing, concave and piecewise-affine map ρ^* as a function of the initial markings N_A and N_P , since ρ^* alone suffices to characterize all the $(\rho_q)_{q \in Q}$. As claimed after the Corollary 4.5, we can associate with this map a polyhedral complex given by the domains on which it is linear. This throughput complex, that from an operational point of view is better thought of as a **congestion diagram**, is represented in Figure 4.7 as well.

We interpret this diagram as follows. The “handling speed” ρ_1 of the medical regulation assistants (MRAs) and ρ_3 of the emergency physicians are always entangled and depend on three key dimensioning parameters: the arrival rate of inbound calls λ , the maximum MRA throughput $\rho_A = N_A/(\tau_1 + \pi\tau_2)$ and the maximum physician throughput $\rho_P = N_P/(\tau_2 + \tau_3)$. These two terms correspond to the number of agents of a given type (MRA or physician) divided by their respective characteristic handling time for a single call.

Based on the solution (4.11), if $N_A \geq N_A^* := \lambda(\tau_1 + \pi\tau_2)$ and $N_P \geq N_P^* := \lambda\pi(\tau_2 + \tau_3)$, we have $\rho_A \geq \lambda$ and $\rho_P/\pi \geq \lambda$, and therefore $\rho^* = \lambda$, which means that all inbound calls are ultimately handled, because both groups of agents work at their maximum speed and are only constrained by the arrivals of calls. These two inequalities delineate the green phase in Figure 4.7, and we may refer to this situation as the **congestion-free** phase. If $N_A \leq N_A^*$ and $N_A/(\tau_1 + \pi\tau_2) \leq N_P/\pi(\tau_2 + \tau_3)$ (equivalently, $\rho_A \leq \rho_P/\pi$) there are too few MRAs to pick up all arriving calls, and as a consequence they impose their maximum handling speed ρ_A to the whole system (recall that emergency physicians wait for MRAs to pass them a fraction π of the calls). Conversely, if $N_P \leq N_P^*$ and $N_P/\pi(\tau_2 + \tau_3) \leq N_A/(\tau_1 + \pi\tau_2)$ (that is, $\rho_P/\pi \leq \rho_A$), there are too few emergency physicians to answer the severe calls, and they impose their handling speed to the whole system again (MRAs are waiting for doctors to take their calls and be released). These two different situations of congestion are shown in red in Figure 4.7.

Observe that there are three different regimes in the Figure 4.7 although there are four policies. As shown in Section 3.2, the policy (λ, P) in fact determines a one-dimensional phase

Policy σ	P^σ	$(P^\sigma)^\star$
$(\lambda, 0)$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$
$(A, 0)$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 - \pi & 0 & \pi & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{\pi+2} & \frac{1-\pi}{\pi+2} & \frac{\pi}{\pi+2} & \frac{\pi}{\pi+2} & 0 \\ 0 & \frac{1}{\pi+2} & \frac{1-\pi}{\pi+2} & \frac{\pi}{\pi+2} & \frac{\pi}{\pi+2} & 0 \\ 0 & \frac{1}{\pi+2} & \frac{1-\pi}{\pi+2} & \frac{\pi}{\pi+2} & \frac{\pi}{\pi+2} & 0 \\ 0 & \frac{1}{\pi+2} & \frac{1-\pi}{\pi+2} & \frac{\pi}{\pi+2} & \frac{\pi}{\pi+2} & 0 \\ 0 & \frac{1}{\pi+2} & \frac{1-\pi}{\pi+2} & \frac{\pi}{\pi+2} & \frac{\pi}{\pi+2} & 0 \end{pmatrix}$
(λ, P)	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}$
(A, P)	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 - \pi & 0 & \pi & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}$

Table 4.8: Transition matrices and spectral projectors associated with the four policies of (EMS-A)

Policy σ	t^σ	c^σ	Final classes		Average-cost
$(\lambda, 0)$	$\begin{pmatrix} 1 \\ 0 \\ \tau_1 \\ \tau_1 \\ \tau_2 \\ \tau_3 \end{pmatrix}$	$\begin{pmatrix} \lambda \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}$	$\{0, 1, 2, 3, 4, 5\}$		$\begin{pmatrix} \lambda \\ \lambda \\ \lambda \\ \lambda \\ \lambda \\ \lambda \end{pmatrix}$
$(A, 0)$	$\begin{pmatrix} 1 \\ 0 \\ \tau_1 \\ \tau_1 \\ \tau_2 \\ \tau_3 \end{pmatrix}$	$\begin{pmatrix} \lambda \\ N_A \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}$	$\{0\}$ $\mu = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \phi = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}$	$\{1, 2, 3, 4\}$ $\mu = \begin{pmatrix} 0 \\ \frac{1}{\pi+2} \\ \frac{1-\pi}{\pi+2} \\ \frac{\pi}{\pi+2} \\ \frac{\pi}{\pi+2} \\ 0 \end{pmatrix}, \phi = \begin{pmatrix} 0 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}$	$\begin{pmatrix} \lambda \\ \rho_A \\ \rho_A \\ \rho_A \\ \rho_A \\ \rho_A \end{pmatrix}$
(λ, P)	$\begin{pmatrix} 1 \\ 0 \\ \tau_1 \\ 0 \\ \tau_2 \\ \tau_3 \end{pmatrix}$	$\begin{pmatrix} \lambda \\ 0 \\ 0 \\ N_P/\pi \\ 0 \\ 0 \end{pmatrix}$	$\{0\}$ $\mu = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \phi = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}$	$\{3, 4, 5\}$ $\mu = \begin{pmatrix} 0 \\ 0 \\ 0 \\ \frac{1}{3} \\ \frac{1}{3} \\ \frac{1}{3} \end{pmatrix}, \phi = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 1 \end{pmatrix}$	$\begin{pmatrix} \lambda \\ \lambda \\ \lambda \\ \rho_P/\pi \\ \rho_P/\pi \\ \rho_P/\pi \end{pmatrix}$
(A, P)	$\begin{pmatrix} 1 \\ 0 \\ \tau_1 \\ 0 \\ \tau_2 \\ \tau_3 \end{pmatrix}$	$\begin{pmatrix} \lambda \\ N_A \\ 0 \\ N_P/\pi \\ 0 \\ 0 \end{pmatrix}$	$\{0\}$ $\mu = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \phi = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}$	$\{3, 4, 5\}$ $\mu = \begin{pmatrix} 0 \\ 0 \\ 0 \\ \frac{1}{3} \\ \frac{1}{3} \\ \frac{1}{3} \end{pmatrix}, \phi = \begin{pmatrix} 0 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}$	$\begin{pmatrix} \lambda \\ \rho_P/\pi \\ \rho_P/\pi \\ \rho_P/\pi \\ \rho_P/\pi \\ \rho_P/\pi \end{pmatrix}$

Table 4.9: Times and cost vectors, final classes and average-costs associated with the four policies of (EMS-A)

in this two-dimensional space, for it achieves the minimum throughput of the complex only if $N_P = \pi\lambda(\tau_2 + \tau_3)$. The policies $(\lambda, 0)$, $(A, 0)$ and (A, P) respectively determine the full-dimensional congestion-free, slowed down by MRAs and slowed down by physicians phases. The throughput is strictly achieved (i.e., by exactly one policy) if (N_A, N_P) is taken in the interior of one of these three polyhedra.

A congestion diagram like Figure 4.7 can be very profitable for a real-life organization whose behaviour can be modeled by means of timed Petri nets. In our application (EMS-A), it can be used to estimate what long-run throughput ρ^* of the system a staffing (N_A, N_P) of the call center will achieve. Conversely and above all, it can be used to determine minimal staffing recommendations. For such an emergency call center and considering that calls arrive with rate λ , at least $\lceil \lambda(\tau_1 + \pi\tau_2) \rceil$ MRAs and $\lceil \lambda\pi(\tau_2 + \tau_3) \rceil$ emergency physicians are needed. We will develop this aspect in the Section 6.

In the light of this case-study, we finally go back to the Figure 4.6 and explain why the different choices of N_A and N_P lead to very different behaviours. Reusing the numerical values that were adopted for these simulations, we obtain $N_A^* = 3.9$ and $N_P^* = 1.65$. As a result, choosing $N_A = 4$ and $N_P = 2$ for the first scenario ensures that the long-run throughput vector is $\rho = \lambda e$, which is what we indeed obtain. In the second case where $N_A = 3$ and $N_P = 2$, we are in the situation where the policy $(A, 0)$ is ultimately active (here as soon as $t \geq 120$ s), and MRAs are limiting. On the contrary when $N_A = 4$ and $N_P = 1$, we have enough MRAs (which is why there is no difference between z_0 and z_1 for quite a long time) but not enough emergency physicians, so that the policy (λ, P) is applied between $t = 283.33$ s and $t = 316.67$ s and the policy (A, P) is applied from $t = 316.67$ s onwards. It is interesting to note that even though the policy (λ, P) is in practice never observed to realize the minimum *ultimate* throughput of the net, it can still be encountered during the execution, before the system reaches a final stable phase. The last situation $N_A = 3.5$ and $N_P = 1.5$ appears to be more chaotic than the others because we are located in the bottom red phase of Figure 4.7 but very close from the frontier $\rho_A = \rho_P/\pi$. Hence, we have many oscillations between the two policies $(A, 0)$ and (A, P) . The system does nevertheless reach a final phase (from $t = 1495.71$ s, which is outside of the charts of Figure 4.6), where (A, P) is active consistently with the congestion-diagram.

- **REMARK 4.7.** It seems important to notice that the minimum values N_A^* and N_P^* we compute here take the form of what is frequently referred to as the **offered load** in queueing theory. These minimum number of agents ensuring stability of the system are indeed the product of the arrival rate λ and an average handling time τ . It is well-known that Little's law (see [Lit61] or [CLo8, p. 457]) predicts that queues with such parameters have an average length of $N^* = \lambda\tau$ customers, and need at least N^* servers to reach stability. It is not very surprising, yet fortunate to retrieve that Petri net models and queues agree on this aspect.
- **REMARK 4.8.** Observe also that the minimum staffing that we obtain are safe lower bounds for the initial discrete Petri nets. Indeed, we computed these values based on continuous-relaxed nets that are typically "more fluid" than the discrete systems. Hence, our minimum staffing recommendations may be insufficient (by a slight amount) but definitely are necessary to prevent congestion in real systems.

3.4 LINEAR PROGRAMMING FORMULATION. – In addition to the Corollaries 4.4 and 4.5 which provide two methods to compute analytically the throughput vector ρ , we now consider the associated computational complexity problem. Indeed, solving the optimality equation (L) using an enumerative technique is in general exponential in the size $|\mathcal{P}| \times |\mathcal{Q}|$ of the net, and the same

complexity is hidden in the expression (4.10) since a minimum is taken on all the policies. The following result guarantees that these enumerative schemes can be avoided, as linear programs can be solved in (weak) polynomial time, for instance by the ellipsoid or interior point methods.

- **COROLLARY 4.9.** *In a continuous-relaxed non-Zeno and conflict-free timed Petri net with affine inputs and that admits a stoichiometric invariant, the throughput ρ can be computed in polynomial time by solving the following linear program:*

$$\max \sum_{q \in Q} \rho_q \quad \text{s.t.} \quad \begin{cases} \rho_q \leq \alpha_{qp}^{-1} \sum_{q' \in Q} \alpha_{pq'} \rho_{q'}, & \forall q \in Q, \quad \forall p \in q^{\text{in}} \\ u_q \leq \alpha_{qp}^{-1} m_p - \rho_q \tau_p + \alpha_{qp}^{-1} \sum_{q' \in Q} \alpha_{pq'} u_{q'}, & \forall q \in Q, \quad \forall p \in q^{\text{in}} \end{cases}$$

in which ρ and u in \mathbb{R}^Q are the variables. More precisely, if (ρ, u) is any optimal solution of this program, then ρ coincides with the throughput vector.

Proof. This is an application of a theorem of Denardo and Fox [DF68, Theorem 2] on undiscounted SMDP. They indeed prove that for any positive vector $v \in \mathbb{R}^S$, the average-cost vector g^* is solution of the linear program whose criterion is $\sum_{i \in S} v_i g_i$ and whose feasibility set is defined by inequalities $g_i \leq \sum_{j \in S} p_j^a g_j$ and $h_i \leq c^a - t^a g_i + \sum_{j \in S} p_j^a h_j$ for all $i \in S$ and $a \in A_i$. Switching to Petri net variables (ρ, u) in view of Correspondence Theorem 2 so that $g^* = \chi(\mathcal{S})_e$ and choosing $v = e$ gives the announced result by Correspondence Theorem 3. \clubsuit

Observe that the constraints of the linear program (4.9) are derived from the optimality equations (L1)–(L2), even though the fact that the minimum in (L2) has to be achieved for upstream places that already achieve minimum in (L1) does not appear explicitly – see to this purpose the discussion of Schäl in [Sch92] on the two forms of the second optimality equation for SMDPs.

In their work [GG04b], Gaujal and Giua derived the same linear program to characterize the throughput of continuous-relaxed Petri nets. What is interesting is that they built the feasible set and its constraints directly by considering instances of Little’s law, i.e., maximizing the flow of fluid-tokens under validity of the semantics rules. The Corollary 4.9 shows on the other hand that this linear program also naturally emerges from the theory of semi-Markov decision processes.

4 Deviation to a congestion-free regime

Now that we have characterized the throughput of monotonic Petri nets, that is according to Corollary 4.2–(ii) the first-order estimate of the long-run trajectories, we want to focus on **second-order estimates**. In the case where the holding times of the places of the Petri net are a multiple of a same positive value, we point out that the results presented in the Section 3.2 of Chapter 3 can be directly applied, owing to Correspondence Theorem 2. In a more general and proper semi-Markov setting however, deriving such results is more difficult. In particular, it involves arithmetic phenomena due to the possible periodic character of certain classes in the SMDP problems.

In what follows, we are going to study a subproblem of this class. We are particularly interested in dealing with Petri net systems that have to meet a known inbound demand. Similarly to our model (EMS-A), it is desirable that the system is **congestion-free** (or input-driven), meaning that only the inputs are ultimately limiting, and not the “internal” system itself. In this section, we are going to focus only on **single-input** Petri nets, i.e., for which $|Q_{\text{input}}| = 1$. Our objective is to characterize the congestion-free phase and to see if finer results than Corollary 4.2–(ii) can be obtained.

4.1 REDUCTION TO A STOCHASTIC SHORTEST PATH PROBLEM. — For the sake of simplicity, let us denote by q_0 the input transition of the generic single-input Petri net considered in this section, and by z_0 its counter function. When this Petri net admits a positive stoichiometric invariant e (element of $\mathbb{R}_{>0}^Q$), it is convenient to use the normalized vector $\bar{e} := e_0^{-1}e$ with unit value for q_0 . On the other hand, not choosing $e_0 = 1$ is a sound way to control the homogeneity of our expressions, which is why we will use $e_0^{-1}e$ instead of \bar{e} in some proofs.

In view of Correspondence Theorem 1 and Proposition 4.1, if z_0 is affine with rate λ , the input transition q_0 can be modeled by an absorbing (or sink) state 0 with cost $\lambda e_0^{-1}\tau_0$ and sojourn time τ_0 in the corresponding semi-Markov decision process (τ_0 is an arbitrary positive time, but the choices $\tau_0 = 1$ or $\tau_0 = \tau_{\max}$ are canonical).

The Assumption D below delineates two conditions under which we can reasonably hope that our single-input Petri net is congestion-free. Let us explain how we derive them. First, the global behaviour of the net shall be driven by the input only if q_0 indirectly constrains all the transitions of the net. This is a so-called “starvation” hypothesis, meaning that if the input flow stops, so will the entire net. Equivalently, and reusing the control theory terminology of the Section 2 of Chapter 2, this a minimal *controllability* assumption. In terms of Petri net structure, it amounts to requiring that q_0 has access to q for all q in $\mathcal{Q} \setminus \{q_0\}$, in the sense that there exists k in \mathbb{N}^* , transitions $q_{i_0}, q_{i_1}, q_{i_2}, \dots, q_{i_k}$ in \mathcal{Q}^{k+1} and places $p_{i_0}, p_{i_1}, \dots, p_{i_{k-1}}$ in \mathcal{P}^k such that $q_{i_0} = q_0, q_{i_k} = q$ and for all $0 \leq \ell \leq k, p_{i_\ell}$ is in $q_{i_\ell}^{\text{out}}$ and $q_{i_{\ell+1}}$ is in p_{i_ℓ} . According to our Proposition 2.16, it guarantees that the Petri net admits at most one positive stoichiometric invariant (up to a multiplicative constant). This justifies the use of \bar{e} , that is the only canonical stoichiometric invariant provided it exists.

In addition, and as shown by (4.5) assuming that there is a positive stoichiometric invariant e , the whole Petri net can reach a congestion-free regime with a throughput $\rho = \lambda \bar{e}$ only if the initial markings of the internal places (or equivalently the internal throughputs of the system) are large enough. Observe that because the state $\{0\}$ of the corresponding SMDP is absorbing, it is always a final class of any policy and it yields an SMDP average-cost $(\lambda e_0^{-1}\tau_0)/\tau_0 = \lambda e_0^{-1}$, or equivalently a Petri net throughput $\rho_0 = \lambda$. Contrarily, the internal throughputs of the system correspond to the average-costs of the policies of \mathfrak{S} on final classes different than $\{0\}$, so it is sound to define the minimal internal throughput ρ_{int} by

$$\rho_{\text{int}} := \min_{\sigma \in \mathfrak{S}} \min_{\substack{F \in \mathcal{F}(\sigma) \\ F \neq \{q_0\}}} \left\{ \frac{\langle \mu_F^\sigma, D^\sigma m^\sigma \rangle}{\langle \mu_F^\sigma, \tau^\sigma \rangle} \right\}, \quad (4.12)$$

where we have reused the notation of Section 3.3, in particular for all policies σ in \mathfrak{S} we have $D^\sigma := \text{diag}((e_0 e_q^{-1} \alpha_{q\sigma(q)}^{-1} \pi_{q\sigma(q)})_{q \in \mathcal{Q}})$. Our second condition to guarantee a congestion-free behaviour of the Petri net will then be $\rho_{\text{int}} \geq \rho_0 = \lambda$.

We can now formally state these two conditions; we will in fact particularly focus on cases where the internal throughput is strictly greater than the input pace to prevent any “quasi-congestion”.

► **ASSUMPTION D.** *The Petri net is continuous-relaxed, non-Zeno, admits a single input transition q_0 with affine profile of rate λ , and a positive stoichiometric invariant e . In addition,*

(1) *The input transition q_0 has access to every other transition q in $\mathcal{Q} \setminus \{q_0\}$,*

(2) *The internal throughput is strictly greater than the input pace, i.e., $\rho_{\text{int}} > \lambda$.*

In view of Correspondence Theorem 2, the Assumption D–(2) also has a clear interpretation in the undiscounted SMDP that corresponds to our Petri net with affine input and positive stoichiometric

invariant. It means that if the decision-maker who plays the game described by the SMDP wants to minimize his/her ultimate average-cost, the best strategy that he/she can find is to look for choices of actions that make him/her end up in state 0, where the average-cost is λ . Any other strategy in which 0 is not reached will for some starting states lead the player in a final class F different than $\{0\}$ where an average-cost strictly bigger than λ is incurred. Conversely, Assumption D-(1) ensures that 0 can be reached in this SMDP, as we are soon going to show.

In order to investigate how the trajectory z of our Petri net behaves under the Assumption D (that we have designed to describe a congestion-free situation) relatively to $t \mapsto \lambda t \bar{e}$ (the candidate function for the ultimate growth of a congestion-free solution of the dynamics), we introduce the reduced counters $(\Delta z)_{q \in Q}$ by $\Delta z(t) := z(t) - \lambda t \bar{e}$ for all $t \geq 0$. We are now ready to state our last correspondence theorem. We encourage the reader to have a look at the first part of the proof which details how this new correspondence arises.

► **CORRESPONDENCE THEOREM 4.** *Suppose that our continuous-relaxed and non-Zeno Petri net admits a single input transition q_0 with affine profile of rate λ , and a positive stoichiometric invariant e . Then, the evolution dynamics of the reduced counters Δz is equivalent to the dynamic programming equations of an undiscounted semi-Markov decision process in stochastic shortest path configuration.*

In addition, this SMDP in SSP configuration satisfies the two conditions of Assumption C if and only if the original Petri net satisfies the two conditions of Assumption D.

Proof. We first show how the SSP configuration emerges.

◇ Recall from the proof of Correspondence Theorem 1 the dynamics verified by the counters $(z_q)_{q \in Q}$ in the case where $Q_{\text{input}} = \{q_0\}$ and still focusing on the simplified conflict-free setting:

$$\forall q \in Q_{\text{sync}}, \quad z_q(t) = \min_{\substack{p \in q_{\text{in}} \\ p \notin \mathcal{P}_{\text{input}}}} \left\{ \alpha_{qp}^{-1} m_p + \alpha_{qp}^1 \sum_{q' \in Q} \alpha_{pq'} z_{q'}(t - \tau_p) \right\} \wedge \min_{q_{\text{in}} \cap q_0^{\text{out}} \neq \emptyset} \{z_0(t)\}. \quad (4.13)$$

Equivalently, reusing the notation of Correspondence Theorem 2, the renormalized counters $(z_q/e_q)_{q \in Q}$ of the Petri net satisfy for all $t \geq 0$ the relations

$$\forall q \in Q_{\text{sync}}, \quad z_q(t)/e_q = \min_{\substack{p \in q_{\text{in}} \\ p \notin \mathcal{P}_{\text{input}}}} \left\{ \alpha_{qp}^{-1} m_p/e_q + \sum_{q' \in Q} \xi_q^p(z_{q'}(t - \tau_p)/e_{q'}) \right\} \wedge \min_{q_{\text{in}} \cap q_0^{\text{out}} \neq \emptyset} \{z_0(t)/e_0\}. \quad (4.14)$$

As already stated by the Correspondence Theorem 2, these are the dynamic programming equations of an undiscounted SMDP with generic dynamic programming equations (DP):

$$\forall t \geq 0, \quad \forall i \in S, \quad v(i, t) = \min_{a \in A_i} \left\{ c^a + \sum_{j \in S} p_j^a v(j, t - t^a) \right\}, \quad (4.15)$$

where consistently with the Proposition 4.1, the input term $z_0(t)/e_0 = \lambda t e_0^{-1}$ distinguished in (4.14) has been integrated in the autonomous dynamics of (4.15), by letting $A_0 := \{a_0\}$, $c^{a_0} := \lambda \tau_{\max} e_0^{-1}$, $t^{a_0} := \tau_{\max}$ and knowing that there is some (unique) state i in $S \setminus \{0\}$ such that A_i contains an action leading deterministically to 0 with null cost and null sojourn time.

Let us suppose that $z_0(t) = \lambda t + \eta_0$. The dynamics of the reduced counters Δz simply write using the definition $\Delta z := z - \lambda t e_0^{-1} e$ and (4.14) as (for all $t \geq 0$ and $q \in Q_{\text{sync}}$):

$$\Delta z_q(t)/e_q = -\lambda t e_0^{-1} + \min_{\substack{p \in q_{\text{in}} \\ p \notin \mathcal{P}_{\text{input}}}} \left\{ \alpha_{qp}^{-1} m_p/e_q + \sum_{q' \in Q} \xi_q^p(\Delta z_{q'}(t - \tau_p)/e_{q'} + \lambda(t - \tau_p) e_0^{-1}) \right\} \wedge \min_{q_{\text{in}} \cap q_0^{\text{out}} \neq \emptyset} \left\{ \frac{\lambda t + \eta_0}{e_0} \right\}.$$

Since for all $p \in \mathcal{P} \setminus \mathcal{P}_{\text{input}}$, we have $\sum_{q' \in Q} \xi_q^p = 1$, all the terms $\lambda t e_0^{-1}$ cancel out and we are left with the dynamic programming equations of an undiscounted SMDP with reduced costs; indeed in SMDP form we

have

$$\forall t \geq 0, \quad \forall i \in S, \quad \Delta v(i, t) = \min_{a \in A_i} \left\{ (c^a - \lambda e_0^{-1} t^a) + \sum_{j \in S} p_j^a \Delta v(j, t - t^a) \right\}, \quad (4.16)$$

where $\Delta v := \Delta z \otimes e$. Most importantly, in this SMDP model, the absorbing state 0 is now cost-free since $c^{a_0} - \lambda e_0^{-1} t^{a_0} = \lambda \tau_{\max} e_0^{-1} - \lambda e_0^{-1} \tau_{\max} = 0$, and therefore this SMDP is exactly in shortest path configuration. To be complete, let us remark that a constant terminal cost $\eta_0 e_0^{-1}$ is incurred in state 0.

- ◇ We now establish that Assumption C and Assumption D are equivalent. We first prove $C(1) \iff D(1)$. Building on Correspondence Theorem 1 and its backward one-to-one correspondence between the above $(\mathcal{P}, \mathcal{Q})$ Petri net and the (S, A) SMDP, the accessibility of any $q \in \mathcal{Q} \setminus \{q_0\}$ from q_0 amounts to the accessibility of 0 from any i in $S \setminus \{0\}$, in the sense of Section 1.1 of Chapter 3. However, recall from this same section that if $i, j \in S$, then the accessibility of j from i amounts to the existence of a particular policy σ in \mathfrak{S} and an integer n in \mathbb{N} such that $(P^\sigma)_{ij}^n > 0$. In this sense, the implication $C(1) \implies D(1)$ is direct, because the existence of proper policy exactly tells us that 0 is accessible from any state. Conversely, suppose that $D(1)$ holds, so that the state 0 is accessible from any other state i of S , and let \mathcal{T} be a directed spanning tree with root 0 of the graph \mathcal{G} , whose nodes set is S and whose arcs set contains $(i, j) \in \mathcal{G}^2$ if there exists $a \in A_i$ such that $p_j^a > 0$ (such a spanning tree can be constructed using depth-first search). We now construct a proper policy state-by-state: for any i in S , there is by assumption a path of the form $i \rightarrow i' \rightarrow \dots \rightarrow 0$ in \mathcal{T} , we then choose for $\sigma(i)$ any action in A_i such that $p_{i'}^a > 0$. The accessibility property and the fact that state 0 is absorbing guarantee that $(P^\sigma|_{(S \setminus \{0\}) \times (S \setminus \{0\})})^{|S|}$ is strictly substochastic, hence $\rho(P^\sigma|_{(S \setminus \{0\}) \times (S \setminus \{0\})}) < 1$ and we have $\lim_{n \rightarrow \infty} (P^\sigma|_{(S \setminus \{0\}) \times (S \setminus \{0\})})^n = 0$, so that σ is proper, hence the implication $D(1) \implies C(1)$.
- ◇ We continue by showing $D(2) \implies C(2)$. Assume $D(2)$ holds, so that $\rho_{\text{int}} > \lambda$. In SMDP notation, it means that

$$\min_{\sigma \in \mathfrak{S}} \min_{\substack{F \in \mathcal{F}(\sigma) \\ F \neq \{0\}}} \left\{ \frac{\langle \mu_F^\sigma, c^\sigma \rangle}{\langle \mu_F^\sigma, t^\sigma \rangle} \right\} > \lambda e_0^{-1}, \quad \text{or equivalently} \quad \underline{\chi} := \min_{\sigma \in \mathfrak{S}} \min_{\substack{F \in \mathcal{F}(\sigma) \\ F \neq \{0\}}} \left\{ \frac{\langle \mu_F^\sigma, c^\sigma - \lambda e_0^{-1} t^\sigma \rangle}{\langle \mu_F^\sigma, t^\sigma \rangle} \right\} > 0,$$

in other words our undiscounted SMDP with reduced costs whose dynamics is given by (4.16) has a positive average-cost if it ends up in a final class different from $\{0\}$. Let us take an improper policy σ in \mathfrak{S} of this SMDP. By definition, because σ is improper, it admits a final class F that is not $\{0\}$. According to Theorem 3.31, playing the stationary strategy associated with this policy starting from a state in F hence forces the player to incur an expected cost with growth rate bigger than $\underline{\chi}$. As a result, applying this policy over an infinite horizon yields positive infinite total cost starting in states of F , that is the requirement $C(2)$.

- ◇ To conclude the proof, we show that $C(1) \wedge \neg D(2) \implies \neg C(2)$, since it brings $C(1) \wedge C(2) \implies C(1) \wedge D(2)$. If $D(2)$ does not hold, it means by (4.12) (converted in SMDP notation with our reduced costs) that there exists a policy σ_1 of \mathfrak{S} and a final class F of σ_1 such that $\langle \mu_F^{\sigma_1}, c^{\sigma_1} - \lambda e_0^{-1} t^{\sigma_1} \rangle / \langle \mu_F^{\sigma_1}, t^{\sigma_1} \rangle \leq 0$. However, if $C(1)$ also holds, it means that we have a proper policy σ_p in \mathfrak{S} . We can therefore introduce the new policy σ in \mathfrak{S} such that for all i in F , $\sigma(i) := \sigma_1(i)$ and for all i in $S \setminus F$, $\sigma(i) := \sigma_p(i)$. Let us denote $S' := S \setminus (\{0\} \cup F)$, then since $P^\sigma|_{S' \times S'}$ is extracted from $P^{\sigma_p}|_{S \setminus \{0\} \times S \setminus \{0\}}$ that have a spectral radius strictly inferior than 1, it cannot contain any stochastic submatrix according to Lemma 3.5, which means that the policy σ admits only $\{0\}$ and F as final classes. The fact that $F \neq \{0\}$ is a final class of policy σ naturally means that σ is improper. Finally, by Theorem 3.31 and Corollary 3.32, we know that for all i in S , the expected cost $v^\sigma(i, t)$ incurred by playing the stationary strategy associated with σ up to horizon t starting from i in our SMDP in SSP configuration with the reduced costs (4.16) is such that $v^\sigma(i, t) = \chi_i t + O(1)$, where χ_i is a nonnegative linear combination of 0 and $\langle \mu_F^{\sigma_1}, c^{\sigma_1} - \lambda e_0^{-1} t^{\sigma_1} \rangle / \langle \mu_F^{\sigma_1}, t^{\sigma_1} \rangle$.

Hence, we have exhibited an improper policy σ which does not yield an infinite total cost, whatever the starting state is, so $C(2)$ is false. \mathfrak{S}

It is enlightening to compare Assumption C and Assumption D in view of Correspondence Theorem 4. Indeed, we see that having an input which imposes starvation in the context of Petri net corresponds to the ability to (virtually) stop the game by reaching the state 0 from any starting state (after some stroll) in the world of SMDPs. Similarly, not being limited by the internal part of the Petri net is equivalent to loose money if we do not stop the game described by the SMDP. This leads to the following important corollary.

- **COROLLARY 4.10.** *Under Assumption D and supposing that the input profile is $z_0(t) = \lambda t + \eta_0$ for all $t \geq 0$, the solution z of the Petri net dynamics satisfies*

$$z(t) \underset{t \rightarrow \infty}{=} \lambda t \bar{e} + \eta + o(1)$$

regardless of the seed function of \mathcal{L} it is initialized with, where the vector η of \mathbb{R}^Q is the unique solution of the equations

$$\forall q \in Q \setminus \{q_0\}, \quad x_q = \min_{p \in q^{\text{in}}} \left\{ \alpha_{qp}^{-1} \pi_{qp} \left(m_p + \sum_{q' \in Q} \alpha_{pq'} x_{q'} \right) - \lambda \tau_p e_q e_0^{-1} \right\} \quad (4.17)$$

and $x_0 = \eta_0$

The function $t \mapsto \lambda t \bar{e} + \eta$ is an affine regime of the Petri net in the sense of Corollary 4.2–(i), and in addition the places achieving the minimum of (4.17) form a proper policy of the corresponding SMDP in SSP configuration.

Proof. Denoting by $\Delta v = z \odot e - (\lambda t + \eta_0) e_0^{-1} \mathbf{1}$, we obtain by computation identical to the beginning of the proof of Correspondence Theorem 4 that Δv is the value function of an undiscounted SMDP in SSP configuration and with null terminal costs, that follows the dynamic programming equations (4.16). We can thus apply Theorem 3.41, saying that under Assumption C (that is here verified for this SMDP since Assumption D is verified for the Petri net), the trajectory of this SMDP converges towards the unique solution u^* of the equations

$$\forall i \in S \setminus \{0\}, \quad u(i) = \min_{a \in A_i} \left\{ c^a - \lambda e_0^{-1} t^a + \sum_{j \in S} p_j^a u(j) \right\} \quad (4.18)$$

and $u(0) = 0$

derived from (3.29) of Theorem 3.38. Multiplying by e , we get that $z(t) - \lambda t \bar{e}$ admits the limit $\eta := u^* \odot e + \eta_0 \bar{e}$ when $t \rightarrow \infty$. Switching back to Petri net notation introduced in the proof of Correspondence Theorem 2 provides the system (4.17) on η .

The fact that $t \mapsto \lambda t \bar{e} + \eta$ is an affine regime of the Petri net follows by the remark subsequent to the Corollary 4.4, since the tuple $(\lambda \bar{e}, \eta)$ is solution of the lexicographic system (L). The final assertion of the theorem on the fact that η is determined by a proper policy is a consequence of Theorem 3.38. \clubsuit

This corollary has major consequences for applications. To begin with, observe that it indeed refines Corollary 4.2–(ii) in which only the asymptotic growth $z(t) \underset{t \rightarrow \infty}{=} \lambda t \bar{e} + O(1)$ was obtained. In this sense, the (large) inequality of Assumption D(2) effectively characterizes the congestion-free regime, in which only the input with rate λ constrains the whole system. But in addition, the strength of Corollary 4.10 is mostly the fact that the trajectory converges towards an affine function that *does not* depend on the initial conditions. This precisely means that **the internal system catches up the input** regardless of its own starting state. This behaviour is enabled by the fact that Assumption D–(2) ensures that the system is strictly in the congestion-free phase, so that the internal part goes strictly faster than the input and can indeed catch it up. The vector η featured in the Corollary 4.10 plays the role of a delay between the internal system and the input, more particularly for all q in $Q \setminus \{q_0\}$, the scalar $\eta_q - \eta_0$ gives an indication of how long it takes the input to access q .

To wrap things up, we let the reader appreciate how in the asymptotic expansion of the Petri net trajectory, we have first characterized the first-order term by an average-cost problem on SMDPs and the second-order term by a stochastic shortest path problem, two classical questions in this literature.

Model	Defining equation	Ultimate behaviour
Single-input monotonic Petri net	(4.13)	$\mathcal{S}_t^{\text{PN}}[r \mapsto \lambda r \bar{e} + \eta](s) = \lambda(t+s)\bar{e} + \eta$ $\forall z^0 \in \mathcal{Z}, \quad \mathcal{S}_t^{\text{PN}}[z^0](s) \underset{t \rightarrow \infty}{=} \lambda(t+s)\bar{e} + \eta + o(1)$
Corresponding undiscounted SMDP	(4.15)	$\mathcal{S}_t^{\text{SMDP}}[r \mapsto (\lambda r + \eta_0)e_0^{-1}\mathbf{1} + u^*](s) = (\lambda(t+s) + \eta_0)e_0^{-1}\mathbf{1} + u^*$ $\forall v^0 \in \mathcal{V}, \quad \mathcal{S}_t^{\text{SMDP}}[v^0](s) \underset{t \rightarrow \infty}{=} (\lambda(t+s) + \eta_0)e_0^{-1}\mathbf{1} + u^* + o(1)$
Reduced SMDP in SSP config.	(4.16)	$\mathcal{S}_t^{\text{SMDP}-\Delta}[r \mapsto u^*](s) = u^*$ $\forall v^0 \in \mathcal{V}_0, \quad \mathcal{S}_t^{\text{SMDP}-\Delta}[v^0](s) \underset{t \rightarrow \infty}{=} u^* + o(1)$

Table 4.10: Affine regimes and long-run behaviour of our three models under Assumption D (or the equivalent Assumption C)

We reckon it can be useful at this point to summarize in the Table 4.10 the journey that we have made from Petri net models to SMDP models, and what it means in terms of their respective dynamics. Recall that the Correspondence Theorem 2 stated that the dynamics of the counter functions of a monotonic fluid Petri net coincide with the dynamic programming equations of an undiscounted SMDP with no stopping costs. Implicitly, it enabled us in the single-input case with $z_0(t) = \lambda t + \eta_0$ for all $t \geq 0$ to describe the dynamics (4.13) of the Petri net by an evolution semigroup $(\mathcal{S}_t^{\text{PN}})_{t \geq 0}$ acting on the seed functions $\mathcal{Z} := \mathcal{F}([- \tau_{\max}, 0), \mathbb{R}^{Q \setminus \{q_0\}})$. This semigroup coincides via scaling by the stoichiometric invariant e with the evolution semigroup $(\mathcal{S}_t^{\text{SMDP}})_{t \geq 0}$ of the corresponding SMDP with dynamics (4.15) and designed to act on the initial conditions $\mathcal{V} := \{v \in \mathcal{F}([- \tau_{\max}, 0), \mathbb{R}^S) \mid \forall s \in [- \tau_{\max}, 0), v(0, s) = \lambda(s + \eta_0)e_0^{-1}\}$. Under the Assumption D that the Petri net behaves in a congestion-free (or input-driven) regime, the Corollary 3.32 and the Correspondence Theorem 3 gave us that $\mathcal{S}_t^{\text{PN}}[z^0](s) \sim \lambda t \bar{e}$ and $\mathcal{S}_t^{\text{SMDP}}[v^0](s) \sim \lambda t e_0^{-1}\mathbf{1}$ when $t \rightarrow \infty$, for all z^0 in \mathcal{Z} and v^0 in \mathcal{V} . But we have in fact successfully found the next term of the asymptotic expansion, because we considered a reduced SMDP with costs $(c^a - \lambda e_0^{-1} t^a)_{a \in A}$, satisfying the dynamics (4.16) and associated with the semigroup $(\mathcal{S}_t^{\text{SMDP}-\Delta})_{t \geq 0}$ acting on the initial conditions of $\mathcal{V}_0 := \{v \in \mathcal{F}([- \tau_{\max}, 0), \mathbb{R}^S) \mid \forall s \in [- \tau_{\max}, 0), v(0, s) = 0\}$. Indeed, this new SMDP in SSP configuration admitted by design an average-cost (or growth-rate) of 0, and the constant function $r \mapsto u^*$ of \mathcal{V}_0 appeared as a fixed-point (special case of affine regime) of $\mathcal{S}^{\text{SMDP}-\Delta}$, where u^* is the solution of the system (4.18). Because the Theorem 3.41 ensured that the dynamics of this reduced-cost SMDP converged towards u^* , we have been able to go back up the chain and prove that the non-reduced SMDP dynamics (resp. Petri net dynamics) converged towards the affine regime $r \mapsto \lambda r e_0^{-1}\mathbf{1} + \eta_0 e_0^{-1}\mathbf{1} + u^*$ (resp. $r \mapsto \lambda r \bar{e} + \eta$ with $\eta := (\eta_0 e_0^{-1}\mathbf{1} + u^*) \odot e$) regardless of the initial conditions.

*
* *

We illustrate the Correspondence Theorem 4 on our (EMS-A) call center model, reusing the positive stoichiometric invariant $e = (1, 1, 1 - \pi, \pi, \pi, \pi)^\top$. It is easy to verify on the Figure 4.5 that the input transition of the Petri net has access to all the transitions of the system (or in the SMDP version of the model, that all the states have access to 0). We verify further that when all the costs $(c^a)_{a \in A}$ of the SMDP are reduced in the form $(c^a - \lambda t^a)_{a \in A}$, the state 0 becomes cost-free. The internal throughput has been computed in Section 3.3 and verifies $\rho_{\text{int}} = \rho_A \wedge (\rho_P / \pi)$, reusing the notation of Section 3.2. Our medical emergency call center therefore is in (strict) congestion-free regime if and only if $\rho_A > \lambda$ and $\rho_P > \pi\lambda$. Moreover, it is clear from Table 4.9 that the policy $(\lambda, 0)$ is the only proper policy of the system (because it is the only one which only admits $\{0\}$ as a final class). Consequently, the vector η featured in Corollary 4.10 is characterized by the equations

$$\begin{cases} \eta_1 = \eta_0 \\ \eta_2 = (1 - \pi)\eta_1 - (1 - \pi)\lambda\tau_1 \\ \eta_3 = \pi\eta_1 - \pi\lambda\tau_1 \\ \eta_4 = \eta_3 - \pi\lambda\tau_2 \\ \eta_5 = \eta_4 - \pi\lambda\tau_3 \end{cases},$$

so that we have $\eta = \eta_0 e - \lambda(0, 0, \tau_1, \tau_1, \tau_1 + \tau_2, \tau_1 + \tau_2 + \tau_3)^\top \odot e$, and whatever the starting conditions of the system (the number of pending calls), the trajectory z of the Petri net verifies

$$z(t) = \begin{pmatrix} z_0(t) \\ z_1(t) \\ z_2(t) \\ z_3(t) \\ z_4(t) \\ z_5(t) \end{pmatrix} \underset{t \rightarrow \infty}{=} \eta_0 e + \begin{pmatrix} \lambda t \\ \lambda t \\ (1 - \pi)\lambda(t - \tau_1) \\ \pi\lambda(t - \tau_1) \\ \pi\lambda(t - (\tau_1 + \tau_2)) \\ \pi\lambda(t - (\tau_1 + \tau_2 + \tau_3)) \end{pmatrix}.$$

We retrieve the ultimate affine regime that we were able to read in the first scenario of Figure 4.6.

4.2 TRANSIENCE TIME NEEDED TO CATCH UP THE INPUT. — As mentioned before, the Correspondence Theorem 4 can be interpreted as a systematic convergence of the trajectory of our Petri net towards the input-driven affine regime $z^{\text{aff}}: t \mapsto \lambda t \bar{e} + \eta$. The next question that comes in mind is naturally... how quick does the convergence occur? Another way to rephrase the problem is the following. Suppose that the system is already in the regime z^{aff} (which entails little loss of generality precisely since this behaviour is always ultimately reached), and that at the instant $\underline{t} > 0$ the input profile suffers a temporary surge of total magnitude $M \geq 0$. Any such perturbation can actually be (crudely) bounded above by an instantaneous extra demand of M , in a step-perturbation (or Heaviside) fashion, so that our total input profile will be $z_0(t) = \lambda t + \eta_0 + M \mathbb{1}_{[\underline{t}, \infty)}(t)$ (see [LGO18] for finer analytical models of bursts in the arrival process). Owing to the input catch-up property, we know that under Assumption D the trajectory z of the system (whose existence is ensured by Theorem 2.5) shall asymptotically verify $z(t) \underset{t \rightarrow \infty}{=} \lambda t \bar{e} + \eta + M \bar{e} + o(1)$. In this context, we introduce the **catch-up times** (or transience times) $(\theta_q)_{q \in Q}$ of the Petri net by

$$\forall q \in Q, \quad \theta_q := \inf \{ t \geq \underline{t} \mid z_q(t) = \lambda t \bar{e}_q + \eta_q + M \bar{e}_q \}, \quad (4.19)$$

that are naturally still independent on the choice of e . Note that we may as well have defined θ_q by the time needed for $z_q(t)$ to approach $(\lambda t + M) \bar{e}_q + \eta_q$ up to a chosen precision ε ; the exact

convergence choice (4.19) will although be justified in Section 5. Still denoting by u^* the solution of (4.18), and denoting by $\mathbb{1}_{S \setminus \{0\}} = \mathbf{1} - \mathbb{1}_{\{0\}}$ the vector of \mathbb{R}^S with all unit coordinates except the first one associated with state 0, the following characterization of the catch-up times holds.

► **THEOREM 4.11.** *Suppose that Assumption D is verified. Let v be the function uniquely determined by the initial condition $s \mapsto u^* - Me_0^{-1}\mathbb{1}_{S \setminus \{0\}}$ in \mathcal{V}_0 and the dynamics (4.16) of the reduced-costs SMDP in SSP configuration featured in Correspondence Theorem 4.*

Then, for all corresponding transition-state pairs (q, i) in $\mathcal{Q} \times S$, we have

$$\theta_q = \underline{t} + \inf \{t \geq 0 \mid v(i, t) = u^*(i)\}.$$

Proof. Let us reuse the notation of the Table 4.10, detailed at the end of Section 4.1. The core equality that we have to justify is that for all $t \geq \underline{t}$ and s in $[-\tau_{\max}, 0)$ such that $t + s \geq \underline{t}$, we have

$$z(t + s) = \mathcal{S}_{t-\underline{t}}^{\text{PN}}[r \mapsto \lambda(r + \underline{t})\bar{e} + \eta + M\mathbb{1}_{\{0\}}](s). \quad (4.20)$$

Indeed, we know that without the perturbation added at $t = \underline{t}$ we would have had $z(t + s) = \mathcal{S}_t^{\text{PN}}[z^{\text{aff}}](s) = \mathcal{S}_{t-\underline{t}}^{\text{PN}}[\mathcal{S}_{\underline{t}}^{\text{PN}}[z^{\text{aff}}]](s) = \mathcal{S}_{t-\underline{t}}^{\text{PN}}[r \mapsto \lambda(r + \underline{t})\bar{e} + \eta](s)$, since the dynamics are initialized with the affine regime z^{aff} : $r \mapsto \lambda r \bar{e} + \eta$. However, from $t = \underline{t}$ the input changes, hence the introduction of the difference $M\mathbb{1}_{\{0\}}$. This correctly models the new input and does not affect the other coordinates, since the input controls the single transition it is directly connected to with no time delay (see (4.2)). Everything actually happens as if the input z_0 was already shifted by M from the instant $\underline{t} - \tau_{\max}$, but this is only revealed to the internal part of the system at the physical time \underline{t} , that prior to this instant had pursued along the affine regime z^{aff} .

Let us define the reduced function $w(t) = z(t + \underline{t}) \odot e - (\lambda(t + \underline{t}) + \eta_0)e_0^{-1}\mathbf{1}$. With the same reductions than from the proof of Correspondence Theorem 4, we have for all $t \geq 0$ and all s in $[-\tau_{\max}, 0)$ such that $t + s \geq 0$ that

$$w(t + s) = \mathcal{S}_{(t+s)-\underline{t}}^{\text{SMDP}-\Delta}[r \mapsto u^* + Me_0^{-1}\mathbb{1}_{\{0\}}](s).$$

Now, the (constant) initial condition $r \mapsto u^* + Me_0^{-1}\mathbb{1}_{\{0\}}$ is not in \mathcal{V}_0 , the natural seed space of $\mathcal{S}^{\text{SMDP}-\Delta}$, because its first coordinate is nonzero. But we can leverage on the additively homogeneous character of this semigroup (see Proposition 3.26-(1)) to transform this “positive shift” in a “negative shift”, and write

$$w(t + s) = \mathcal{S}_t^{\text{SMDP}-\Delta}[r \mapsto u^* - Me_0^{-1}\mathbb{1}_{S \setminus \{0\}}](s) + Me_0^{-1}\mathbf{1} = v(t + s) + Me_0^{-1}\mathbf{1},$$

where we have now recognized the function v defined in the claim of the theorem. Because Assumption D is satisfied for the Petri net, the Assumption C is satisfied for the reduced SMDP in SSP configuration, and we know by Theorem 3.41 that $\lim_{t \rightarrow \infty} w(t) = u^* + Me_0^{-1}\mathbf{1}$.

To sum it up, we have for all $t \geq \underline{t}$ that $z(t) = \lambda t \bar{e} + \eta_0 \bar{e} + M \bar{e} + w(t - \underline{t}) \odot e$, we can confirm using this equality that $z(t)_{t \rightarrow \infty} = (\lambda t + M)\bar{e} + \eta + o(1)$, but also that for all $q \in \mathcal{Q}$, $z_q(t) = (\lambda t + M)\bar{e}_q + \eta_q$ if and only if $v(i, t - \underline{t}) = u^*(i)$, where i is the state of S corresponding to q ; hence the result. \mathcal{E}

In other words, the Theorem 4.11 tells us that the catch-up times associated with a step perturbation of magnitude M affecting the input of a monotonic fluid Petri net correspond to the (state-by-state) convergence times of the **value iteration** algorithm in an SMDP in SSP configuration, to reach its value u^* starting from the initial condition $u^* - Me_0^{-1}\mathbb{1}_{S \setminus \{0\}}$.

*
* *

We illustrate the Theorem 4.11 on the Figure 4.11, choosing a number $N_A = 6$ of medical regulation assistants and $N_P = 2$ of emergency physicians for the model (EMS-A). Consistently with the analysis in Section 3.3, this ensures that we are in the congestion-free regime. The inbound demand begins with $z_0(t) = \lambda t$ still with $\lambda = 0.025 \text{ s}^{-1}$ (one call every 40 seconds).

First have a look at the Petri net counters z in Figure 4.11a. The system is not initialized in the affine regime exhibited in Section 4.1 but just like shown in the first scenario of Figure 4.6, this

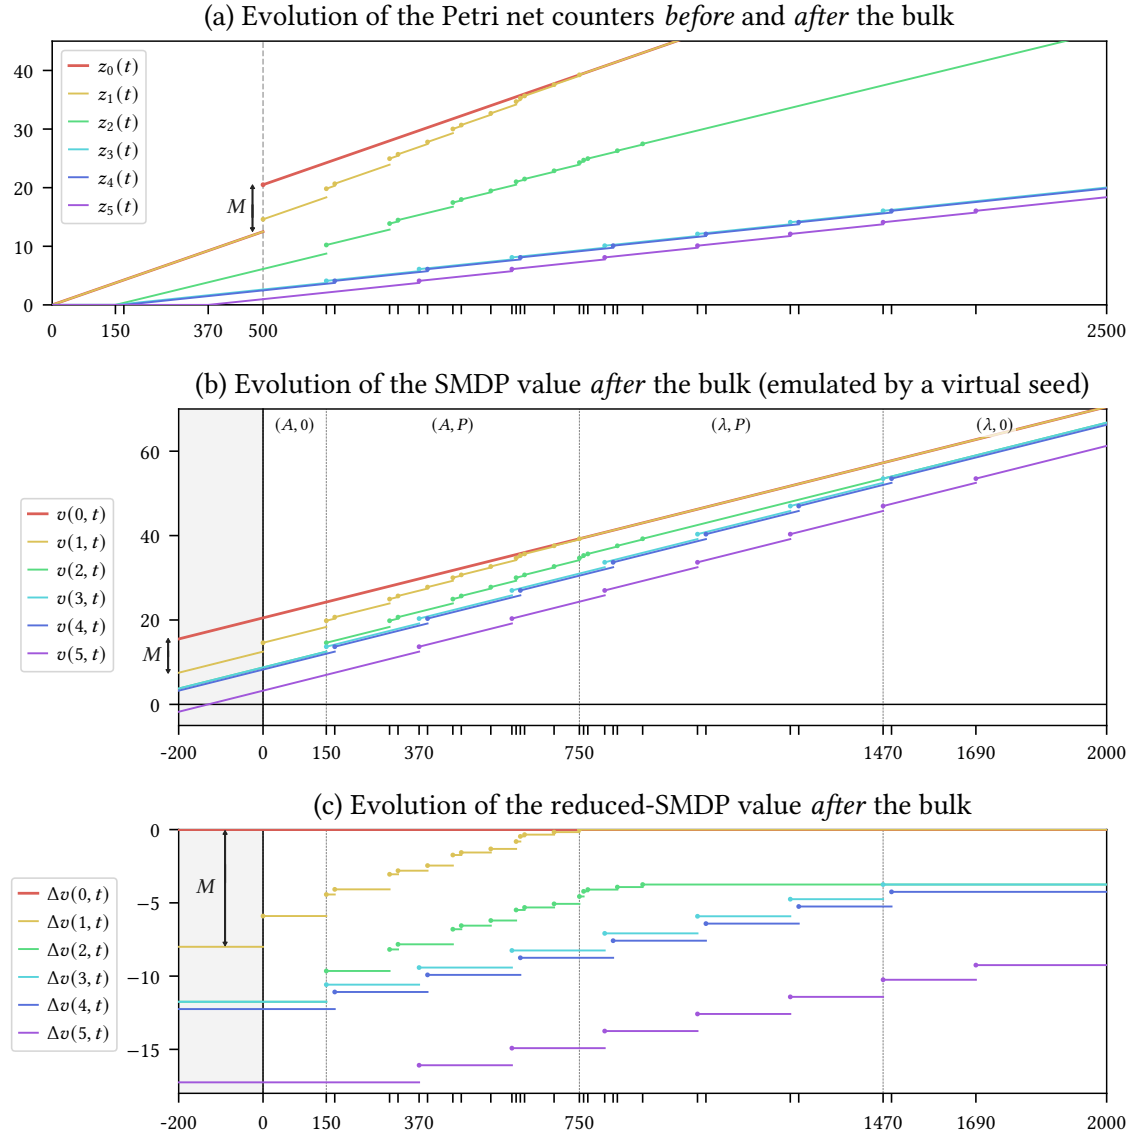


Figure 4.11: Illustration of the input catch-up after a perturbation, in (EMS-A).

regime is reached from the physical time $t = 370$ s onwards. The same evolution than the latter scenario would actually continue along if it was not for the bulk of $M = 8$ calls arising at $\underline{t} = 500$ s, so that z_0 is discontinuous at \underline{t} (we have $z_0(\underline{t}^-) = 12.5$ and $z_0(\underline{t}) = 20.5$). Before the bulk, all the calls were picked up with no delay and there was even a number $N_A - N_A^* = N_A - \lambda(\tau_1 + \pi\tau_2) = 2.1$ of idle agents, which explains why z_1 can instantaneously jump to the value $z_1(\underline{t}) = 12.5 + 2.1 = 13.6$. But then, all the MRAs are busy and only the progressive returns of previously occupied MRAs can make z_1 grow. All the further jumps correspond to call pick ups by MRAs or emergency doctors that were idle before \underline{t} .

We depict in Figure 4.11b the evolution of the value function v of the (non-reduced) SMDP corresponding to the Petri net (recall that it is represented in Figure 4.5) from the initial condition $r \mapsto \lambda(r + \underline{t})e_0^{-1}\mathbf{1} + u^* + M\mathbb{1}_{\{0\}}$ in \mathcal{V} . This chart is easier to read since all the slopes are equal. Although this initial condition is purely *virtual*, it does successfully emulate the behaviour of the renormalized counters $(z_q/e_q)_{q \in Q}$ for $t \geq \underline{t}$ as shown in the proof of Theorem 4.11. Remark

that since $u^*(1) = u^*(0) = 0$ in this example, the value $v(0, -\tau_{\max}) - v(1, -\tau_{\max})$ reduces to $-M$. The input catch-up can be effectively seen in this figure; the transience regime occurs as long as the “value function beam” shrinks and is over when the size of the beam only amounts to u^* . This arises from an extra time of 1690 s after the bulk, corresponding to $t = 2190$ s for the original system. Observe that the catch-up occurs in multiple phases, because the MRAs have absorbed all the extra inbound calls at the instant $t = 1250$ s (or 750 s after the bulk). The rest of the transient regime is only caused by the emergency doctors, for which the margin $N_P - N_P^* = N_P - \pi\lambda(\tau_2 + \tau_3) = 0.35$ was thin. At the time $t = 1970$ s though, the last call of the bulk that arose at \underline{t} is taken for transfer by an emergency doctor, so that z_3 is all caught up. The associated consultation indeed ends at $t = 2190$ s and the whole peak has been handled.

The Figure 4.11c provides the same information than Figure 4.11b however normalized by the quantity $\lambda(t + \underline{t})$. In other words, it exactly depicts the convergence towards u^* of the value function Δv of the SMDP with reduced costs $(c^a - \lambda t^a)_{a \in A}$ that is in SSP configuration when the terminal costs are $u^* - M\mathbb{1}_{S \setminus \{0\}}$ – this seed can be seen in the shaded area $[-\tau_{\max}, 0)$.

5 Convergence time of the value iteration algorithm for semi-Markov SSP problems

As suggested by Theorem 4.11 and its subsequent remark, we are going to focus on the convergence time of the value iteration procedure for SMDPs in SSP configuration, in the hope of getting more insight on the catch-up times of Petri nets that are in congestion-free phases. We shall be particularly interested in knowing if a **finite-time** convergence can be obtained, because it is a very desirable property of the systems we study in our applications – notably emergency call centers in which we would like the peak of calls to be absorbed in a short duration.

In all this section, we let a bit aside the Petri net models that correspond to SMDPs and we mainly focus on a generic non-Zeno undiscounted SMDP in SSP configuration, meaning that Assumption A and Assumption C hold. We suppose with no loss of generality that a single action a_0 is playable from the absorbing and cost-free state 0.

5.1 THE DIFFERENT CONVERGENCE PHASES. — Recall as stated in Theorem 3.41 that a function determined by an initial condition v^0 in \mathcal{V}_0 and satisfying the dynamic programming equations of an SMDP in SSP configuration converges towards u^* , when the planning horizon approaches infinity. We first provide a quantitative description of the associated speed of convergence.

For all i in $S \setminus \{0\}$, we denote by A_i^* the nonempty subset of A_i composed of **optimal actions**, that is to say those achieving minimality in (3.29). Similarly, we introduce $\mathfrak{S}^* := A_0 \times \prod_{i \in S \setminus \{0\}} A_i^*$ the set of **optimal policies**, included in \mathfrak{S} . Observe that optimal policies must be proper (for improper policies yield infinite costs for some starting states). In accordance with the Remark 3.27, optimal actions induce an evolution semigroup \mathcal{S}^* which corresponds to applying only optimal policies. The following theorem shows that there is always an instant t^* such that: (i) before t^* , either optimal or non-optimal actions can be picked, and thus the evolution of the dynamics amounts to applying the semigroup \mathcal{S} ; (ii) after t^* , only optimal actions are chosen, so that the finer dynamics associated with the semigroup \mathcal{S}^* are actually used. Moreover, the second phase induces a geometric speed of convergence controlled by the spectral radii of the probability matrices of optimal policies. Once again, we leverage on techniques from non-linear Perron–Frobenius theory to control the contraction aspects of the dynamics.

- **THEOREM 4.12.** *Suppose Assumptions A and C hold. Let v^0 in \mathcal{V}_0 and denote by v the solution of (DP) it determines. Then, there exists $t^* \geq 0$ such that for all $t \geq t^*$ and all s in $[-t_{\max}, 0)$, we have*

$$v(t+s) = \mathcal{S}_{t-t^*}^* [\mathcal{S}_{t^*}^* [v^0]](s).$$

Moreover, if we define

$$\nu := \max_{\sigma \in \mathfrak{S}^*} \rho(P^\sigma|_{(S \setminus \{0\}) \times (S \setminus \{0\})}) < 1,$$

then for all $\varepsilon > 0$ small enough, there is a weighted sup-norm in which $\mathcal{S}_{t_{\max}}^*$ is a contraction of rate $\nu + \varepsilon$.

Proof. We start by showing the ultimate restriction to the optimal dynamics.

- ◊ We know that v satisfies the equation (DP), which means that

$$\forall i \in S, \quad \forall t \geq 0, \quad v(i, t) = \min_{a \in A_i^*} \left\{ c^a + \sum_{j \in S} p_j^a v(j, t - t^a) \right\} \wedge \min_{a \in A_i \setminus A_i^*} \left\{ c^a + \sum_{j \in S} p_j^a v(j, t - t^a) \right\}.$$

In addition, in view of Theorem 3.41, $\lim_{t \rightarrow \infty} v(t) = u^*$, where u^* is characterized by the equations $u(0) = 0$ and $\forall i \in S, u(i) = \min_{a \in A_i^*} \{c^a + \sum_{j \in S} p_j^a u(j)\}$ (see (3.29)). By definition of the optimal actions, we know that

$$\forall i \in S, \quad \forall a \in A_i \setminus A_i^*, \quad u^*(i) < c^a + \sum_{j \in S} p_j^a u^*(j).$$

Let $\varepsilon^* := \min_{i \in S} \min_{a \in A_i \setminus A_i^*} \{c^a + \sum_{j \in S} p_j^a u^*(j) - u^*(i)\} > 0$. We know that there exists $t_0 \geq 0$ such that $\forall t \geq t_0, \forall i \in S, |v(i, t) - u^*(i)| \leq \varepsilon^*/2$. In particular

$$\forall t \geq t_0 + t_{\max}, \quad \forall i \in S, \quad \forall a \in A_i \setminus A_i^*, \quad c^a + \sum_{j \in S} p_j^a v(j, t - t^a) \geq c^a + \sum_{j \in S} p_j^a (u^*(j) - \varepsilon^*/2) \geq \varepsilon^*/2 + u^*(i)$$

which proves that only optimal actions are played from the time $t^* := t_0 + t_{\max}$.

- ◊ For the second part of the claim, let us consider the self-map T^* of $\mathbb{R}^{S \setminus \{0\}}$ that is defined by $T^*(x) = \max_{\sigma \in \mathfrak{S}^*} \{P^\sigma|_{S \setminus \{0\} \times S \setminus \{0\}} x\}$ (that can be seen as the Bellman operator T of MDPs restricted to optimal policies and non-terminal states). Recall as detailed in the proof of Proposition 3.13 that the generalized Collatz-Wielandt theorem (see [AGN11, Theorem 1.1 and Proposition 8.1]) provides

$$\inf \{ \ell > 0 \mid \exists x \in \mathbb{R}_{>0}^{S \setminus \{0\}}, T^*(x) \leq \ell x \} = \max_{\sigma \in \mathfrak{S}^*} \{ \rho(P^\sigma|_{S \setminus \{0\} \times S \setminus \{0\}}) \} = \nu.$$

Taking $\varepsilon > 0$ such that $\nu + \varepsilon < 1$, we obtain the existence of a vector w in $\mathbb{R}^{S \setminus \{0\}}$ such that $T^*(w) \leq (\nu + \varepsilon)w$. We still denote by w the augmented vector $(0, w)$ in \mathbb{R}^S , and we denote by \tilde{w} the constant function of \mathcal{V}_0 such that $\tilde{w}(s) = w$ for all $s \in [-t_{\max}, 0)$. Let us now take an initial condition v^0 in \mathcal{V}_0 , and $\alpha > 0$. We want to show that $\mathcal{S}_{t_{\max}}^*$ is subhomogeneous with respect to the function \tilde{w} and the factor $\nu + \varepsilon$, which means that $\mathcal{S}_{t_{\max}}^* [v^0 + \alpha \tilde{w}] \leq \mathcal{S}_{t_{\max}}^* [v^0] + \alpha(\nu + \varepsilon)\tilde{w}$. Denoting by v (resp. v_α) the solution of the optimal dynamics (equation (DP) replacing A_i by A_i^*) determined by the seed v^0 (resp. $v^0 + \alpha \tilde{w}$), we recall that for $s \in [-t_{\max}, 0)$, the vectors $\mathcal{S}_{t_{\max}}^* [v^0 + \alpha \tilde{w}](s)$ and $\mathcal{S}_{t_{\max}}^* [v^0](s)$ are respectively equal to $v_\alpha(t_{\max} + s)$ and $v(t_{\max} + s)$. We are going to show that $v_\alpha(t_{\max} + s) \leq v(t_{\max} + s) + \alpha(\nu + \varepsilon)w$ by finite induction over s . First suppose that all the $(t^a)_{a \in A}$ are positive and let $K := \lfloor t_{\max}/t_{\min} \rfloor$.

Let $s \in [-t_{\max}, -t_{\max} + t_{\min})$ (initialization of the finite induction). In this case, we have using the optimal dynamics that for all i in $S \setminus \{0\}$

$$v_\alpha(i, t_{\max} + s) = \min_{a \in A_i^*} \left\{ c^a + \sum_{j \in S} p_j^a v_\alpha(j, t_{\max} + s - t^a) \right\} = \min_{a \in A_i^*} \left\{ c^a + \sum_{j \in S} p_j^a (v^0(j, t_{\max} + s - t^a) + \alpha w(j)) \right\},$$

where we have used the fact that $t_{\max} + s - t^a < 0$ to make the initial condition $v^0 + \alpha \tilde{w}$ appear. We then observe that for all a in A_i^* , it holds that $\sum_{j \in S} p_j^a w(j) \leq \max_{\sigma \in \mathfrak{S}^*} [P^\sigma|_{S \setminus \{0\} \times S \setminus \{0\}} w]_i \leq (\nu + \varepsilon)w(i)$. The remaining term $\min_{a \in A_i^*} \{c^a + \sum_{j \in S} p_j^a v^0(j, t_{\max} + s - t^a)\}$ coincides with $v(i, t)$. Hence, we indeed have $v_\alpha(t_{\max} + s) \leq v(t_{\max} + s) + \alpha(\nu + \varepsilon)w$ for all $s \in [-t_{\max}, -t_{\max} + t_{\min})$.

Let us now assume that $v_\alpha(t_{\max} + s) \leq v(t_{\max} + s) + \alpha(\nu + \varepsilon)w$ holds for all $s \in [-t_{\max}, -t_{\max} + kt_{\min})$ where k is an integer such that $1 \leq k \leq K - 1$. We choose s fixed in $[-t_{\max} + kt_{\min}, -t_{\max} + (k+1)t_{\min})$

(and $s < 0$). We can still write thanks to the optimal dynamics that for all i in $S \setminus \{0\}$

$$v_\alpha(i, t_{\max} + s) = \min_{a \in A_i^*} \left\{ c^a + \sum_{j \in S} p_j^a v_\alpha(j, t_{\max} + s - t^a) \right\}. \quad (4.21)$$

We distinguish to sub-cases. First, if $a \in A_i^*$ is such that $t^a > t_{\max} + s$ (recall that $t_{\max} + s$ is always nonnegative), then the term $v_\alpha(j, t_{\max} + s - t^a)$ in (4.21) can be replaced by $v^0(j, t_{\max} + s - t^a) + \alpha w(j) = v(j, t_{\max} + s - t^a) + \alpha w(j)$. Otherwise, if $a \in A_i^*$ is such that $t^a \leq t_{\max} + s$, then we use the fact that $s \leq -t_{\max} + (k+1)t_{\min}$ and $t^a \geq t_{\min}$ to obtain that $s - t^a$ lies in $[-t_{\max}, -t_{\max} + kt_{\min}]$. Our induction hypothesis guarantees that the term $v_\alpha(j, t_{\max} + s - t^a)$ in (4.21) can be bounded above by $v(j, t_{\max} + s - t^a) + \alpha(v + \varepsilon)w(j)$ or further by $v(j, t_{\max} + s - t^a) + \alpha w(j)$ (recall that $v + \varepsilon < 1$). In both cases, we have

$$\begin{aligned} c^a + \sum_{j \in S} p_j^a v_\alpha(j, t_{\max} + s - t^a) &\leq c^a + \sum_{j \in S} p_j^a (v(j, t_{\max} + s - t^a) + \alpha w(j)) \\ &\leq c^a + \sum_{j \in S} p_j^a v(j, t_{\max} + s - t^a) + \alpha(v + \varepsilon)w(i), \end{aligned}$$

and we get $v_\alpha(i, t_{\max} + s) \leq v(i, t_{\max} + s) + \alpha(v + \varepsilon)w(i)$, regardless of s in $[-t_{\max} + kt_{\min}, -t_{\max} + (k+1)t_{\min}]$, which completes the induction scheme. If Zeno phenomena were at stake, we would have use the induction principle (\mathcal{I}_Z), first handling at each step the necessarily non-Zeno states and then the Zeno-states, recovering the above inequality in the same way as the proof of Proposition 3.26.

◇ Now that we have $\mathcal{S}_{t_{\max}}^* [v^0 + \alpha \tilde{w}] \leq \mathcal{S}_{t_{\max}}^* [v^0] + \alpha(v + \varepsilon)\tilde{w}$ (regardless of v^0 and α), let us recall how to derive the contraction result. We choose two functions v_1 and v_2 in \mathcal{V}_0 , and define

$$\|v_1 - v_2\|_{\infty, w} := \sup_{\substack{j \in S \setminus \{0\} \\ s \in [-t_{\max}, 0]}} \left| \frac{v_1(j, s)}{w(j)} - \frac{v_2(j, s)}{w(j)} \right| \quad \text{so that} \quad v_1 \leq v_2 + \|v_1 - v_2\|_{\infty, w} \tilde{w},$$

where the last inequality stands in \mathcal{V}_0 . Applying the operator $\mathcal{S}_{t_{\max}}^*$ that is order-preserving (as shown in Proposition 3.26-(2)), we obtain

$$\mathcal{S}_{t_{\max}}^* [v_1] \leq \mathcal{S}_{t_{\max}}^* [v_2] + (v + \varepsilon)\|v_1 - v_2\|_{\infty, w} \tilde{w}.$$

The converse inequality (with v_1 and v_2 swapped) holds as well, and it follows that regardless of v_1 and v_2 , we have

$$\|\mathcal{S}_{t_{\max}}^* [v_1] - \mathcal{S}_{t_{\max}}^* [v_2]\|_{\infty, w} \leq (v + \varepsilon)\|v_1 - v_2\|_{\infty, w},$$

which is the contraction result claimed. \clubsuit

We point out that this proposition is reminiscent of the different phases of convergence of the deviation $v(t) - \chi(\mathcal{S})t$ in the case of MDPs when this quantity admits a limit, established by Schweitzer and Federgruen in their work [SF79]. They showed in particular that after elimination of non-optimal actions, the convergence is geometric. It is worth noting that in Theorem 4.12, we could only control the contraction rate of $\mathcal{S}_{t_{\max}}^*$, and not \mathcal{S}_t^* for $t \leq t_{\max}$. This is characteristic of the semi-Markov setting, since in the worst case we must wait for a time t_{\max} until the full dynamics switch on.

We now come back to the question of the finite time convergence of the value iteration. The next theorem tells us that it cannot be expected unless all probability matrices associated with optimal policies restricted to states of $S \setminus \{0\}$ are nilpotent, hence making the rate of geometric convergence ν featured in Theorem 4.12 null and effectively calling off the last phase. This nilpotency condition is also equivalent to some restrictions on the SMDP topology, that we may interpret as requiring a form of **hierarchy** within the set of states. It must be compatible with the moves made by the optimal policies; these should always “descend” in the hierarchy until finally reaching 0, the minimal element. To our knowledge, these characterizations of finite-time convergence of the value iteration in SSP configurations are new even in the simpler MDP setting.

► **THEOREM 4.13.** *Suppose Assumptions A and C hold. Then, the following are equivalent:*

- (i) *for all v^0 in \mathcal{V}_0 and associated solution v of the dynamics (DP), there exists t^* in \mathbb{R} such that for all $t \geq t^*$, $v(t) = u^*$;*
- (ii) *for all optimal policies σ in \mathfrak{S}^* , we have $\rho(P^\sigma|_{(S \setminus \{0\}) \times (S \setminus \{0\})}) = 0$;*
- (iii) *there exists a partial ordering ($<$) on S such that for all σ in \mathfrak{S}^* and for all i in $S \setminus \{0\}$, $\{j \in S \mid p_j^{\sigma(i)} > 0\} \subset \{j \in S, j < i\}$.*

Proof. We will prove (i) \implies (ii) \implies (iii) \implies (i), but we begin by stating a first-order decomposition lemma.

- ◊ Let us reuse the positive scalar $\varepsilon^* := \min_{i \in S} \min_{a \in A_i \setminus A_i^*} \{c^a + \sum_{j \in S} p_j^a u^*(j) - u^*(i)\} > 0$ introduced in the proof of Theorem 4.12, and let $\alpha := \varepsilon^*/3$. We claim that

$$\forall h^0 \in \mathcal{V}_0, \quad \|h^0\|_\infty \leq \alpha \implies \forall t \geq 0, \quad \mathcal{S}_t[\tilde{u}^* + h^0] = \mathcal{S}_t[\tilde{u}^*] + \mathcal{S}'_t[h^0] = \tilde{u}^* + \mathcal{S}'_t[h^0], \quad (4.22)$$

where $(\mathcal{S}'_t)_{t \geq 0}$ denotes the *semidifferentiate* semigroup of $(\mathcal{S}_t)_{t \geq 0}$ at point \tilde{u}^* (constant function of \mathcal{V}_0 equal to u^*), defined as in the Definition 3.23 but according to the dynamical equations

$$\forall i \in S, \quad h(i, t) = \min_{a \in A_i^*} \left\{ \sum_{j \in S} p_j^a h(j, t - t^a) \right\}. \quad (4.23)$$

To establish (4.22), it suffices to apply the induction principle (\mathcal{I}_Z) to work step-by-step with the dynamic programming equations, to use the specific value of α to show that only optimal actions can achieve the minimum (as in Theorem 4.12), and finally the fixed-point property satisfied by \tilde{u}^* (see Lemma 3.39). Similarly to \mathcal{S} , the semidifferentiate semigroup \mathcal{S}' is a nonlinear self-map of \mathcal{V}_0 , however it is positively homogeneous (but still non-linear). Note that (4.22) is an asymptotic expansion that is actually exact owing to the fact that the dynamics (DP) are piecewise-affine.

- ◊ We now show (i) \implies (ii). Because (i) holds, we know that convergence of the dynamics towards u^* for the specific initial condition $v^0(s) := u^* - \alpha \mathbb{1}_{S \setminus \{0\}}$ ($s \in [-t_{\max}, 0)$) occurs in finite time, so we have some $t_0 \geq 0$ such that $\mathcal{S}_{t_0}[v^0] = \tilde{u}^*$. According to (4.22), we have $\mathcal{S}'_{t_0}[r \mapsto -\alpha \mathbb{1}_{S \setminus \{0\}}] = \tilde{\mathbf{0}}$, or even $\mathcal{S}'_{t_0}[r \mapsto -\mathbb{1}_{S \setminus \{0\}}] = \tilde{\mathbf{0}}$ via positive homogeneity of \mathcal{S}' .

Denote by h the solution of the optimal semidifferentiate dynamics (4.23) uniquely determined by the initial condition $s \mapsto -\mathbb{1}_{S \setminus \{0\}}$, take an optimal policy $\sigma \in \mathfrak{S}^*$ and let us show that for all $t \geq 0$, we have $h(t) \leq (P^\sigma)^{\lceil t/t_{\min} \rceil}(-\mathbb{1}_{S \setminus \{0\}})$, by induction on t supposing that all the sojourn times are positive (otherwise we would still use the induction principle (\mathcal{I}_Z)). If $t \in [0, t_{\min})$, the result is direct by using the dynamics (4.23) restricted to the policy σ and the initial condition. Suppose the result holds for $t \in [0, kt_{\min})$ for some $k \in \mathbb{N}^*$ and take $t \in [kt_{\min}, (k+1)t_{\min})$. If i is in $S \setminus \{0\}$, we have by (4.23) that $h(i, t) \leq \sum_{j \in S} p_j^{\sigma(i)} h(j, t - t^{\sigma(i)})$. However, $t - t^{\sigma(i)}$ lies in $[-t_{\max}, kt_{\min})$, hence we can bound above the term $h(j, t - t^{\sigma(i)})$ by $[(P^\sigma)^k(-\mathbb{1}_{S \setminus \{0\}})]_j$ (either using the induction assumption if $t - t^{\sigma(i)} \geq 0$ or the initial condition otherwise). It follows that $h(i, t) \leq [(P^\sigma)^{k+1}(-\mathbb{1}_{S \setminus \{0\}})]_i$ and the induction scheme is completed.

In particular, for all optimal policies $\sigma \in \mathfrak{S}^*$, for all $t \geq 0$ and $s \in [-t_{\max}, 0)$ such that $t + s > 0$, we have $\mathcal{S}'_t[-\mathbb{1}_{S \setminus \{0\}}](s) \leq (P^\sigma)^{\lceil \frac{t+s}{t_{\min}} \rceil}(-\mathbb{1}_{S \setminus \{0\}})$. Applying this result to $t = t_0$, we know that there exists an integer $k_0 \in \mathbb{N}^*$ such that $\mathbf{0} = \mathcal{S}'_{t_0}[-\mathbb{1}_{S \setminus \{0\}}](-t_{\max}) \leq (P^\sigma)^{k_0}(-\mathbb{1}_{S \setminus \{0\}}) \leq \mathbf{0}$ (the last inequality has been obtained by order-preserving character of stochastic matrices). We deduce that $(P^\sigma)^{k_0} \mathbb{1}_{S \setminus \{0\}} = \mathbf{0}$, or equivalently forgetting about the component associated with state 0 that $(P^\sigma|_{S \setminus \{0\} \times S \setminus \{0\}})^{k_0} \mathbf{1} = \mathbf{0}$. Still by the order-preserving character of the linear map associated with this matrix (to show that $(P^\sigma|_{S \setminus \{0\} \times S \setminus \{0\}})^{k_0}$ maps some open ball to $\mathbf{0}$), it entails that $P^\sigma|_{S \setminus \{0\} \times S \setminus \{0\}}$ is nilpotent or equivalently $\rho(P^\sigma|_{S \setminus \{0\} \times S \setminus \{0\}}) = 0$. This holds for any $\sigma \in \mathfrak{S}^*$, which is our item (ii).

- ◊ Let us prove the implication (ii) \implies (iii). Observe to this purpose that we can interpret the assumption $\max_{\sigma \in \mathfrak{S}^*} \rho(P^\sigma|_{S \setminus \{0\} \times S \setminus \{0\}}) = 0$ by the fact that the (generalized) spectral radius of the application T^* defined in the proof of Theorem 4.12 is zero as well, which in terms of the Collatz–Wielandt property gives us $\inf \{\ell > 0 \mid \exists x \in \mathbb{R}_{>0}^{S \setminus \{0\}}, T^*(x) \leq \ell x\} = 0$. As a result we know that for all $\ell > 0$, there is some

$w_\ell \in \mathbb{R}_{>0}^{S \setminus \{0\}}$ such that $T^*(w_\ell) \leq \ell w_\ell$, or equivalently $P^\sigma|_{S \setminus \{0\} \times S \setminus \{0\}} w_\ell \leq \ell w_\ell$ for all $\sigma \in \mathfrak{S}^*$. Introducing $\bar{P} := (1/|\mathfrak{S}^*|) \sum_{\sigma \in \mathfrak{S}^*} P^\sigma|_{S \setminus \{0\} \times S \setminus \{0\}}$, we naturally still have $\bar{P} w_\ell \leq \ell w_\ell$. Because such an inequality can be written for all $\ell > 0$, we can reuse the Collatz–Wielandt characterization of the spectral radius and conclude that the (linear) operator \bar{P} is also nilpotent. This entails that the accessibility graph \bar{G} associated with \bar{P} (with nodes set included in $S \setminus \{0\}$) does not contain any circuit. We can thus associate a partial order ($<$) with this graph, such that each node $i \in S \setminus \{0\}$ has only access to nodes $\{j \in S, j < i\}$. However by definition of \bar{P} (that can be thought of as a randomization of all policies), the accessibility graph G^σ of $P^\sigma|_{S \setminus \{0\} \times S \setminus \{0\}}$ is included in \bar{G} for all $\sigma \in \mathfrak{S}^*$. It follows that for all optimal policies and states $i \in S \setminus \{0\}$, we have $\{j \in S \mid p_j^{\sigma(i)} > 0\} \subset \{j \in S, j < i\}$, that is to say (iii).

- ◇ To finish, we focus on (iii) \implies (i), that is easier to prove. Let $v^0 \in \mathcal{V}_0$. Recall from Theorem 4.12 that there is some t^* from which only optimal actions are used. According to the item (iii), each time an optimal action is pulled from a state in $S \setminus \{0\}$, the player goes to a state strictly inferior in the hierarchy, as a result after $t^* + |S|t_{\max}$ we know that the state 0 is reached, and that the value function has stabilized to u^* . Remark that the time t^* needed to approach \tilde{u}^* from v^0 to a precision α is in fact only controlled by $\|v^0 - \tilde{u}^*\|_\infty$. \clubsuit

Observe that the condition in item (ii) cannot be strengthened to require that all proper policies induce nilpotent transition probability matrices on $S \setminus \{0\}$, as shown in the next example.

Example 4.14. Consider the simple SMDP of Figure 4.12 that is actually an MDP since all the sojourn times are taken equal to one.

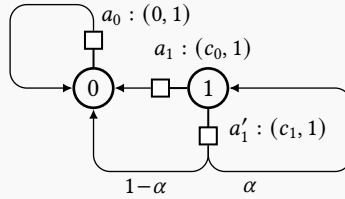


Figure 4.12: A 2-state SMDP with two proper policies but only one optimal policy

The state 0 is cost-free and absorbing, and from the state 1, the player can either choose to pay the cost c_0 and go to state 0 (hence the game stops), or to pay the cost c_1 and come back to 1 with the probability $0 < \alpha < 1$. The fact that $\alpha < 1$ guarantees that both policies (a_0, a_1) and (a_0, a'_1) are proper, and we easily compute that starting from state 1, they respectively induce total costs equal to c_0 and $c_1/(1 - \alpha)$.

Let us suppose that $c_0 < c_1/(1 - \alpha)$. In this MDP context, the value iteration procedure consists in computing the sequence $(v(1, t))_{t \in \mathbb{N}}$ defined by $v(1, t) = c_0 \wedge (c_1 + \alpha v(1, t - 1))$ for all t in \mathbb{N} and the initial condition $v(1, t - 1) = M$ with M in \mathbb{R} . Clearly, if for some t_0 in \mathbb{N} the action a_1 becomes minimizing, it remains minimizing for all $t \geq t_0$ and the convergence of the value towards c_0 happens in finite time. Suppose for contradiction that a_1 is never minimizing, and therefore that $v(1, t) = c_1 + \alpha v(1, t - 1)$ for all t in \mathbb{N} . This is a classic arithmetico-geometric sequence, and we obtain $v(1, t) = M\alpha^{t+1} + c_1(1 - \alpha^{t+1})/(1 - \alpha)$. But with $\delta = (c_1/(1 - \alpha) - c_0)/2$, and $t^* := (\log(\delta) - \log(|M - c_1/(1 - \alpha)|))/\log(\alpha)$, we get $v(1, t) > c_0$ for $t + 1 \geq t^*$, which contradicts the dynamic programming equation. Hence a_0 becomes minimizing for some $t \leq t^*$ and convergence towards the SSP value always happens in finite time for this Markov decision process.

Since the matrix $P^\sigma|_{S \setminus \{0\} \times S \setminus \{0\}}$ associated with the (proper but not optimal) policy (a_0, a'_1) reduces to the scalar $\alpha > 0$, we see that it is not nilpotent. Hence, although it is sufficient, it is not necessary that all proper policies give nilpotent transition matrices over $S \setminus \{0\}$ to obtain a finite-time convergence of the value iteration regardless of the initial condition.

5.2 HIERARCHICAL SSP CONFIGURATIONS. — According to the Theorem 4.12, the set \mathfrak{S}^* of optimal policies controls the ultimate rate of convergence. It is in general not known, and depends on the costs and sojourn times attached to the actions of the SMDP – as illustrated for instance in Example 4.14. It can be desirable to identify conditions of a more topological nature under which finite time convergence occurs regardless of the optimal character of some policies, that however need to remain proper. Building on the statement of Theorem 4.13–(iii), we choose to enforce the existence of a partial ordering (\leq) on S such that applying some actions necessarily makes the order strictly decrease. This is formalized in the next assumption, where whenever a is an action of A , we denote its **support** $\{i \in S \mid p_i^a > 0\}$ by $\text{supp}(a)$.

► **ASSUMPTION E.** *There is a partial ordering (\leq) on S , such that for all states i in S , the set A_i of playable actions from i admits a partition of the form $A_i = A_i^- \uplus A_i^+$, with the condition that*

$$\begin{aligned} \text{if } a \in A_i^-, \text{ then } p_j^a > 0 &\implies j < i \quad (\text{equivalently } \text{supp}(a) \subset \{j < i\}) \quad \text{and} \\ \text{if } a \in A_i^+, \text{ then } p_j^a > 0 &\implies i \leq j \quad (\text{equivalently } \text{supp}(a) \subset \{i \leq j\}), \end{aligned}$$

where $j < i$ means $j \leq i$ and $j \neq i$. In addition, $A_i^- \neq \emptyset$ for all i in $S \setminus \{0\}$.

In other words, Assumption E requires that actions either strictly “descend” or weakly “ascend” relatively to the states hierarchy, the first case being always possible. In what follows, we shall refer to a in A_i^- (resp. a in A_i^+) as a “descending action” (resp. an “ascending action”). Remark that imposing in Assumption E that $A_i^- \neq \emptyset$ for all i in $S \setminus \{0\}$ implies the Assumption C(1), since any policy which only makes use of descending actions is proper. This implication turns into an equivalence (i.e., any proper policy must only make use of descending actions) if the next condition is also met.

► **ASSUMPTION F.** *For all σ in \mathfrak{S} , if there exists i in S such that $\sigma(i)$ is in A_i^+ , then σ is improper.*

Under Assumptions E and F (stronger than the condition (iii) of Theorem 4.13), we obtain that $A_i^* \subset A_i^-$ for all i in S . It implies that optimal policies, that are proper, must also use only strictly descending actions, and according to Theorem 4.13, convergence towards u^* arises in finite time.

*
* *

In the case of the medical emergency call center first depicted on Figure 4.5 and reproduced in Figure 4.13a, the two hierarchical assumptions E and F are satisfied with the partial ordering (\leq) such that $0 \leq 1 \leq 2$ and $0 \leq 1 \leq 3 \leq 4 \leq 5$ (the other pairs of states do not need to be comparable). Observe that this states hierarchy is the one naturally given by the (partial) order in which the different tasks are performed in the Petri net of Figure 4.5. We verify that (strictly) descending actions (labeled by means of the “−” symbol) can be played for all states in $S \setminus \{0\}$. There are only three ascending actions here (labeled by means of the “+” symbol), the action from 0 that makes the decision-maker stay in 0, the action from 1 that sends him/her stochastically to 2 or 4, and the action from 3 that forces to evolve in 5. Consistently with the discussion based on the Table 4.9, the policy that picks the descending actions in A_1 and A_3 is the only proper policy (and thus it is optimal).

Observe how a slight modification of the SMDP topology can make the Assumption F violated. Indeed, in the Figure 4.13b, the policy that picks the ascending action from state 1 and the descending action from state 3 becomes proper (the player will necessarily end up in state 0 by applying this policy starting from any state). It may be optimal for suitable choices of costs and sojourn times as in Example 4.14, but convergence towards u^* will happen geometrically.

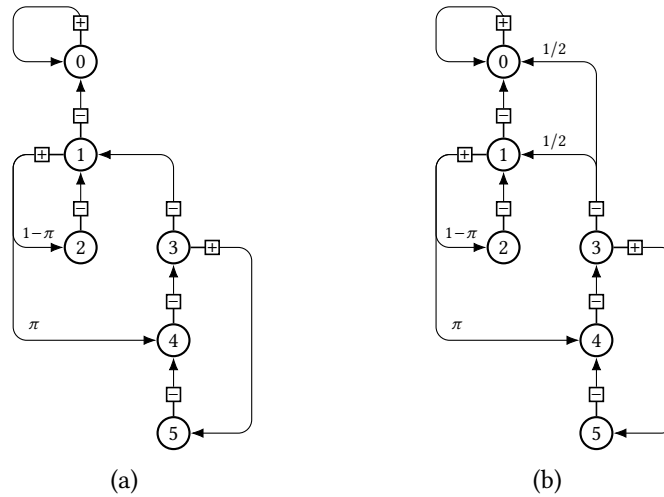


Figure 4.13: Two SMDPs in SSP configuration verifying Assumption E, but only the first one verifies Assumption F

5.3 CONVERGENCE BY ABOVE. — Now that we have exhibited conditions under which the convergence of the dynamics of an SMDP in SSP configuration towards its value u^* occurs in finite time, we are interested in finding an upper bound of the (state-by-state) convergence times $(\theta_i)_{i \in S}$ defined by

$$\text{for all } i \text{ in } S, \quad \theta_i := \inf\{t \geq 0 \mid v(i, t) = u^*(i)\}, \quad (4.24)$$

where v is the solution of (DP) determined by an initial condition v^0 in \mathcal{V}_0 . Although we shall not write it explicitly, the $(\theta_i)_{i \in S}$ depend on v^0 .

A somehow easy case to analyze occurs when the initial condition v^0 is **above** the limit u^* , i.e., $v^0(s) \geq u^*$ for all s in $[-t_{\max}, 0)$. Recall from Section 2.3 of Chapter 3 that $v(i, t)$ gives the minimum expected cost by starting the game in state i and playing during t units of time; knowing that terminating the game in state i with a horizon “overshoot” of s costs $v^0(i, -s) \geq u^*(i)$ (and this terminal cost is null in state 0 because v^0 is in \mathcal{V}_0). In that case, the key is that terminating in state $j \neq 0$ starting from state i is more costly than rushing directly towards 0 (the cost of travelling from i to j plus incurring at least $u^*(j)$ at the end exceeds the cost of going from i to 0 and stop playing there). As a result, each state has interest to join state 0 using an optimal policy as soon as possible.

Bertsekas has already observed this fact for Markov decision processes (SMDPs with sojourn times equal to one) in SSP configuration in his study of the value iteration algorithm [Ber11, Section 3.4.1] when the transition probability graph of some optimal policy is acyclic, taking in particular an initial condition with $v^0(i) = \infty$ for all i in $S \setminus \{0\}$. In the same spirit, we are going to take advantage of the states hierarchy to propose an inductive scheme to bound the $(\theta_i)_{i \in S}$. Indeed, the state 0 is the bottom element of the order (\leq), and the states that lie in low layers of the hierarchy shall have “quicker” access to 0 than those located higher in the hierarchy (since the latter may need to pass through the former). From this perspective, if i is a state of S , it is natural to focus on the convergence time θ_i after the states located lower than i in the hierarchy already converged themselves. It amounts to determining the $(\theta_i)_{i \in S}$ by following an inductive scheme given by the partial ordering (\leq).

- **PROPOSITION 4.15.** *Suppose that the SMDP satisfies the hierarchical assumptions E and F. Suppose in addition that $v^0(s) \geq u^*$ for all s in $[-t_{\max}, 0)$. Then we have $\theta_0 = 0$ and for all i in $S \setminus \{0\}$, θ_i consistently verifies by induction:*

$$\theta_i \leq \min_{a \in A_i^-} \left\{ t^a + \max_{j \in \text{supp}(a)} \theta_j \right\}.$$

Proof. We first deal with a particular case where v^0 is of the form $v^0(s) = u^* + M \mathbb{1}_{S \setminus \{0\}}$ for all s in $[-t_{\max}, 0)$ and some $M \geq 0$ (so that we indeed have $v^0 \geq \tilde{u}^*$). We still denote by v the function determined by this particular initial condition, and as a first step we show below that v is nonincreasing. Indeed, for $t \in [0, t_{\min})$ and $s \in [-t_{\max}, 0)$, we have (supposing for the sake of simplicity that all the sojourn times are positive) for all $i \in S \setminus \{0\}$ that

$$\begin{aligned} v(i, t+s) = \mathcal{S}_t[v^0](s)_i &= \min_{a \in A_i} \left\{ c^a + \sum_{j \in S} p_j^a v^0(j, t+s-t^a) \right\} \\ &= \min_{a \in A_i} \left\{ c^a + \sum_{j \in S} p_j^a (u^*(j) + M \mathbb{1}_{j \neq 0}) \right\} \\ &\leq u^*(i) + M \end{aligned}$$

and in addition $v(0, t+s) = 0$. Therefore for all $t \in [0, t_{\min})$, we have $\tilde{u}^* \leq \mathcal{S}_t[v^0] \leq v^0$. This double inequality can be propagated for all $t \geq 0$ by applying the induction principle (\mathcal{I}_Z); which we also resort to in the case where some sojourn times are null, still with similar schemes than those of Chapter 3. We can now apply the Lemma 3.40 whose proof guarantees that v is nonincreasing, bounded below by u^* and converges towards u^* . In particular, we necessarily have $t \geq \theta_i \implies v(i, t) = u^*(i)$ for all $i \in S$, meaning that as soon as $v(i, t)$ has reached once its final value $u^*(i)$, it stays constant.

- ◇ We show the result of the theorem by proving by induction that the property $\mathcal{P}(i)$: “the bound on θ_i in the theorem is valid” holds for all states i in S . To this purpose, it suffices to show that $\mathcal{P}(0)$ is true (initialization), and that if $\mathcal{P}(j)$ holds for all $j < i$, then $\mathcal{P}(i)$ holds as well (heredity). Indeed, although not all states are comparable relatively to the partial ordering (\leq), the state 0 is a minimal element and any state can reach 0 using only descending actions (following from Assumption E). Finally, all states are comparable within such a path.

It is immediate to obtain that $\mathcal{P}(0)$ is true, i.e. $\theta_0 = 0$, since for all $t \geq 0$, we have $v(0, t) = 0$ (recall that \mathcal{V}_0 is stable under evolution equations associated with SSP configurations).

In what follows, we let $i \in S \setminus \{0\}$ and suppose that $\mathcal{P}(j)$ is true for all $j \in S$ such that $j < i$. Cases are distinguished based on the existence or absence of ascending actions from i .

- ◇ Let $i \in S \setminus \{0\}$ such that $A_i^+ = \emptyset$, so that by (DP), for all $t \geq 0$,

$$v(i, t) = \min_{a \in A_i^-} \left\{ c^a + \sum_{j \in S} p_j^a v(j, t-t^a) \right\}.$$

Let us take $t = \max_{a \in A_i^-} \{t^a + \max_{j \in \text{supp}(a)} \theta_j\}$. For all a in A_i^- , we have for all j in $\text{supp}(a)$ that $v(j, t-t^a) = u^*(j)$, by the monotonic character of v shown in the beginning of the demonstration. As a result, we have

$$v(i, t) = \min_{a \in A_i^-} \left\{ c^a + \sum_{j \in S} p_j^a u^*(j) \right\}$$

where we recognize the right-hand side of (3.29), which provides $v(i, t) = u^*(i)$. As a result we must have by minimality $\theta_i \leq \max_{a \in A_i^-} \{t^a + \max_{j \in \text{supp}(a)} \theta_j\}$.

- ◇ Let us now take $i \in S \setminus \{0\}$ such that $A_i^+ \neq \emptyset$. The (DP) equation here gives us for all $t \geq 0$ and $i \in S \setminus \{0\}$ that

$$v(i, t) = \min \left(\min_{a \in A_i^-} \left\{ c^a + \sum_{j \in S} p_j^a v(j, t-t^a) \right\}, \min_{a \in A_i^+} \left\{ c^a + \sum_{j \in S} p_j^a v(j, t-t^a) \right\} \right). \quad (4.25)$$

Denote for short $\Theta_i := \max_{a \in A_i^-} \{t^a + \max_{j \in \text{supp}(a)} \theta_j\}$, and choose $t \geq \Theta_i$. Just like the previous part of the proof, we know that $v(j, t-t^a) = u^*(j)$ for all $a \in A_i^-$ and $j \in \text{supp}(a)$. In addition, by monotonicity of v again and the fact that $v(t) \geq u^*$ for all $t \geq -t_{\max}$, we have $v(j, t-t^a) \geq u^*(j)$ if $j \in S$ and $a \in A_i^+$.

Hence, (4.25) becomes

$$v(i, t) = \min \left(\underbrace{\min_{a \in A_i^-} \left\{ c^a + \sum_{j \in S} p_j^a u^*(i) \right\}}_{= u^*(i)}, \underbrace{\min_{a \in A_i^+} \left\{ c^a + \sum_{j \in S} p_j^a v(j, t - t^a) \right\}}_{\geq u^*(i)} \right). \quad (4.26)$$

As a result, if $t \geq \Theta_i$, the first term of the minimum in (4.26) is active and we get $v(i, t) = u^*(i)$, which proves that $\theta_i \leq \Theta_i$, and achieves the induction scheme on the states.

- ◇ To finish, let us come back to the general case where v^0 is any bounded function of \mathcal{V}_0 such that $v^0 \geq u^*$. Denoting $M := \sup_{s \in [-t_{\max}, 0)} \{v^0(s) - u^*\}$, we obtain that $u^* \leq v^0(s) \leq u^* + M \mathbb{1}_{S \setminus \{0\}}$ with $M \geq 0$. This double inequality extends to the dynamics continuation of these three initial conditions, and when the i -th coordinate of $t \mapsto \mathcal{S}_{t+\varepsilon}[r \mapsto u^* + M \mathbb{1}_{S \setminus \{0\}}](-\varepsilon)$ has reached its final value $u^*(i)$, so has the i -th coordinate of $t \mapsto \mathcal{S}_{t+\varepsilon}[v^0](-\varepsilon)$, hence the conclusion. Remark that the bounds on the $(\theta_i)_{i \in S}$ actually do not depend upon M . \clubsuit

The following corollary that focuses on the total convergence time $t^* := \max_{i \in S} \theta_i$ is a coarser version of Proposition 4.15, but it highlights the tree structure of hierarchical configurations by featuring its depth d . Its proof is immediate.

- **COROLLARY 4.16.** *Under the conditions of Proposition 4.15, denote by d the maximal length of a descending path in S relatively to the ordering (\leq) . Then, finite-convergence of v towards u^* occurs in a time t^* such that*

$$t^* \leq d \times t_{\max}.$$

This claim is consistent with the result of Bertsekas who showed still in [Ber11, Section 3.4.1] that the value u^* could be computed in $|S|$ steps of the value iteration algorithm for MDPs by starting with a positive infinite initial condition (since in this case $d \leq |S|$ and $t_{\max} = 1$).

*
* *

The top chart of the Figure 4.14 illustrates the Proposition 4.15 for the SMDP in SSP configuration associated with our running example (EMS-A), for quite an eccentric choice of initial condition but that respects the requirement of being above u^* . We can verify that after the time $t^* = \tau_1 + \tau_2 + \tau_3 = 370$ s, corresponding with the time-length of the longest branch of the system, the final regime (constant equal to u^*) is reached. During this transient regime, only the single proper policy $(\lambda, 0)$ of the SMDP has been played.

- 5.4 CONVERGENCE BY BELOW.** — Bounding the $(\theta_i)_{i \in S}$ in the converse case where the initial condition is **below** the limit, (meaning that $v^0(s) \leq u^*$ for all s in $[-t_{\max}, 0)$) is a different kettle of fish. Indeed, joining state 0 from state i as soon as it becomes possible (and thus accumulating a total cost of $u^*(i)$) may not be the best choice, since terminating in another state with low terminal cost can compensate the costly decision of not joining 0. For bigger horizons though, we know by Assumption C(2) that not leaving $S \setminus \{0\}$ results in accumulating positive costs, which will eventually mitigate the benefit of terminating in $S \setminus \{0\}$. For instance on the bottom graph of Figure 4.14, a simulation of (EMS-A) with an initial condition below u^* shows that we have no interest in pulling all the optimal actions until $t = 724.78$ s. It is only after this instant that not leaving $S \setminus \{0\}$ is always more costly than going towards 0. This figure confirms the intuition that convergence by below is typically more complex and takes more time to happen than convergence by above.

The possible adverse behaviour that we have just described prompts us to focus on the evolution of v if one decides not to join state 0 from some state. Let us consider the collection of

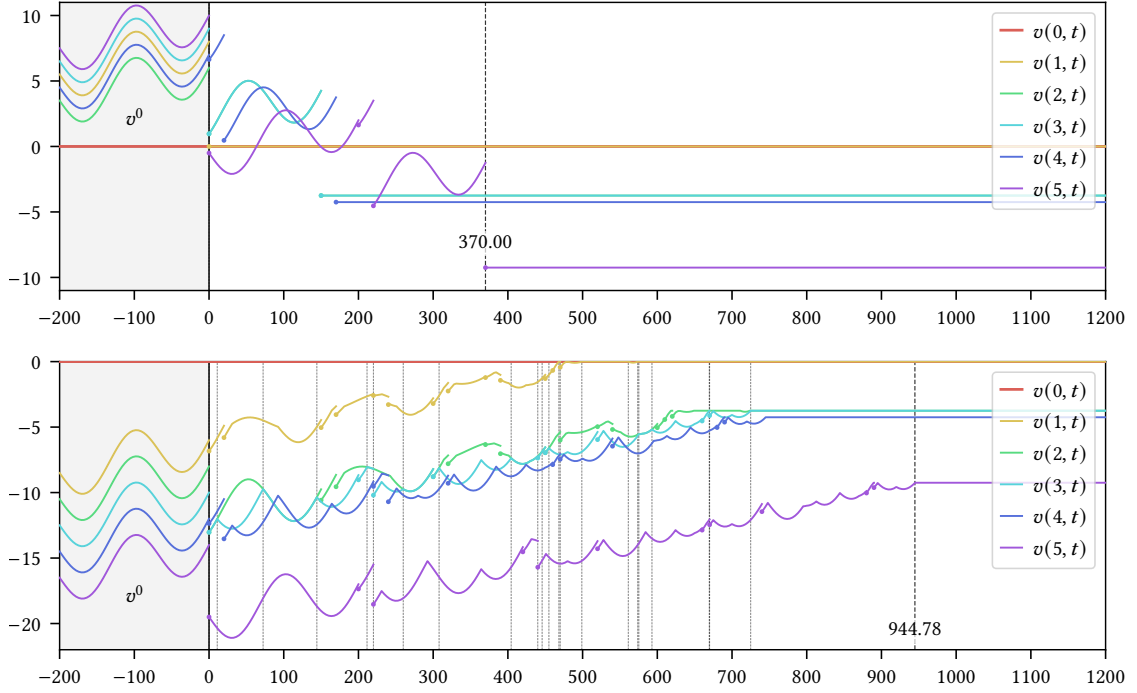


Figure 4.14: Convergence of the dynamics of the SMDP in SSP configuration associated with (EMS-A), for two different initial conditions

state-accessibility graphs of our SMDP in which a particular state cannot use descending actions. More precisely, given i in $S \setminus \{0\}$ such that $A_i^+ \neq \emptyset$, we denote by $\mathcal{G}_S^{(i)}$ the graph with nodes S and with edges set $\{(k, \ell) \in S^2 \mid \exists a \in A_k \setminus A_i^- \text{ s.t. } p_\ell^a > 0\}$. In other words, $\mathcal{G}_S^{(i)}$ describes the alternative SMDP in which i is forced to use actions in A_i^+ , all the descending actions playable from *other* states being still available. We denote by $C(i)$ the strongly connected component of i in the graph $\mathcal{G}_S^{(i)}$. The next lemma proves that evolving in $\mathcal{G}_S^{(i)}$, that is to say not selecting actions in A_i^- , amounts to staying within states of $C(i)$ with no way of joining 0.

► **LEMMA 4.17.** *Suppose that Assumptions C, E and F hold. For all i in $S \setminus \{0\}$ such that $A_i^+ \neq \emptyset$, the following properties hold:*

- (i) *0 is not accessible from i in $\mathcal{G}_S^{(i)}$,*
- (ii) *every state accessible from i in $\mathcal{G}_S^{(i)}$ also has access to i via descending actions.*

Proof. Remember that relatively to the graph \mathcal{G}_S with no removed actions, the state 0 is accessible from any other state, according to the Assumption E, and in addition, an accessibility walk can be achieved using only actions in $\bigcup_{i \in S} A_i^-$, the whole set of descending actions.

The subgraph $\mathcal{G}_S^{(i)}$ represents only improper policies by Assumption F. Thus, there is a state j that has not access to 0. This entails that any path connecting j to 0 in the original graph \mathcal{G}_S (possibly obtained by only descending actions) had to go through state i (the only state from which actions have been ruled out), otherwise these paths would still be available. It follows that i itself has not access to 0, otherwise j could reach 0 again, which proves (i). Property (ii) follows since if j is a state accessible from i in $\mathcal{G}_S^{(i)}$, then j has no longer access to 0 in $\mathcal{G}_S^{(i)}$, and the above argument shows that j has access to i . \clubsuit

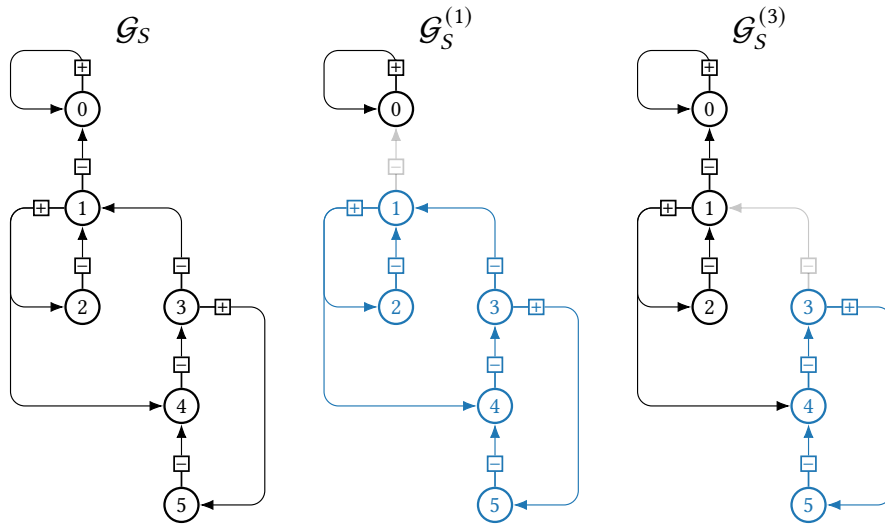


Figure 4.15: The SMDP of (EMS-A) in hierarchical SSP configuration, and the remaining sub-graphs when ruling out actions of A_1^- or A_3^- , with strongly connected components $C(1)$ and $C(3)$ outlined in blue.

We illustrate on the Figure 4.15 the construction of the graphs $(\mathcal{G}_S^{(i)})_{i \in S, A_i^+ \neq \emptyset}$ and the subsets $(C(i))_{i \in S, A_i^+ \neq \emptyset}$ for the SMDP associated with (EMS-A), recalled in Figure 4.13a. We let the reader check that the two items of Lemma 4.17 apply on these examples.

After having “topologically” described what happens to a decision-maker who decides to make use only of descending actions in a state i such that $A_i^+ \neq \emptyset$, we now specify what it implies cost-wisely. Let us first denote by $\underline{\chi}$ the minimum average-cost of all the final classes (except $\{0\}$) of improper policies, defined by

$$\underline{\chi} := \min_{\sigma \in \mathfrak{S}} \min_{\substack{F \in \mathcal{F}(\sigma) \\ F \neq \{0\}}} \left\{ \frac{\langle \mu_F^\sigma, c^\sigma \rangle}{\langle \mu_F^\sigma, t^\sigma \rangle} \right\}. \quad (4.27)$$

As showed by the proof of Correspondence Theorem 4, the Assumption C(2) is equivalent to $\underline{\chi} > 0$. In addition, observe that the pair $(C(i), A_{C(i)})$, with $A_{C(i)} := A_i^+ \uplus \biguplus_{j \in C(i) \setminus \{i\}} A_j$ is a consistent subset of states and actions in the sense of Remark 3.27. We thus introduce the evolution semigroup $(\mathcal{S}_t|_{C(i), A_{C(i)}})_{t \geq 0}$ portraying the associated reduced dynamics, denoted for short by $(\mathcal{S}_t^{(i)})_{t \geq 0}$. This evolution semigroup describes the costs collected by a player over time when he/she starts in some state of $C(i)$ and cannot play actions of A_i^- .

► **LEMMA 4.18.** *Suppose Assumptions E and F hold. Let i in S such that $A_i^+ \neq \emptyset$. Then, the growth rate of $\mathcal{S}^{(i)}$ (defined by (3.23)) is uniform on $C(i)$, i.e., of the form $\chi^{(i)} \mathbf{1}$, where $\chi^{(i)} \geq \underline{\chi}$.*

Proof. Let σ be a policy of

$$\mathfrak{S}_i := A_i^+ \times \prod_{j \in C(i) \setminus \{i\}} A_j$$

that realizes the minimal average-cost vector $\chi(\mathcal{S}^{(i)})$ of $\mathcal{S}^{(i)}$, and let F be a final class associated with σ included in $C(i)$ such that $\langle \mu_F^\sigma, c^\sigma \rangle / \langle \mu_F^\sigma, t^\sigma \rangle = \chi^{(i)} := \min_{j \in C(i)} \chi(\mathcal{S}^{(i)})_j$.

Observe that $\chi^{(i)} \geq \underline{\chi}$ by definition of the latter in (4.27). If for all $j \in C(i)$, we have $\phi_{F,j}^\sigma = 1$ (so that F is the only final class of σ in $C(i)$), we directly obtain $\chi(\mathcal{S}^{(i)}) = \chi^{(i)} \mathbf{1}$. Otherwise, since $C(i)$ is strongly connected, we may construct a policy $\sigma' \in \mathfrak{S}_i$ (identical to σ on F) such that σ' admits F as its

unique final class and as a consequence $\chi^{(i)}\mathbf{1}$ as optimal cost vector. By minimality of $\chi(\mathcal{S}^{(i)})$, we obtain $\chi(\mathcal{S}^{(i)}) \leq \chi^{(i)}\mathbf{1}$ and it follows that $\chi(\mathcal{S}^{(i)}) = \chi^{(i)}\mathbf{1}$. \clubsuit

We are now ready to compute quantitative upper bounds on the state-by-state convergence times $(\theta_i)_{i \in S}$ of the value iteration algorithm for SMDPs in SSP configuration, when the initial condition v^0 is below u^* . To our knowledge, this result is also new in the context of MDP in SSP configuration. The crux of the proof is to bound below the evolution of v (the trajectory determined by v^0 and the dynamic programming equations (DP)) by an affine regime of the dynamics associated with improper policies, characterized by a positive average-cost.

- **THEOREM 4.19.** *Suppose that the SMDP satisfies the hierarchical assumptions E and F. Suppose in addition that $v^0(s) \leq u^*$ for all s in $[-t_{\max}, 0)$ and denote $M := \sup_{s \in [-t_{\max}, 0)} \{u^* - v^0(s)\}$. Then we have $\theta_0 = 0$ and for all i in $S \setminus \{0\}$, θ_i consistently verifies by induction:*

$$\theta_i \leq \max_{a \in A_i^-} \{t^a + \max_{j \in \text{supp}(a)} \theta_j\} + \frac{M}{\chi^{(i)}} \mathbb{1}_{A_i^+ \neq \emptyset},$$

where for all i in S such that $A_i^+ \neq \emptyset$, $\chi^{(i)}\mathbf{1}$ denotes the growth rate vector of the evolution semigroup $\mathcal{S}^{(i)}$ introduced in Lemma 4.18.

Proof. The proof of this theorem greatly follows the lines of the one of Proposition 4.15.

- ◊ We first consider the particular case where $v^0(s) := u^* - M\mathbb{1}_{S \setminus \{0\}}$ for all $s \in [-t_{\max}, 0)$, for some given $M \geq 0$. By the same arguments than the proof of Proposition 4.15, this initial condition of \mathcal{V}_0 determines a solution v that is nondecreasing, bounded above by u^* and that converges towards u^* . For all $i \in S$, the property $t \geq \theta_i \implies v(i, t) = u^*(i)$ also survives.
- ◊ We implement the same induction scheme on the states than in the proof of Proposition 4.15, that is to say, we let $i \in S \setminus \{0\}$ and suppose that $\mathcal{P}(j)$ is true for all $j \in S$ such that $j < i$. Our goal is to prove that $\mathcal{P}(i)$ is true as well. If $A_i^+ = \emptyset$, we are in the “easy” case and $\mathcal{P}(i)$ is established exactly as in the proof of Proposition 4.15.
- ◊ The novelty (and the difficulty!) arises if $A_i^+ \neq \emptyset$. Indeed, we can apply the same techniques than the proof of Proposition 4.15 until the equation (4.26), which as a reminder claims that

$$\forall t \geq \Theta_i, \quad v(i, t) = \min \left(u^*(i), \min_{a \in A_i^+} \left\{ c^a + \sum_{j \in S} p_j^a v(j, t - t^a) \right\} \right)$$

where $\Theta_i := \max_{a \in A_i^-} \{t^a + \max_{j \in \text{supp}(a)} \theta_j\}$. This time, the second term in the minimum is no longer always greater than $u^*(i)$. However, still by monotonicity of v , we know that the minimum is attained by the rightmost term if $t < \theta_i$. Suppose then that $\theta_i > \Theta_i$ (the opposite leads to the same upper bound $\theta_i \leq \Theta_i$ as before).

For all t such that $\Theta_i \leq t < \theta_i$, only actions of A_i^+ are used from state i to determine $v(t)$ based on the values of v at prior times. In addition, it is necessary and sufficient to focus on the states of $C(i)$ to analyze the evolution of $v(i, t)$ for the particular state i , since by Lemma 4.17–(ii), this class is fully autonomous under the action of the dynamics when it is forbidden to play actions in A_i^- . In particular, reusing the semigroup $\mathcal{S}^{(i)}$ studied in the Lemma 4.18, we have

$$\forall s \in [-t_{\max}, 0), \quad v(t+s)|_{C(i)} = \mathcal{S}_{t-\Theta_i}^{(i)}[w](s). \quad (4.28)$$

where w is the function of $[-t_{\max}, 0)$ to $\mathbb{R}^{C(i)}$ such that $\forall s \in [-t_{\max}, 0)$, $w(s) := v(s + \Theta_i)|_{C(i)}$.

Invoking monotonicity of v , we can bound w below using our initial condition:

$$\forall s \in [-t_{\max}, 0), \quad w(s) \geq -M\mathbf{1} + u^*|_{C(i)} \quad (4.29)$$

According to Corollary 3.32–(1), the semigroup $\mathcal{S}^{(i)}$ admits an affine regime, with the throughput $\chi^{(i)}\mathbf{1}$ mentioned in Lemma 4.18. Let us denote by $h^{(i)}$ a bias vector of $\mathbb{R}^{C(i)}$ so that $t \mapsto \chi^{(i)}t\mathbf{1} + h^{(i)}$ is an affine regime of $\mathcal{S}^{(i)}$. Using the fact given by Lemma 4.18 and Assumption C(2) that $\chi^{(i)} > 0$, we can also bound

w below by this affine regime, up to shifting it by a constant. Formally, we can write for all s in $[-t_{\max}, 0)$:

$$w(s) \geq \chi^{(i)} s \mathbf{1} + h^{(i)} - \underbrace{\left(M + \max_{k \in C(i)} \{h_k^{(i)} - u^*(k)\} \right)}_{\alpha} \mathbf{1}$$

Unravelling the dynamics by applying the operator $\mathcal{S}_{t-\Theta_i}^{(i)}$ as done in (4.28) has just the effect of a shift in time of magnitude $t - \Theta_i$ for this last affine function, as explained just before Proposition 3.30 and owing to the additively homogeneous character of the evolution semi-group (see Proposition 3.26–(1)). Building on the order-preserving character of the latter (see Proposition 3.26–(2)), we obtain from (4.28) that for all s in $[-t_{\max}, 0)$:

$$v(t+s)|_{C(i)} \geq (t+s-\Theta_i)\chi^{(i)}\mathbf{1} + h^{(i)} - \alpha\mathbf{1}. \quad (4.30)$$

Now suppose by contradiction that $\theta_i > \Theta_i + M/\chi^{(i)}$ so that for $\varepsilon > 0$ sufficiently small, we can let $t = \Theta_i + M/\chi^{(i)} + \varepsilon$ in (4.28) and then evaluate the result in $s = -\varepsilon$, to get using the inequality (4.30) that

$$v(\Theta_i + M/\chi^{(i)})|_{C(i)} \geq h^{(i)} - \max_{k \in C(i)} \{h_k^{(i)} - u^*(k)\} \mathbf{1}.$$

But by denoting j a state of $C(i)$ such that $\max_{k \in C(i)} \{h_k^{(i)} - u^*(k)\} = h_j^{(i)} - u^*(j)$, we end up with

$$v(j, \Theta_i + M/\chi^{(i)}) \geq u^*(j), \quad \text{from which we derive } \theta_j \leq \Theta_i + M/\chi^{(i)} \quad (4.31)$$

◇ We are left with showing that for all j in $C(i)$, we have $\theta_j \geq \theta_i$. Indeed for such a j , the dynamic programming equation (DP) evaluated in θ_j implies:

$$u^*(j) = v(j, \theta_j) = \min_{a \in A_j} \left\{ c^a + \sum_{k \in S} p_k^a v(k, \theta_j - t^a) \right\} \leq \min_{a \in A_j^-} \left\{ c^a + \sum_{k \in S} p_k^a v(k, \theta_j - t^a) \right\}.$$

However by monotonicity of v which is nondecreasing and tends towards u^* , we have by using the fixed-point property in (3.29) verified by u^* that

$$\min_{a \in A_j^-} \left\{ c^a + \sum_{k \in S} p_k^a v(k, \theta_j - t^a) \right\} \leq \min_{a \in A_j^-} \left\{ c^a + \sum_{k \in S} p_k^a u^*(k) \right\} = u^*(j).$$

As a result we obtain the equality

$$\min_{a \in A_j^-} \left\{ c^a + \sum_{k \in S} p_k^a v(k, \theta_j - t^a) \right\} = \min_{a \in A_j^-} \left\{ c^a + \sum_{k \in S} p_k^a u^*(k) \right\}. \quad (4.32)$$

Invoking monotonicity still, it is easily seen that there exists a particular a in A_j^- in both argmins above, i.e. such that

$$c^a + \sum_{k \in S} p_k^a v(k, \theta_j - t^a) = c^a + \sum_{k \in S} p_k^a u^*(k) = u^*(j)$$

(otherwise we could find an action in the right-hand side argmin of (4.32) and not in the left-hand side argmin, which would contradict nondecreasingness of v). The bounds on v are saturated, hence for all k in $\text{supp}(a)$, we have $v(k, \theta_j - t^a) = u^*(k)$. We deduce that for all k in $\text{supp}(a)$, we have $\theta_k \leq \theta_j - t^a \leq \theta_j$. Recall in addition that since a is in A_j^- , $\text{supp}(a)$ is included in $\{k \in S \mid k < j\}$, and if $j \neq i$, then $\text{supp}(a)$ is also included in $C(i)$ (indeed according to Lemma 4.17, going through state i via a descending action is the only manner of leaving $C(i)$). Therefore, if i is in $\text{supp}(a)$, it is clear that $\theta_i \leq \theta_j$, otherwise it suffices to reapply this reasoning (finitely many times) for all k in $\text{supp}(a)$ until eventually hitting i .

◇ We can therefore safely conclude from (4.31) that $\theta_i \leq \Theta_i + M/\chi^{(i)}$, which contradicts our hypothesis $\theta_i > \Theta_i + M/\chi^{(i)}$. We have consequently the bound announced by the theorem, which achieves to show that $\mathcal{P}(i)$ is true, making the induction scheme on the states hierarchy complete.

Dealing with the general case where v^0 is not of the form $u^* - M\mathbb{1}_{S \setminus \{0\}}$ is immediate and done just like in the proof of Proposition 4.15. In particular, our transience bound is only governed by the distance $M := \sup_{s \in [-t_{\max}, 0)} \{v^0(s) - u^*\}$ and the specific values of v^0 do not matter much. \clubsuit

As done after the Proposition 4.15, we propose in the following corollary a coarser version of Theorem 4.19 on the global convergence time $t^* := \max_{i \in S} \theta_i$ that emphasizes the tree structure of hierarchical SMDPs. We still denote by d the maximal length of a descending path in S relatively to the ordering (\leq), and we denote by d^+ the maximal number of states with non-empty set of ascending playable actions along a descending path.

► **COROLLARY 4.20.** *Under the conditions of Theorem 4.19, finite-convergence of v towards u^* occurs in a time t^* such that*

$$t^* \leq d \times t_{\max} + d^+ \times (M/\underline{\chi}).$$

We highlight the fact that both upper bounds stated in Theorem 4.19 or Corollary 4.20 are essentially governed by the ratio $M/\underline{\chi}$, i.e., the convergence time of the value iteration algorithm is proportional to the distance of the initial condition relatively to u^* , and inversely proportional to the (positive) average-cost of the inner part of the system. Such forms echo the expressions derived in the field of network calculus for similar problems, see for instance [LBT01, Chapter 2]

These bounds are generally not tight. In addition of bounding the initial condition v^0 by a function of the form $s \mapsto u^* - M\mathbb{1}_{S \setminus \{0\}}$, an important error can arise in the inequality (4.29) since this lower bound uses the initial values of v (that is to say v^0) over sets of the form $(C(i))_{i \in S, A_i^+ \neq \emptyset}$, not taking into account the fact that during the previous steps of the induction, the function v has grown. In particular, we reckon that the degree d^+ featured in Corollary 4.20 is highly overestimated. Another way to improve the result of Theorem 4.19 would be to consider the “numerical hierarchy” of the average-costs $(\chi^{(i)})_{i \in S, A_i^+ \neq \emptyset}$ in addition of the states “topological hierarchy”. We will propose in the particular case of (EMS-A) below a refined bound in this spirit. However, dealing with these aspects in general is a challenge left for future work.

*
* *

Consistently with Table 4.9, we have for our SMDP with *reduced costs* associated with the system (EMS-A) that $\chi^{(1)} = \rho_A \wedge \rho_P/\pi - \lambda$ and $\chi^{(3)} = \rho_P/\pi - \lambda$. Consequently, we can state building on Theorem 4.11 and Theorem 4.19 that the total absorption time of a peak of M calls in the congestion-free regime (i.e., $\rho_A > \lambda$ and $\rho_P/\pi > \lambda$) is less than t^* , with

$$t^* = \frac{M}{(\rho_A \wedge \rho_P/\pi) - \lambda} + \tau_1 + \frac{M}{\rho_P/\pi - \lambda} + \tau_2 + \tau_3. \quad (4.33)$$

Remark again how the quantity $\rho_{\text{int}} - \lambda$ with $\rho_{\text{int}} := \rho_A \wedge \rho_P/\pi$ expresses the difference between an intrinsic throughput of the system and the input throughput. From a dynamical systems point of view, the quantity $\rho_{\text{int}} - \lambda$ plays the role of a second lowest eigenvalue. From the practical point of view though, it can be interpreted as a margin of security to overcome a sudden peak of calls, that is null if $(N_A, N_P) = (N_A^*, N_P^*)$ and that increases when taking extra agents.

More specifically in the transience bound (4.33), we identify three terms coming from the mere communication delays between states (the times τ_1 , τ_2 and τ_3), and two terms originating from the states that could play ascending actions. The term $M/(\rho_A \wedge \rho_P/\pi - \lambda)$ corresponds to a maximum time needed for the MRAs to catch-up the input and pick up all the extra calls arrived since the bulk. We check that it is proportional to the size of the bulk and is governed by the minimum residual handling speed of agents after they performed all the usual tasks, i.e. the throughput on class $C(1)$ of the semigroup $(\mathcal{S}_t^{(1)})_{t \geq 0}$, given by either $\rho_A - \lambda$ or $\rho_P/\pi - \lambda$ depending on which policy realizes the minimal average-cost between $(A, 0)$ and (A, P) . Similarly, the term

$M/(\rho_P/\pi - \lambda)$ bounds the time needed for the emergency physicians to handle all the calls passed by the MRAs after the latter have caught up with the input. The denominator corresponds to the throughput on the class $C(3) = \{3, 4, 5\}$ of the semigroup $(\mathcal{S}_t^{(3)})_{t \geq 0}$ (applying policy (λ, P)).

A numerical application of the Theorem 4.19 in the case $N_A = 6$, $N_P = 2$ and $M = 8$ which gave the trajectory depicted in Figure 4.13 leads to $t^* = 3386$ s. This is much more than the real convergence time value 1690 s observed in the chart. As explained in the discussion after the Theorem 4.19, having two terms depending on M in (4.33) is too conservative. The next proposition is specific to our example (EMS-A) but presents a refined bound t^* . It still uses an inductive scheme on the states hierarchy, but contrarily to Theorem 4.19, the growth of the value function is retained between the first and the second step.

► **PROPOSITION 4.21.** *The convergence times $(\theta_i)_{i \in S, A_1^*}$ of the value iteration algorithm for the SMDP in SSP configuration associated with (EMS-A) can be bounded by*

$$\theta_1 \leq \frac{M}{(\rho_A \wedge \rho_P/\pi) - \lambda} \quad \text{and} \quad \begin{cases} \theta_3 \leq \theta_1 + \tau_1 & \text{if } \rho_A \leq \rho_P/\pi \\ \theta_3 \leq \theta_1 + \tau_1 + \frac{\rho_A - \rho_P/\pi}{\rho_P/\pi - \lambda} \left(\frac{\tau_1 + \pi\tau_2}{\pi} \right) & \text{if } \rho_A \geq \rho_P/\pi \end{cases}$$

So that $t^* \leq \frac{M}{\rho_{\text{int}} - \lambda} + \tau_1 + \frac{\rho_A - \rho_{\text{int}}}{\rho_P/\pi - \lambda} \left(\frac{\tau_1 + \pi\tau_2}{\pi} \right) + \tau_2 + \tau_3$, with $\rho_{\text{int}} := \rho_A \wedge \rho_P/\pi$.

Proof. Recall that we work in the SMDP of Figure 4.5 with reduced costs (so that it is in SSP configuration). The policy previously denoted by $(\lambda, 0)$ is the unique proper policy in this SSP, in particular it provides a null average-cost, and applying the stationary strategy associated with $(\lambda, 0)$ gives the player an expected total cost u^* . As mentioned before, the constant function $t \mapsto u^*$ is an affine regime of the dynamics semigroup \mathcal{S} of the *complete* SMDP, which we indicate below in the average-cost and bias form.

Policy σ	Average-cost χ	Bias h
$(\lambda, 0)$	$\mathbf{0}$	$-\lambda \begin{pmatrix} 0 \\ 0 \\ \tau_1 \\ \tau_1 \\ \tau_2 \\ \tau_3 \end{pmatrix} (= u^*)$

◇ To control the growth of the function v determined by v^0 and (DP), let us introduce the dynamics semigroup $\mathcal{S}^{(1)}$, which does not use the descending action from 1. In other words, $\mathcal{S}^{(1)}$ uses only the policies $(A, 0)$ and (A, P) . When restricted to $C(1) = \{1, 2, 3, 4, 5\}$, these two policies admit only one final class (see Table 4.9), hence the average-cost on $C(i)$ is uniform, and the bias vector of the affine regime has a single additive degree of freedom over $C(i)$ (as shown by [SF78]). The average-cost and the bias of affine regimes of $\mathcal{S}^{(i)}$ naturally depend on which policy is average-cost optimal between $(A, 0)$ and (A, P) , which amounts to comparing ρ_A and ρ_P/π .

Regardless of which one of ρ_A or ρ_P/π is minimum, let us denote by $\chi^{(1)}\mathbf{1}$ and $h^{(1)}$ the pair of vectors in the table below that gives an affine regime of $\mathcal{S}^{(1)}$. We denote by w^0 the affine initial condition of \mathcal{V}_0 satisfying $w^0(0, s) = 0$ and $w^0(s)|_{C(1)} := \chi^{(1)}s\mathbf{1} + h^{(1)} - M\mathbf{1}$ for all $s \in [-\tau_{\max}, 0)$ and by w the function from $[-\tau_{\max}, \infty)$ to \mathbb{R}^S uniquely determined by w^0 and the complete dynamics \mathcal{S} . Using the property that the dynamics $\mathcal{S}^{(1)}$ are in fact applied (and therefore that the affine regime $(\chi^{(1)}, h^{(1)})$ carries over) until the value for state 1 catches-up $u^*(1)$, we derive the exact expression $w(t) = (\chi^{(1)}t\mathbf{1} + h^{(1)} - M\mathbf{1}) \wedge u^*$ for all $t \geq -\tau_{\max}$.

Policy σ	Average-cost χ restricted to $\{1, 2, 3, 4, 5\}$	Bias h restricted to $\{1, 2, 3, 4, 5\}$
$(A, 0)$	$(\rho_A - \lambda)\mathbf{1}$	$-\rho_A \begin{pmatrix} 0 \\ \tau_1 \\ \tau_1 \\ \tau_2 \\ \tau_3 \end{pmatrix}$
(A, P)	$(\rho_P/\pi - \lambda)\mathbf{1}$	$-\rho_P/\pi \begin{pmatrix} 0 \\ \tau_1 \\ \tau_1 \\ \tau_2 \\ \tau_3 \end{pmatrix} + (\rho_P/\pi - \rho_A) \left(\frac{\tau_1 + \pi\tau_2}{\pi} \right) \begin{pmatrix} 0 \\ 0 \\ 1 \\ 1 \\ 1 \end{pmatrix}$

Since $u^* - M\mathbb{1}_{S \setminus \{0\}} \geq w^0(s)$ for all $s \in [-\tau_{\max}, 0)$, we know that $v(t) \geq w(t)$ for all $t \geq -\tau_{\max}$. In particular, denoting by $\bar{\theta}_1 := M/\chi^{(1)}$, we have $v(\bar{\theta}_1) \geq h^{(1)} \wedge u^* = h^{(1)}$ (indeed, whether ρ_A or ρ_P/π is minimum, we always have $h^{(1)} \leq u^*$). However, because $v(1, \bar{\theta}_1) \geq h_1^{(1)} = 0 = u^*(1)$ and that on the other hand we knew that $v(1, \cdot) \leq u^*(1)$, it means that $v(1, \bar{\theta}_1) = u^*(1)$ and as a result $\theta_1 \leq \bar{\theta}_1$, which is the bound promised by the proposition claim.

- ◇ We now study the system evolution for $t \geq \theta_1$. It is sufficient to only focus on the states of $C(3) = \{3, 4, 5\}$, since $v(1, t) = u^*(1)$ already and $v(2, t) = u^*(2)$ as soon as $t \geq \theta_1 + \tau_1$. Recall that $v(3, t) = (v(1, t - \tau_1) - \lambda\tau_1) \wedge (N_P/\pi + v(5, t))$, and depending on which term is minimizing, either the policy $(\lambda, 0)$ or the policy (λ, P) is applied. Let us suppose that (λ, P) and the associated dynamics $\mathcal{S}^{(3)}$ are used. The following table provides an affine regime of this policy over $C(3)$ (once again, this class is strongly connected so there is only one degree of freedom)

Policy σ	Average-cost χ restricted to $\{3, 4, 5\}$	Bias h restricted to $\{3, 4, 5\}$
(λ, P)	$(\rho_P/\pi - \lambda)\mathbf{1}$	$-\rho_P/\pi \begin{pmatrix} 0 \\ \tau_2 \\ \tau_2 + \tau_3 \end{pmatrix}$

The first part of the proof gives us $v(s)|_{C(3)} \geq \chi^{(1)}(s - \bar{\theta}_1)\mathbf{1} + h^{(1)}|_{C(3)}$ for all s in $[-\tau_{\max} + \bar{\theta}_1, \bar{\theta}_1]$ (replacing in the expression of w the term $-M\mathbf{1}$ by $-\chi^{(1)}\bar{\theta}_1\mathbf{1}$). If $\chi^{(3)}\mathbf{1}$ and $h^{(3)}$ are the average-cost and bias vectors of (λ, P) in our last table and if $\alpha \geq 0$ is such that $\chi^{(1)}(s - \bar{\theta}_1)\mathbf{1} + h^{(1)}|_{C(3)} \geq \chi^{(3)}(s - \bar{\theta}_1)\mathbf{1} + h^{(3)} - \alpha\mathbf{1}$ for all s in $[-\tau_{\max} + \bar{\theta}_1, \bar{\theta}_1]$, then by propagation of the dynamics to this lower affine bound we obtain for all i in $\{3, 4, 5\}$ and $t \geq \theta_1$ that $v(i, t) \geq (\chi^{(3)}(t - \bar{\theta}_1) + h^{(3)}(i) - \alpha) \wedge u^*(i)$. Actually since $\chi^{(3)} \geq \chi^{(1)}$ always holds, it suffices to choose α such that $h^{(1)}|_{C(3)} \geq h^{(3)} - \alpha\mathbf{1}$.

- ◇ Suppose first that $\rho_A \leq \rho_P/\pi$. Comparing the expressions of $h^{(1)}|_{C(3)}$ and $h^{(3)}$, we find that $\alpha := \rho_A\tau_1$ convenes. Since $h^{(3)}(3) = 0$, we obtain $v(3, \bar{\theta}_3) = u^*(3)$ for the instant $\bar{\theta}_3 := \bar{\theta}_1 + (\rho_A - \lambda)\tau_1/(\rho_P/\pi - \lambda)$. However, $\rho_A \leq \rho_P/\pi$ implies $\bar{\theta}_3 \leq \bar{\theta}_1 + \tau_1$, so that we were wrong about assuming that (λ, P) was played, and we had $v(3, t) = v(1, t - \tau_1) - \lambda\tau_1$ for $t \geq \theta_1$. In particular, $\theta_3 \leq \bar{\theta}_1 + \tau_1$.
- ◇ Suppose now that $\rho_P/\pi \leq \rho_A$. This time, we find that $\alpha := (\rho_P/\pi)\tau_1 + (\rho_A - \rho_P/\pi)(\tau_1 + \pi\tau_2)/\pi$ convenes. As a result we are ensured that $v(3, \bar{\theta}_3) = u^*(3)$ for the instant $\bar{\theta}_3 := \bar{\theta}_1 + \tau_1 + (\rho_A - \rho_P/\pi)(\tau_1 + \pi\tau_2)/(\rho_P/\pi - \lambda)$, which is indeed greater than τ_1 . By definition of θ_3 , we have $\theta_3 \leq \bar{\theta}_3$.
- ◇ It suffices to add the times τ_2 (resp. $\tau_2 + \tau_3$) to $\bar{\theta}_3$ to derive an upper bound on θ_4 and θ_5 , and in this case we have $t^* = \theta_5$. The final expression of t^* in the claim covers all the cases regardless on which of ρ_A and ρ_P/π attains the minimum. ♣

This new bound leads to a value of $t^* = 2678$ s for the experiment depicted in Figure 4.13, that is much better than the previous one, but still far from the real value 1670 s. In fact, the bound

of Proposition 4.21 is tight whenever the initial condition is $h^{(1)} - M\mathbb{1}_{S \setminus \{0\}}$ instead of $u^* - M\mathbb{1}_{S \setminus \{0\}}$, where $h^{(1)}$ is the bias vector of the dynamics semigroup $\mathcal{S}^{(1)}$ featured in the proof (such that $h^{(1)}(1) = 0$). Unfortunately, bounding u^* below by $h^{(1)}$ can bring a significant starting error.

The Figure 4.16 illustrates the tight bounds computed by Proposition 4.21 when the initial condition is of the form $h^{(1)} - M\mathbb{1}_{S \setminus \{0\}}$, for a first case where $N_A = 6$ and $N_P = 2$ (so that $\rho_{\text{int}} = \rho_A$) and for a second case where $N_A = 5$ and $N_P = 3$ (which gives $\rho_{\text{int}} = \rho_P/\pi$), with still $M = 8$. The upper bounds on the $(\theta_i)_{1 \leq i \leq 5}$ are represented by the five vertical lines on each graph, and the lower dashed lines give the affine regime $t \mapsto \chi^{(1)}t + h^{(1)}$.

Remark that the expression of the bound of Proposition 4.21 depends on the sign of $\rho_A - \rho_P/\pi$. We retrieve the frontier equation $\rho_A = \rho_P/\pi$ between the two slowed down phases of the congestion diagram of Figure 4.7. This suggests that just like the throughput vector, we can also associate to the catch-up times a polyhedral complex, however constrained to be a subdivision of the congestion-free phase of the throughput complex.

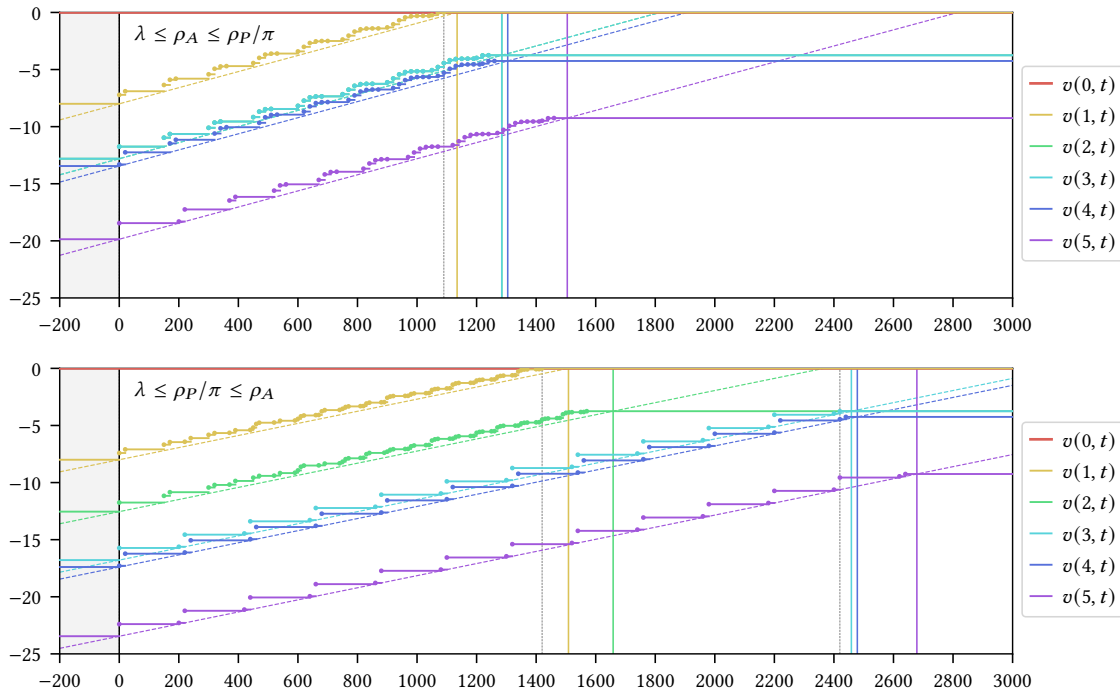


Figure 4.16: Visualization of the upper bounds computed in Proposition 4.21 that are tight if the initial condition is of the form $h^{(1)} - M\mathbb{1}_{S \setminus \{0\}}$

6 Application to a staffing problem of the SAMU

We finish this chapter by illustrating what operational recommendations can be drawn from our analytical formulas. As previously mentioned, the model (EMS-A) was indeed designed with the physicians of the four SAMU (medical emergency services) of Paris area. A key question of theirs was to know if they were understaffed and what staffing levels would arise from the mathematical modeling.

To address this problem, we computed the throughput diagram of Figure 4.7 and deduced minimal numbers N_A^* of MRAs and N_P^* of emergency physicians needed to be in the congestion-free phase. However, we had to amend these theoretical formulas by an **occupancy factor**. Indeed, our model assumes not only that the agents are infinitely divisible, but also that they can work without any interruption. In actual fact, health professionals recommend that occupancy rates of MRAs (fraction of their time spent on the phone) should not go over $\eta_A := 40\%$ (see [MP19]) to guarantee satisfying work conditions (sufficient level of vigilance, more time to pick up new calls immediately...). Emergency physicians as well must keep low online occupancy. It is considered that their occupancy level for the handling of inbound calls should stay below $\eta_P := 20\%$, in order to account for the fact that they have back-office tasks to do (for instance finding the most adapted hospital facilities to receive the patient) and also other calls to handle (reports from dispatched teams, etc.).

In order to determine staffing levels N_A and N_P for MRAs and physicians that achieve occupancy rates η_A and η_P , we choose to divide our minimal frontiers N_A^* and N_P^* by η_A and η_P , thus obtaining

$$N_A = \frac{N_A^*}{\eta_A} = \frac{\lambda(\tau_1 + \pi\tau_2)}{\eta_A} \quad \text{and} \quad N_P = \frac{N_P^*}{\eta_P} = \frac{\pi\lambda(\tau_2 + \tau_3)}{\eta_P}. \quad (4.34)$$

The logic is that taking $N_A = N_A^*$ would represent a 100% occupancy level for MRAs, that choosing twice as more agents $N_A = 2N_A^*$ would make them work under a 50% occupancy, and so on. This rule is empirical but we show in the Section 3.1 of Chapter 8 that it is in practice a very good approximation. Notice that this choice falls in the category of the “quality-driven” staffing rules (see [GMR02]), where the number of agents N equals the offered load N^* plus a safety percentage, so that $N = (1 + \gamma)N^*$, with γ to tune (recall from Section 3.3 that our congestion-free frontiers coincide with the offered load of queueing theory). Less conservative staffing rules exist, like the so-called “square-root principle” (see [KMo2]) where $N = N^* + \beta\sqrt{N^*}$ (and β is a parameter to adjust), but are better suited for commercial call centers and not emergency response.

We verify that the recommended staffings in (4.34) are proportional to the arrival rate of inbound calls and inversely proportional to the occupancy factors. These formulas can in fact be used in two manners:

- either we know the throughput λ of inbound calls that we want to cope with, and the target occupancy factor η that a pool of agents has to remain under, and we use (4.34) to determine minimal number of agents in this pool;
- either we do not have the choice on N that is already given, and since λ is exogenous as well we use (4.34) to compute the actual occupancy undergone by our N agents.

The first approach allows us to obtain **staffing maps**, indicating for all possible values of λ and η the minimal number of agents to take in the continuous-relaxed model. The iso-staffing values of these maps are lines going through the origin, according to the expression (4.34). We depict on the Figure 4.17 such a map for both the medical regulation assistants and the emergency

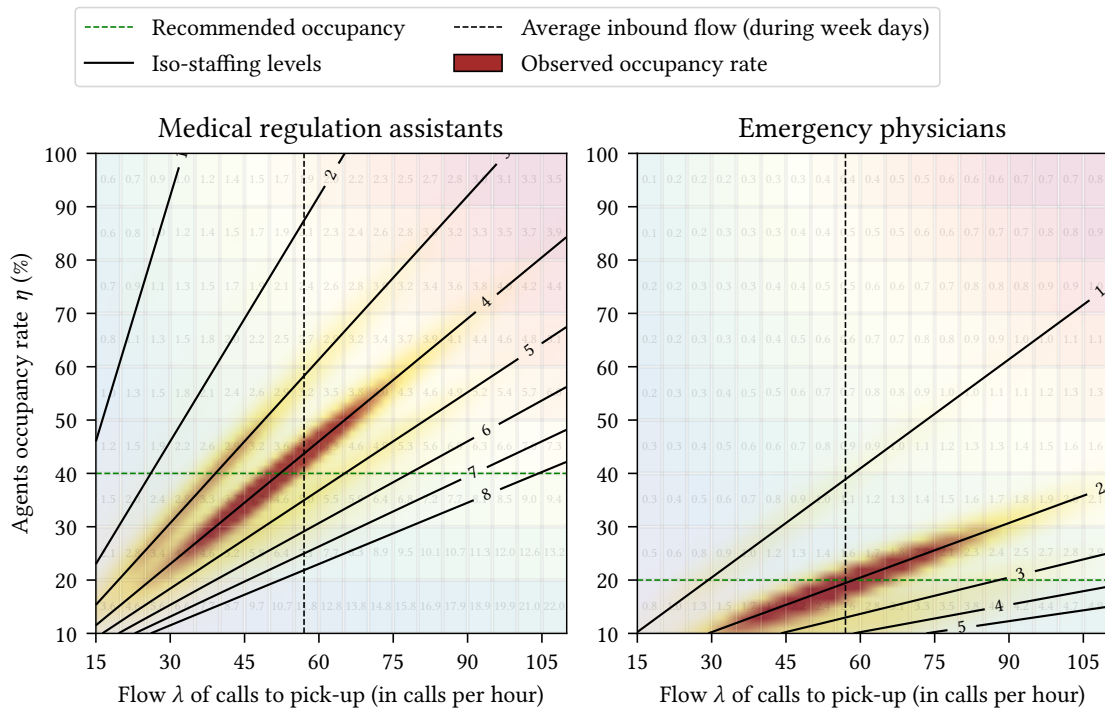


Figure 4.17: Staffing maps N_A and N_P as functions of the rate of arrivals of inbound calls and occupancy rate (the example of the SAMU 94)

physicians, using real parameters measured in the SAMU 94 call center using the phone call logs of more than 10,000 emergency calls received in February 2019. Conversely, we have measured the number of logged agents in the SAMU 94 call center (averaged over windows of 15 minutes). Using the corresponding values of the arrival rate λ (averaged on the same timespans), we also depict in Figure 4.17 the estimated occupancy level of agents based on the formula (4.34).

The first map for MRAs tells us that to handle an average throughput λ of 57 calls per hour and not exceed a 40 % occupancy rate, there should be at least 4.3 MRAs in the call center. This is close from the most typical staffing of four agents, with a little understaffing that therefore yields an undergone occupancy rate of 44 % when facing the average demand. However, it is insufficient to only look at this mean arrival rate (taking into account all slots of the day from 7 AM to 9 PM), and we can see that there are sometimes up to 75 calls per hour. For this high solicitation, 4 MRAs will experience a 58 % occupancy, which is too much with respect to the official recommendations. To stay below the 40 % limit with 75 calls per hour, we find that 6 answering agents should be present in the call center. As far as emergency physicians are concerned, the typical number of two doctors handling the most severe calls is sufficient in average to stay below the 20 % limit of occupancy, but an additional specialist is needed for the periods of high demand.

These results, that we have adapted in [ABG20b] for each SAMU of Paris area taking into account their specificities, were taken into account by the emergency doctors we work with to raise the staffing policies. They confirmed that our computations matched their own operational observations.

5

NON-MONOTONIC PETRI NET MODELS OF EMERGENCY CALL CENTERS AND THEIR CONGESTION DIAGRAMS

Contents

1	Stationary regimes of Petri nets with priority rules	190
2	A bilevel emergency call center with a priority mechanism	194
2.1	Description of the model	194
2.2	Congestion diagram	196
3	An emergency call center for medical needs with a monitored reservoir . .	201
3.1	Description of the model	201
3.2	Congestion diagram	203
4	Variations of the model EMS-B	208
4.1	Support of the reservoir assistant to pick up inbound calls	208
4.2	Assigning another extra flow to the reservoir assistants	211
4.3	Mutual support of two first-answering agent pools	215

The previous chapters have given us a deep understanding of monotonic fluid Petri nets. The analysis conducted there allowed us to characterize several measures of performance of such nets, with a great level of confidence. Indeed, and to name just these two positive results, recall from Section 3 of Chapter 4 that under good assumptions, monotonic fluid Petri nets admit ultimate throughputs whose dependence upon the resources is well-described by a polyhedral complex (the congestion diagram), and that in addition the behaviour of such nets is a fair approximation of more realistic discrete Petri nets – even subject to non-deterministic events (see the Section 3 of Chapter 2).

In this chapter, we come back to a modeling aspect discussed in Chapter 1 and mostly left aside afterwards, the priority routing. As discussed there, priorities are essentially non-monotonic, and

loosing this good property makes the associated dynamics much harder to study. In particular, the correspondence results with semi-Markov decision processes established in Chapter 4 do not hold anymore. However, in view of the rich modeling interest of priority rules to describe real-life organizations by means of Petri nets, and especially emergency call centers, we want to evaluate to which extent the results and techniques derived for the monotonic case carry over the non-monotonic one.

This chapter builds very much on the work of Allamigeon, Bœuf and Gaubert [ABG15] who introduced the counter equations for two-level priority routings in Petri nets, already featured in Chapter 1. The study of these mechanisms was one of the main aspects of the dissertation of Bœuf [Bœu17]. The approach that these authors developed was inspired by a work of Farhi, Goursat and Quadrat [FGQ11], who built a min-plus model for a road traffic network and used the idea of representing priorities by rational min-plus dynamics in the discrete-time case. The treatment of the continuous time case was addressed in [ABG15] using a symbolic perturbation technique, working with a semifield of germs. This method has also been used in [GG98] for algorithmic purposes, see also [AFG⁺14].

In Section 1, we recall the approach used by Allamigeon, Bœuf and Gaubert to compute stationary regimes of non-monotonic fluid Petri nets. In contrast with [ABG15], we deal with a case allowing an arbitrary number of priority levels and also under general arc multipliers; we also briefly discuss the associated open problems. The rest of the chapter is much more applied and focuses on various systems of emergency call centers. In Section 2, we recall the simple model of a so-called “bilevel” emergency call center studied in [ABG15, Bœu17], designed and used nowadays by the police and rescue forces of the Paris area to handle the calls of three different emergency numbers. Although the analysis of this particular system is not new from the theoretical point of view, it comes convenient in order to illustrate the computation techniques used in this chapter. Above all, the results derived here shall be reused in the Part II of the thesis where we conduct the extensive case-study of this Parisian emergency call center. We dedicate the Section 3 to the analysis of the non-monotonic Petri net model (EMS-B), which is a complexified version of the medical emergency call center model (EMS-A) broadly studied in the Chapter 4. The aim of (EMS-B) is in particular to alleviate the global slowdown undergone by (EMS-A) during too important peaks of calls. We show that when the system (EMS-B) is understaffed, some counter-intuitive behaviours can arise owing to the priority rules. Finally, the Section 4 discusses several variations of the model (EMS-B) that can be implemented to extend the congestion-free regime.

1 Stationary regimes of Petri nets with priority rules

Let us recall in the opposite table the dynamical equations followed by the transitions counter functions $(z_q)_{q \in Q}$ of a continuous-relaxed Petri net, as they were presented in the Table 2.3. The notation of Chapters 1 and 2 are reused, in particular recall that the priority routing governing the downstream transitions of a place p in \mathcal{P} is described by a total order (\prec_p) on p^{out} . We readily remark that the generic equation for transitions in Q_{prio} do not correspond anymore to a dynamic programming equation of semi-Markov models. This is due to the fact that counter equations for priority rules involve differences, therefore somehow “negative probabilities” in the world of SMDPs...

Recall from Theorem 2.5 that these dynamical equations, alongwith the knowledge of the input profiles $(z_q)_{q \in Q_{\text{input}}}$, determine a unique trajectory of the fluid Petri net in the monotonic setting. It also holds with priority rules provided that the input profiles are regular enough. We

Type	Counter equation in the continuous-relaxed model
$q \in \mathcal{Q}_{\text{input}}$	$z_q(t)$ is prescribed externally
$q \in \mathcal{Q}_{\text{sync}}$	$z_q(t) = \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}(t - \tau_p) \right)$
$q \in \mathcal{Q}_{\text{psel}}$	$z_q(t) = \alpha_{qp}^{-1} \left(\pi_{qp}^0 m_p + \pi_{qp} \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}(t - \tau_p) \right)$
$q \in \mathcal{Q}_{\text{prio}}$	$z_q(t) = \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}(t - \tau_p) - \sum_{q' <_{\rho} q} \alpha_{q'p} z_{q'}(t) - \sum_{q' >_{\rho} q} \alpha_{q'p} z_{q'}(t^-) \right)$

Table 2.3 reminded: dynamical equations followed by the counter functions of the transitions in continuous-relaxed Petri nets

have seen in Chapter 4 that whenever $\mathcal{P}_{\text{prio}} = \emptyset$ and the input profiles are affine, this trajectory is satisfactorily described by an affine function up to a bounded error term (see Corollary 4.2). This result for monotonic Petri nets motivates us to also look for solutions of the dynamics of Table 2.3 (where in general $\mathcal{P}_{\text{prio}} \neq \emptyset$) that are ultimately affine, i.e., of the form $z_q(t) \underset{t \rightarrow \infty}{=} \rho_q t + u_q + o(1)$ for all q in \mathcal{Q} , with nonnegative throughputs $(\rho_q)_{q \in \mathcal{Q}}$ and real intercepts $(u_q)_{q \in \mathcal{Q}}$.

In the Section 3.2 of Chapter 4 (hence for monotonic Petri nets), we showed building on the correspondence results with SMDPs that the tuples $(\rho_q, u_q)_{q \in \mathcal{Q}}$ verified min-plus equations with no more recursive aspect but involving minimums relatively to a lexicographic ordering. Recall for example the generic optimality equation (L) for transitions in $\mathcal{Q}_{\text{sync}}$

$$\forall q \in \mathcal{Q}_{\text{sync}}, \quad (\rho_q, u_q) = \min_{p \in q^{\text{in}}}^{\text{LEX}} \left\{ \left(\alpha_{qp}^{-1} \sum_{q' \in Q} \alpha_{pq'} \rho_{q'}, \alpha_{qp}^{-1} m_p + \alpha_{qp}^{-1} \sum_{q' \in Q} \alpha_{pq'} (u_{q'} - \rho_{q'} \tau_p) \right) \right\}, \quad (\text{L})$$

where in contrast with the optimality equation featured in Chapter 4 we transformed the term ρ_q in the second term of the right-hand side by $\alpha_{qp}^{-1} \sum_{q' \in p^{\text{in}}} \alpha_{q'p} \rho_{q'}$, as permitted by the rules of the lexicographic ordering.

It is of course no accident that our detour via semi-Markov decision processes provided an equation (L) on the tuples $(\rho_q, u_q)_{q \in \mathcal{Q}}$ that bears much resemblance with the original counter equations of Table 2.3, and this can be algebraically justified by working with **germs of affine functions**. In general, a germ at infinity of a function (later simply referred to as a germ for brevity) is an equivalence class for the relation which identifies two functions that coincide for sufficiently large values of their argument. In what follows, we shall choose a positive scalar δ and work with the equivalence relation \sim such that for two functions f and g , we have $f \sim g$ if and only if $f(n\delta) = g(n\delta)$ for all n in \mathbb{N} large enough.

The tuple (ρ, u) in \mathbb{R}^2 will represent the germ of the affine function $t \mapsto \rho t + u$. For instance if $\delta = 1$, the two usual functions $t \mapsto t$ and $t \mapsto \lfloor t \rfloor$ are described by the same germ $(1, 0)$, even though the second one is not properly affine. The pointwise order on functions induces a total order on germs of affine functions, which coincides with the already introduced lexicographic order on the coordinates (ρ, u) , the ρ coordinate being considered first. As before, the lexicographic

minimum of two germs will indifferently be denoted by \min^{LEX} or the symbol \wedge . Since \wedge is an associative operation on germs of affine functions, (\mathbb{R}^2, \wedge) is a semigroup. In fact, we can complete it in a monoid (i.e., a semigroup with a neutral element) by introducing a greatest element \top with respect to the lexicographic minimum, so that $(\rho, u) \wedge \top := (\rho, u)$. Observe further that the usual sum of functions induces a sum over the set of germs, with $(\rho, u) + (\rho', u') := (\rho + \rho', u + u')$, and by definition $(\rho, u) + \top := \top$. Since this addition is also associative and distributive with respect to \wedge , this makes $\mathbb{G} := (\mathbb{R}^2 \cup \{\top\}, \wedge, +)$ a **min-plus semifield** (every germ (ρ, u) indeed admits the inverse $(-\rho, -u)$ relatively to $+$). Finally, a notion of exponentiation can also be considered in \mathbb{G} since the multiplicative group $\mathbb{R}_{>0}$ acts on \mathbb{G} by setting $a(\rho, u) := (a\rho, au)$, for all $a > 0$ and (ρ, u) in \mathbb{R}^2 , and $a\top := \top$.

The technique developed in [ABG15] is to look for ultimately affine solutions of the equations of Table 2.3 by considering the associated germs. If (ρ, u) is the germ of the function f , it is immediate to see that $(\rho, u - \rho\tau)$ is the germ of $t \mapsto f(t - \tau)$, and as a result one can see that the equation (L) of Chapter 4 is a “germ” version of the equations for transitions in $\mathcal{Q}_{\text{sync}}$ and $\mathcal{Q}_{\text{psel}}$ of Table 2.3. The remaining challenge is to derive a germ equation for transitions in $\mathcal{Q}_{\text{prio}}$. The problem may seem ill-posed due to the presence of terms of the form $z_q(t^-)$ in the corresponding counter equation, since if z is affine, $z(t^-)$ coincides with $z(t)$. Allamigeon et al. tackled this issue by showing that it is well-posed to look for ultimately affine solutions on the **δ -discretization** of the dynamics instead. It means that the holding times of the places are all multiples of $\delta > 0$ and thus trajectories that are constant on the intervals $([k\delta, (k+1)\delta))_{k \in \mathbb{N}}$ are computed. In this discretized model, the term $z(t^-)$ is replaced by $z(t - \delta)$. The detour via the discretized dynamics enables one to prove that, regardless of the choice of δ small enough, some terms cannot achieve the minimum in some dynamical equations associated with transitions ruled by priority, and thus can be removed.

Applying the above method, we arrive at the equations of Table 5.1 giving the min-plus system of equations verified by the germs of affine solutions of the dynamics of Table 2.3, the correctness being stated in Theorem 5.1. Compared with the result stated in [ABG15], our framework covers the case of multiple levels of priorities and not the mere two-level priority rules. We also take into account the arc multipliers of the net although this aspect is pretty straightforward.

Type	Germ equation in stationary regime
$q \in \mathcal{Q}_{\text{sync}}$	$(\rho_q, u_q) = \min_{p \in q^{\text{in}}}^{\text{LEX}} \alpha_{qp}^{-1} \left((0, m_p) + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} (\rho_{q'}, u_{q'} - \rho_{q'} \tau_p) \right)$
$q \in \mathcal{Q}_{\text{psel}}$	$(\rho_q, u_q) = \pi_{qp} \cdot \alpha_{qp}^{-1} \left((0, m_p) + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} (\rho_{q'}, u_{q'} - \rho_{q'} \tau_p) \right)$
$q \in \mathcal{Q}_{\text{prio}}$	$(\rho_q, u_q) = \min_{\substack{p \in q^{\text{in}} \\ \sum_{q' > p q} \rho_{q'} = 0}}^{\text{LEX}} \alpha_{qp}^{-1} \left((0, m_p) + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} (\rho_{q'}, u_{q'} - \rho_{q'} \tau_p) - \sum_{q' \in p^{\text{out}} \setminus \{q\}} \alpha_{q'p} (\rho_{q'}, u_{q'}) \right)$

Table 5.1: System of equations followed by germs of transitions counter functions

Observe that for a transition q in $\mathcal{Q}_{\text{prio}}$, the equation featured in Table 5.1 claims that the germ (ρ_q, u_q) is only determined by places p of q^{in} such that $\sum_{q' > p q} \rho_{q'} = 0$. Although this condition first looks involved, it actually encodes the fact that q can be constrained only by an upstream place p that does not provide fluid-tokens to another transition q' with lesser priority level than

q , and therefore it is a natural requirement. We retrieve that the behaviour of the transitions in $\mathcal{Q}_{\text{prio}}$ does not depend only on their “parent” transitions but also on their siblings.

- **THEOREM 5.1.** *The germs of ultimately affine solutions of the dynamics of Table 2.3 are solutions of the system of equations of Table 5.1, with the condition that $\rho_q \geq 0$ for all q in \mathcal{Q} .*

Proof. We prove the result for transitions ruled by priority, the proof being direct for other patterns. Let us first write the generic counter equation followed by a transition $q \in \mathcal{Q}_{\text{prio}}$:

$$z_q(t) = \min_{p \in q^{\text{in}}} \alpha_{qp}^{-1} \left(m_p + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} z_{q'}(t - \tau_p) - \sum_{q' <_p q} \alpha_{q'p} z_{q'}(t) - \sum_{q' >_p q} \alpha_{q'p} z_{q'}(t^-) \right).$$

Our claim is that due to the priority mechanism, some terms in the above minimum cannot realize minimality and thus can be removed.

Let $q_1 \in \mathcal{Q}_{\text{prio}}$, $p \in q_1^{\text{in}}$ and $q_2 \in p^{\text{out}}$, with $q_2 \neq q_1$. Let us suppose that $q_1 <_p q_2$. Substituting counters by their corresponding germs and replacing t^- by $t - \delta$ with $\delta > 0$, we have

$$(\rho_{q_1}, u_{q_1}) \leq \alpha_{q_1 p}^{-1} \left((\rho_{\Sigma}^p, u_{\Sigma}^p) - \sum_{q' <_p q_1} \alpha_{q' p} (\rho_{q'}, u_{q'}) - \sum_{q' >_p q_1} \alpha_{q' p} (\rho_{q'}, u_{q'} - \rho_{q'} \delta) \right) \quad (5.1)$$

$$\leq \alpha_{q_1 p}^{-1} \left((\rho_{\Sigma}^p, u_{\Sigma}^p) - \sum_{q' \neq q_1} \alpha_{q' p} (\rho_{q'}, u_{q'}) + \sum_{q' >_p q_1} \alpha_{q' p} (0, \rho_{q'} \delta) \right) \quad (5.2)$$

where $(\rho_{\Sigma}^p, u_{\Sigma}^p)$ stands for the germ $(0, m_p) + \sum_{q' \in p^{\text{in}}} \alpha_{pq'} (\rho_{q'}, u_{q'} - \rho_{q'} \tau_p)$. Similarly, we have

$$(\rho_{q_2}, u_{q_2}) \leq \alpha_{q_2 p}^{-1} \left((\rho_{\Sigma}^p, u_{\Sigma}^p) - \sum_{q' \neq q_2} \alpha_{q' p} (\rho_{q'}, u_{q'}) + \sum_{q' >_p q_2} \alpha_{q' p} (0, \rho_{q'} \delta) \right) \quad (5.3)$$

Let us apply the nondecreasing mapping of $\mathbb{G} \rightarrow \mathbb{G}$: $g \mapsto \alpha_{q_2 p} \alpha_{q_1 p}^{-1} (g - (\rho_{q_2}, u_{q_2})) + (\rho_{q_1}, u_{q_1})$ in both sides of the latter equation. We obtain:

$$(\rho_{q_1}, u_{q_1}) \leq \alpha_{q_1 p}^{-1} \left((\rho_{\Sigma}^p, u_{\Sigma}^p) - \sum_{q' \neq q_1} \alpha_{q' p} (\rho_{q'}, u_{q'}) + \sum_{q' >_p q_2} \alpha_{q' p} (0, \rho_{q'} \delta) \right) \quad (5.4)$$

Comparing (5.4) to (5.2), one can observe that if $\rho_{q'} > 0$ for some $q_1 <_p q' \leq_p q_2$, then the right-hand side of (5.4) strictly bounds below the one of (5.2), thus the equality in (5.2) cannot be achieved, and the corresponding germ can be removed from the original minimum. This reasoning can be applied when q_2 is the transition of p^{out} with the least priority, so that the inequality (5.2) is strict whenever the sum $\sum_{q' >_p q_1} \rho_{q'}$ is positive (the $(\rho_q)_{q \in \mathcal{Q}}$ variables are nonnegative). Conversely, only the contributions of upstream places $p \in q_1^{\text{in}}$ such that $\sum_{q' >_p q_1} \rho_{q'} = 0$ remain in the minimum. \clubsuit

The Theorem 5.1 is one of the few theoretical results on non-monotonic Petri nets. In general, we do not know if a system of equations over \mathbb{G} determined by the Table 5.1 admits solutions, nor can we conclude if we have found a solution that the trajectory of the Petri net indeed admits the ρ -part of this solution as a throughput vector! This negative result was highlighted by numerical simulations in [ABG15, Figure 4–(d)]. However, such pathological behaviours appear to be related to the existence of arithmetical relationships between the holding times of places.

Establishing under which conditions the germ equations associated with non-monotonic Petri nets admit solutions that satisfactorily provide the throughput of the trajectories is an important question that remains open and is left for future work.

2 A bilevel emergency call center with a priority mechanism

We now start studying the properties of several non-monotonic emergency call center organizations. To begin with, and following our presentation of the different French emergency services in the [Introduction](#), we discuss the case of the Emergency Calls Platform of Paris (in French “Plateforme d’appels d’urgence”, abbreviated PFAU), which shows many specificities. We shall discuss them very thoroughly in the Part II of the thesis with additional elements of context. For now, we mainly focus on the organizational aspects. First, this platform receives several millions of calls each year, since it covers the whole city of Paris but also its inner suburbs (roughly seven million inhabitants in total). Second, it receives emergency calls from three different emergency numbers (17–18–112) and is operated by both firefighters and police officers, hence gathering calls about thefts, aggressions, fires, gas leaks, rescue needs, etc. For medical needs, the population most often dials the number 15 to reach the SAMU (and its system (EMS-A)) but can also call the firefighters in some situations (traffic accidents, first-aid support...). Last but not least, it handles the calls in a two-level process, contrarily to the majority of other police or rescue call centers in France that operate with the single-tier logic introduced in the [Example 2.6](#).

2.1 DESCRIPTION OF THE MODEL. — In this section, we study an idealized model of the Parisian 17–18–12 PFAU, in particular we only consider a one-job system answering a single inbound flow of emergency calls. The general case of three different numbers handled by a two-job organization will be tackled in Part II. The call center functions as follows. Emergency calls arrive via the uppermost transition z_0 and have to be picked up by a **first-level operator**, that we also refer to as a LVL 1 operator. These agents have the mission to quickly filter out and qualify the demand of the caller by asking basic questions (this takes a characteristic time τ_1) with three different possible outcomes:

- (i) the call is **very urgent (VU)**, (for instance a fire, a heart attack, an ongoing robbery, etc.) and should be handled with **priority** by a qualified policeman of firefighter to dispatch response teams as soon as possible;
- (ii) the call is **urgent (U)** and should also be transferred to a skilled police officer/firefighter in order to answer the caller demand;
- (iii) the call is **non urgent (NU)** (car impoundment, caller locked himself out, or simply a dial error) and there is no need to transfer the call to another agent, the LVL 1 operator may give an advice to the caller or tell him/her to call a more suitable and non urgent service.

We depict in [Figure 5.2](#) a timed Petri net representing our bilevel call center. As done in [Section 1](#) of [Chapter 4](#), the transitions are referred to directly by their counters. We denote by π^{VU} , π^U and π^{NU} the proportions of very urgent, urgent and non urgent calls in the inbound flow (naturally with $\pi^{VU} + \pi^U + \pi^{NU} = 1$), so that the conversation place with holding time τ_1 is ruled by proportional preselection routing. If the call is non urgent, the LVL 1 agent can terminate the call and come back to his/her availability pool by firing the transition z_2'' . Otherwise, the caller is transferred to a **second-level operator** (also referred to as a LVL 2 operator), whose mission is to handle the request, ask the caller for more information on his/her location and the situation, potentially dispatch vehicles, or demand support from other services.

In case the call is very urgent, the situation is considered too serious for the caller to be left on hold. Consequently, after the call is qualified, the transition z_2 is fired but the LVL 1 agent stays on line with the caller, waiting for a LVL 2 operator to pick up the call at transition z_3 , with

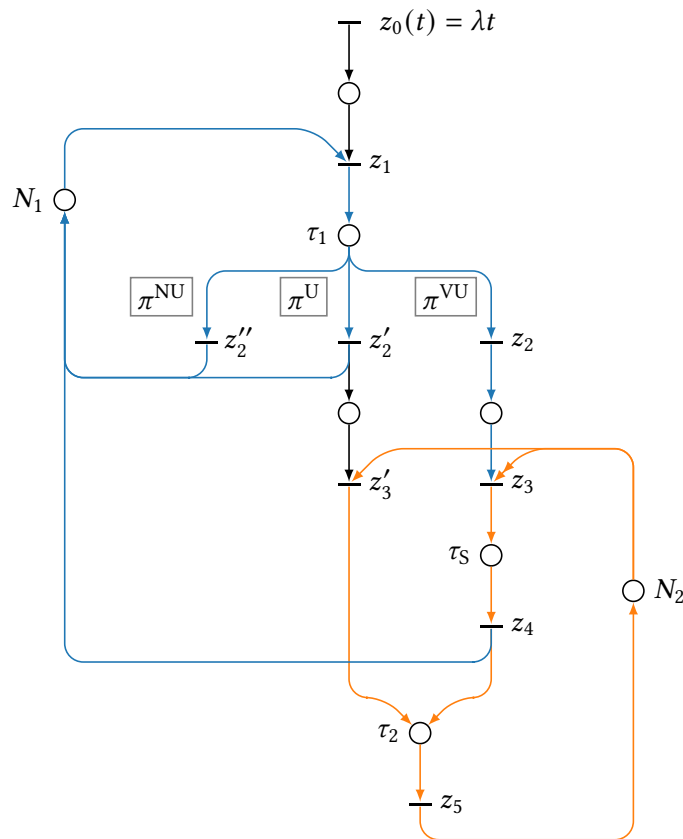


Figure 5.2: A Petri net model of a bilevel call center with priority

priority over urgent calls. This triggers the beginning of a synchronization step (also referred to as accompaniment or escort step), in which the LVL 1 operator hands over the case for a duration τ_S . At the end of this synchronization, z_4 is fired and releases the first level agent to the LVL 1 pool while the proper LVL 2 instruction of the VU call by the second-level operator begins. Once it is finished (after a characteristic time τ_2), the LVL 2 agent comes back to his/her pool.

In case the call is qualified as urgent by the LVL 1 operator, the transition z_2' is fired and the level-one agent comes back directly to the LVL 1 pool, because it is considered acceptable in this case that the caller is placed in a waiting room. Once a LVL 2 agent is available and no very urgent call is pending, the transition z_2' can be fired and the real instruction begins.

We denote by N_1 (resp. N_2) the total number of first-level (resp. second-level) operators, all idle at the start of the execution of the net and therefore giving the initial marking of the LVL 1 and LVL 2 agent pools. To keep the model simple, we assume that it takes a LVL 1 operator the same (constant) time τ_1 to qualify and treat an inbound call whether it is very urgent, urgent or non-urgent. Similarly, we suppose that the instruction of a call by a LVL 2 operator always takes a time τ_2 , regardless of the urgent or very urgent character. The strongest assumption is once again that the callers have an infinite patience, and that they will not hang up the call themselves even though they wait for a long time. In Part II, we will consider non-deterministic and gravity-dependent conversation times with very realistic distributions, and we will also use realistic non-infinite patience functions.

According to this model description and the Table 2.3, we are able to write the dynamical recursive equations satisfied by the counter functions $(z_q)_{q \in Q}$:

$$\left\{ \begin{array}{l} z_1(t) = z_0(t) \wedge (N_1 + z'_2(t) + z''_2(t) + z_4(t)) \\ z_2(t) = \pi^{\text{VU}} z_1(t - \tau_1) \\ z'_2(t) = \pi^{\text{U}} z_1(t - \tau_1) \\ z''_2(t) = \pi^{\text{NU}} z_1(t - \tau_1) \\ z_3(t) = z_2(t) \wedge (N_2 + z_5(t) - z'_3(t^-)) \\ z'_3(t) = z'_2(t) \wedge (N_2 + z_5(t) - z_3(t)) \\ z_4(t) = z_3(t - \tau_2) \\ z_5(t) = z'_3(t - \tau_2) + z_4(t - \tau_5) \end{array} \right. \quad (\text{BiLVL})$$

Observe that the “non-trivial” equations are those on z_1 , z_3 and z'_3 with synchronizations phenomena, respectively giving the (cumulated) number of picked up calls by the LVL 1 operators, the number of very urgent calls picked up by LVL 2 operators and the number of urgent calls picked up by LVL 2 operators. The other equations can be easily derived from these three fundamental counters, which by substitution satisfy the reduced dynamics

$$\left\{ \begin{array}{l} z_1(t) = z_0(t) \wedge (N_1 + (1 - \pi^{\text{VU}})z_1(t - \tau_1) + z_3(t - \tau_2)) \\ z_3(t) = \pi^{\text{VU}} z_1(t - \tau_1) \wedge (N_2 + z'_3(t - \tau_2) + z_3(t - \tau_2 - \tau_5) - z'_3(t^-)) \\ z'_3(t) = \pi^{\text{U}} z_1(t - \tau_1) \wedge (N_2 + z'_3(t - \tau_2) + z_3(t - \tau_2 - \tau_5) - z_3(t)) \end{array} \right. \quad (\text{BiLVL})$$

2.2 CONGESTION DIAGRAM. — In order to find the affine regimes of the net, we use the germs equations presented in the Table 5.1. We can either write the germs of the reduced dynamics above, or write the complete system on germs and then combine the resulting equations together to keep only the germs (ρ_1, u_1) , (ρ_3, u_3) and (ρ'_3, u'_3) . We also suppose that the input z_0 admits an affine profile with throughput λ , and we replace the term z_0 in (BiLVL) by the germ $(\lambda, 0)$ (the value 0 here is arbitrary but does not change the results to come). We obtain the following system

$$\left\{ \begin{array}{l} (\rho_1, u_1) = (\lambda, 0) \wedge ((1 - \pi^{\text{VU}})\rho_1 + \rho_3, N_1 + (1 - \pi^{\text{VU}})(u_1 - \rho_1\tau_1) + u_3 - \rho_3\tau_2) \\ (\rho_3, u_3) = \begin{cases} (\pi^{\text{VU}}\rho_1, \pi^{\text{VU}}(u_1 - \rho_1\tau_1)) \wedge (\rho_3, N_2 + u_3 - \rho_3(\tau_2 + \tau_5)) & \text{if } \rho'_3 = 0 \\ (\pi^{\text{VU}}\rho_1, \pi^{\text{VU}}(u_1 - \rho_1\tau_1)) & \text{otherwise} \end{cases} \\ (\rho'_3, u'_3) = (\pi^{\text{U}}\rho_1, \pi^{\text{U}}(u_1 - \rho_1\tau_1)) \wedge (\rho'_3, N_2 + u'_3 - \rho_3(\tau_2 + \tau_5) - \rho'_3\tau_2) \end{array} \right. \quad (5.5)$$

The major difference with the generic optimality system (L) that we obtained in the monotonic case is that due to the priority rules, some equations are now conditional on the values of other transitions' throughputs. This is the case here for the transition z_3 , that cannot be constrained by the second-level agents if the concurrent transition z'_3 with inferior priority is fired with a positive throughput. Nevertheless, we can still look for solutions of these equations by adopting an enumerative scheme similar to the one we used in Section 3.2 of Chapter 4. This amounts to

making a choice of minimizing terms in each synchronization equation (or equivalently choosing a bottleneck policy). This determines a value of the throughputs $(\rho_q)_{q \in Q}$, while the non-minimizing terms provide inequalities that are necessary conditions for the chosen policy to be effectively minimizing.

The Proposition 5.2 gives the solutions of the ρ -part of the system (5.5). It makes use of the two nonnegative ratios

$$r' := \frac{\pi^{\text{VU}}(\tau_2 + \tau_S)}{\tau_1 + \pi^{\text{VU}}\tau_S} \quad \text{and} \quad r := \frac{(\pi^{\text{U}} + \pi^{\text{VU}})\tau_2 + \pi^{\text{VU}}\tau_S}{\tau_1 + \pi^{\text{VU}}\tau_S} \quad (5.6)$$

that are such that $r' \leq r$, and $r = r'$ if it happens that $\pi^{\text{U}} = 0$.

► **PROPOSITION 5.2.** *The solutions $(\rho_1, \rho_3, \rho'_3)$ in $\mathbb{R}_{\geq 0}^3$ of the system (5.5) are given by*

$$\rho_1 = \lambda \wedge \left(\frac{N_1}{\tau_1 + \pi^{\text{VU}}\tau_S} \right) \wedge \left(\frac{N_2}{\pi^{\text{VU}}(\tau_2 + \tau_S)} \right), \quad \rho_3 = \pi^{\text{VU}}\rho_1, \quad \text{and}$$

$$\rho'_3 = \begin{cases} \pi^{\text{U}}\rho_1 & \text{if } N_2 \geq r \min(N_1, N_1^*) \\ \frac{N_2 - \pi^{\text{VU}}\rho_1(\tau_S + \tau_2)}{\tau_2} & \text{if } r' \min(N_1, N_1^*) \leq N_2 \leq r \min(N_1, N_1^*) \\ 0 & \text{if } N_2 \leq r' \min(N_1, N_1^*) \end{cases}$$

where $N_1^* := \lambda(\tau_1 + \pi^{\text{VU}}\tau_S)$.

Proof. We distinguish two main cases, whether $\rho'_3 > 0$ or not.

◇ If $\rho'_3 > 0$, the system (5.5) provides the following inequalities:

$$(\rho_1, u_1) \leq (\lambda, 0) \quad (\text{E1.}\lambda)$$

$$(\rho_1, u_1) \leq ((1 - \pi^{\text{VU}})\rho_1 + \rho_3, N_1 + (1 - \pi^{\text{VU}})(u_1 - \rho_1\tau_1) + u_3 - \rho_3\tau_2) \quad (\text{E1.1})$$

$$(\rho_3, u_3) = (\pi^{\text{VU}}\rho_1, \pi^{\text{VU}}(u_1 - \rho_1\tau_1)) \quad (\text{E3.0})$$

$$(\rho'_3, u'_3) \leq (\pi^{\text{U}}\rho_1, \pi^{\text{U}}(u_1 - \rho_1\tau_1)) \quad (\text{E3'.0})$$

$$(\rho'_3, u'_3) \leq (\rho'_3, N_2 + u'_3 - \rho_3(\tau_2 + \tau_S) - \rho'_3\tau_2) \quad (\text{E3'.2})$$

Minimum reached by (E1.1) and (E3'.0). — We derive from the ρ -parts of (E1.1) and (E3'.0) that $\rho_1 = \lambda$ and $\rho'_3 = \pi^{\text{U}}\lambda$, and in addition $\rho_3 = \pi^{\text{VU}}\lambda$ according to the ρ -part of (E3.0). The u -part of (E3.0) gives the equality $u_3 = \pi^{\text{VU}}(u_1 - \rho_1\tau_1) = \pi^{\text{VU}}(u_1 - \lambda\tau_1)$, that we can re-inject in the inequality $u_1 \leq N_1 + (1 - \pi^{\text{U}})(u_1 - \lambda\tau_1) + u_3 - \pi^{\text{VU}}\lambda\tau_2$ brought by the u -part of (E1.1) to conclude that $N_1 \geq \lambda(\tau_1 + \pi^{\text{U}}\tau_2)$. Finally, the u -part of (E3'.2) gives the inequality $u'_3 \leq N_2 + u'_3 - \pi^{\text{VU}}\lambda(\tau_S + \tau_2) - \pi^{\text{U}}\lambda\tau_2$ or equivalently $N_2 \geq \pi^{\text{VU}}\lambda(\tau_S + \tau_2) + \pi^{\text{U}}\lambda\tau_2$.

Minimum reached by (E1.1) and (E3'.0). — Using the ρ -part of (E3'.0), we know that $\rho'_3 = \pi^{\text{U}}\rho_1$, while $\rho_3 = \pi^{\text{VU}}\rho_1$ still by (E3.0). We can use the u -part of (E1.1) to obtain the equality $u_1 = N_1 + (1 - \pi^{\text{VU}})u_1 - (1 - \pi^{\text{VU}})\rho_1\tau_1 + u_3 - \pi^{\text{VU}}\rho_1\tau_2$. On the other hand, we have $u_3 = \pi^{\text{VU}}u_1 - \pi^{\text{VU}}\rho_1\tau_1$ according to the u -part of (E3.0), thus we obtain $\rho_1 = N_1/(\tau_1 + \pi^{\text{VU}}\tau_2)$, we deduce $\rho_3 = \pi^{\text{VU}}N_1/(\tau_1 + \pi^{\text{VU}}\tau_2)$ and $\rho'_3 = \pi^{\text{U}}N_1/(\tau_1 + \pi^{\text{VU}}\tau_2)$.

The ρ -part of (E1.1) gives the inequality $\rho_1 \leq \lambda$, that we can write in the form $N_1 \leq \lambda(\tau_1 + \pi^{\text{VU}}\tau_2)$. The u -part of (E3'.2) leads to $N_2 \geq \rho_3(\tau_S + \tau_2) + \rho'_3\tau_2$, and using the values of ρ_3 and ρ'_3 we find $N_2 \geq (\pi^{\text{VU}}(\tau_S + \tau_2) + \pi^{\text{U}}\tau_2)N_1/(\tau_1 + \pi^{\text{VU}}\tau_2)$, or equivalently $N_2 \geq rN_1$.

Minimum reached by (E1.1) and (E3'.2). — The ρ -part of (E1.1) ensures that $\rho_1 = \lambda$, the ρ -part of (E3.0) gives $\rho_3 = \pi^{\text{VU}}\lambda$ and the u -part of (E3'.2) provides $u_3 = N_2 + u_3 - \rho_3(\tau_S + \tau_2) - \rho'_3\tau_2$ so $\rho'_3 = (N_2 - \pi^{\text{VU}}\lambda(\tau_S +$

$\tau_2)/\tau_2$. The ρ -part of (E3'.0) tells us $\rho_3' \leq \pi^U \rho_1$, and in addition we are in a case where by assumption $\rho_3' \geq 0$, so we have the double inequality $\pi^{VU} \lambda (\tau_S + \tau_2) \leq N_2 \leq \pi^{VU} \lambda (\tau_S + \tau_2) + \pi^U \lambda \tau_2$. Finally, the u -part of (E1.1) leads by combination with the u -part of (E3.0) to the inequality $N_1 \geq \lambda (\tau_1 + \pi^{VU} \tau_2)$ once again.

Minimum reached by (E1.1) and (E3'.2). – The ρ -part of (E3.0) gives $\rho_3 = \pi^{VU} \rho_1$, the u -part of the same equation gives $u_3 = \pi^{VU} (u_1 - \rho_1 \tau_1)$, which combined with the u -part of (E1.1) provides $\rho_1 = N_1 / (\tau_1 + \pi^{VU} \tau_2)$. The u -part of (E3'.2) yields $u_3 = N_2 + u_3 - \rho_3 (\tau_S + \tau_2) - \rho_3' \tau_2$ so $\rho_3' = (N_2 - \pi^{VU} \rho_1 (\tau_S + \tau_2)) / \tau_2 = (N_2 - \pi^{VU} N_1 (\tau_S + \tau_2)) / (\tau_1 + \pi^{VU} \tau_2) / \tau_2$.

From the ρ -inequality of (E1.1), we have $\rho_1 \leq \lambda$ hence $N_1 \leq \lambda (\tau_1 + \pi^{VU} \tau_S)$. The ρ -inequality of (E3'.0) gives $\rho_3' \leq \pi^U \rho_1$ and using the assumption $\rho_3' \geq 0$, we can write the double inequality $\pi^{VU} N_1 (\tau_S + \tau_2) / (\tau_1 + \pi^{VU} \tau_2) \leq N_2 \leq \pi^{VU} N_1 (\tau_S + \tau_2) / (\tau_1 + \pi^{VU} \tau_2) + \tau_2 \pi^U N_1 / (\tau_1 + \pi^{VU} \tau_2)$, which in compact notation is equivalent to $r' N_1 \leq N_2 \leq r N_1$.

◇ In the other case $\rho_3' = 0$, we have the inequalities

$$(\rho_1, u_1) \leq (\lambda, 0) \quad (\text{E1.1})$$

$$(\rho_1, u_1) \leq ((1 - \pi^{VU}) \rho_1 + \rho_3, N_1 + (1 - \pi^{VU})(u_1 - \rho_1 \tau_1) + u_3 - \rho_3 \tau_2) \quad (\text{E1.1})$$

$$(\rho_3, u_3) \leq (\pi^{VU} \rho_1, \pi^{VU} (u_1 - \rho_1 \tau_1)) \quad (\text{E3.0})$$

$$(\rho_3, u_3) \leq (\rho_3, N_2 + u_3 - \rho_3 (\tau_2 + \tau_S)) \quad (\text{E3.2})$$

$$(0, u_3') \leq (\pi^U \rho_1, \pi^U (u_1 - \rho_1 \tau_1)) \quad (\text{E3'.0})$$

$$(0, u_3') \leq (0, N_2 + u_3' - \rho_3 (\tau_2 + \tau_S)) \quad (\text{E3'.2})$$

Minimum reached by (E1.1) and (E3.0). – We derive from the ρ -parts of (E1.1) and (E3.0) that $\rho_1 = \lambda$ and $\rho_3 = \pi^{VU} \lambda$. The fact that $\rho_1 > 0$ and by assumption $\rho_3' = 0$ implies that (E3'.2) attains the minimum, and as a result $u_3' = N_2 + u_3' - \rho_3 (\tau_2 + \tau_S)$, and we obtain that $N_2 = \pi^{VU} \lambda (\tau_2 + \tau_S)$. We can stop the computation here since our initial assumptions has led to a phase that is not full-dimensional.

Minimum reached by (E1.1) and (E3.0). – Using the ρ -part of (E3.0), we know that $\rho_3 = \pi^{VU} \rho_1$. Therefore, the ρ -part of (E1.1) just amounts to the trivial relation $\rho_1 = \rho_1$. However, we can use the u -part of the same (E1.1) to obtain the equality $u_1 = N_1 + (1 - \pi^{VU}) u_1 - (1 - \pi^{VU}) \rho_1 \tau_1 + u_3 - \pi^{VU} \rho_1 \tau_2$. On the other hand, we have $u_3 = \pi^{VU} u_1 - \pi^{VU} \rho_1 \tau_1$ according to the u -part of (E3.0), thus by combination we obtain $\rho_1 = N_1 / (\tau_1 + \pi^{VU} \tau_2)$, and $\rho_3 = \pi^{VU} N_1 / (\tau_1 + \pi^{VU} \tau_2)$. But the u -part of (E3'.2) still gives $\rho_3 = N_2 / (\tau_S + \tau_2)$, hence we have $N_2 = r' N_1$ and the phase is not full-dimensional.

Minimum reached by (E1.1) and (E3.2). – The ρ -part of (E1.1) ensures that $\rho_1 = \lambda$, while the u -part of (E3.2) provides $u_3 = N_2 + u_3 - \rho_3 (\tau_S + \tau_2)$ so $\rho_3 = N_2 / (\tau_S + \tau_2)$. On one hand, remark that we have the inequality $\rho_3 \leq \pi^{VU} \rho_1 \leq \pi^{VU} \lambda$ by the ρ -part of (E3.0), but on the other hand we know from the ρ -part of (E1.1) that $\lambda \leq (1 - \pi^{VU}) \lambda + \rho_3$, so $\rho_3 \geq \pi^{VU} \lambda$. This is again not a full-dimensional case.

Minimum reached by (E1.1) and (E3.2). – The u -part of (E3.2) still gives $\rho_3 = N_2 / (\tau_S + \tau_2)$, and because the ρ -part of (E1.1) gives $\rho_3 = \pi^{VU} \rho_1$, we obtain $\rho_1 = N_2 / (\pi^{VU} (\tau_S + \tau_2))$. The ρ -inequality of (E1.1) gives $\rho_1 \leq \lambda$ hence $N_2 \leq \pi^{VU} (\tau_S + \tau_2)$. In addition, the u -part of (E1.1) provides $N_1 - \rho_1 (\tau_1 + \pi^{VU} \tau_2) = \pi^{VU} (u_1 - \rho_1 \tau_1) - u_3$, but we know that the latter quantity is nonnegative by the u -inequality of (E3.0), hence $N_1 \geq N_2 (\tau_1 + \pi^{VU} \tau_2) / (\pi^{VU} (\tau_S + \tau_2))$, or equivalently $N_2 \leq r' N_1$. \clubsuit

The Proposition 5.2 shows that for the system (BiLvl), we can still compute a **throughput complex**, that is to say a polyhedral complex covering $\mathbb{R}_{\geq 0}^P$ such that the throughputs are linear functions of the resources over each cells of the complex. We represent this diagram and its five full-dimensional phases in the Figure 5.3.

The Figure 5.4 further gives the profile of the three throughputs ρ_1 , ρ_3 and ρ_3' characterized by Proposition 5.2, as functions of N_1 and N_2 . We verify that the linearity regions of these maps induce the polyhedral complex of Figure 5.3.

We interpret the Figures 5.3 and 5.4 as follows. If $N_1 \geq N_1^*$ and $N_2 \geq N_2^* := r N_1^*$ (the green phase), we have $\rho_1 = \lambda$, $\rho_3 = \pi^{VU} \lambda$ and $\rho_3' = \pi^U \lambda$, so that all the calls are handled with no delay, and the system is **congestion-free**. If $N_1 \geq N_1^*$ but $r' N_1^* \leq N_2 \leq r N_1^*$ (the yellow phase), we still

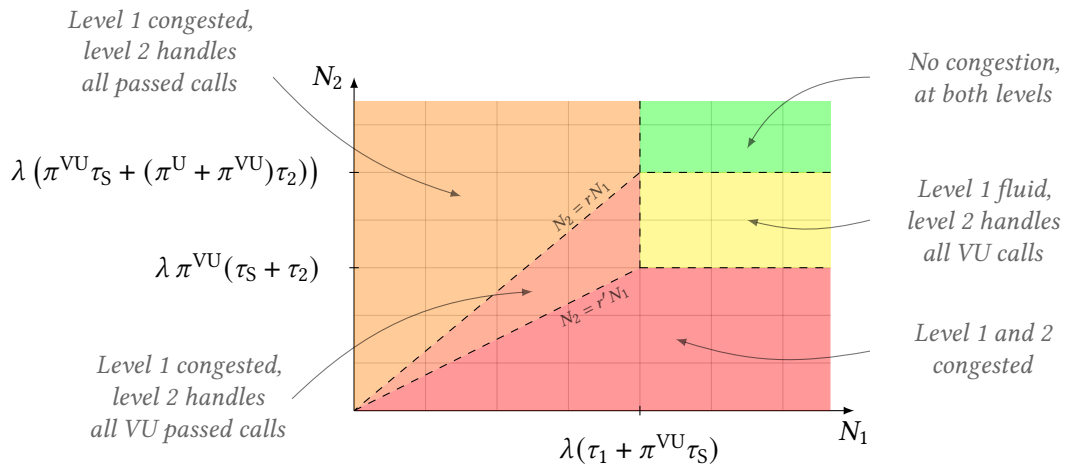


Figure 5.3: Phase diagram of the throughput of the system (BiLVL)

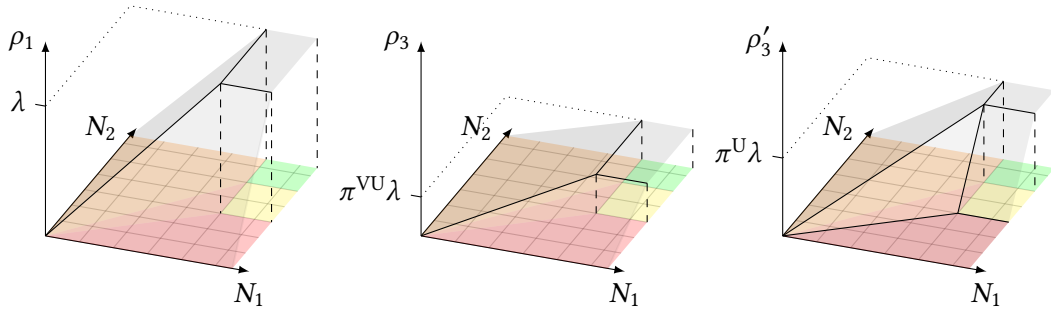


Figure 5.4: Contrary to ρ_1 and ρ_3 , the throughput ρ_3' is not concave, and can decrease as N_1 increases

have $\rho_1 = \lambda$ and $\rho_3 = \pi^{\text{VU}}\lambda$, hence the LVL 1 operators pick up all the inbound calls and the LVL 2 operators handle all the VU calls handed over by the LVL 1 without delay, however there are not enough second-level agents to handle all the U calls. In this case, the urgent calls pile up indefinitely because LVL 2 operators pick them up with more and more delay.

If $N_1 \leq N_1^*$, there are not enough LVL 1 operators to pick up all the inbound calls as soon as they arrive. Nevertheless, the picked up calls are still qualified in the three categories NU, U and VU. If in addition $N_2 \geq rN_1$ (light orange phase), it means that the LVL 2 is sufficiently staffed to handle all the U and VU calls qualified as such by the LVL 1, but there remains pending U and VU calls not yet picked up by the first level. If $r'N_1 \leq N_2 \leq rN_1$ (dark orange phase), not only is the LVL 1 slowed down due to an insufficient first-level staffing, but so is the LVL 2 with not all urgent calls handled. All the VU calls detected and handed over by the LVL 1 are however still handled with no delay by the LVL 2.

Finally, whenever $N_2 \leq r'N_1$ and $N_2 \leq N_1^*$ (red phase) there are so few LVL 2 agents that not all VU calls picked up by the first level can be treated without delay, in particular no U calls at all are handled. Due to the synchronization step between LVL 1 and LVL 2 agents in the VU calls procedure and the fact that these VU calls are picked up with more and more delay by second-level operators, the LVL 1 agents end up being slowed down by the LVL 2 regardless of their own staffing. Indeed, they are all ultimately kept on line in the synchronization steps waiting for available LVL 2 agents to release them. In a more realistic model where the caller's patience is not infinite, this phenomenon would not arise, although the behaviour of the call center would still be very hampered in this phase.

In spite of the bad behaviour of the organization in this most congested phase, let us observe that the priority mechanism fully fulfills its purpose, as notably showed by the two **intermediary phases** (or buffer phases). In these, even though there are not LVL 2 agents to cope with the whole demand, all the very urgent calls are still protected and none of them suffers a wait. In the real PFAU call center implementing this bilevel layout, many procedure are in practice adapted to this specific feature, in particular the LVL 1 operators have the mission to maintain a short conversation time τ_1 (between 20 s and 30 s) and not to delve too much into details with the caller. This enables the first level to pick up quickly all the arriving calls and detect the minority of calls that are very urgent. These calls can then receive an in-depth treatment by second-level operators. Compared with a usual single-tier call center, the bilevel system requires a little more agents but very urgent calls are much better handled, as we will also see in Part II.

The Figure 5.5 gives an alternative illustration of the throughputs ρ_3 and ρ'_3 in the congestion-free, semi-congested (at LVL 2) and congested phases as a function of N_2 only, when the LVL 1 is sufficiently staffed (this is therefore a vertical cross-section of the Figure 5.4 for $N_1 \geq N_1^*$).

Compared with the result of Corollary 4.5 in the monotonic setting, observe in the presence of priority rules that the throughputs of the transitions of the net are not necessarily concave functions of the resources anylonger, as it appears for ρ'_3 in Figures 5.4 and 5.5. Moreover, although all the throughputs in this example are nondecreasing functions of the headcount at LVL 1, this is not the case relatively to the number of agents at LVL 2. Indeed, as stated by Proposition 5.2 and illustrated by Figure 5.4, the throughput ρ'_3 decreases if N_2 increases when (N_1, N_2) stays in

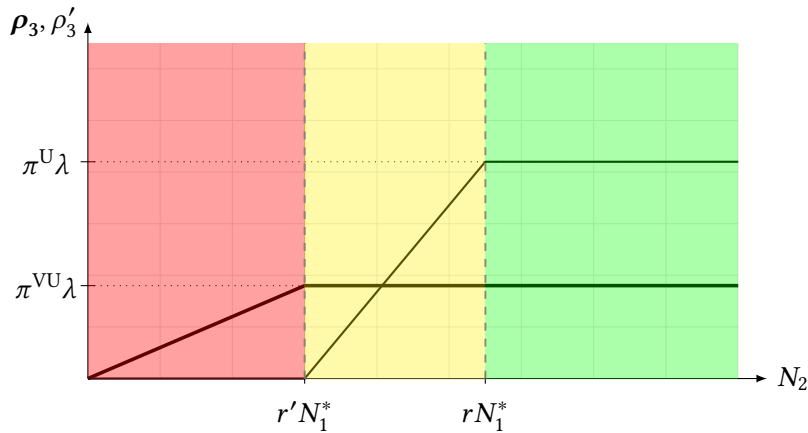


Figure 5.5: A cross-section of the throughputs ρ_3 (bold) and ρ'_3 from Figure 5.4 for a constant value of $N_1 \geq N_1^*$

$\{N_1 \leq N_1^*, r'N_1 \leq N_2 \leq rN_1\}$ (our dark orange phase). This is not counter-intuitive, because in this phase not all inbound calls are handled by the LVL 1 (not enough first-level agents) and the LVL 2 agents handle all the handed over VU calls, in addition with some U calls. A small increase of N_1 shall increase the number of VU calls transferred to the LVL 2 agents, who have to treat less U calls in order to handle these new VU calls first, as prescribed by the priority mechanism.

Observe nevertheless that the transition z_3 with the highest priority level still admits a throughput ρ_3 that is a concave and nondecreasing piecewise-affine function of the resources N_1 and N_2 . We invite the reader to keep this fact in mind for comparison with the next case-study.

3 An emergency call center for medical needs with a monitored reservoir

We now come back to the case of emergency call centers for medical needs, still drawn from the French SAMU organizations and specifically discussed in the Section 1 of Chapter 4 with the system (EMS-A). If needed, we encourage the reader to go over this previous section again, since we are now about to introduce the alternative model (EMS-B).

- 3.1 DESCRIPTION OF THE MODEL.** — Recall from the Section 3.3 of Chapter 4 in which the throughput of the system (EMS-A) is computed that the handling speed of medical regulation assistants (abb. MRAs, the first-answering agents) and the handling speed of the emergency physicians are always entangled. In particular for this organization, a slowdown arising either in the MRAs circuit or in the physicians circuit causes a slowdown of the whole system. This phenomenon is due to the synchronization step at transition z_3 , that can as a consequence be identified as the “weakest link” of the treatment chain.

To alleviate this issue and still maintain the presence of an interlocutor with the patient and the brief oral summary told to physician, emergency doctors we were in touch with¹ as part of our collaboration proposed to consider another model. One may create a new type of MRA, the **reservoir assistant**, who after a brief discussion with the MRA having picked up the call, places the patient in a monitored reservoir. The first-answering MRA is released to pick up other inbound calls. When an emergency physician becomes available, the reservoir assistant passes on the short briefing to the doctor and transfers the patient. While the queue of patients in the reservoir is non empty, the reservoir assistant checks on the patients in the reservoir, and can call patients back in case they hang up. This replaces the synchronizations between physicians and answering MRAs, enabling the latter to pick up new calls more quickly.

Another advantage of the reservoir mechanism is that if a single reservoir assistant is sufficient to handle all the calls, this agent can have a consolidated vision of all the patients waiting for emergency physicians and revise in real time their priority level if more severe cases arrive. In the system with no reservoir, the emergency physicians may have to ask each of the waiting MRAs which call he/she should take first. In what follows, we are going to explicitly consider that the reservoir assistant distinguishes high priority calls and low priority calls, and transfers them to the doctors accordingly.

¹ We would like to especially thank DR Éric Lecarpentier from Hôpital Henri Mondor in Créteil for his leading role on this topic

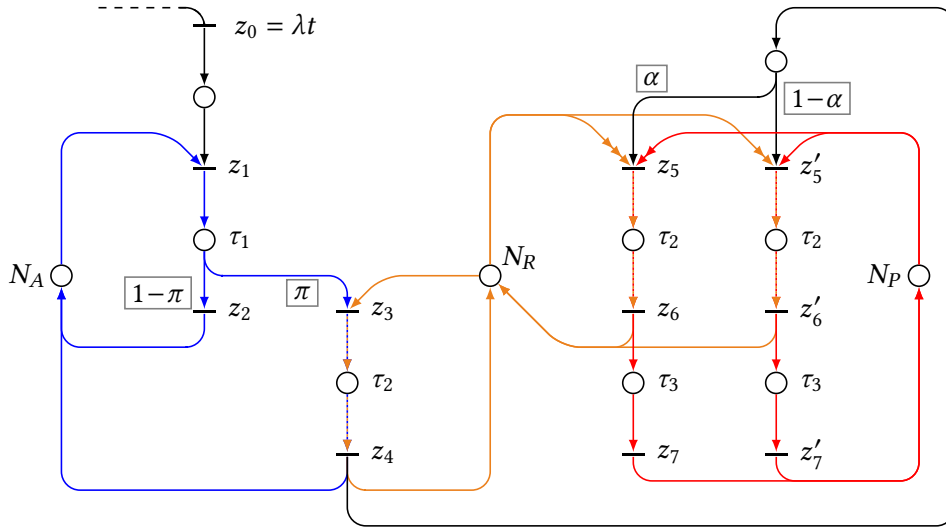


Figure 5.6: Medical emergency call center with a monitored reservoir (EMS-B)

The model (EMS-B) depicted in Figure 5.6 implements these modifications, to be compared with the model (EMS-A) depicted in Figure 4.3. The reservoir assistant pool is a new place with initial marking N_R (not necessarily equal to 1). Reservoir assistants receive patients from the answering MRAs at transition z_3 and transfer them to physicians at transitions z_5 and z'_5 , depending on the severity of the case. We still denote by τ_2 the characteristic time of the short briefing steps between several agents. We denote by α the proportion of very urgent calls among patients who need to talk to an emergency physician. In case of conflict, that is to say if a same reservoir assistant can either welcome a new patient in the reservoir or transfer a patient already in the reservoir to an emergency doctor, we decide that this reservoir assistant must first empty the reservoir before filling it, and should first escort out the very urgent calls. Note that this induces a three-level priority rule for reservoir assistant tokens. The three possible release transitions of the reservoir assistants are z_4 , z_6 and z'_6 (respectively after taking a new patient in the reservoir, transferring a very urgent call from the reservoir to the emergency doctor, and transferring a less urgent call from the reservoir to the doctor). Consultations with a physician still take a time τ_3 after which transitions z_7 and z'_7 can be fired. The circuits involving the reservoir assistant are depicted with color orange. It can be verified that the places standing for the pool of reservoir assistants and physicians have compatible priority rules.

Using Table 2.3, we can write the dynamics of the counter variables of the net. As before, we present below a reduced system of equations where z_2 , z_4 , z_6 , z'_6 , z_7 and z'_7 have been substituted by expressions depending on z_1 , z_3 , z_5 and z'_5 only. For the sake of readability, we denote $z|_{t_1}^{t_2} := z(t_2) - z(t_1)$, and $z|_t^t := z(t)$.

$$\begin{aligned}
 z_1(t) &= z_0|_t^t \wedge \left(N_A + (1 - \pi) z_1|_t^{t-\tau_1} + z_3|_t^{t-\tau_2} \right) & \text{(EMS-B)} \\
 z_3(t) &= \pi z_1|_t^{t-\tau_1} \wedge \left(N_R + z_3|_t^{t-\tau_2} + z_5|_t^{t-\tau_2} + z'_5|_t^{t-\tau_2} \right) \\
 z_5(t) &= \alpha z_3|_t^{t-\tau_2} \wedge \left(N_P + z_5|_t^{t-\tau_2-\tau_3} + z'_5|_t^{t-\tau_2-\tau_3} \right) \wedge \left(N_R + z_3|_t^{t-\tau_2} + z_5|_t^{t-\tau_2} + z'_5|_t^{t-\tau_2} \right) \\
 z'_5(t) &= (1 - \alpha) z_3|_t^{t-\tau_2} \wedge \left(N_P + z_5|_t^{t-\tau_2-\tau_3} + z'_5|_t^{t-\tau_2-\tau_3} \right) \wedge \left(N_R + z_3|_t^{t-\tau_2} + z_5|_t^{t-\tau_2} + z'_5|_t^{t-\tau_2} \right)
 \end{aligned}$$

3.2 CONGESTION DIAGRAM. — We suppose that the input z_0 admits an affine profile with rate λ . Applying Theorem 5.1 and the equations of Table 5.1 to the model (EMS-B) provides the following system on the affine germs of counter variables (again after substitutions of some germs easily expressed in terms of those associated with counters z_1, z_3, z_5 and z'_5):

$$\begin{aligned}
(\rho_1, u_1) &= (\lambda, 0) \wedge ((1 - \pi)\rho_1 + \rho_3, N_A + (1 - \pi)(u_1 - \rho_1\tau_1) + u_3 - \rho_3\tau_2) \\
(\rho_3, u_3) &= (\pi\rho_1, \pi(u_1 - \rho_1\tau_1)) \wedge (\rho_3, u_3 + N_R - \rho_3\tau_2 - (\rho_5 + \rho'_5)\tau_2) \\
(\rho_5, u_5) &= \begin{cases} \alpha(\rho_3, u_3 - \rho_3\tau_2) \wedge (\rho_5, u_5 + N_P - \rho_5(\tau_2 + \tau_3)) \wedge (\rho_3, u_3 + N_R - (\rho_3 + \rho_5)\tau_2) & \text{if } \rho'_5 = 0 \text{ and } \rho_3 = 0 \\ \alpha(\rho_3, u_3 - \rho_3\tau_2) \wedge (\rho_5, u_5 + N_P - \rho_5(\tau_2 + \tau_3)) & \text{if } \rho'_5 = 0 \text{ and } \rho_3 > 0 \\ \alpha(\rho_3, u_3 - \rho_3\tau_2) & \text{if } \rho'_5 > 0 \end{cases} \\
(\rho'_5, u'_5) &= \begin{cases} (1 - \alpha)(\rho_3, u_3 - \rho_3\tau_2) \wedge (\rho'_5, u'_5 + N_P - (\rho_5 + \rho'_5)(\tau_2 + \tau_3)) \wedge (\rho'_5, N_R + u'_5 - (\rho_5 + \rho'_5)\tau_2) & \text{if } \rho_3 = 0 \\ (1 - \alpha)(\rho_3, u_3 - \rho_3\tau_2) \wedge (\rho'_5, u'_5 + N_P - (\rho_5 + \rho'_5)(\tau_2 + \tau_3)) & \text{if } \rho_3 > 0 \end{cases}
\end{aligned}$$

Similarly to our study of the (BlLvl) model, we now have to distinguish cases on the possible bottleneck upstream places in $\mathcal{P}_{\text{prio}}$ depending on the throughputs of the other transitions they rule, which brings a much more complex system of equations on germs than (4.9), the one derived for (EMS-A) in Chapter 4.

We point out that the cases where $\rho_3 = 0$ in germs equations on (ρ_5, u_5) and (ρ'_5, u'_5) can reasonably be neglected for further analysis. Indeed, the first two equations above always ensure $\rho_3 = \pi\rho_1$, and supposing $\rho_3 = 0$ and $\lambda > 0$ leads to $\min(N_A, N_R) = 0$ by combination of the germs that achieve the minimum. Therefore, the throughput ρ_3 is positive as soon as we suppose $\lambda > 0$, $N_A > 0$ and $N_R > 0$, i.e., positive inflow of calls and positive number of agents to pick them up. As a result, when this condition is met, the priority routing that rules tokens from the pool of reservoir assistants does not appear on the affine germs of z_5 and z'_5 anylonger. This is an expected outcome since transitions z_5 and z'_5 (very high and high level of priority for the reservoir pool) can only receive tokens that have passed through transition z_3 (low level of priority for the reservoir pool) before, thus z_5 and z'_5 cannot ultimately inhibit themselves. Such a layout of priorities does remain appropriate to perform arbitration of tokens orientation in case of conflicts and we show below that it still produces effects in the scope of long-run analysis of the system.

As in Section 3.2 of Chapter 4 and as in Section 2 of this current chapter, a choice of policy (minimizing terms in the lexicographic system) provides affine equalities determining the throughput as an affine function of the resources of the model, the validity region of this expression being obtained by inequalities derived from the remaining (non-minimizing) terms of the system. This leads to nine full-dimensional congestion phases (maximal cells of the throughput complex) covering $(\mathbb{R}_{\geq 0})^3$, that we depict and number on Figure 5.7 and whose complete polyhedral form is given in Table 5.11 later, with expressions of the throughputs.

We now describe this congestion diagram. As expected, the introduction of a new type of resource agent (the reservoir assistant) introduces more slowdown phases if its initial marking N_R is too small. Therefore, to ensure the good behaviour of the (EMS-B) model whose design relies substantially on the reservoir, one needs to take $N_R \geq N_R^* := 2\pi\lambda\tau_2$. Note that the minimal number of MRAs (resp. emergency physicians) to answer all the calls is not affected by the presence of the reservoir by comparison with (EMS-A) model, and is still equal to $N_A^* := \lambda(\tau_1 + \pi\tau_2)$ (resp. $N_P^* := \pi\lambda(\tau_2 + \tau_3)$). These three lower bounds on N_A, N_R and N_P define the phase $\boxed{1}$, that we can still refer to as the congestion-free phase.

Compared with (EMS-A), the duplication of the physician's lane and the fact that very urgent calls (in proportion α among all calls transferred to doctors) are handled in priority has the effect

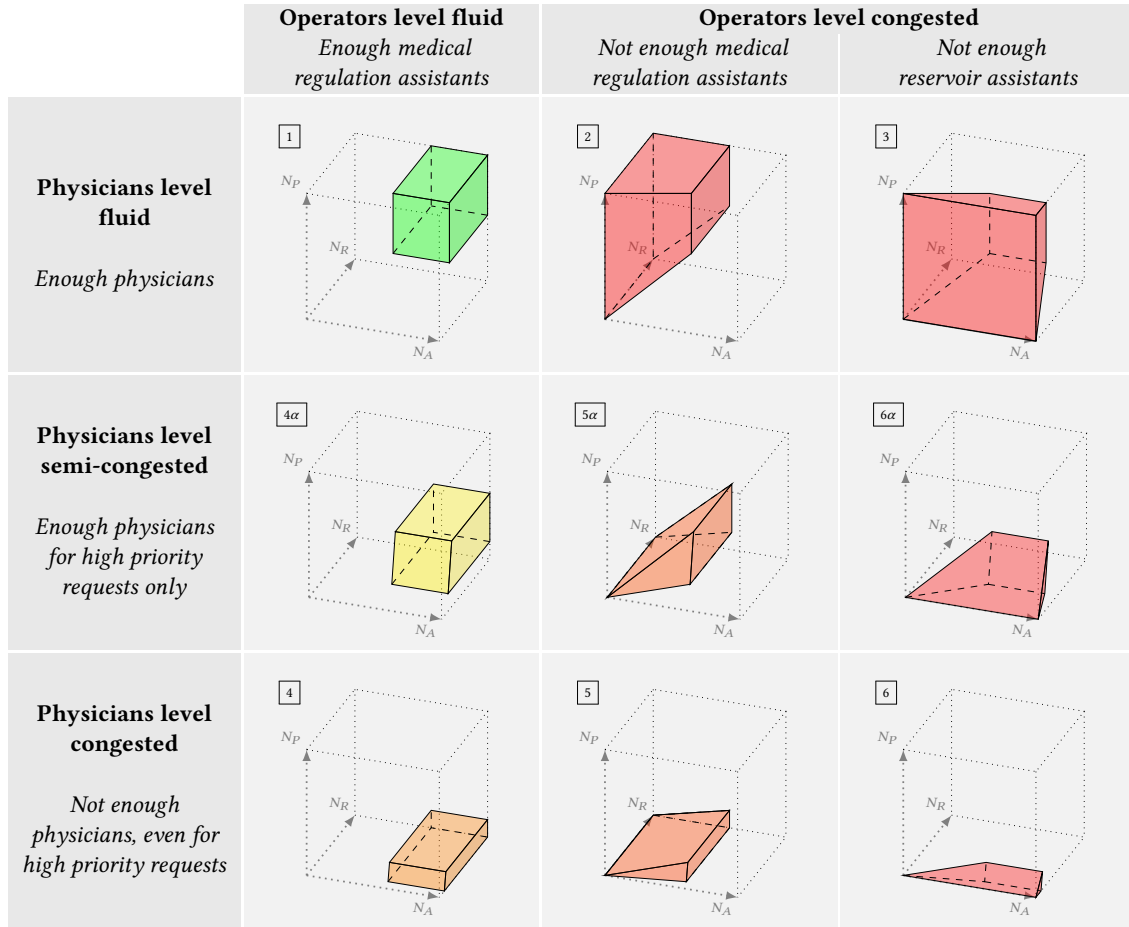


Figure 5.7: Congestion phases of the model (EMS-B)

of splitting each congestion phase associated with a lack of emergency physicians in two parts, exactly like in the model (Bi Lvl). Given an MRA throughput ρ_1 , we can define the two functions \underline{N}_P and \overline{N}_P by

$$\underline{N}_P(\rho_1) := \pi\alpha(\tau_2 + \tau_3)\rho_1 \quad \text{and} \quad \overline{N}_P(\rho_1) := \pi(\tau_2 + \tau_3)\rho_1,$$

so that a minimum number of \overline{N}_P physicians is needed to handle all the calls passed by the MRAs via the reservoir assistant. However, in case of a lack of physicians, the priority mechanism ensures that the very urgent calls remain handled as long as $N_P \geq \underline{N}_P$ (phases 4α , 5α and 6α). Below the latter threshold, there are too few physicians to handle these very urgent calls (phases 4 , 5 and 6), and *a fortiori* the less urgent ones as well.

A major qualitative advantage of the system (EMS-B) is that contrary to the model (EMS-A), we observe that a slowdown in the emergency physician circuit does not affect the throughput of the MRAs, as an effect of their *desynchronization* by the reservoir buffer. It may still happen that we encounter both a lack of MRAs and physicians, but the latter do not prevent the former to pick up inbound calls at their maximal possible throughput (phases 5α and 5 where the throughputs ρ_1 and $\rho_5 + \rho'_5$ are disentangled).

This improvement compared to (EMS-A) can be visualized in terms of phase diagrams, by letting $\alpha = 1$ in the Figure 5.7 so that as in this original model, all the calls requiring the

intervention of the emergency physician have the same degree of priority. This makes the phases 4α , 5α and 6α collapse. The Figure 5.8 represents the four remaining phases 1 , 2 , 4 and 5 where there are enough reservoir assistants; a two-dimensional phase diagram in the space (N_A, N_P) can thus be recovered by taking a cross-section of these polyhedra, with $N_R \geq N_R^*$.

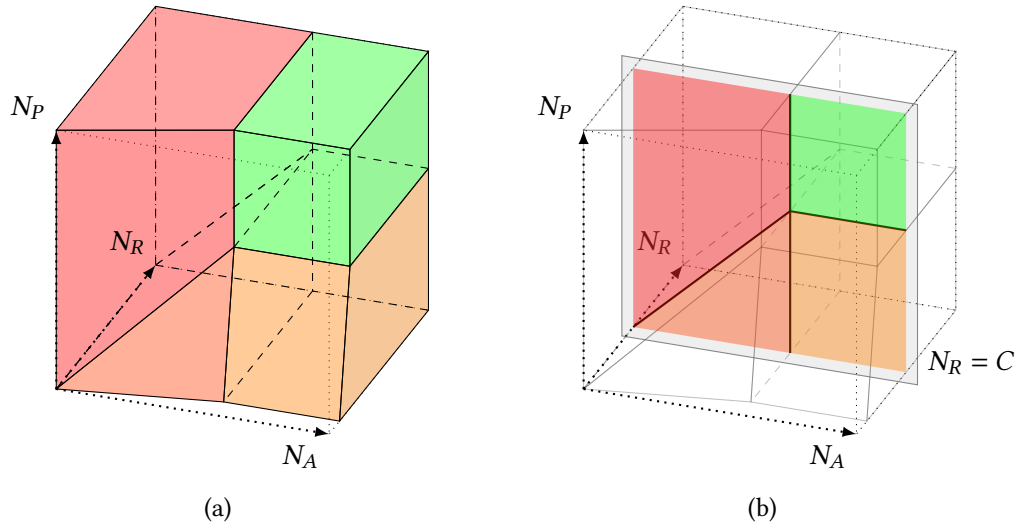


Figure 5.8: Phases of (EMS-B) whenever $\alpha = 1$ and the reservoir is not congested (left); and a cross-section for a constant value of $N_R \geq N_R^*$ (right)

In the resulting congestion diagram depicted in Figure 5.9b, we can verify that the throughput ρ_1 is now indeed independent of the emergency physicians' staffing N_P . Compare with the phase diagram of (EMS-A) first presented in Figure 4.7 and reproduced in Figure 5.9a. In view of these two throughput maps, we can confirm that the reservoir mechanism satisfactorily serves its initial purpose.

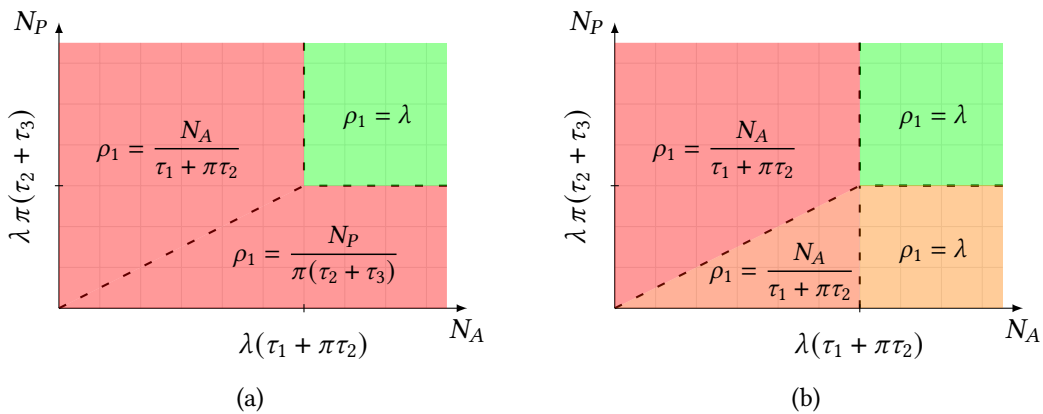


Figure 5.9: The (N_A, N_P) congestion diagrams of (EMS-A) (left) and (EMS-B) (right) compared

It is also instructive to study the situations in which the reservoir assistants are understaffed (phases 3 , 6α and 6), although they are not desirable in practice. We read in particular from the Table 5.11 that the throughputs ρ_3 and ρ_5 are proportional to $N_R/\tau_2 - N_P/(\tau_2 + \tau_3)$ in phase 6α

(it remains true for ρ_3 in phase $\boxed{6}$ as well), which means that increasing the number of emergency physicians slows down the handling of top priority calls! Compare with the model (B1LVL) in Section 2.2, where *only* the throughput with low priority was subject to this effect. This establishes the following seemingly paradoxical property that must be highlighted to insist on the difference with respect to monotonic models:

In presence of priority routings, an increase of the resources may result in a decrease of the asymptotic throughputs of some transitions, even with the highest degree of priority.

We depict this phenomenon on Figure 5.10 (supposing that first-answering MRAs are non-limiting so $N_A \geq N_A^*$). The red curve shows that ρ_5 is nonmonotonic as N_P grows and as we go through phases $\boxed{6}$, then $\boxed{6\alpha}$ and finally $\boxed{3}$. This counter-intuitive situation can be explained as follows. Suppose for the sake of simplicity that $N_R < \pi\lambda\tau_2$, so that there are not enough reservoir assistants to even fill the reservoir room (while twice this amount of agents would be needed to fill it and empty it). At $N_P = 0$, because there are no physicians, both very urgent and urgent calls queues build up in the reservoir at throughput $N_R/\pi\tau_2$ (maximum filling speed of reservoir agents). As N_P increases (phase $\boxed{6}$), some very urgent cases can now be handled by emergency physicians at rate $N_P/(\tau_2 + \tau_3)$, however this task requires a second accompaniment step with reservoir assistants and is prioritized for them. Hence, they spend less time filling the reservoir and ρ_3 decreases. As N_P increases again (phase $\boxed{6\alpha}$), there are enough emergency physicians to also pick up calls from the second-priority room, requiring again the intervention of reservoir assistants (before admitting new patients), as a result ρ_3 decreases again and so does $\rho_5 = \alpha\rho_3$. The reservoir assistants have “less time” to admit and detect very urgent calls as they must escort already admitted very urgent and urgent calls to doctors before, and eventually less very urgent calls are handled.

We insist on the fact that this unusual phenomenon arises because top-priority transitions are served downstream of some inferior-priority ones. Recall that the same type of “diamond” pattern involving priorities already led to difficulties in the Chapter 1, see in particular the discussion about the Figure 1.14. We also note that the above paradoxical principle echoes a similar pathological behaviour observed in [Yeo7].

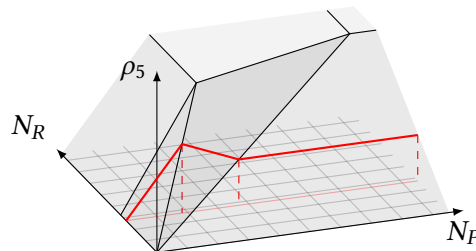


Figure 5.10: The throughput ρ_5 does not increase with respect to N_P

Phase	Bounding inequalities	ρ_1	ρ_5	ρ'_5
1	$\frac{N_A}{\tau_1 + \pi\tau_2} \geq \lambda$; $\frac{N_R}{2\tau_2} \geq \pi\lambda$ $\frac{N_P}{\tau_2 + \tau_3} \geq \pi\lambda$	λ	$\pi\alpha\lambda$	$\pi(1 - \alpha)\lambda$
4 α	$\frac{N_A}{\tau_1 + \pi\tau_2} \geq \lambda$; $\frac{N_R}{\tau_2} \geq \pi\lambda + \frac{N_P}{\tau_2 + \tau_3}$ $\pi\alpha\lambda \leq \frac{N_P}{\tau_2 + \tau_3} \leq \pi\lambda$	λ	$\pi\alpha\lambda$	$\frac{N_P}{\tau_2 + \tau_3} - \pi\alpha\lambda$
4	$\frac{N_A}{\tau_1 + \pi\tau_2} \geq \lambda$; $\frac{N_R}{\tau_2} \geq \pi\lambda + \frac{N_P}{\tau_2 + \tau_3}$ $\frac{N_P}{\tau_2 + \tau_3} \leq \pi\alpha\lambda$	λ	$\frac{N_P}{\tau_2 + \tau_3}$	0
2	$\frac{N_A}{\tau_1 + \pi\tau_2} < \lambda$; $\frac{N_R}{2\tau_2} \geq \frac{\pi N_A}{\tau_1 + \pi\tau_2}$ $\frac{N_P}{\tau_2 + \tau_3} \geq \frac{\pi N_A}{\tau_1 + \pi\tau_2}$	$\frac{N_A}{\tau_1 + \pi\tau_2}$	$\frac{\pi\alpha N_A}{\tau_1 + \pi\tau_2}$	$\frac{\pi(1 - \alpha)N_A}{\tau_1 + \pi\tau_2}$
5 α	$\frac{N_A}{\tau_1 + \pi\tau_2} < \lambda$; $\frac{N_R}{\tau_2} \geq \frac{\pi N_A}{\tau_1 + \pi\tau_2} + \frac{N_P}{\tau_3 + \tau_2}$ $\frac{\pi\alpha N_A}{\tau_1 + \pi\tau_2} \leq \frac{N_P}{\tau_2 + \tau_3} \leq \frac{\pi N_A}{\tau_1 + \pi\tau_2}$	$\frac{N_A}{\tau_1 + \pi\tau_2}$	$\frac{\pi\alpha N_A}{\tau_1 + \pi\tau_2}$	$\frac{N_P}{\tau_2 + \tau_3} - \frac{\pi\alpha N_A}{\tau_1 + \pi\tau_2}$
5	$\frac{N_A}{\tau_1 + \pi\tau_2} < \lambda$; $\frac{N_R}{\tau_2} \geq \frac{\pi N_A}{\tau_1 + \pi\tau_2} + \frac{N_P}{\tau_3 + \tau_2}$ $\frac{N_P}{\tau_3 + \tau_2} \leq \frac{\pi\alpha N_A}{\tau_1 + \pi\tau_2}$	$\frac{N_A}{\tau_1 + \pi\tau_2}$	$\frac{N_P}{\tau_2 + \tau_3}$	0
3	$\frac{N_R}{2\tau_2} < \pi\lambda$; $\frac{N_R}{2\tau_2} \leq \frac{\pi N_A}{\tau_1 + \pi\tau_2}$ $\frac{N_P}{\tau_2 + \tau_3} \geq \frac{N_R}{2\tau_2}$	$\frac{N_R}{2\pi\tau_2}$	$\frac{\alpha N_R}{2\tau_2}$	$\frac{(1 - \alpha)N_R}{2\tau_2}$
6 α	$\frac{N_R}{\tau_2} - \frac{N_P}{\tau_2 + \tau_3} < \pi\lambda$; $\frac{\pi N_A}{\tau_1 + \pi\tau_2} \geq \frac{N_R}{\tau_2} - \frac{N_P}{\tau_2 + \tau_3}$ $\frac{\alpha N_R}{(1 + \alpha)\tau_2} \leq \frac{N_P}{\tau_2 + \tau_3} \leq \frac{N_R}{2\tau_2}$	$\frac{N_R}{\pi\tau_2} - \frac{N_P}{\pi(\tau_2 + \tau_3)}$	$\frac{\alpha N_R}{\tau_2} - \frac{\alpha N_P}{\tau_2 + \tau_3}$	$\frac{(1 + \alpha)N_P}{\tau_2 + \tau_3} - \frac{\alpha N_R}{\tau_2}$
6	$\frac{N_R}{\tau_2} - \frac{N_P}{\tau_2 + \tau_3} < \pi\lambda$; $\frac{\pi N_A}{\tau_1 + \pi\tau_2} \geq \frac{N_R}{\tau_2} - \frac{N_P}{\tau_2 + \tau_3}$ $\frac{N_P}{\tau_2 + \tau_3} \leq \frac{\alpha N_R}{(1 + \alpha)\tau_2}$	$\frac{N_R}{\pi\tau_2} - \frac{N_P}{\pi(\tau_2 + \tau_3)}$	$\frac{N_P}{\tau_2 + \tau_3}$	0

Table 5.11: Bounding inequalities of the nine full-dimensional phases of the throughput complex of the model (EMS-B); along with the expression of the throughputs ρ_1 , ρ_5 and ρ'_5 . The throughput ρ_3 always satisfies $\rho_3 = \pi\rho_1$.

4 Variations of the model EMS-B

In this section, we study how slight modifications of the model (EMS-B) can improve its performance and make it more robust. The three alternative models presented here are first and foremost tools for reflection, and illustrations of our systematic method of analysis for the throughput of non-monotonic Petri nets. From the operational point of a view of managing a medical emergency call center, these new models can come with “qualitative” drawbacks, that we also discuss.

4.1 SUPPORT OF THE RESERVOIR ASSISTANT TO PICK UP INBOUND CALLS. — A first idea to bring more fluidity to the system is to also allow the reservoir assistants to pick up inbound calls in case they are idle. If both a standard MRA and a reservoir assistant are available and an inbound call arrives to the platform, we decide that the MRA should pick it up, and keep the occupancy of the reservoir assistants as low as possible – for them to be ready for their other tasks with higher priority degrees. This transforms the transition z_1 served by MRAs in the model (EMS-B) of Figure 5.6 into two transitions z_1^+ and z_1^- , respectively served by MRAs and reservoir assistants. We illustrate the organization implementing this modification in the Figure 5.12, and for the sake of the simplicity we have not depicted the duplication of the physicians’ lane here.

If a reservoir assistant picks up an inbound call, he/she will handle it with the same characteristic time τ_1 than the MRAs, and a fraction π of these calls shall still require a consultation with the emergency physician. However, no synchronization step here is needed.

We do not write the counter dynamics nor the associated system on germs of affine functions, as they become cumbersome, but we present in the Table 5.13 the twelve full-dimensional phases that the resolution of the throughput problem yields (with the two different levels of priority to speak to the emergency physician).

It can then be seen that the model of Figure 5.12 has the effect of dividing in two each one of the phases $[2]$, $[5\alpha]$ and $[5]$ of the Figure 5.7 that were valid for the model (EMS-B). Within the newly created phases $[2']$, $[5'\alpha]$ and $[5']$, there are not enough MRAs to absorb all the inbound flow of calls and reservoir assistants come in support to fill the breach, without altering their other tasks. In the phases $[2]$, $[5\alpha]$ and $[5]$ however, there are already not enough reservoir assistants to handle the monitored reservoir and as a consequence no idle resource to lend a hand to MRAs.

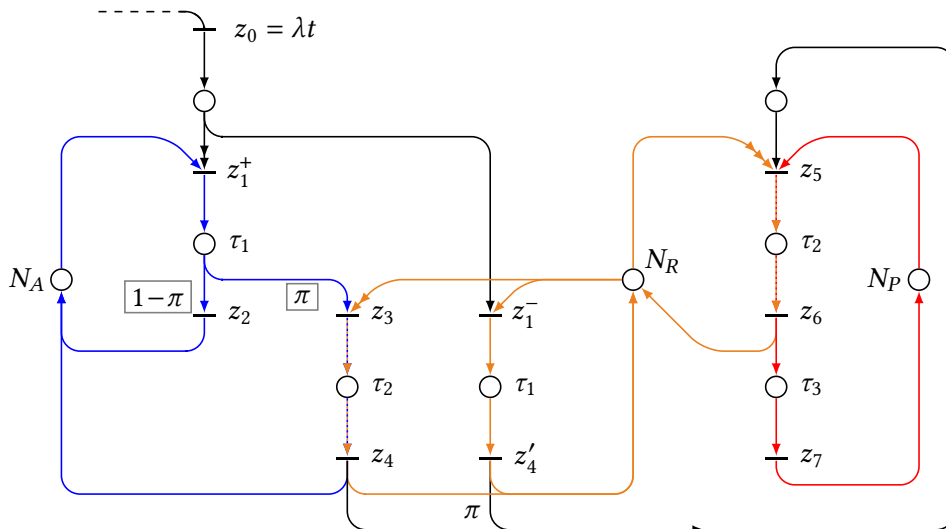


Figure 5.12: (EMS-B) model in which reservoir agents can pick up inbound calls

	Operators level fluid		Operators level congested	
	Enough medical regulation assistants	MRAs supported by the reservoir assistants	Not enough medical regulation assistants	Not enough reservoir assistants
Physician level fluid				
Physician level semi-congested				
Physician level congested				

Table 5.13: Congestion phases of the model (EMS-B) where the reservoir assistants can also pick up inbound calls

Granting reservoir assistants the right to pick up inbound calls therefore introduces a **second congestion-free phase**, since the system is fluid in both phases $\boxed{1}$ and $\boxed{2'}$. Another way to see it is to consider that the congestion-free regime has been extended. We can still take a cross-section of the complete throughput diagram to recover a two-dimensional (N_A, N_P) diagram, as done in Section 3.2. However, unlike the original (EMS-B) model, the shape of some phases change depending on $N_R \geq N_R^* = 2\pi\lambda\tau_2$.

We depict in the Figure 5.14a only the “back” phases $\boxed{1}$, $\boxed{2'}$, $\boxed{4}$ and $\boxed{5'}$ in the simplified case where $\alpha = 1$. These four phases are those for which the “operator level” is fluid (with or without the support of reservoir assistants) and no inbound call is handled with delay. The complete equation of the front surface in Figure 5.14a, that separates the four depicted phases from the other phases $\boxed{2}$, $\boxed{3}$, $\boxed{5}$ and $\boxed{6}$, writes

$$N_R - \pi(\rho_A \wedge \lambda)\tau_2 - (\rho_P \wedge (\pi\lambda))\tau_2 \geq (\lambda - (\rho_A \wedge \lambda))\tau_1. \quad (5.7)$$

This inequality describes the fact that there are enough **remaining** reservoir agents, after welcoming in the reservoir the calls handed over by the MRAs (this consumes $\pi(\rho_A \wedge \lambda)\tau_2$ reservoir resource) and escorting these calls out of the reservoir to the doctors (this consumes another $(\rho_P \wedge (\pi\lambda))\tau_2$ reservoir assistants) to treat the inbound calls not picked up by the MRAs, with rate $\lambda - (\rho_A \wedge \lambda)$ and average handling time τ_1 .

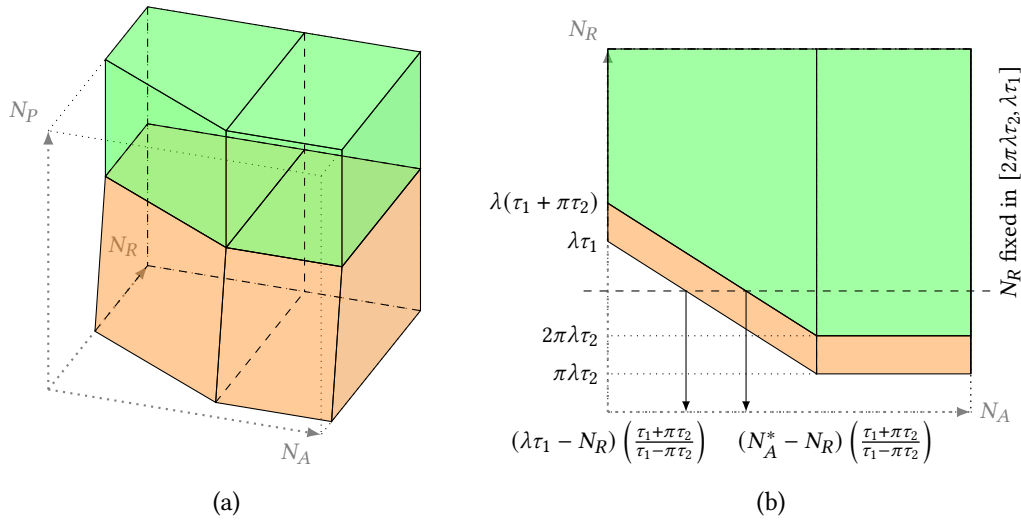


Figure 5.14: Left: The phases of Table 5.13 for which the operator level is fluid whenever $\alpha = 1$. Right: The same phases seen from above with a cross-section for N_R fixed

The Figure 5.14b shows the same phases than Figure 5.14a but seen from above. Since in practice $\pi\tau_2$ is a short duration (less than 10 s) and τ_1 is relatively high (around two minutes), we shall have $\pi\lambda\tau_2 \leq 2\pi\lambda\tau_2 \ll \lambda\tau_1 \leq N_A^* = \lambda(\tau_1 + \pi\tau_2)$, and it is reasonable to take a cross-section of the throughput diagram with a value of N_R in $[N_R^*, \lambda\tau_1]$. We obtain using the equation (5.7) the (non-trivial) values of N_A delineating the phases created by our slicing operation. This 2D map is shown in Figure 5.15.

Remark how there are now two congestion-free phases in Figure 5.15. Although it happens in a very degraded mode, note also that the point $(N_A = 0, N_P = 0)$ is no longer in the frontier of the phases [2] and [5], which was bound to be the case for the original model (EMS-B). By comparison with the latter, allowing the reservoir assistants to pick up inbound calls indeed enable the emergency physicians to treat some requests even if $N_A = 0$!

An interesting computation to make is to check if with the same total number of MRAs, the model of Figure 5.12 does better than the model (EMS-A) or not. It is rather intuitive that a staffing (N_A, N_R, N_P) in a model with the monitored reservoir will do better than the organization (EMS-A) with a staffing (N_A, N_P) , because the former uses strictly more regulation assistants than the latter – $N_A + N_R$ versus N_A . Since a natural choice for the model with reservoir is $N_R = 1$, we can rather compare the extended (EMS-B) model with $(N_A - 1, 1, N_P)$ agents – where $N_A > 1$ – with the original (EMS-A) model with (N_A, N_P) agents. Both systems are equivalent in the intersection of the fluid phases, which happens if $N_A \geq N_A^* + (2\pi\tau_2 / (\tau_1 - \pi\tau_2))(N_A^* - 1)$ and $N_P \geq N_P^*$. Outside this fluide regime, we can show that it is more interesting to adopt the model with extended reservoir (staffed with one agent) if and only if

$$\frac{N_P}{\pi(\tau_2 + \tau_3)} \leq \min \left(\frac{1}{\tau_1 + \pi\tau_2} + \left(\frac{\tau_1 - \pi\tau_2}{\tau_1 + \pi\tau_2} \right) \frac{N_A - 1}{\tau_1 + \pi\tau_2}, \lambda \right). \quad (5.8)$$

We verify that due to the small value of $\pi\tau_2$ in practice, the above condition is very close from the boundary $\rho_P \leq \min(\rho_A, \pi\lambda)$ of (EMS-A). As a result, when the model (EMS-A) is slowed down by its MRAs, taking off one of them to adopt the reservoir model decreases the overall performance (however, with real-life parameters, the performance loss is no more than 3 %). Conversely, if the model (EMS-A) is slowed down due to a lack of physicians, implementing the

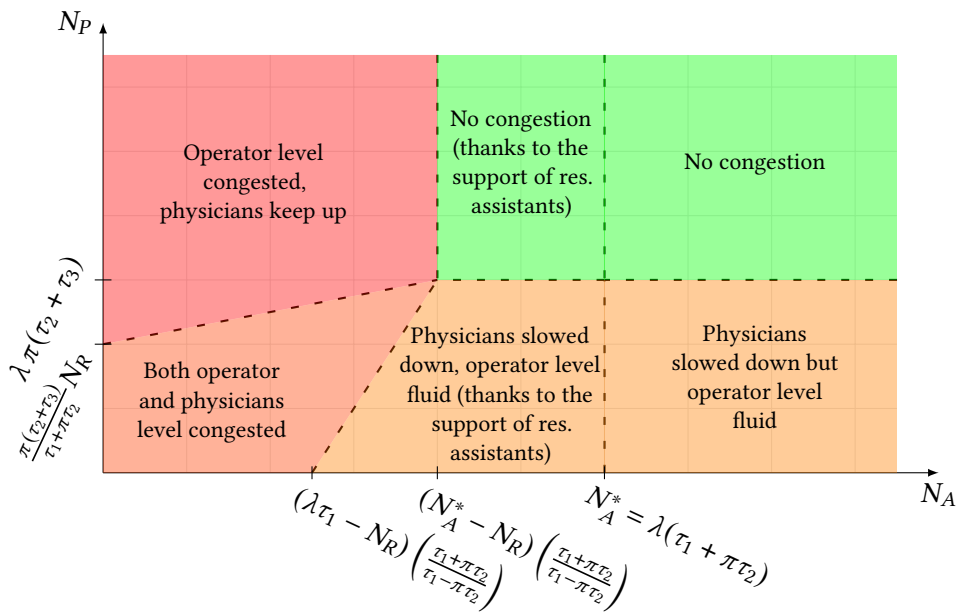


Figure 5.15: Phases diagram (N_A, N_P) of the modified system (EMS-B), where the reservoir assistants can support the MRAs

(extended) reservoir strategy by asking one classical MRA to handle the reservoir and possibly extra inbound calls does improve the throughput of the system (up to 10 % with our realistic parameters). This is consistent with the initial *raison d'être* of the reservoir mechanism. Note that although we did not write it explicitly to simplify the discussion, numerical evaluations in practice of minimum staffings – or equations of type (5.8) – need to also feature the occupancy rates of the agents, as introduced in the Section 6 of Chapter 4. This can be done by replacing each term N_X by a ratio N_X/η_X , where η_X is the target occupancy rate of the pool of agents X .

4.2 ASSIGNING ANOTHER EXTRA FLOW TO THE RESERVOIR ASSISTANTS. – In practice, the previous system was not considered to be a desirable design by the emergency doctors we work with. Indeed, even though it extends the congestion-free regime, it puts more weight on the reservoir assistant. As the centrepiece of the organization, this agent should on the contrary remain as much available as possible to manage the reservoir. In fact, because they work hand in hand with the emergency physicians (to indicate them which pending calls ought to be treated first), it is also the mission of reservoir assistants to realize other tasks on behalf of the doctors, like handling urgent calls transferred by partner services, or dispatching a mobile intensive care unit, etc.

To represent these extra tasks, that are different from the management of the reservoir, we introduce a new exogenous flow with rate μ , accounting either for partner services' inbound calls to pick up or outbound calls to make. In order to protect the usual calls of the population, among which there can be very urgent needs that are not yet acknowledged by an emergency physician, the extra flow of tasks has to be handled by the reservoir assistant only if no entrance or exit task is pending in the reservoir. The Figure 5.16 depicts a Petri net implementing this modification (to keep things simple, we let aside here the distinction of calls prioritized to physicians). Note that some of these extra tasks performed by the reservoir assistants may further need the intervention

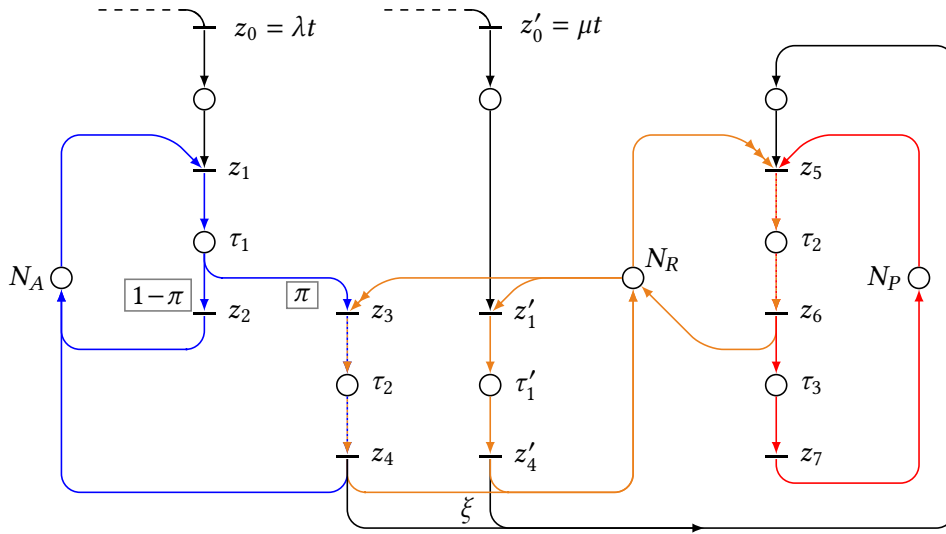


Figure 5.16: Medical emergency call center with a monitored reservoir (EMS-B) and where reservoir assistants can treat an extra flow of calls

of the emergency doctor (for instance, medical calls transferred by firefighters). We denote by ξ the fraction of the new flow which adds to the doctors' jobs.

The Table 5.17 depicts the ten full-dimensional phases of this new system. The six phases $\boxed{1}$ – $\boxed{6}$ initially present in the Figure 5.7 for the model (EMS-B) in which $\alpha = 1$ are still there. The phases $\boxed{1}$, $\boxed{2}$, $\boxed{4}$ and $\boxed{5}$ correspond to a sufficient amount of reservoir assistants to achieve all their tasks. The phases $\boxed{3}$ and $\boxed{6}$ stand for an insufficient amount of them to manage the mere reservoir itself. In addition, the four intermediary (or buffer) phases $\boxed{1\mu}$, $\boxed{2\mu}$, $\boxed{4\mu}$ and $\boxed{5\mu}$ have appeared along the dimension N_R , indicating when the extra flow is only partially handled. This phenomenon of buffer phases is similar to the one encountered along the dimension N_P when the emergency physicians could handle two types of flows with different priority degrees.

It is important to note that the phase $\boxed{1}$ in this new model is congestion-free in the sense that not only the usual calls from the general public are handled without delay at all steps, but also the new extra flow picked up by reservoir assistants. In particular, because a fraction ξ of this flow has to go through the physicians as well (with no distinguished priority), the greater the new flow μ is, the more physicians are needed. Computing the equation of the frontier between phases $\boxed{1}$ and $\boxed{2}$ on one hand, and phases $\boxed{4}$ and $\boxed{5}$ on the other hand provides

$$\rho_P \geq \pi(\rho_A \wedge \lambda) + \xi\mu, \quad (5.9)$$

where as introduced in the Section 3.2 of Chapter 4, ρ_A (resp. ρ_P) stands for the MRA intrinsic throughput $N_A/(\tau_1 + \pi\tau_2)$ (resp. the physician intrinsic throughput $N_P/(\tau_2 + \tau_3)$). We recognize with the first term that there must be enough emergency doctors to handle the usual calls of the population first picked up by the MRAs and transferred by the reservoir assistants (whose constraints does not appear here because it is not active within these specific phases). The second term corresponds to the number of extra tasks per unit of time requiring a discussion with the emergency doctor. Note again that even if $N_A = 0$, so that no usual call is picked up, the physicians can still treat the demands brought by the reservoir assistants.

In order to evaluate the maximum number of tasks that the reservoir assistants can handle, we must first fix a staffing N_R of reservoir agents that is such that $N_R > 2\pi\lambda\tau_2$, as found in Section 3.2,

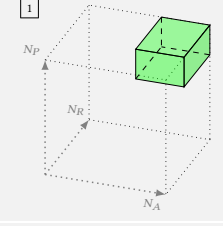
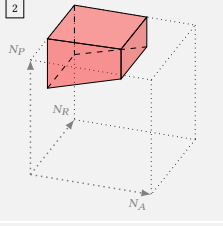
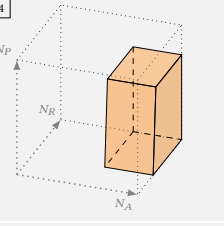
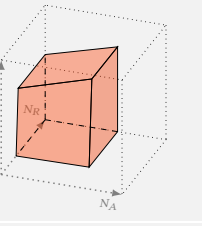
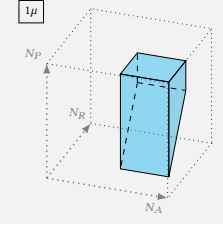
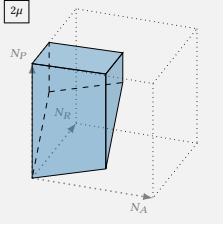
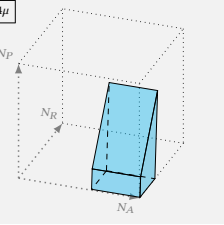
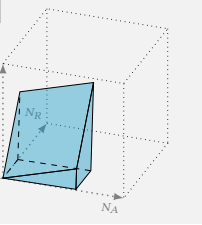
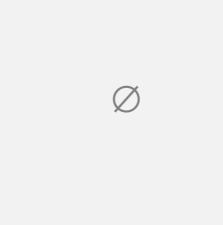
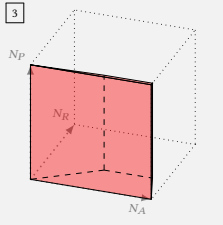
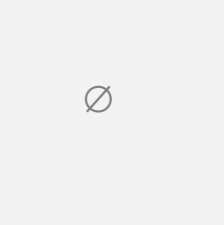
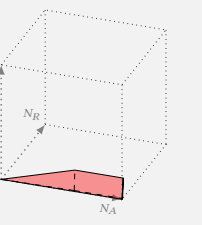
	Physician level fluid		Physician level congested	
	MRA level fluid	MRA level congested	MRA level fluid	MRA level congested
Reservoir fluid and extra flow handled				
Reservoir fluid but extra flow not fully handled				
Reservoir congested, extra flow not picked-up				

Table 5.17: Congestion phases of the model (EMS-B) where the reservoir assistants can also treat an extra flow of calls

and start with a null extra flow $\mu = 0$. This ensures that we are located in one of the four phases $\boxed{1}$, $\boxed{2}$, $\boxed{4}$ or $\boxed{5}$ of the Table 5.17 (because $\mu = 0$, the phases $\boxed{1\mu}$, $\boxed{2\mu}$, $\boxed{4\mu}$ and $\boxed{5\mu}$ are flat). These are also the phases $\boxed{1}$, $\boxed{2}$, $\boxed{4}$ and $\boxed{5}$ of the model (EMS-B) in which no extra flow was considered, as depicted in Figure 5.7. Now, for any staffing choice (N_A, N_P) , we can start increasing the parameter μ , which make the phases $\boxed{1\mu}$, $\boxed{2\mu}$, $\boxed{4\mu}$ and $\boxed{5\mu}$ grow thicker and thicker. For some value μ_{\max} of μ , the point (N_A, N_R, N_P) becomes overwhelmed by one of the blue phases, which is therefore the maximum extra flow intensity that the reservoir assistants can handle before putting the system in a regime where these extra tasks are only partially handled. We illustrate this idea on the Figure 5.18, for two different values of μ . Think of the blue phases as a tidal wave that is going to submerge the fixed points M and M' as μ increases, and as the distance to the wavefront as the remaining margin of calls that it is possible to take. We let the reader verify that to be within the congestion-free phase $\boxed{1}$, more and more emergency physicians are needed as μ increases, consistently with (5.9).

Computing the equation of the “wavefront” of Figure 5.18 enables us to determine that for any fixed value of $N_R > 2\pi\lambda\tau_2$ and any (N_A, N_P) in $\mathbb{R}_{\geq 0}^2$, we have

$$\mu_{\max} = \max\left(\frac{N_R - 2\pi(\rho_A \wedge \lambda)\tau_2}{\tau_1 + \xi\tau_2}, \frac{N_R - (\pi(\rho_A \wedge \lambda)\tau_2 + \rho_P\tau_2)}{\tau_1}\right). \quad (5.10)$$

Note how the two main terms of the right-hand side in (5.10) do not have the same denominator. Indeed, the first one corresponds to a frontier of phases $\boxed{1}$ and $\boxed{2}$, where there are sufficient

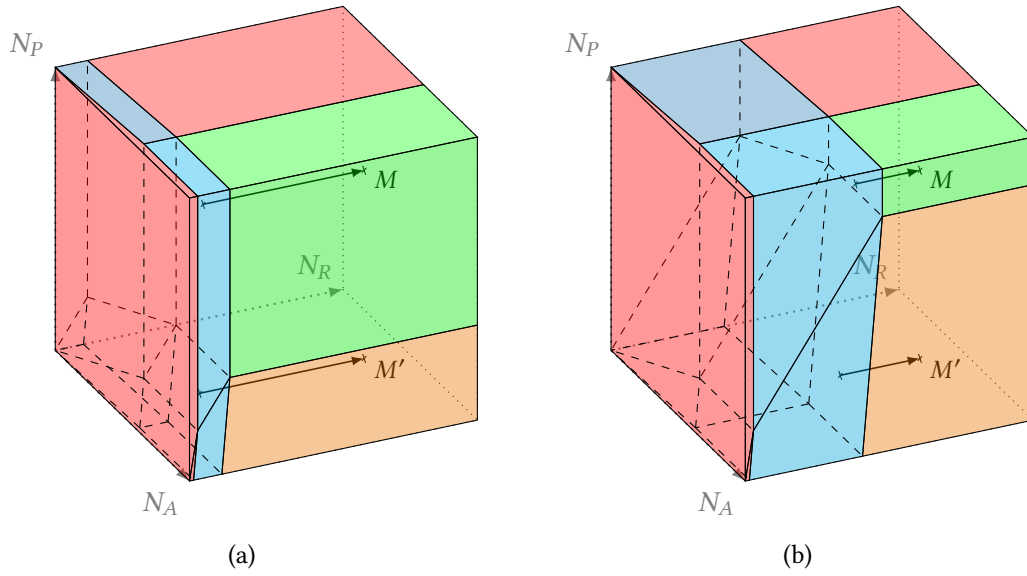


Figure 5.18: Visualization of the evolution of the throughput complex of Table 5.17 when μ increases, going from a small value (left) to a higher value (right). The distance from the blue phases to two arbitrary points M and M' is shown, accounting for a remaining quantity of manageable extra calls.

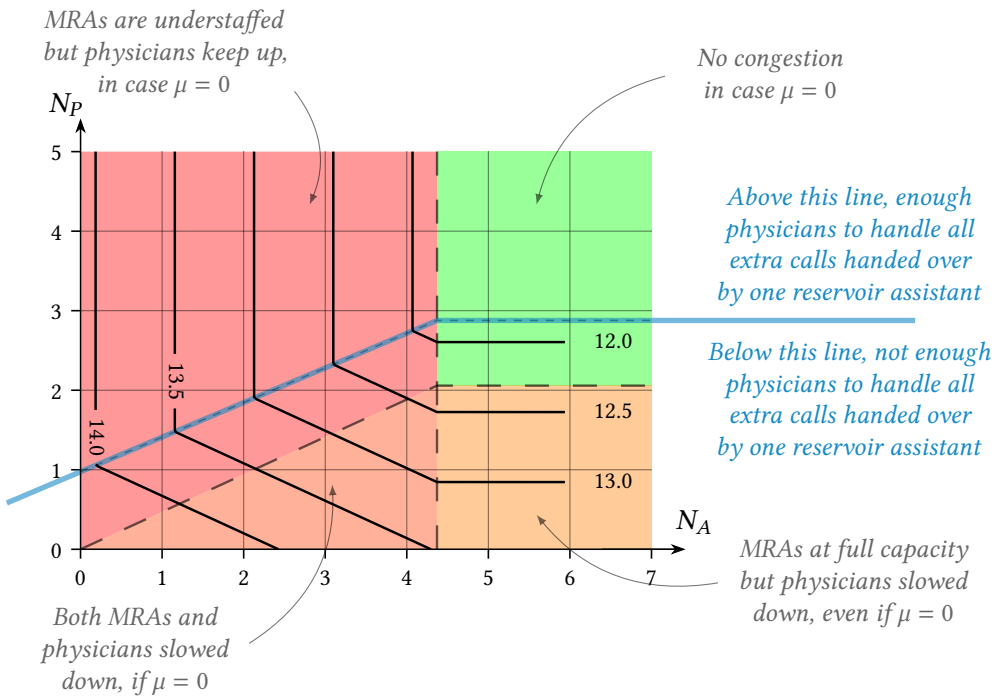


Figure 5.19: Maximum absorbable extra flow of calls by a single reservoir assistant ($N_R = 1$), in calls per hour

emergency physicians to take the extra tasks handed over by the reservoir assistants, therefore the latter handle each extra task with a mean time of $\tau_1 + \xi\tau_2$ (the initial time τ_1 is mandatory, and the transfer time τ_2 occurs for a fraction ξ of the calls). However, if there are not enough emergency doctors (which corresponds to the second term of (5.10) and a frontier of phases [4] and [5]), then everything happens as if there were $\pi(\rho_A \wedge \lambda)\tau_2 + \rho_P\tau_2$ reservoir agents managing entries and exits in the reservoir (with an average time τ_2 to treat each task). This already overwhelms the doctors, so a remaining number $N_R - \pi(\rho_A \wedge \lambda)\tau_2 + \rho_P\tau_2$ can effectively pick up calls from the extra flow but the synchronization with physicians does not happen, hence the mean treatment time of τ_1 .

The Figure 5.19 provides the level sets of the map μ_{\max} satisfying (5.10) in the space (N_A, N_P) , for the cross-section $N_R = 1$. We have superimposed the four congestion phases of (EMS-B) already depicted in Figure 5.9b (that are also those of the current system with $\mu = 0$). Indeed, the reasonable question to ask, starting from the model (EMS-B), is the following. Given our staffing (N_A, N_P) , if we were to ask our reservoir assistants to pick up and handle extra tasks, how many of them could they take? Interestingly enough, since the evolution of the wavefront in Figure 5.18 is not a mere translation along the N_R axis due to the dependence of the N_P frontier (5.9) upon μ , the break of slopes in μ_{\max} and hence those of its level sets are not “lined up” with the original phases of (EMS-B). Instead, we must depict the frontier associated with the inequality

$$\rho_P \geq \pi(\rho_A \wedge \lambda) + \frac{\xi(N_R - 2\pi(\rho_A \wedge \lambda)\tau_2)}{\tau_1 + \xi\tau_2}, \quad \text{with } N_R = 1, \quad (5.11)$$

that is the inequality (5.9) in which μ has been replaced by $(N_R - 2\pi(\rho_A \wedge \lambda)\tau_2)/(\tau_1 + \xi\tau_2)$ (with $N_R = 1$), that is to say the maximum number of extra tasks that a single reservoir agent can handle after first ensuring fluidity of the reservoir. Had we taken $\xi = 0$, the N_P minimal staffing would have been independent of μ , and the level sets of μ_{\max} would have completed the phases of (EMS-B) into a neat spider web (as the one arising in Figure 5.19 once the blue frontier is added).

The Figure 5.19 actually uses realistic values of the parameters for a medical emergency call center of the Paris area, and it also takes into account the occupancy rates of the agents. It numerically reads as follows. Taking 5 MRAs, 1 reservoir assistant and 3 emergency physicians guarantee that the system is congestion-free for the chosen value of λ (here 57 calls per hour). In addition, the single reservoir assistant can also handle about 12 extra tasks per hour (given $\tau_1 = 98.3$ s, $\tau_2 = 16.8$ s and $\xi = 23$ %). As N_A or N_P decrease, the reservoir assistant can actually treat more and more extra tasks, because he/she is less busy with managing the reservoir...which is not desirable in practice! This figure illustrates the practical interest of visualizations based on the congestion diagram, to compare it with certain notions of security margins (for instance the extra flow we have studied here).

4.3 MUTUAL SUPPORT OF TWO FIRST-ANSWERING AGENT POOLS. — We finish this section by discussing a last alternative feature for the model (EMS-B), that could actually also be applied for (EMS-A) because it does not revolve around the reservoir. Instead, it is a different mechanism to welcome inbound calls to the platform, where several flows are at stake. Indeed, medical emergency call centers also receive calls that do not come from the **primary flow** dialed by the general public. Most of time, this **secondary flow** is made of calls transferred by partners (traffic agents, underground or rail services), other emergency units (like the firefighters in France) but also phone reports passed by teams dispatched on the ground. In some French departments, the SAMU call centers also centralize non-urgent out-of-hours services phone calls.

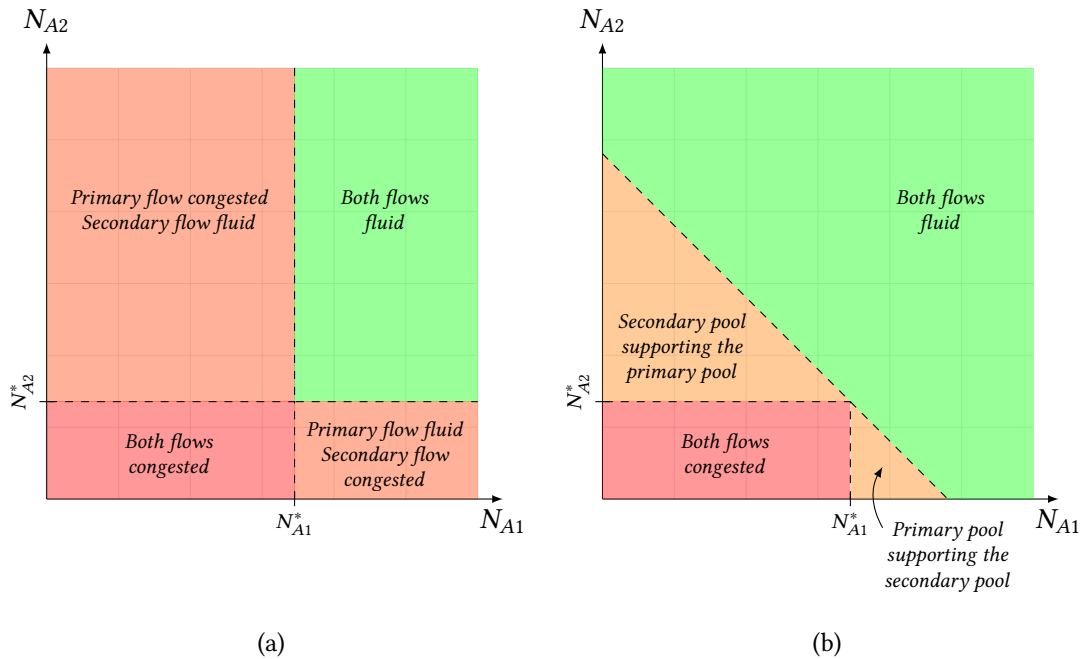


Figure 5.21: The MRA phase diagram of a call center without mutual support (left) and with mutual support (right)

Using the Table 5.1 and the Theorem 5.1 causes no particular difficulty but still yields a pretty cumbersome system that we do not write here. We represent directly on the Figure 5.21b the effect of the mutual support mechanism in the space (N_{A1}, N_{A2}) , which is a two-dimensional cross-section (taking N_R and N_P large enough) of a global four-dimensional throughput complex. We compare this design with a simpler system where no mutual support is allowed (agents of each pool can pick up only their main flow), and thus no priority rules are used; its phase diagram is depicted in Figure 5.21a.

When no mutual support of the two pools of first-answering agents is possible, the Figure 5.21a shows that the arrivals circuit of the organization is congestion-free if and only if headcounts are such that $N_{A1} \geq N_{A1}^* := \lambda(\tau_1 + \pi\tau_2)$ and $N_{A2} \geq N_{A2}^* := \lambda'(\tau'_1 + \pi'\tau'_2)$. Outside this area, either one of the two flows is handled with no congestion while the other one is slowed down, or both of them are congested with maximal handling speeds $\rho_1 = N_{A1}/(\tau_1 + \pi\tau_2)$ and $\rho_2 = N_{A2}/(\tau'_1 + \pi'\tau'_2)$.

However, as shown by the Figure 5.21b, resorting to the mutual support mechanism allows to greatly extend the domain of the congestion-free phase, as soon as $N_{A1} + N_{A2} \geq N_{A1}^* + N_{A2}^*$. In fact, it suffices to have $N_{A1} \geq N_{A1}^*$ or $N_{A2} \geq N_{A2}^*$ for the support of one pool to the other one to switch on. In these phases **the global performance is strictly improved** compared to the organization in silos. This is consistent with the intuition that merging the resources brings more fluidity. We also observe that such priority schemes can in effect transform “and” conditions into “or” conditions.

6

POSYNOMIAL SYSTEMS AND UNDERLYING GEOMETRIC PROBLEMS

Contents

1	Real and tropical posynomial systems	220
1.1	Definition and illustrations	220
1.2	Solving posynomial systems is NP-hard	222
1.3	A linear programming approach for tropical posynomial systems	223
1.4	The relationship between tropical posynomial systems and Markov decision processes	226
1.5	Geometric programming approach of posynomials systems	230
2	Properties of the colorful interior of convex sets	232
2.1	Structure of the colorful interior of pointed cones	233
2.2	Nonvacuity conditions	235

This chapter takes a more abstract and algebraical look on the objects studied in the rest of Part I. We focus here on posynomials, that is to say multivariate polynomials on positive variables but with arbitrary exponents. These generalized polynomials can be considered on any semiring, and we shall in fact study both classical and tropical posynomials – in order words whether the ground semiring $(\mathbb{R}, +, \times)$ or $(\mathbb{R} \cup \{-\infty\}, \max, +)$ is involved. Posynomials arise in convex optimization, especially in geometric and entropic programming [CS16] and in polynomial optimization [DIdW17]. They also arise in the theory of nonnegative tensors [Lim05, FG18], in risk sensitive control [AB17] and game theory [AGGCG19].

Our main motivation to study posynomials comes from the fact that several piecewise affine mappings featured in the previous chapters, whose fixed-points are estimates of interest of Petri net models (like the throughput or the catch-up delays), have in fact tropical posynomial components. Understanding posynomial systems and figuring out resolution methods is therefore a way to gain insight on the underlying structure of piecewise affine dynamical systems, even

though we shall only work here on “static” problems derived from these time-delay dynamics. To our knowledge, maxima or minima of finitely many affine functions were not studied before under the tropical posynomial perspective.

In Section 1.2, we first show that solving (square) posynomial systems is in general NP-hard, for the classical instances but for the tropical ones as well. Then, we identify a tractable subclass. We show in Section 1.3 that the tropical version can be solved exactly in polynomial time by reduction to a linear program, whereas the classical version can be solved approximately by reduction to a geometric program (Section 1.5). Our approach is based on a notion of colorful vector relatively to the Newton polytopes of the posynomial exponents. We show that the class of Markov decision processes is a special case of tropical posynomial systems, and we recover the linear programming formulation of the associated value problem – and the geometric programming formulation of risk sensitive problems as well. In Section 2, we focus on the subproblem of characterizing the set of colorful vectors relatively to a collection of convex bodies, that is somehow of independent interest. We explain how this question hinges on notions of tangent hyperplanes to several convex sets, and we provide a necessary nonvacuity condition for the set of colorful vectors, that we also show to be sufficient up to dimension 3.

1 Real and tropical posynomial systems

In this section, we study systems of equations that abstract and generalize some of the problems met in Chapters 3 and 4.

- 1.1 DEFINITION AND ILLUSTRATIONS.** — We first introduce the main objects of this section. A **posynomial** is a function P of the form

$$P(x) = \sum_{a \in A} c_a x_1^{a_1} x_2^{a_2} \cdots x_n^{a_n} \quad (6.1)$$

where the variable $x := (x_1, \dots, x_n)$ is a vector with real positive entries, A is a finite subset of vectors of \mathbb{R}^n , and the $(c_a)_{a \in A}$ are positive real numbers. The set A is called the **support** of P , also denoted by S_P , its elements are called the **exponents** of the posynomial and the $(c_a)_{a \in A}$ its **coefficients**. Posynomials generalize polynomials given that the exponents are not necessarily nonnegative integers, but as a result one needs to consider only positive variables. Note that one talks about *signomials* if in addition the $(c_a)_{a \in A}$ can take arbitrary real values.

We are interested in solving (square) **posynomial systems**, that is, given n classical posynomials P_1, \dots, P_n , finding x in $\mathbb{R}_{>0}^n$ such that

$$P_i(x) = 1 \quad \text{for all } i \text{ in } \{1, \dots, n\}. \quad (6.2)$$

The form of the equation (6.2) can be thought of as the normalization of a two-sided system that writes $P_i(x) = Q_i(x)$ and where Q_i is a monomial (i.e., a posynomial for which $|A| = 1$). Indeed, dividing both sides of the i -th equation still gives a posynomial on the left-hand side, and the “unit” posynomial 1 on the right-hand side.

Example 6.1. Many usual curves that are beloved by mathematicians can be represented by posynomial equations. For instance, the lemniscate of Bernoulli with polar equation $\rho^2 = a^2 \cos(2\theta)$ also admits the Cartesian form $(x^2 + y^2)^2 = a^2(x^2 - y^2)$, that we can rewrite as the equality of two posynomials $x^4 + 2x^2y^2 + y^4 + a^2y^2 = a^2x^2$, or equivalently in normalized

form $a^{-2}x^2 + 2a^{-2}y^2 + a^{-2}x^{-2}y^4 + x^{-2}y^2 = 1$. We recognize in the left-hand side a posynomial with exponents $(2, 0)$, $(0, 2)$, $(-2, 4)$ and $(-2, 2)$, and respective coefficients a^{-2} , $2a^{-2}$, a^{-2} and 1. Similarly, the usual astroid has the Cartesian equation $\sqrt[3]{x^2} + \sqrt[3]{y^2} = b$, which after normalization gives a posynomial with exponents $(2/3, 0)$ and $(0, 2/3)$, with same coefficients b^{-1} .

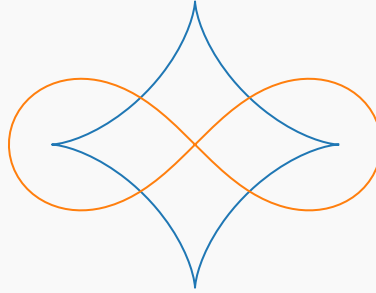


Figure 6.1: Two notorious curves that can be represented by posynomial

The Figure 6.1 shows the level-sets associated with the value 1 for our two posynomials (that are only defined over $\mathbb{R}_{>0}^n$ so we have here mirrored the positive orthant). The intersections of the lemniscate and the astroid can thus be thought of as solutions of a posynomial system on two variables.

Placing ourselves in the max-plus (or tropical) semifield $\mathbb{R}_{\max} := (\mathbb{R} \cup \{-\infty\}, \max, +)$ where the maximum (resp. the usual sum) plays the role of the additive law (resp. the multiplicative law), we can bring up a similar notion. We define a **tropical posynomial** to be a function P^{trop} of the form

$$P^{\text{trop}}(x) = \max_{a \in A} (c_a + \langle a, x \rangle) \quad (6.3)$$

where $\langle \cdot, \cdot \rangle$ is the usual dot product of \mathbb{R}^n , the $(c_a)_{a \in A}$ are now real coefficients, and $x = (x_1, \dots, x_n)$ is a vector of \mathbb{R}^n . Compared with (6.1), the usual sum has been replaced in (6.3) by the maximum, the usual product by the sum, and the usual exponentiation by the multiplication by a scalar – consistently with the max-plus structure. The fact that the $(c_a)_{a \in A}$ are here arbitrary real numbers is the correct counterpart of the positivity constraint on the coefficients of real posynomials, since the additive neutral of \mathbb{R}_{\max} is $-\infty$. Hence, the inequalities $c_a > -\infty$ for all a in A play the role of positivity constraints (in other words, there are no negative numbers in \mathbb{R}_{\max}).

We shall also study tropical posynomial systems, which means that given the n tropical posynomials $(P_i^{\text{trop}})_{1 \leq i \leq n}$, we will look for x in \mathbb{R}^n such that

$$P_i^{\text{trop}}(x) = 0 \quad \text{for all } i \text{ in } \{1, \dots, n\}. \quad (6.4)$$

This is also a normalized form of the two-sided system $P_i^{\text{trop}}(x) = Q_i^{\text{trop}}(x)$ in case Q_i^{trop} is a monomial. Our interest in posynomial systems is in fact motivated by the latter form. Indeed, the counter equations of Table 2.3 featured in Chapter 2 for Petri nets or the dynamic programming equations (DP) that we have met for semi-Markov decision processes in Chapter 3 reveal that the trajectories of monotonic Petri nets or SMDPs arise as fixed-points of min-plus systems with delays. Focusing on the average-cost or the throughput yield optimality systems of the form (L), that are square systems involving tropical posynomials in the right-hand side and simple monomials in the left-hand side (however these systems are obtained over a field of real tuples).

As a result, studying tropical posynomial systems is a manner to gain insight on the problem of computing the throughput of Petri nets (or the average-cost vector of SMDPs).

The definition (6.3) shows that tropical posynomials are nothing more than maximum of finitely many affine functions, and thus are convex piecewise-linear functions. Yet, thinking in terms of posynomials is fruitful from the algebraic point of view, and is indicative of the difficulty of the problem we deal with. Typically, the fact that our optimality systems of the form (L) involve posynomial functions and not mere linear maps (in the tropical sense) is a direct consequence of the use of proportional routings in the context of Petri nets, or the use of stochastic moves in the context of SMDPs. In this sense, it translates the increased complexity of these models compared to deterministic cases (see our discussion on event graphs in the introduction of Chapter 4).

Example 6.2. We illustrate in the Figure 6.2 a graphical resolution of a tropical posynomial system on two variables x_1 and x_2 . We consider a first posynomial P_1 defined by $P_1(x_1, x_2) := \max(-2x_1 + x_2 + 1, -3x_1 + 4x_2 + 1)$, with exponents $(-2, 1)$ and $(-3, 4)$, and a second polynomial $P_2(x_1, x_2) := \max(2x_1 + 4x_2 - 5, 3x_1 + 2x_2 - 6)$ with exponents $(2, 4)$ and $(3, 2)$.

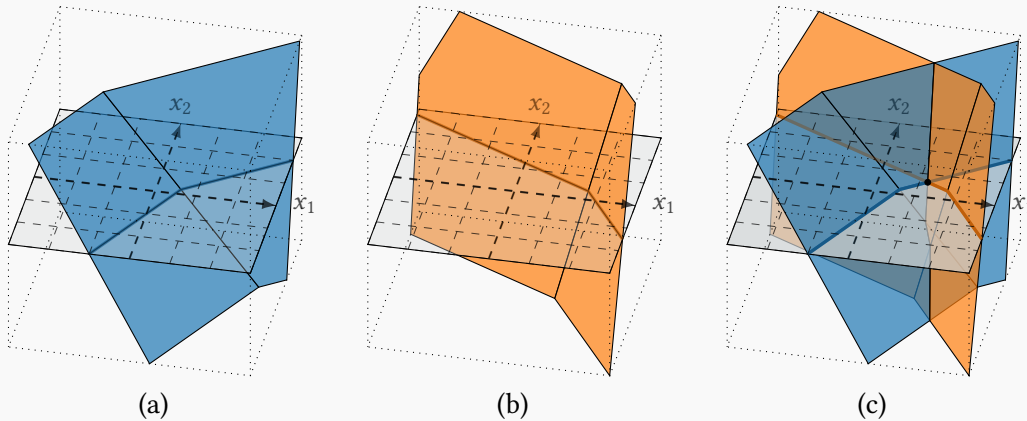


Figure 6.2: Visualizing two tropical posynomials on two variables and the associated square system

The Figure 6.2a (resp. Figure 6.2b) shows the surface with equation $z = P_1(x_1, x_2)$ (resp. $z = P_2(x_1, x_2)$), in addition with their intersection with the plane $z = 0$ (giving the 0-level set of both posynomials). We verify that both P_1 and P_2 are piecewise-affine convex functions. The Figure 6.2c combines the two surfaces, and it can be seen that the tropical posynomial system $P_1(x) = 0$ and $P_2(x) = 0$ admits a unique solution represented by the black dot.

In what follows and as done in this example, we shall no longer write P^{trop} to denote a generic tropical posynomial, but write P instead for brevity.

1.2 SOLVING POSYNOMIAL SYSTEMS IS NP-HARD. — Following up on the discussion about the increased difficulty to deal with posynomial systems compared to fixed-points of piecewise linear maps, we state two formal results showing that the feasibility problems for classical or tropical posynomial systems are NP-hard, even with integer exponents. We refer to the book of Arora and Barak for in-depth presentation of complexity questions [AB09].

➤ **PROPOSITION 6.3.** *Solving a square tropical posynomial system is NP-hard.*

Proof. We reduce 3-SAT to the problem (6.4). Let us consider a Boolean formula in conjunctive normal form $C_1 \wedge \cdots \wedge C_p$ made of p clauses, each one of them using three out of n real variables x_1, \dots, x_n ($p, n \in \mathbb{N}$).

We consider $n+2p$ additional unknown $(y_i)_{1 \leq i \leq n}$, $(z_k)_{1 \leq k \leq p}$ and $(s_k)_{1 \leq k \leq p}$ and we introduce the following (square) tropical posynomial system in the $2n + 2p$ variables $(x_1, \dots, x_n, y_1, \dots, y_n, z_1, \dots, z_p, s_1, \dots, s_p)$:

$$\begin{aligned} \forall i \in \{1, \dots, n\} \quad \max(x_i - 1, y_i - 1) = 0, & \quad x_i + y_i - 1 = 0, \\ \forall j \in \{1, \dots, p\} \quad \max\left(\max_{x_i \in C_j} (x_i - z_j), \max_{\neg x_i \in C_j} (y_i - z_j)\right) = 0, & \quad \max\left(\frac{1}{2} - z_j, s_j - z_j\right) = 0. \end{aligned}$$

This system can be constructed in polynomial time from the Boolean formula. The first $2n$ equations ensure that for all $i \in \{1, \dots, n\}$, $x_i \in \{0, 1\}$ and that x_i and y_i have opposite logical values. The next p equations express that for all $j \in \{1, \dots, p\}$, the variable z_j has the same Boolean value as the clause C_j , with the notation $x_i \in C_j$ (resp. $\neg x_i \in C_j$) if the variable x_i occurs positively (resp. negatively) in the clause C_j . The last (dummy) equations ensure that $z_j = 1$ for all $j \in \{1, \dots, p\}$. The instance $C_1 \wedge \cdots \wedge C_p$ is satisfiable if and only if this system admits a solution. \clubsuit

➤ **THEOREM 6.4.** *Solving a square classical posynomial system is NP-hard.*

Proof. We modify the previous construction to obtain a square posynomial system over $\mathbb{R}_{>0}^{2n+2p}$, along the lines of Maslov's dequantization principle [Lito7] or Viro's method [Viro1]:

$$\begin{aligned} \forall i \in \{1, \dots, n\} \quad \frac{2}{5}x_i + \frac{2}{5}y_i = 1, & \quad x_i y_i = 1, \\ \forall j \in \{1, \dots, p\} \quad \sum_{x_i \in C_j} \frac{1}{6}x_i z_j^{-1} + \sum_{\neg x_i \in C_j} \frac{1}{6}y_i z_j^{-1} = 1, & \quad \frac{1}{3}z_j^{-1} + s_j z_j^{-1} = 1. \end{aligned}$$

From the first $2n$ equations, the variables x_i and y_i range over $\{2, 1/2\}$, the values 2 and $1/2$ respectively encoding the true and false Boolean values. The variable $y_i = 1/x_i$ corresponds to the Boolean negation of x_i . Since each clause has precisely three literals, using the p next equations, we deduce that the variable z_j takes one of the values $\{1/2, 3/4, 1\}$ if the clause C_j is satisfied, and that it takes the value $1/4$ otherwise. The last p equations impose that z_j can take any value in $(1/3, \infty)$. We deduce that the formula $C_1 \wedge \cdots \wedge C_p$ is satisfied if and only if the posynomial system that we have obtained in this way admits a solution in $\mathbb{R}_{>0}^{2n+2p}$, hence the reduction to 3-SAT. \clubsuit

1.3 A LINEAR PROGRAMMING APPROACH FOR TROPICAL POSYNOMIAL SYSTEMS. — Given tropical posynomials P_1, \dots, P_n , we write the system (6.4) as $P(x) = 0$, where $P := (P_1, \dots, P_n)$. The **support** of this system, denoted by S , is defined as the disjoint union $\bigsqcup_{1 \leq i \leq n} S_{P_i}$ of the supports of the posynomials $(P_i)_{1 \leq i \leq n}$ (by disjoint union, we mean the coproduct in the category of sets, since these supports may have non-empty intersections or they may even coincide).

Our objective is to study if *some* tropical posynomial systems can be solved using linear programming (according to Proposition 6.3, we cannot hope such an approach to work in general). This is motivated by the fact that the zero-sublevel set (or feasibility set) \mathcal{P} associated with a tropical posynomial system is a closed polyhedron. Formally, we have

$$\mathcal{P} := \{x \in \mathbb{R}^n \mid P(x) \leq 0\} = \{x \in \mathbb{R}^n \mid \max_{a \in S} (c_a + \langle a, x \rangle) \leq 0\}.$$

Given a vector y in \mathbb{R}^n , we will denote by $(\text{LP}(y))$ the following linear program:

$$\text{Maximize } \langle y, x \rangle \quad \text{subject to } x \in \mathcal{P}. \quad (\text{LP}(y))$$

If the posynomial system $P(x) = 0$ admits a solution x^* that is a vertex of \mathcal{P} , then we can always construct a suitable vector y such that x^* arises as the solution of $(\text{LP}(y))$, however this vector will depend upon the coefficients of the posynomials. In comparison, the following definition

introduces candidates of interest for the cost vector of $(LP(y))$ only in terms of the posynomials' exponents.

► **DEFINITION 6.5.** We say that a vector y in the (convex) conic hull $\text{cone}(S)$ is *colorful* if for all μ in $(\mathbb{R}_{\geq 0})^S$,

$$y = \sum_{a \in S} \mu_a a \implies \forall i \in \{1, \dots, n\}, \exists a \in S_{P_i}, \mu_a > 0.$$

In other words, a vector y in \mathbb{R}^n is colorful if it arises as a nonnegative combination of the exponents of P , but also if all such combinations make use of at least one exponent of each one of the tropical posynomials P_1, \dots, P_n . Our terminology is justified by the fact that if we think of $S_{P_1}, S_{P_2}, \dots, S_{P_n}$ as colored sets, then we need all the colors to decompose a colorful vector y over these. Recall that by Carathéodory's theorem, every vector in the conic hull $\text{cone}(S)$ can be written as a positive linear combination of an independent family of vectors of S . Hence, when y is a colorful vector, it is obtained as a positive linear combination of precisely one vector a_i in each color class S_{P_i} , and the family a_1, \dots, a_n must be a basis (if not, Carathéodory's theorem would imply that y is a positive linear combination of a proper subset of $\{a_1, \dots, a_n\}$, so that y could not be a colorful vector.)

Example 6.6. We illustrate the Definition 6.5 in the Figure 6.3a, reusing the exponents of the two tropical posynomials of Example 6.2, i.e., $S_{P_1} := \{a_1, a_2\}$ and $S_{P_2} := \{a'_1, a'_2\}$, with $a_1 := (-2, 1)$, $a_2 := (-3, 4)$, $a'_1 := (2, 4)$ and $a'_2 := (3, 2)$. We represent a first vector y_0 that is indeed colorful, because writing it as a linear combination of vectors in S_{P_1} and S_{P_2} necessarily makes use of at least one vector in each support. This is not the case for the vector y_1 , because it lies in the blue span (that is, $\text{cone}(S_{P_1})$). The vector y_2 cannot be colorful because it is not in the conic hull $\text{cone}(S_{P_1} \cup S_{P_2})$. It can be seen on this figure that the set of colorful vectors is precisely $\text{int}(\text{cone}(a_2, a'_1))$, the dark-shaded region of the plane.

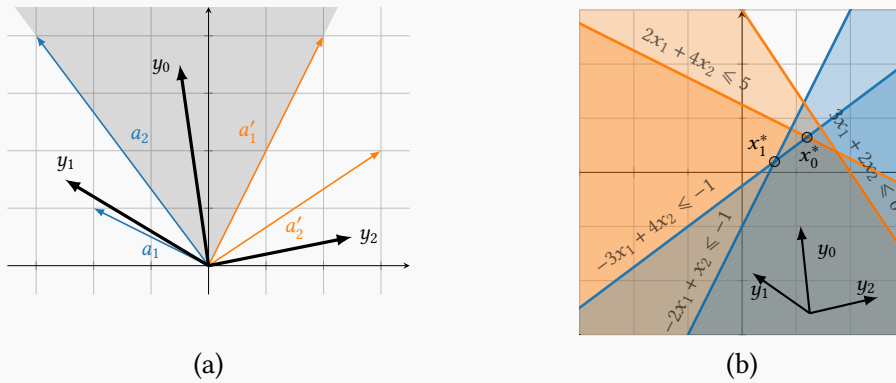


Figure 6.3: One colorful and two non-colorful vectors relative to the exponents of posynomials from Example 6.2

We represent in the Figure 6.3b the feasibility set \mathcal{P} of our posynomial system $P_1(x_1, x_2) \leq 0$ and $P_2(x_1, x_2) \leq 0$, the one that we could already see in Figure 6.2c. We also reproduce the three vectors y_0, y_1 and y_2 . It appears here that maximizing the cost y_0 over the polyhedron \mathcal{P} yields a point x_0^* that is a solution of the posynomial system, since it lies in the zero-level-sets of both P_1 and P_2 . The cost y_1 also determines an optimal vertex x_1^* of \mathcal{P} but that does not achieve $P_2(x_1^*) = 0$, and therefore is not a solution of the system. Finally, the cost y_2 yields an infinite solution of the linear program $(LP(y))$ since \mathcal{P} is unbounded in this direction.

The following theorem guarantees that the phenomenon observed in the Example 6.6 is general, meaning that the existence of colorful vectors is a sufficient condition for feasible tropical posynomial systems to have a solution.

- **THEOREM 6.7.** Assume that y is a colorful vector, and that the linear program $(LP(y))$ is feasible. Then, the linear program $(LP(y))$ has an optimal solution, and any optimal solution x satisfies $P(x) = 0$.

Proof. Since the feasibility set \mathcal{P} of $(LP(y))$ is nonempty, we can consider its recession cone, which is given by $C = \{x \in \mathbb{R}^n \mid \forall a \in S, \langle a, x \rangle \leq 0\}$. As a colorful vector, y belongs to the polyhedral cone generated by the vectors $a \in S$, so $\langle y, x \rangle \leq 0$ for all $x \in C$. By the Minkowski–Weyl theorem, \mathcal{P} is a Minkowski sum of the form $\mathcal{T} + C$ where \mathcal{T} is a polytope, i.e., every feasible point x can be written as $x = x' + x''$ with $x' \in \mathcal{T}$ and $x'' \in C$. Since $\langle y, x'' \rangle \leq 0$, the maximum of the objective function $x \mapsto \langle y, x \rangle$ over the polyhedron \mathcal{P} is attained (by an element of \mathcal{T}).

- ◇ Let $x^* \in \mathbb{R}^n$ be an optimal solution of $(LP(y))$. From the strong duality theorem, the dual linear program admits an optimal solution $(\mu_a^*)_{a \in S} \in (\mathbb{R}_{\geq 0})^S$ which satisfies $y = \sum_{a \in S} \mu_a^* a$ and $\mu_a^* (c_a + \langle a, x^* \rangle) = 0$ for all $a \in S$. Since y is a colorful vector, for all $i \in \{1, \dots, n\}$, there is some $a_i \in S_{P_i}$ such that $\mu_{a_i}^* > 0$. We then get that, for all $i \in \{1, \dots, n\}$, $P_i(x^*) \geq c_{a_i} + \langle a_i, x^* \rangle = 0$. As a result, $P(x^*) = 0$. ♣

We now want to provide a geometric condition ensuring that the linear program $(LP(y))$ is feasible, still regardless of the coefficients $(c_a)_{a \in S}$. We say that the tropical posynomial map P associated with our system has **pointed** exponents if its support is contained in an open halfspace, i.e., there exists $z \in \mathbb{R}^n$ such that for all a in S , we have $\langle a, z \rangle < 0$.

The interest of pointed systems comes from the following property:

- **PROPOSITION 6.8.** The inequality problem $P(x) \leq 0$ has a solution x in \mathbb{R}^n regardless of the coefficients of P if and only if P has pointed exponents.

Proof. Suppose that for all values of $(c_a)_{a \in S}$, there exists $x \in \mathbb{R}^n$ such that $P(x) \leq 0$. By choosing $c_a \equiv 1$, there exists $x_0 \in \mathbb{R}^n$ that satisfies $\forall a \in S, 1 + \langle a, x_0 \rangle \leq 0$. Hence, for all $i \in \{1, \dots, n\}$, the exponents of P_i lie in the open halfspace $\{a \in \mathbb{R}^n \mid \langle a, x_0 \rangle < 0\}$.

- ◇ Suppose now that P has pointed exponents. Then there is some $z \in \mathbb{R}^n$ such that for all $a \in S$, we have $\langle a, z \rangle < 0$. We define $\lambda := \max_{a \in S} (-c_a) / \langle a, z \rangle$ so that $\forall a \in S, c_a + \langle a, \lambda z \rangle \leq 0$ and therefore for all $i \in \{1, \dots, n\}$, $P_i(\lambda z) \leq 0$. ♣

Example 6.9. Let us illustrate the Proposition 6.8, by reusing again the exponents of the tropical posynomial system in Example 6.2, and visualized in Figure 6.3a.

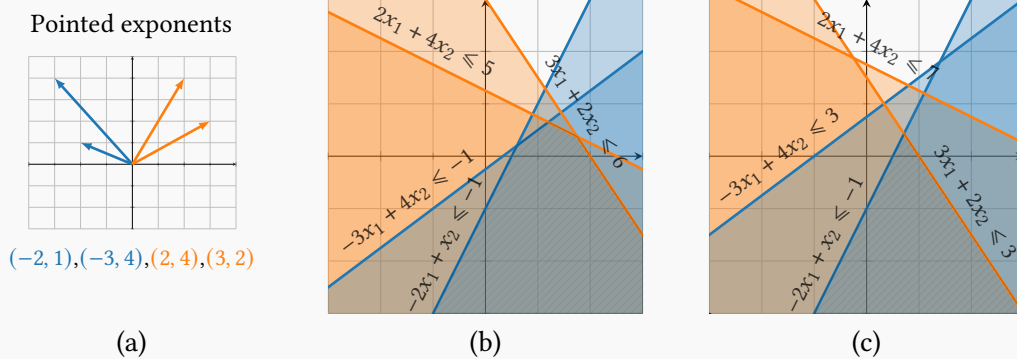
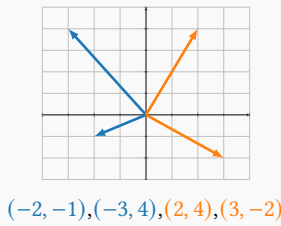


Figure 6.4: When the exponents are pointed, the polyhedron \mathcal{P} is nonempty regardless of the coefficients $(c_a)_{a \in S}$

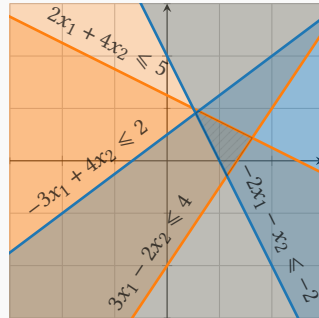
These exponents are pointed since they all have a positive x_2 coordinate. The Figure 6.4b and the Figure 6.4c show that we may modify the coefficients of P as we want, thus translating the defining inequalities of \mathcal{P} , the polyhedron \mathcal{P} remains nonempty.

This is not the case for a feasibility set associated with a tropical posynomial with non-pointed exponents. For instance, taking $S_{P_1} := \{(-2, -1), (-3, 4)\}$ and $S_{P_2} := \{(2, 4), (3, -2)\}$, we can see on the Figure 6.5a that S_{P_1} and S_{P_2} (positively) span the whole two-dimensional space, thus the configuration of exponents is not pointed. The Figures 6.5b and 6.5c show that the corresponding polyhedron \mathcal{P} can be either empty or nonempty depending on the choice of coefficients.

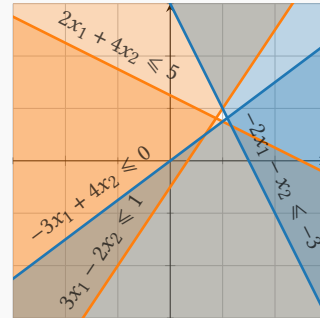
Non-pointed exponents



(a)



(b)



(c)

Figure 6.5: When the exponents are not pointed, the polyhedron \mathcal{P} may be empty for some coefficients $(c_a)_{a \in S}$

As a consequence of Theorem 6.7 and Proposition 6.8, if a tropical posynomial system $P(x) = 0$ has pointed exponents and there exists a colorful vector, then the system admits a solution which can be found by linear programming (and whatever the coefficients of the posynomials!).

1.4 THE RELATIONSHIP BETWEEN TROPICAL POSYNOMIAL SYSTEMS AND MARKOV DECISION PROCESSES. – Let us discuss the remarkable special case of tropical posynomial systems made of Markov decision processes (MDPs). We refer the reader to our Chapter 3 for a detailed presentation of Markov and semi-Markov decision processes. The results presented there allow us to establish that the discounted value function v of MDPs (when we want to maximize a reward instead of minimizing a cost) satisfy a Bellman fixed-point equation of the form

$$\text{for all } i \text{ in } \{1, \dots, n\}, \quad v_i = \max_{a \in A_i} (c^a + \gamma^a \langle p^a, v \rangle), \quad (6.5)$$

where $\{1, \dots, n\}$ is the state space of the MDP, A_i is the set of playable actions from i , the scalar c^a (resp. γ^a) is the reward (resp. discount factor) collected by the player as a result of choosing action a , and $p^a := (p_j^a)_{1 \leq j \leq n}$ is the (stochastic) transition probability vector determining in which state the player goes after playing a .

The equation (6.5) reduces to the form (6.4) with $S_{P_i} := \{\gamma^a p^a - e_i \mid a \in A_i\}$, where e_i denotes the i -th element of the canonical basis of \mathbb{R}^n . In what follows, when we say that an MDP is discounted, we mean that all the discount factors are such that $0 \leq \gamma^a < 1$ (in the usual sense, and not the generalized framework of Chapter 3). The following holds:

► **PROPOSITION 6.10.** Any negative vector is colorful with respect to the tropical posynomial system associated with a discounted Markov decision process.

Proof. Let x be a vector with negative entries. We first claim that x is in the cone spanned by the vectors $(\gamma^{a_i} p^{a_i} - e_i)_{1 \leq i \leq n}$, for all choice of (a_1, a_2, \dots, a_n) in $A_1 \times A_2 \times \dots \times A_n$. Indeed, the system $x = \sum_{j=1}^n \lambda_j (\gamma^{a_j} p^{a_j} - e_j)$ can be rewritten as $\lambda = -x + \Gamma P \lambda$, where P is the stochastic matrix with columns $(p^{a_i})_{1 \leq i \leq n}$, Γ the diagonal matrix $\text{diag}((\gamma^{a_i})_{1 \leq i \leq n})$ and λ the vector $(\lambda_i)_{1 \leq i \leq n}$. Since the spectral radius $\rho(\Gamma P)$ is strictly less than 1, we have $\lambda = (I - \Gamma P)^{-1}(-x)$, and $(I - \Gamma P)^{-1} = \sum_{k=0}^{\infty} (\Gamma P)^k$ has only nonnegative entries. Thus, we have $\lambda \geq 0$ and x lies in the conic hull of $(\gamma^{a_i} p^{a_i} - e_i)_{1 \leq i \leq n}$. This is true for any choice of actions, hence x is in $\text{cone}(S)$.

◇ Suppose by contradiction that x is not a colorful vector. Then, according to the Definition 6.5, there exists a vector μ with nonnegative entries such that $x = \sum_{j=1}^n \sum_{a \in A_j} \mu_a (\gamma^a p^a - e_j)$ and $\exists i \in \{1, \dots, n\}$ such that for all $a \in A_i$, $\mu_a = 0$, meaning that the i -th color is not used in the decomposition. It follows that

$$x_i = \langle x, e_i \rangle = \sum_{j=1}^n \sum_{a \in A_j} \mu_a \langle \gamma^a p^a - e_j, e_i \rangle = \sum_{j=1}^n \sum_{a \in A_j} \mu_a \gamma^a p_j^a - \sum_{p \in A_i} \mu_a = \sum_{j=1}^n \sum_{a \in A_i} \mu_a \gamma^a p_j^a,$$

By nonnegativity of the multipliers $(\mu_a)_{a \in S}$, of the discount factors and of the entries of the probability vectors, we end up with $x_i \geq 0$, hence the contradiction. ♣

Example 6.11. In order to illustrate what precedes, let us depict in Figure 6.6a a 2-state MDP (with the same graphical conventions than Chapter 3) with $|A_1| = 3$ and $|A_2| = 2$. The rewards are not depicted because not relevant here, however, the transition probabilities are given along the arcs and the discount factors written in bold next to their parent actions. In the Figure 6.6b, we show the five vectors $(\gamma^a p^a)_{a \in S}$, that all lie in the half-open simplex $\{x \geq 0\} \cap \{y \geq 0\} \cap \{x + y < 1\}$. Remark that one of these vectors (associated with state 1) is on the y -axis because it corresponds to a deterministic action making the player go from state 1 to state 2. All the other actions are properly stochastic, hence not on the axes.

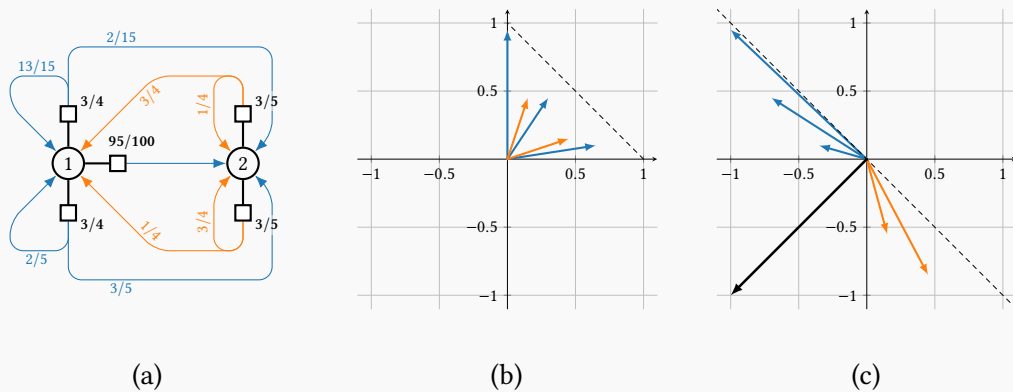


Figure 6.6: An MDP with two states and five actions (a), its probability vectors (b) and its posynomial exponents (c)

In the Figure 6.6c, we represent the vectors $(\gamma^a p^a - e_i)_{a \in S}$. One can then see that all the vectors of $(\mathbb{R}_{<0})^2$ are colorful and we highlight in particular the vector -1 . Remark that the configuration is pointed as well, because all the exponents lie in the half-space $x + y < 0$.

The Proposition 6.10 together with the Theorem 6.7 allow us to recover the usual linear programming approach to find the value of discounted MDPs, (see [Put14, p. 223]), sometimes called

Perron's method, stating that the value can be obtained by minimizing the function $v \mapsto \sum_{i=1}^n v_i$ subject to the constraints $v_i \geq c^a + \gamma^a \langle p^a, v \rangle$ for all i in $\{1, \dots, n\}$ and a in A_i .

The fact that discounted MDPs yield tropical posynomial systems with pointed configuration of exponents, as showed by Example 6.11, is true in general. Conversely, the following theorem tells us that tropical posynomial systems with pointed exponents admitting a colorful vector are actually not more general than discounted MDPs. Note that the proof makes use of an important property of the set of colorful vectors that will only be shown in Section 2.

► **THEOREM 6.12.** *Tropical posynomial systems with pointed exponents and for which a colorful vector exists are equivalent to discounted Markov decision processes up to a change of variables.*

Proof. Let us consider the general system (6.4) $P_i(x) = \max_{a \in S_{P_i}} (c_a + \langle a, x \rangle) = 0$ for all $i \in \{1, \dots, n\}$. Following the claim of the theorem, we suppose that (i) the configuration of exponents is pointed, as a result we have a linear form $z \in (\mathbb{R}^n)^\star$ such that $z(a) > 0$ for all $i \in \{1, \dots, n\}$ and for all $a \in A_i$; (ii) in addition, the configuration admits a colorful vector, hence by Theorem 6.20, we know that the set of colorful vectors is an n -dimensional simplicial cone. We represent the hyperplanes supporting the maximal faces of this cone by their normal vectors $(h_i)_{1 \leq i \leq n}$, also identified to vectors of the dual space $(\mathbb{R}^n)^\star$ of \mathbb{R}^n , i.e., linear forms. It is shown in Section 2 that the $(h_i)_{1 \leq i \leq n}$ are such that for all $1 \leq i, j \leq n$ with $j \neq i$, if $a \in S_{P_i}$ then $h_i(a) > 0$ while if $a \in S_{P_j}$ then $h_i(a) \leq 0$.

The collection $(h_i)_{1 \leq i \leq n}$ is a basis of $(\mathbb{R}^n)^\star$ (for they are the normal vectors to the faces of a full-dimensional simplicial cone). We denote by $(s_i)_{1 \leq i \leq n}$ its dual basis, whose elements belong to \mathbb{R}^n and generate (by positive scaling) the edges of the set of colorful vectors. From Lemma 6.18, we deduce that for all $1 \leq i \leq n$, we have $s_i \in \text{cone}(S)$ and therefore $z(s_i) > 0$. We also define the positive scalar λ by

$$\lambda := \min_{1 \leq i \leq n} \min_{a \in S_{P_i}} [z(s_i)h_i(a)]^{-1}.$$

◊ We now introduce the automorphism f of \mathbb{R}^n that maps s_i to $-\lambda z(s_i)e_i$ for all $1 \leq i \leq n$ (where e_i denotes the i -th vector of the standard basis). From the duality of bases $(h_i)_{1 \leq i \leq n}$ and $(s_i)_{1 \leq i \leq n}$, for any $a \in \mathbb{R}^n$, we have the expression

$$f(a) = -\lambda \sum_{j=1}^n z(s_j)h_j(a)e_j.$$

For all $i \in \{1, \dots, n\}$ and for all $a \in A_i$, we define $\gamma^a := \|f(a) + e_i\|_1$. We can verify that $0 \leq \gamma^a < 1$, indeed

$$0 \leq \|f(a) + e_i\|_1 = \lambda \sum_{j \neq i} z(s_j)(-h_j(a)) + (1 - \lambda z(s_i)h_i(a)) = 1 - \lambda z(a) < 1,$$

since $z(a) = z(\sum_{j \in \{1, \dots, n\}} h_j(a)s_j) > 0$ from the pointed character of the exponents' configuration. We can therefore introduce the stochastic vector $p^a := (\gamma^a)^{-1}(f(a) + e_i)$ for all $a \in S$ (as shown in the above expression, it has nonnegative coordinates from our choice of λ).

Eventually, denoting by f^\star the adjoint of the isomorphism f and letting $y := (f^\star)^{-1}(x)$, we have for all $i \in \{1, \dots, n\}$:

$$\max_{a \in S_{P_i}} (c_a + \langle a, x \rangle) = 0 \iff \max_{a \in S_{P_i}} (c_a + \gamma^a \langle p^a, y \rangle) = y_i$$

where we recognize Bellman's optimality equation of a MDP with rewards $(c_a)_{a \in S}$, discount factors $(\gamma^a)_{a \in S}$ and transitions probabilities $(p^a)_{a \in S}$. \clubsuit

In a sense, the Theorem 6.12 is at the same time a negative and a positive result. On one hand, we learn that the class of interest made of tropical posynomial systems with pointed exponents and admitting a colorful vector, for which we can easily find solutions, is not more general than the widely studied class of discounted MDPs. On the other hand, it allows us to use well-known algorithms for MDPs to solve this class of posynomial systems.

We understand alternatively from the proof of Theorem 6.12 that we can distort tropical posynomial systems in general onto MDPs by means of a linear transformation, but there are multiple degrees of freedom in this map. Indeed, we could have taken for f the automorphism

which maps s_i to $-\mu_i e_i$, where $(\mu_i)_{1 \leq i \leq n}$ are positive coefficients to tune, chosen such that for all $1 \leq i \leq n$ and a in A_i , the vector $f(a) + e_i$ is nonnegative and sub-stochastic (such coefficients do exist since the previous result shows that $\mu_i = -\lambda z(s_i)$ satisfies these requirements). Let us denote by \mathcal{X} the polyhedron of admissible choices of μ :

$$\mathcal{X} = \left\{ \mu \in (\mathbb{R}_{>0})^n \mid \forall i \in \{1, \dots, n\}, \forall a \in A_i, 1 - \mu_i h_i(a) \geq 0 \text{ and } \sum_{j=1}^n \mu_j h_j(a) \geq 0 \right\}.$$

Tweaking the values of the $(\mu_i)_{1 \leq i \leq n}$ by remaining in \mathcal{X} allows one to have a little control on the largest discount factor $\gamma(\mu)$ featured in the resulting MDP, that we can compute by taking the largest ℓ_1 norm of the vectors $f(a) + e_i$

$$\gamma(\mu) = \max_{1 \leq i \leq n} \max_{a \in A_i} \left\{ 1 - \sum_{j=1}^n \mu_j h_j(a) \right\}.$$

For the $(A_i)_{1 \leq i \leq n}$ are finite, a minimal discount factor $\gamma^* := \min_{\mu \in \mathcal{X}} \gamma(\mu)$ can be determined by solving $\sum_{i=1}^n |A_i|$ linear programs.

Since this approach is applicable to any posynomial tropical system, we can use it on... systems that already arise from MDPs. This amounts to geometrically distort an MDP by a linear change of variables into another MDP, to obtain reduced discount factors. Since it is customary to obtain that the largest discount factor controls the complexity of algorithms used in the framework of MDPs (see for instance the complexity result from [Ye11] for the simplex algorithm), this “discount boosting” technique could turn advantageous in some applications. It is also prone to transform undiscounted configurations into discounted configurations for some instances. Remark that this approach echoes other reduction techniques involving diagonal scalings [AG13, AGQS19].

We illustrate the above discussion with the following example.

Example 6.13. We come back to the Markov decision process depicted in Figure 6.6a, and with vectors $(\gamma^a p^a)_{a \in S}$ depicted in Figure 6.6b. Associating with each action a of A_i with discount factor γ and transition vector p the vector $a := \gamma p - e_i$ ($i = 1$ or $i = 2$), we have (referring to the three actions of A_1 and the two actions A_2 by the indices 11, 12, 13, 21 and 22):

$$\begin{array}{ccccc} \gamma_{11} = 0.75 & \gamma_{12} = 0.75 & \gamma_{13} = 0.95 & \gamma_{21} = 0.60 & \gamma_{22} = 0.60 \\ p_{11} = (0.866\dots, 0.133\dots) & p_{12} = (0.4, 0.6) & p_{13} = (0, 1) & p_{21} = (0.25, .75) & p_{22} = (0.75, 0.25) \\ a_{11} = (-0.35, 0.10) & a_{12} = (-0.70, 0.45) & a_{13} = (-1, 0.95) & a_{21} = (0.15, -0.55) & a_{22} = (0.45, -0.85) \end{array}$$

The Figure 6.7a shows in solid lines the (pointed) configuration of exponents in the half-plane $x + y < 0$, and the vectors $(\gamma^a p^a)_{a \in S}$ as well in dashed lines (this figure hence combines the Figure 6.6b and the Figure 6.6c). We can graphically read that the maximum discount factor of the problem is 0.95, either by taking the norm of $\gamma_{13} p_{13}$, or by looking at the (normal) distance from the tip of the vector a_{13} to the hyperplane $x + y = 0$. Our objective is to increase this margin by applying a linear transformation to the vectors a_{11}, \dots, a_{22} .

As done in the proof of Theorem 6.12, we can map via the function Ψ the set of colorful vectors $\text{cone}(a_{11}, a_{21})$ onto the negative orthant $(\mathbb{R}_{\leq 0})^2$, but this still leaves us with too degrees of freedom μ_1 and μ_2 such that $\Psi(a_{11}) = -\mu_1 e_1$ and $\Psi(a_{21}) = -\mu_2 e_2$. We let the reader check that the function Ψ admits the expression

$$\Psi(x, y) = \left(\mu_1 \frac{(11x + 3y)}{355/100}, \mu_2 \frac{(2x + 7y)}{355/100} \right).$$

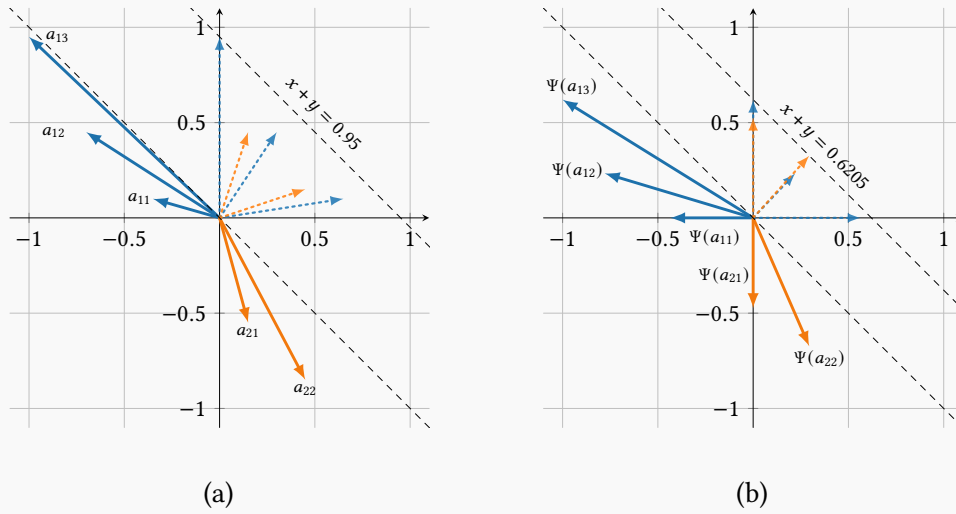


Figure 6.7: The exponents of the posynomial system associated with the MDP of Example 6.11 before (left) and after (right) a change of variables to reduce the discount factors

The saturated constraints of \mathcal{X} that determine μ_1 and μ_2 in this example are on one hand the fact that $\Psi(a_{13}) + e_1$ needs to have a nonnegative x -coordinate, and on the other hand the fact that $\Psi(a_{13})$ and $\Psi(a_{22})$ shall yield the same discount factor. We are thus able to compute $\mu_1 = 71/163$ and $\mu_2 = 14981/31622$.

We can now go back to a MDP framework by shifting our new exponents of $\Psi(A_1)$ (resp. $\Psi(A_2)$) by e_1 (resp. e_2). An initial discount γ and transition vector p determine with our transformation another “improved” tuple $\gamma'p' := \Psi(\gamma p - e_i) + e_i$ (it suffices to take the ℓ_1 norm to isolate γ' in this expression). The maximum new discounts are $\gamma'_{13} = \gamma'_{22} = 19623/31622 \simeq 0.6205$, which represents a 65 % shrink compared to the initial value $\gamma_{13} = 0.95$.

To sum it up, we started from an MDP system with value function v satisfying

$$\begin{aligned} v_1 &= \max \left(c_{11} + \frac{3}{4} \left(\frac{13}{15} v_1 + \frac{2}{15} v_2 \right), c_{12} + \frac{3}{4} \left(\frac{3}{5} v_1 + \frac{2}{5} v_2 \right), c_{13} + \frac{95}{100} v_2 \right) \\ v_2 &= \max \left(c_{21} + \frac{3}{5} \left(\frac{1}{4} v_1 + \frac{3}{4} v_2 \right), c_{22} + \frac{3}{5} \left(\frac{3}{4} v_1 + \frac{1}{4} v_2 \right) \right) \end{aligned}$$

and we have built an equivalent MDP system with value function $w = (\Psi^*)^{-1}(v)$, identical rewards, but different transition probabilities and reduced discount factors verifying

$$\begin{aligned} w_1 &= \max \left(c_{11} + 0.5644w_1, c_{12} + 0.4544 \left(0.4860w_1 + 0.5140w_2 \right), c_{13} + 0.6205w_2 \right) \\ w_2 &= \max \left(c_{21} + 0.5262w_2, c_{22} + 0.6205 \left(0.4745w_1 + 0.5255w_2 \right) \right). \end{aligned}$$

1.5 GEOMETRIC PROGRAMMING APPROACH OF POSYNOMIALS SYSTEMS. — In this section, we come back to the problem of solving “classical” posynomial systems of the form (6.2), i.e., given a collection $P = (P_1, \dots, P_n)$ of posynomials, we want to find x in $(\mathbb{R}_{>0})^n$ such that $P(x) = 1$. We are going to leverage on the results proved for the tropical case to tackle the classical case. Indeed, remark that the definitions of colorful vectors and pointed exponents introduced in Section 1.3

still make sense for usual posynomials, since they depend only upon the disjoint union S of S_{P_1}, \dots, S_{P_n} , and these remain subsets of \mathbb{R}^n . We shall need the following lemma:

► **LEMMA 6.14.** *If y is a colorful vector relatively to the exponents of S , the polyhedron \mathcal{P} defined by*

$$\mathcal{P} := \{x \in \mathbb{R}^n \mid \forall a \in S, \quad \log c_a + \langle a, x \rangle \leq 0 \quad \text{and} \quad \langle y, x \rangle \geq \mu\}$$

is bounded (possibly empty), regardless of our choice of positive $(c_a)_{a \in S}$ or μ in \mathbb{R} .

Proof. If \mathcal{P} is nonempty, let $C := \{x \in \mathbb{R}^n \mid \forall a \in S, \langle a, x \rangle \leq 0, \langle y, x \rangle \geq 0\}$ denote its recession cone, and let $x \in C$. Since y is a colorful vector, there exists $(\lambda_1, \dots, \lambda_n) \in \mathbb{R}_{>0}^n$ and a basis $(a_1, \dots, a_n) \in \prod_{1 \leq i \leq n} S_{P_i}$ such that $y = \sum_{i=1}^n \lambda_i a_i$. Thus, $\langle y, x \rangle \leq 0$, and so $\langle y, x \rangle = \sum_{i=1}^n \lambda_i \langle a_i, x \rangle = 0$. As a consequence, since $\lambda_i > 0$ for all $i \in \{1, \dots, n\}$, we have $\langle a_i, x \rangle = 0$, but since (a_1, \dots, a_n) is a basis, we get $x = 0$. Thus, $C = \{0\}$, and according to the Minkowski–Weyl theorem, \mathcal{P} is bounded. \clubsuit

We now exhibit a class of posynomial systems that we can solve using geometric programs (we refer the reader to [CS16] for background on this topic), still using the notion of colorful vector introduced by Definition 6.5, and the notion of pointed exponents discussed in Section 1.3 as well. Given a vector X in \mathbb{R}^n , we denote $\exp(X) := (\exp(X_i))_{1 \leq i \leq n}$.

► **THEOREM 6.15.** *Let $P(x) = 1$ be a posynomial system with pointed exponents S , and y be a colorful vector relatively to S . Then, the system has a solution $x = \exp(X^*)$ in $(\mathbb{R}_{>0})^n$, where X^* is an arbitrary solution of the following geometric program:*

$$\text{Maximize } \langle y, X \rangle \quad \text{subject to} \quad \forall i \in \{1, \dots, n\} \quad g_i(X) \leq 0, \quad (\text{G})$$

where $g_i(X) := \log \left(\sum_{a \in S_{P_i}} c_a e^{\langle a, X \rangle} \right)$.

Proof. For $x \in \mathbb{R}_{>0}^n$, we define $X := \log(x)$ (componentwise) so that $P(x) = 1$ is equivalent to solving $g_i(X) = 0$ for all $i \in \{1, \dots, n\}$. By Hölder’s inequality, the functions $(g_i)_{1 \leq i \leq n}$ are convex. We define $h_i: X \mapsto \max_{a \in S_{P_i}} (\log(c_a) + \langle a, X \rangle)$ for $i \in \{1, \dots, n\}$ and we observe that $h_i(X) \leq g_i(X) \leq h_i(X) + \log(|S_{P_i}|)$ (see in particular the Example 6.16 to visualize this inequality).

Since the system $P(x) = 1$ has pointed exponents, by Proposition 6.8, the polyhedron $\{X \in \mathbb{R}^n \mid \forall i \in \{1, \dots, n\}, h_i(X) + \log(|S_{P_i}|) \leq 0\}$ is nonempty. A fortiori, the feasible set of (G) is nonempty.

◊ Let us now prove that the maximum of (G) is finite and attained, by proving that the μ -superlevel set $S_\mu = \{X \in \mathbb{R}^n \mid \langle y, X \rangle \geq \mu \text{ and } \forall i \in \{1, \dots, n\}, g_i(X) \leq 0\}$ of the objective function (included in the feasible set) is compact for all $\mu \in \mathbb{R}$. Closedness is direct, and we observe that for $\mu \in \mathbb{R}$, S_μ is contained in $\{X \in \mathbb{R}^n \mid \langle y, X \rangle \geq \mu \text{ and } \forall i \in \{1, \dots, n\}, h_i(X) \leq 0\}$, but by Lemma 6.14, this polyhedron is bounded. Hence, (G) admits an optimal solution X^* .

Furthermore, again by Proposition 6.8, there exists \bar{X} such that for all $i \in \{1, \dots, n\}$, $h_i(\bar{X}) + \log(|S_{P_i}|) + 1 \leq 0$. Therefore, we have $g_i(\bar{X}) < 0$ for all $1 \leq i \leq n$, which means that (G) satisfies Slater’s condition. Problem (G) being convex, optimality of X^* is characterized by the Karush–Kuhn–Tucker conditions (see [BV04]). Hence, there is a vector of nonnegative multipliers $\lambda^* = (\lambda_1^*, \dots, \lambda_n^*)$ such that (X^*, λ^*) is a stationary point of the Lagrangian of (G), and the complementarity slackness conditions hold, i.e. for all $i \in \{1, \dots, n\}$, $\lambda_i^* g_i(X^*) = 0$. Defining $Z_i := \sum_{a \in S_{P_i}} c_a e^{\langle a, X^* \rangle} > 0$ for $i \in \{1, \dots, n\}$, the stationarity conditions give

$$y = \sum_{i=1}^n \frac{\lambda_i^*}{Z_i} \sum_{a \in S_{P_i}} c_a e^{\langle a, X^* \rangle} a.$$

Since y is colorful, λ_i^* is positive for all $1 \leq i \leq n$. The complementarity slackness conditions yield $g_i(X^*) = 0$ for all $1 \leq i \leq n$. Finally, $x^* := \exp(X^*)$ satisfies $P(x^*) = 1$. \clubsuit

Example 6.16. The Theorem 6.15 for classical posynomials uses our result Theorem 6.7 dealing with tropical posynomials, thanks to the inequalities $h_i(X) \leq g_i(X) \leq h_i(X) + \log(|S_{P_i}|)$ for all $1 \leq i \leq n$ on the constraints. We reckon it is interesting to illustrate this aspect, that is a customary tool when going from the field $(\mathbb{R}, +, \times)$ to the semifield $(\mathbb{R} \cup \{-\infty\}, \max, +)$.

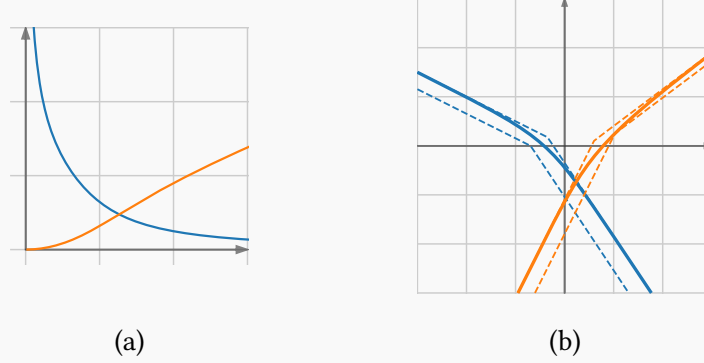


Figure 6.8: Seeing classical posynomials through log-glasses shows max-plus boundaries

The Figure 6.8a depicts the 1-level-sets of the two posynomials $P_1(x_1, x_2) = x_1^2 x_2^4 + 2x_1^3 x_2^2$ and $P_2(x_1, x_2) = 3x_1^{-2} x_2 + 4x_1^{-3} x_2^4$ over $(\mathbb{R}_{>0})^2$, while the Figure 6.8b gives the same plot once we wear “log-glasses”, i.e., $(X_1, X_2) \mapsto \log(P_i(\exp(X_1), \exp(X_2)))$, with $i = 1$ or $i = 2$. We can then see that each one of these new level-sets (i.e., g_1 and g_2) is bounded between a pair of “parallel” tropical posynomials. In this example, these posynomials are those featured in Examples 6.2, 6.6 and 6.9. In particular the configuration of exponents is pointed and admits a colorful vector, so that the associated geometric program can be solved via linear programming.

2 Properties of the colorful interior of convex sets

The Theorems 6.7 and 6.15 rely on the existence of a colorful vector (see Definition 6.5). The purpose of this section is to study the properties of the set of such vectors. In fact, colorful vectors can be defined more generally from a family of n closed convex cones, independently of posynomial systems.

► **DEFINITION 6.17.** Let $C = (C_1, \dots, C_n)$ be a collection of n closed convex cones of \mathbb{R}^n . We say that a vector y in \mathbb{R}^n is *colorful* if it belongs to the set

$$\text{cone}(C_1 \cup \dots \cup C_n) \setminus \bigcup_{1 \leq i \leq n} \text{cone}\left(\bigcup_{j \neq i} C_j\right).$$

The latter set is referred to as the *colorful interior* of C .

The name of the set still comes from the idea that if the initial cones C_1, \dots, C_n are thought of as different colored sets, then a colorful vector necessarily makes use of all the colors to be decomposed over these cones, since it is in the whole conic hull $\text{cone}(C_1 \cup \dots \cup C_n)$ but not in the conic hull $(n - 1)$ colors. The earlier Definition 6.5, that is more specific to posynomials, can be recovered by taking $C_i := \text{cone}(S_{P_i})$ for all i in $\{1, \dots, n\}$.

2.1 STRUCTURE OF THE COLORFUL INTERIOR OF POINTED CONES. — In what follows, we shall restrict to the case where the collection \mathcal{C} is **pointed**, i.e., $\text{cone}(C_1 \cup \dots \cup C_n) \setminus \{0\}$ is included in an open halfspace – we will explain at the end of the section what properties can be lost otherwise. Remark that we do not know so far if the colorful interior of \mathcal{C} is convex, or even connected, and using the Definition 6.17 head first does not suffice to answer these questions.

Suppose that $\{x \in \mathbb{R}^n \mid \langle z, x \rangle > 0\}$ is an open halfspace containing the $(C_i \setminus \{0\})_{1 \leq i \leq n}$. Then, as a cone, the colorful interior of \mathcal{C} can be more simply studied from its cross-section with $\{x \in \mathbb{R}^n \mid \langle x, z \rangle = 1\}$. The latter can be shown to coincide with the set

$$\text{conv}(S_1 \cup \dots \cup S_n) \setminus \bigcup_{1 \leq i \leq n} \text{conv}\left(\bigcup_{j \neq i} S_j\right) \tag{6.6}$$

where for all i in $\{1, \dots, n\}$, S_i is the cross-section of the cone C_i by $\{x \in \mathbb{R}^n \mid \langle x, z \rangle = 1\}$. In the setting of posynomial systems, S_i is isomorphic to the convex hull of S_{p_i} , hence we recognize the role played by the Newton polytopes of the different posynomials (see [DLRS10]). Given a collection $\mathcal{S} = (S_1, \dots, S_n)$ of n closed convex sets of \mathbb{R}^{n-1} , we still refer to the set (6.6) as the **colorful interior** of \mathcal{S} , and denote it by $\text{colint } \mathcal{S}$. The Figure 6.9 illustrates this dimensional reduction with three colored sets, hence going from three conic hulls in \mathbb{R}^3 to an equivalent configuration of three convex hulls in \mathbb{R}^2 .

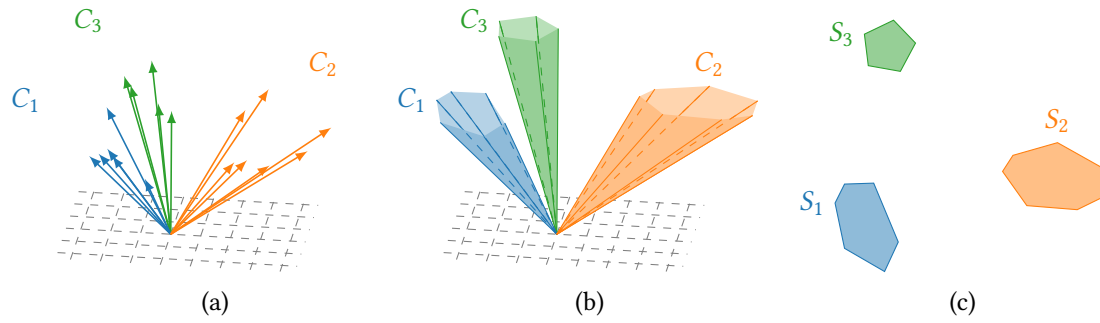


Figure 6.9: Taking three finite subsets of \mathbb{R}^n (a), their conic hulls that are all included in a common halfspace (b), and switching to convex hulls in a cross-section parallel with the ground plane (c)

We start with a lemma justifying why we have chosen the “interior” terminology.

► **LEMMA 6.18.** *Let $\mathcal{S} = (S_1, \dots, S_n)$ be a collection of n closed convex sets of \mathbb{R}^{n-1} . Then $\text{colint } \mathcal{S}$ is an open set included in $\text{int conv}(S_1 \cup \dots \cup S_n)$.*

Proof. Since $\text{colint } \mathcal{S} \subset \text{conv}(\mathcal{S})$, it suffices to prove that $\text{colint } \mathcal{S} \cap \partial \text{conv}(\mathcal{S}) = \emptyset$ (where ∂ denotes the frontier of a set) to show $\text{colint } \mathcal{S} \subset \text{int conv}(\mathcal{S})$. Let $x \in \partial \text{conv}(\mathcal{S})$, and let H be a supporting hyperplane of $\text{conv}(\mathcal{S})$ at x , so that $x \in \text{conv}(\mathcal{S}) \cap H = \text{conv}(\bigcup_{1 \leq i \leq n} (S_i \cap H))$. By Carathéodory’s theorem, x can be written using at most $n - 1$ vectors of $(S_i \cap H)_{1 \leq i \leq n}$, so it is in $\text{conv}(\bigcup_{j \neq i} S_j)$ for some $i \in \{1, \dots, n\}$ (the color not used in the decomposition). Therefore $x \notin \text{colint } \mathcal{S}$.

◇ Now we show that $\text{colint } \mathcal{S}$ is open. Let $x \in \text{colint } \mathcal{S}$. Let $i \in \{1, \dots, n\}$, x is not in $\text{conv}(\bigcup_{j \neq i} S_j)$ which is closed, so $d_i := d(x, \text{conv}(\bigcup_{j \neq i} S_j)) > 0$. Similarly, $\partial \text{conv}(\mathcal{S})$ is closed and does not contain x by the above argument, so $d' := d(x, \partial \text{conv}(\mathcal{S})) > 0$. The open ball centered in x with radius $\min((d_i)_{1 \leq i \leq n}, d')$ is clearly included in $\text{conv}(\mathcal{S})$. ♠

The set $\text{colint } \mathcal{S}$ has appeared before in a work of Lawrence and Soltan [LS09], in one of the proof where they characterize the intersection of convex transversals to a collection of sets. In more details, Lemma 6.18 and [LS09, Lemma 6] imply:

- **PROPOSITION 6.19.** *Let $\mathcal{S} = (S_1, \dots, S_n)$ be a collection of n closed convex sets of \mathbb{R}^{n-1} . Define the set \mathcal{D} of colorful simplices by $\mathcal{D} := \{\text{conv}(\{x_1, \dots, x_n\}) \mid x_1 \in S_1, \dots, x_n \in S_n\}$, i.e., with one vertex in each colored set. Then we have*

$$\text{colint } \mathcal{S} = \bigcap_{\Delta \in \mathcal{D}} \text{int } \Delta = \text{int } \bigcap_{\Delta \in \mathcal{D}} \Delta .$$

Remark that the Proposition 6.19 still holds if the colorful simplices Δ in \mathcal{D} are replaced by the convex transversals to the sets S_1, \dots, S_n . We illustrate in the Figure 6.10 our two ways so far to delineate the colorful interior $\text{colint } \mathcal{S}$, reusing the three convex sets of Figure 6.9c. The Figure 6.10a uses our definition (6.6), where for all $1 \leq i \leq n$, the convex hull $\text{conv}(\cup_{j \neq i} S_j)$ of all colors but the i -th one is denoted for short by \widehat{S}_i . The Figure 6.10b shows that $\text{colint } \mathcal{S}$ arises as the intersection of colored simplices (three of them suffice here), as claimed by Proposition 6.19.

The Figure 6.10 for this precise example with three polyhedra in \mathbb{R}^2 shows that the set $\text{colint } \mathcal{S}$ of colorful vectors is a triangle, that does not arise as any colored simplex of \mathcal{D} though. This property is in fact general, as showed by the following characterization theorem that is a corollary of [LS09, Theorem 2]. Given an affine (and oriented) hyperplane $H := \{x \in \mathbb{R}^{n-1} : \langle h, x \rangle = b\}$, we shall denote by $H^>$ (resp. H^\leq) the open (resp. closed) halfspace $\{x \in \mathbb{R}^{n-1} : \langle h, x \rangle > b\}$ (resp. $\{x \in \mathbb{R}^{n-1} : \langle h, x \rangle \leq b\}$).

- **THEOREM 6.20.** *Let $\mathcal{S} = (S_1, \dots, S_n)$ be a collection of n closed convex sets of \mathbb{R}^{n-1} , and assume that $\text{colint } \mathcal{S}$ is nonempty. Then, $\text{colint } \mathcal{S}$ is the interior of a $(n - 1)$ -dimensional simplex. Moreover, if the sets $(S_i)_{1 \leq i \leq n}$ are bounded, then there are n unique hyperplanes $(H_i)_{1 \leq i \leq n}$ such that for all i in $\{1, \dots, n\}$, $S_i \subset H_i^>$, and for all $j \neq i$, $S_j \subset H_i^\leq$ and $S_j \cap H_i \neq \emptyset$. In this case, we have $\text{colint } \mathcal{S} = \bigcap_{1 \leq i \leq n} H_i^>$.*

Geometrically, every H_i in Theorem 6.20 is a tangent hyperplane to the convex sets $(S_j)_{j \neq i}$ which separates them from the set S_i , and in particular we have $\widehat{S}_i \subset H_i^\leq$ for all $1 \leq i \leq n$. The existence (and uniqueness) of the tangent hyperplanes follows from the work of Cappell et al. [CGP⁺94], revisited by Lewis, Klee and von Hohenbalken [LvHK96] who gave a simpler and

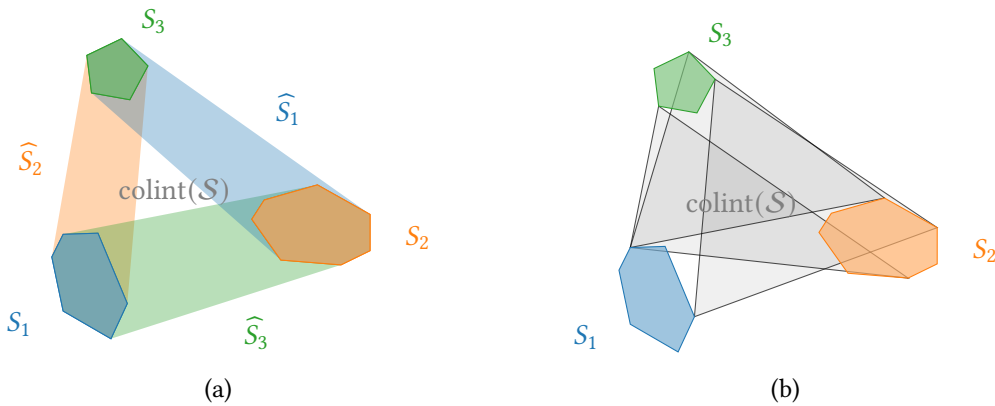


Figure 6.10: Visualizing the colorful interior of three polyhedral in \mathbb{R}^2 using the definition (6.6) (left) or the Proposition 6.19 (right).

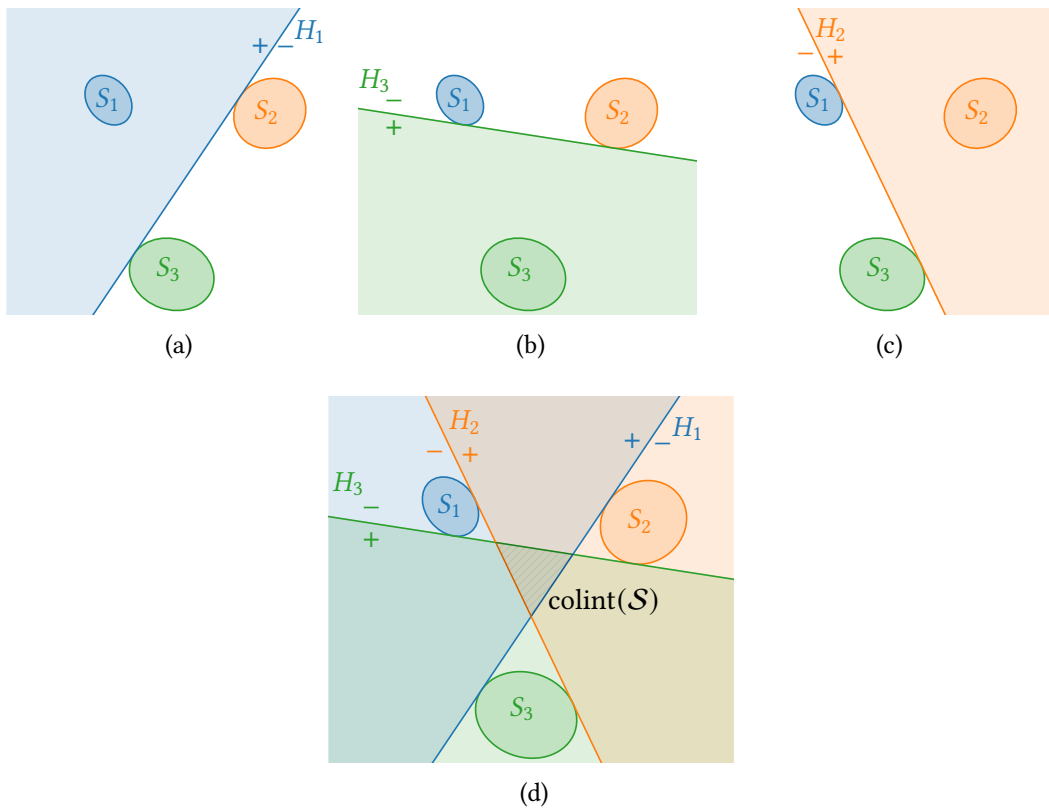


Figure 6.11: The colorful interior of S_1, \dots, S_n convex bodies of \mathbb{R}^{n-1} is a simplex whose faces are supported by $\widehat{S}_1, \dots, \widehat{S}_n$

constructive proof. We depict in the Figure 6.11 three convex sets S_1, S_2 and S_3 in \mathbb{R}^2 that this are time are not polyhedra. Their nonempty colorful interior $\text{colint}(S_1, S_2, S_3)$ remains a triangle (two-dimensional simplex) as claimed by Theorem 6.20. Observe how the maximal faces of this simplex (edges of the triangle) are supported by tangent hyperplanes to all but one bodies.

2.2 NONVACUITY CONDITIONS. — Given a collection $\mathcal{S} = (S_1, \dots, S_n)$ of n closed convex sets of \mathbb{R}^{n-1} , we now discuss necessary and sufficient conditions for $\text{colint}(\mathcal{S})$ to be nonempty. Indeed, we have so far only represented configurations of convex bodies for which the colorful simplex exists, but it is likely to be empty in general. We can easily understand that in order to have $\text{colint}(\mathcal{S}) \neq \emptyset$, it is necessary that the $(S_i)_{1 \leq i \leq n}$ are pairwise disjoint. However, as shown by the Figure 6.12b, this is way not sufficient. Even imposing that $S_i \cap \widehat{S}_i \neq \emptyset$ for all $1 \leq i \leq n$ does not ensure the nonvacuity of $\text{colint} S$, so a stronger notion on the relative positions of the $(S_i)_{1 \leq i \leq n}$ is needed. It looks from the Figure 6.12d that it is worth focusing on the “pointy” convex bodies delineated by $n - 1$ convex hulls among the $(\widehat{S}_i)_{1 \leq i \leq n}$, i.e., the sets $(\overline{S}_i)_{1 \leq i \leq n}$ where for all i in $\{1, \dots, n\}$, $\overline{S}_i := \bigcap_{j \neq i} \widehat{S}_j$. Indeed in Figure 6.12, the “spikes” of the $(\overline{S}_i)_{1 \leq i \leq n}$ seem to control the size of the colorful interior, and if they overlap, $\text{colint}(\mathcal{S})$ collapses and becomes empty.

Recall (for instance from [GP88]) that the collection \mathcal{S} is **separated** if for any choice of $k \leq n$ points x_1, \dots, x_k in $S_{i_1} \times \dots \times S_{i_k}$ (and pairwise distinct i_1, \dots, i_k), the points x_1, \dots, x_k are in general position (spanning a $(k - 1)$ -dimensional affine space). We have the following:

► **PROPOSITION 6.21.** *Let S_1, \dots, S_n be a collection of n compact convex sets of \mathbb{R}^{n-1} . The collection $(\overline{S}_i)_{1 \leq i \leq n}$ is separated if and only if $\bigcap_{1 \leq i \leq n} \widehat{S}_i = \emptyset$.*

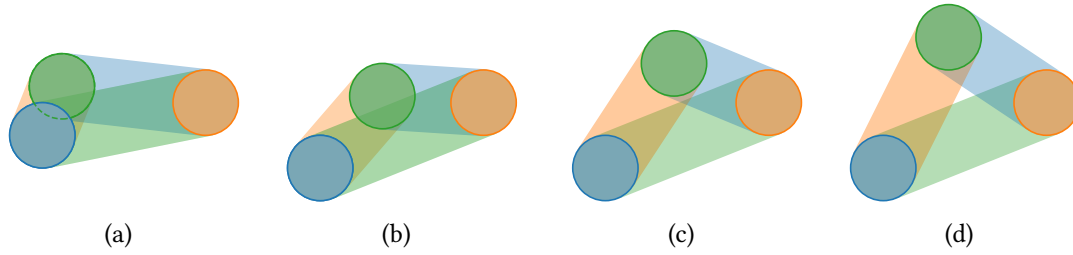


Figure 6.12: If the convex bodies are too close to each other, their colorful interior is empty

Proof. Recall that the collection \mathcal{S} is separated if and only if for all $k+l \leq n$, any two disjoint k -subcollection and l -subcollection of \mathcal{S} can be separated by an affine hyperplane (this result, already stated in [BHJ08] for example, is easy to show).

- ◇ We first suppose that $\bigcap_{1 \leq i \leq n} \widehat{S}_i = \emptyset$, then for all partition (I, J) of $\{1, \dots, n\}$, we have $(\bigcap_{i \in I} \widehat{S}_i) \cap (\bigcap_{j \in J} \widehat{S}_j) = \emptyset$. The convex sets separation theorem ensures that there exists an affine hyperplane H separating the convex sets $\bigcap_{i \in I} \widehat{S}_i$ and $\bigcap_{j \in J} \widehat{S}_j$. However, for all $j \in J$, $\bar{S}_j := \text{conv}(\{\bar{S}_i\}_{i \in I}) \subset \widehat{S}_j$, as a result $\bar{S}_j \subset \bigcap_{j \in J} \widehat{S}_j$. In particular, H separates the two subcollections $(\bar{S}_i)_{i \in I}$ and $(\bar{S}_j)_{j \in J}$. If (I', J') are two disjoint subcollections of $\{1, \dots, n\}$ such that $I' \cup J' \neq \{1, \dots, n\}$, we can still complete (I', J') in a partition of the form (I, J) to get as above a hyperplane H separating $(\bar{S}_i)_{i \in I'}$ and $(\bar{S}_j)_{j \in J'}$. From the proposition we have recalled on the separation property, the collection $(\bar{S}_i)_{1 \leq i \leq n}$ is separated.
- ◇ Conversely, the separation of $(\bar{S}_i)_{1 \leq i \leq n}$ implies that for all $i \in \{1, \dots, n\}$, there exists some affine hyperplane separating \bar{S}_i from the collection $(\bar{S}_k)_{k \neq i}$. By convexity, this hyperplanes separates as well \bar{S}_i and $\text{conv}(\bigcup_{k \neq i} \bar{S}_k) = \widehat{S}_i$. As a result $\bar{S}_i \cap \widehat{S}_i = \emptyset$, or equivalently $\bigcap_{1 \leq i \leq n} \widehat{S}_i = \emptyset$. \clubsuit

Note that the criterion $\bigcap_{1 \leq i \leq n} \widehat{S}_i = \emptyset$ is advantageous in practice because it can be assessed in (weak) polynomial time via linear programming. The interest of the Proposition 6.21 comes from the fact that the separation of the $(\bar{S}_i)_{1 \leq i \leq n}$ is a strong necessary condition for nonvacuity of the colorful interior.

- **PROPOSITION 6.22.** *Let S_1, \dots, S_n be a collection of n compact convex sets of \mathbb{R}^{n-1} . Then, if $\text{colint } \mathcal{S}$ is nonempty, the collection $(\bar{S}_i)_{1 \leq i \leq n}$ is separated.*

Proof. We will show by contraposition that $\text{colint } \mathcal{S} \neq \emptyset \implies \bigcap_{1 \leq i \leq n} \widehat{S}_i = \emptyset$. Suppose that we have x_0 such that for all $i \in \{1, \dots, n\}$, $x_0 \in \widehat{S}_i$. Let $x \in \text{conv}(\mathcal{S})$. If $x = x_0$, it is clear that $x \notin \text{colint } \mathcal{S}$. Otherwise, consider the halfline $[x_0, x)$, which is not entirely included in $\text{conv}(\mathcal{S})$ because all the $(S_i)_{1 \leq i \leq n}$ are bounded. We define $x' := \max_{z \in [x_0, x)} \{z \in \text{conv}(\mathcal{S})\}$, in the sense of the order induced on $[x_0, x)$. Note that the maximum is indeed attained by closedness of the $(S_i)_{1 \leq i \leq n}$. By definition, we have $x \in [x_0, x')$ and $x' \in \partial \text{conv}(\mathcal{S})$. The latter ensures that $x \in \text{conv}(\bigcup_{1 \leq i \leq n} (S_i \cap H))$ where H is a supporting hyperplane to $\text{conv}(\mathcal{S})$. Hence, from Carathéodory's theorem, there is $k \in \{1, \dots, n\}$ such that $x' \in \widehat{S}_k$ (see the proof of Lemma 6.18 where it is also done). But by assumption, x_0 is in \widehat{S}_k as well and so does x by convexity. We have thus just proved that for all $x \in \text{conv}(\mathcal{S})$, $x \in \widehat{S}_k$ for some $k \in \{1, \dots, n\}$, which implies that $\text{colint } \mathcal{S} = \emptyset$. Proposition 6.21 terminates the proof. \clubsuit

We conjecture that the necessary condition of Proposition 6.22 is actually sufficient, so that the following would hold:

- **CONJECTURE 6.23.** *Let S_1, \dots, S_n be a collection of n compact convex sets of \mathbb{R}^{n-1} . Then $\text{colint } \mathcal{S}$ is nonempty if and only if the family $(\bar{S}_i)_{1 \leq i \leq n}$ is separated.*

This conjecture is easy to show for $n = 2$, either in the conic setting by describing the cones C_1 and C_2 by means of polar angles, or in the affine setting where convex hulls become mere segments. We prove the conjecture in the case where $n = 3$, taking advantage of linear algebra formulas involving cross products.

► **PROPOSITION 6.24.** *Let $\mathcal{S} = (S_1, S_2, S_3)$ be a collection of three convex compact sets of \mathbb{R}^2 . Then, colint \mathcal{S} is nonempty if and only if $(\bar{S}_1, \bar{S}_2, \bar{S}_3)$ is separated.*

Proof. Suppose that $(\bar{S}_1, \bar{S}_2, \bar{S}_3)$ is separated. We know from [LvHK96] that for all $i \in \{1, 2, 3\}$ we have two hyperplanes (in this case affine lines) tangent to sets of the collection $(\bar{S}_j)_{j \neq i}$ and inducing opposite orientation on these. Such lines cannot meet \bar{S}_i by separation property, so one of them, denoted H_i , is such that $\bar{S}_i \subset H_i^>$ and $\bar{S}_j \subset H_i^{\leq}$ for $j \neq i$. In particular, note that $\text{conv}((S_j)_{j \neq i}) \subset H_i^{\leq}$. For $i, j \in \{1, 2, 3\}$ and $j \neq i$, the hyperplane H_i is not only tangent to \bar{S}_j but also to S_j . Indeed, take a support y_i^j of H_i in \bar{S}_j , it arises as a convex combination $y_i^j = \sum_{k \neq i} \lambda_k x_k$ with $x_k \in S_k$ for $y_i^j \in \bar{S}_i$. By $S_i \subset \bar{S}_i$, we derive for all $k \neq i$, $x_k \in H_k$ or $\lambda_k = 0$, the latter being ruled out by separation. Hence, let us denote by x_i^j a support of hyperplane H_i in S_j . Note that once again from the separation of $(\bar{S}_1, \bar{S}_2, \bar{S}_3)$, two supports of a tangent line in two different colors cannot be equal.

If $x := (a, b)^\top$ and $y := (a', b')^\top$ are two distinct vectors, we denote $x \wedge y := (ab' - a'b)^{-1}(b - b', a' - a)^\top$ their usual cross-product. As is customary, $h_1 := x_1^2 \wedge x_1^3$ (resp. $h_2 := x_2^3 \wedge x_2^1$ and $h_3 := x_3^1 \wedge x_3^2$) is a normal vector to H_1 (resp. H_2 and H_3), and $\langle h_i, x \rangle + 1 = 0$ is an equation defining H_i . Furthermore, the intersection of H_1 and H_2 is given by $s_3 := h_1 \wedge h_2$, or using the triple product formula, by

$$s_3 = h_1 \wedge (x_2^3 \wedge x_2^1) = \frac{(\langle h_1, x_2^1 \rangle + 1) x_2^3 - (\langle h_1, x_2^3 \rangle + 1) x_2^1}{(\langle h_1, x_2^1 \rangle + 1) - (\langle h_1, x_2^3 \rangle + 1)}. \quad (6.7)$$

Because $x_2^1 \in \bar{S}_1 \subset H_1^>$ and $x_2^3 \in \bar{S}_3 \subset H_1^{\leq}$, we get $\langle h_1, x_2^1 \rangle + 1 \neq 0$ and $(\langle h_1, x_2^1 \rangle + 1)(\langle h_1, x_2^3 \rangle + 1) \leq 0$. As a result of (6.7), s_3 indeed exists and arises as a convex combination of x_2^3 and x_2^1 , so $s_3 \in \text{conv}(S_1 \cup S_3)$. By writing $s_3 = (x_1^2 \wedge x_1^3) \wedge h_2$ as in (6.7), we show likewise that s_3 is a convex combination of x_1^2 and x_1^3 , thus $s_3 \in \text{conv}(S_2 \cup S_3)$. This finally entails that $s_3 \in \bar{S}_1 \cap \bar{S}_2 = \bar{S}_3$ and therefore $s_3 \in H_3^>$. It now suffices to define $s_1 := h_2 \wedge h_3$ and $s_2 := h_3 \wedge h_1$ in a similar way and consider $y = (s_1 + s_2 + s_3)/3$. It is clear that $y \in \text{conv}(S_1 \cup S_2 \cup S_3)$, and for all $i \in \{1, 2, 3\}$, $y \in H_i^>$, in particular $y \notin \text{conv}((S_j)_{j \neq i})$. As a consequence, y is a colorful vector for S_1, S_2 and S_3 . \clubsuit

Proving this conjecture for $n \geq 4$ remains an open problem. Indeed, even though multiple-cross-product identities exist in higher dimension, they feature determinants whose sign is not easy to control. More generally, the sets $(\bar{S}_i)_{1 \leq i \leq n}$ are difficult to work with as they do not admit straightforward description in terms of the initial sets $(S_i)_{1 \leq i \leq n}$. We point out that another interesting problem is the computational complexity of determining whether the colorful interior is empty or not, in the case where the sets $(S_i)_{1 \leq i \leq n}$ are polytopes. The problem could be tackled by studying the complexity of separating a point from the colorful interior. This is tightly linked with the computation of the tangent hyperplanes of Theorem 6.20, for which the complexity status is not well understood.

We also mention that when we do not assume that $\text{cone}(C_1 \cup \dots \cup C_n)$ is pointed, the results highlighted in this section do not hold anymore, in particular the colorful interior of the cones $(C_i)_{1 \leq i \leq n}$ can have several connected components. See for instance our Example 6.9, where the non-pointed configuration of vectors of Figure 6.5a turns out to yield a colorful interior that is the union of two open cones.

Part II

**A COMPREHENSIVE STUDY OF
A REAL-LIFE AND LARGE-SCALE
EMERGENCY CALL-CENTER**

FOREWORD

AND CONTEXT OF THE STUDY

This second part of the thesis is drawn from the report “*Optimisation de la performance de centres de traitement d’appels d’urgence*” [ABCG21], a work funded by the French *Institut des Hautes Études du Ministère de l’Intérieur* (IHEMI) in 2020. This project relied on the collaboration of our research team with the *Préfecture de Police de Paris* (PP). The latter supervises the *Direction de la Sécurité de Proximité de l’Agglomération de Paris* (DSPAP) that operates Parisian police stations, and also the *Brigade de Sapeurs Pompiers de Paris* (BSPP) who is in charge of fires and rescue operations. We were asked to carry out a quantitative evaluation of the *Plateforme d’Appels d’Urgence* (PFAU). This emergency call platform receives all police and rescue emergency calls across four French departments (the city of Paris and its inner suburbs) through the emergency lines 17, 18 and 112. We refer to the [Introduction](#) of the thesis for detailed discussion on the different actors of emergency in France and the Parisian geography.

We recall below the different purposes served by the initial study [ABCG21]. Although we already presented a bit the design of the PFAU in the Section 2 of Chapter 5, we provide in this introductory section a much more in-depth description of the handling chain of calls. Being familiar with this layout (and some acronyms!) is essential before reading Chapters 7 and 8. Because the study is quite dense, we also propose a summary of our results at the end of this foreword.

1 Object of the study

The goal of the study [ABCG21] is to develop decision-support and optimization tools for operations management of emergency call centers. These tools must cover “nominal” (or standard) situations but may also be applied in response to an atypical event, planned or not.

Our approach builds on (extended) timed Petri net models of call centers, similar to those studied in the Part I of the thesis – but even more flexible. As already mentioned, Petri nets come in handy to account for synchronization and concurrency phenomena (wait of a resource, shared resource between several tasks,...). The models considered here capture additional key features of emergency systems, such as priority mechanisms (for the most urgent calls), or the limited patience of people calling.

All of our models first take advantage of an in-depth analysis of the PFAU phone logs data. This step enables us to get a clear understanding of the organization – the different circuits of the calls depending on their characteristics, but it also yields precise numerical values of all the call center parameters: handling times by the operators, patience of callers, intensity of arrivals, etc. This statistical analysis of emergency calls constitutes our Chapter 7. Beyond its interest

for modeling purposes, it already shows several strengths of the current PFAU design. These important results are summarized in the Section 3 below.

Then, we use a numerical simulation approach, running a software developed in our TROPICAL research team (Inria and CMAP, École polytechnique), specifically for this application. This software, whose design is presented in the Section 1 of the Chapter 8, reproduces *in silico* the response given by the call center to emergency demands. It takes into account all the characteristics of the model, and yields a performance analysis whose precision level is directly linked to the one of the statistical analysis of calls data. This is an independent and complementary approach from the theoretical formulas derived in the first part of the thesis.

We use these two techniques to evaluate the performance of the current “two-tier” (or bilevel) PFAU organization, and to compare the latter with alternative designs, in particular to a more classical “single-tier” layout. We also study the impact on the performance of several modifications of the organization, such as the unification of previously separated agents pools. We also evaluate the response of bilevel call centers to unusual peaks of calls. All these results are presented in the Chapter 8 and summed up in the Section 4 below.

2 Presentation of the PFAU emergency platform

The *Plateforme d’Appels d’Urgence* (PFAU) is located in the Parisian firebrigade headquarters (Caserne Champerret, 17th district of Paris). It is operated by both the BSPP and the DSPAP, under the authority of the PP. It receives emergency calls from three emergency numbers: line 17 for police requests, line 18 for fires and rescue needs, and line 112 (European number for all emergency needs). All calls emitted from Paris and its inner suburbs (French departments named Hauts-de-Seine, Seine-Saint-Denis and Val-de-Marne) arrive at this same PFAU site, which represents more than 8,000 calls per day – or during the day up to 400 calls per hour. Before this common PFAU call center existed, the calls were dispatched to five physically different call centers – one for the fire brigade for the whole area, and one in each department for the police.

The PFAU was officially inaugurated in 2016, in a context where there is an increasing demand on the 17–18–112 lines with very varied call motives – a significant part of them is actually just about giving advice to the population and does not trigger an intervention. In addition, the Paris area experienced a series of one-off events (demonstrations, attacks...) where massive peaks of calls were received. The PFAU was therefore designed to address a triple goal:

- improve the ability to filter out the less urgent calls (but still offering them an appropriate answer) to handle more efficiently life-threatening emergencies and other prioritized calls
- enhance the robustness and the resilience of the instruction chain, in particular in response to exceptional bursts in arrivals
- take advantage of the mutualization of sites and services.

To fulfill these objectives, an innovative design in which **calls are picked up in a two-step process** has been chosen for the PFAU:

- callers are first welcomed by a voice prompt (which lasts for less than 10 seconds), explaining that an emergency operator is going to pick up their call
- at the first level (Lvl 1), a first agent picks up the call with the mission to swiftly qualify the demand of the caller and identify its gravity. To this purpose, these first-level answering

agents use a comprehensive decision tree. The calls can either be qualified as urgent (U), very urgent (VU) or non urgent (NU). The urgent and very urgent calls are transferred to the second level of instruction, with a higher **priority** degree for VU calls. The non urgent calls are given advice (or rerouting suggestions) by the first-level operator, and are not handed over to the second level.

- at the second level (LVL 2), a service expert (either a police officer, either a firefighter) handles the call transmitted by the LVL 1 agent. This LVL 2 operator can dispatch on-site response teams and can also reach out to partners (SAMU, rail or road traffic services...). Second-level agents pick up in priority the very urgent calls, labeled as such by operators at LVL 1. These calls are transferred to LVL 2 with a special protocol, hereafter referred to as the *accompagnement* or *escorting mechanism*. Indeed, when they identify a very urgent call, the agents at LVL 1 stay on line with the caller until a second-level operator is available. The call is then transferred in a short synchronization step, where a “threeway” conversation can take place. In contrast, the calls tagged U (and not VU) by the first level are not escorted to the LVL 2, and instead placed in a virtual waiting room until picked up by second-level operators.

This architecture is referred to as **bilevel**, in opposition with a “monolevel” or **single-tier** system where all the operators pick up and handle all calls from start to end. In addition of this two-level organization, the PFAU has the specificity to bring together two different services – police and rescue missions:

- first-level police operators (agents of the DSPAP service) are instructed to pick up calls from emergency number 18 and also line 112. Then, second-level police officers handle the demands handed over by the LVL 1 that fall within their missions
- first-level firebrigade operators (agents of the BSPP service) are instructed to pick up calls from emergency number 17 and also line 112. Then, second-level firefighters handle the demands handed over by the LVL 1 that here as well fall within their missions.

The two types of first-level agents thus share the pick up of calls from line 112, and work side-by-side in the same room. We talk about a **multi-job** organization. The DSPAP and BSPP chains are actually not symmetrical, because currently, the instruction of police calls at LVL 2 by the DSPAP is *departmentalized* – there is a single first level, but four second levels, one for each department. On the contrary, BSPP uses a common first level and a common second level for the four departments it receives calls from. We shall explain in the Section 4.3 of Chapter 8 how this structural difference affects the performance. A last specificity of the BSPP that we will briefly discuss in the Chapter 7 is to have a “*coordination médicale*” staffed with physicians. The second-level firefighters or on-site dispatched teams can request the help on the phone of these physicians for health-related demands. However, such requests are most of time the area of expertise of the SAMU organization (emergency number 15) and therefore transferred to them.

3 Summary of the data analysis step

We recap below the main takeaways from the analysis of phone calls data¹ of our Chapter 7.

- We first point out that this analysis benefited from a huge amount of very precise data, which brought a whole development in itself. Having a complete year of emergency calls and consolidated datasets gives to this study a very high degree of confidence.
- The analysis of the throughputs of call arrivals in the Section 2 confirms the high level of “pollution” on line 112, well-known by the operators. It also shows that many wrong number errors are filtered out by the welcoming voice prompt. This protects the first-level agents from an addition pollution.
- A great attention was brought to the exploration of dependence relations between the characteristics of the PFAU. We summarize them in the dependence matrix depicted in Figure II.1. Beyond slight territorial variation, this matrix recalls that the different handling times at the different steps of instruction of a same call are correlated. This is an essential phenomenon to take into account to ensure the realism of simulations.

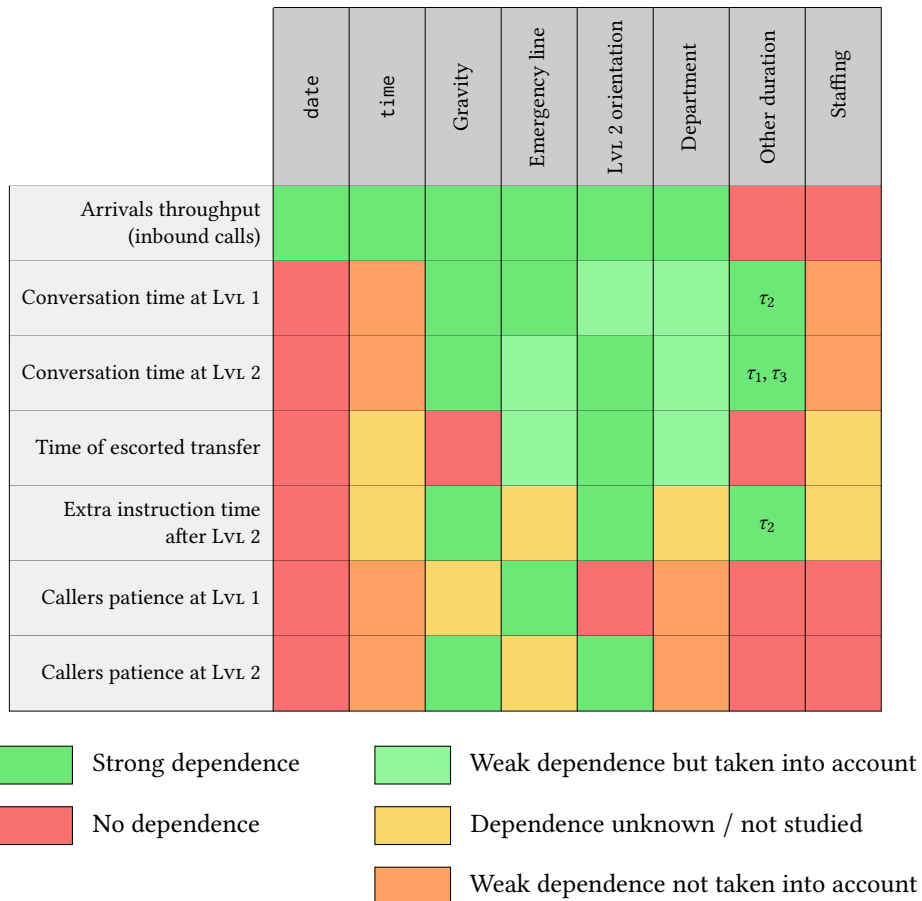


Figure II.1: Paramaters dependences matrix

¹ Performed with the pandas [pDT20] and SciPy [WM10] libraries for Python. We mention by the way that like the rest of the thesis, all the charts were made with Matplotlib [Hun07].

- While the handling times of the non urgent calls are not identically distributed across the four departments, they are very similar for urgent or very urgent calls (see Section 4). This shows that the notion of emergency or serious distress is felt and dealt with evenly in the whole Paris area.
- The measurement in Section 4.4 of an “extra instruction time” at LVL 2 after the caller hangs up (when contacting another partner is needed) is an originality of our study. It allows us to represent more faithfully the workload of LVL 2 agents. The strong uniformity of this additional treatment for very urgent calls is characteristic of a standardized and dedicated protocol implemented for those calls.
- We show in Section 2.2 that assuming in simulations that call arrivals follow (nonhomogeneous) Poisson laws is an excellent approximation of the reality.
- The study of “cross-flows” $17 \rightarrow$ BSPP and $18 \rightarrow$ DSPAP shows the great resilience of the PFAU relatively to calls for which the initially dialed number was not the most appropriate one. Indeed, the rerouting of these calls happens very quickly.
- The patience analysis that is presented in Section 5 proves that the patience of callers is greatly enhanced after the first pick up at LVL 1; this is an intrinsic advantage of the bilevel layout.
- The analysis of the accompagnement / escorting mechanism for very urgent calls demonstrates its major interest to prevent inter-level abandonments. We show that there are seven times less such abandonments thanks to this feature.
- The study of the number of logged agents (see Section 3) illustrates how quick the PFAU can adapt its staffing to cope with events. It raises a focus point on the transitions between operator shifts – that managers of the call center know how to handle.

4 Summary of the comparative study of organizations

Our Chapter 8 aims at providing **quantitative** performance evaluation for emergency call centers such as the PFAU. To this end, we resort to two complementary approaches:

- **numerical simulation**, which mimics the behaviour of the real system and makes as few assumptions as possible. It allows us to **empirically** compute a minimal staffing or performance levels (average waiting time, percentage of picked up calls...) for a given architecture. This is done by running simulations of scenarios that are very close and faithful to call arrivals of the PFAU – that we can accurately determine thanks to the data analysis step presented in the Chapter 7.
- **analytical computations**, that were extensively presented in the first part of the thesis. This approach yields mathematical formulas that also determine minimal staffing levels of a call center in terms of its parameters. It can also evaluate several performance estimates, but sometimes requires simplifications to effectively do the computation.

These two methods inherit the modularity of the Petri net formalism on which they hinge. They enable us to benchmark varied and concrete design of emergency call centers, in order to compare their respective performance.

Note that the **qualitative** performance of the organization are not in the scope of this work. This for instance includes the improvement of work practices made possible by bringing two

different services under the same roof, the working comfort of agents brought by the organization design, or even the perception of the two-tier handling chain by the callers. All these aspects are indeed not measurable by mathematical or statistical tools. In the case of the PFAU, the analysis of such questions is left to their own teams and to the *Préfecture de Police*. In particular, a couple of quantitative conclusions should not be reproduced without the complementary enlightenment of qualitative comments.

SINGLE-TIER / BILEVEL COMPARISON (SINGLE-JOB ORGANIZATIONS). — We first compare the performance of single-tier and bilevel systems in a fictitious but interesting configuration, where the call center receives only the calls corresponding to one type of job (either police or rescue matters). This single-job comparison is represented in the Figure II.2 in the case of rescue calls (emergency number 18 and most calls from 112). The Section 2 of Chapter 8 is entirely devoted to this simplified case, and plays the role of “scientific demonstrator” of our methodology.

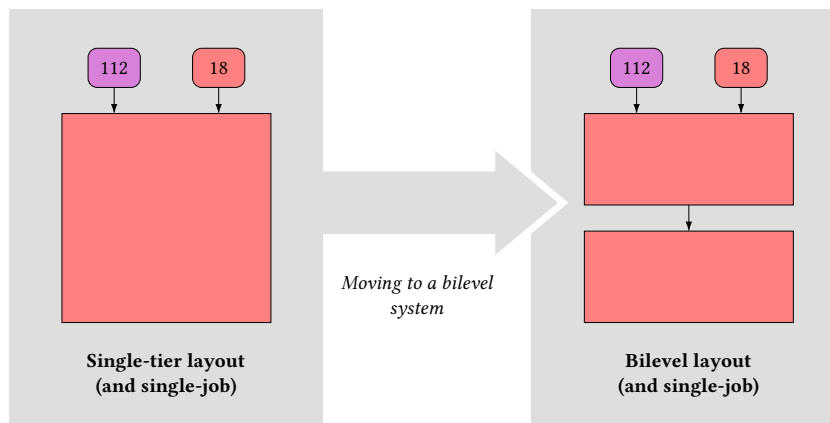


Figure II.2: Comparison between single-tier and bilevel systems

We point out that this first comparison is also of academic interest. For instance, the EUROCALL study [NCM⁺₁₇] also focuses on this aspect, by comparing the total time needed to access a dispatcher agent in 21 medical emergency call centers across 11 countries. The authors conclude that from this perspective, single-tier architectures are more advantageous. To the contrary, we show that **well-tuned** bilevel layouts yield lower abandonment rates than single-tier systems, without increasing too much the access time. We show in fact that several necessary conditions need to be met for the bilevel design to outperform its single-tier counterpart, and we reckon that they may explain why our global results diverge from those of [NCM⁺₁₇].

For each service (police or rescue), we first carry out **iso-staffing comparisons**. This means that the bilevel system is compared to a single-tier organization with the same total number of agents (see Sections 2.3 and 2.4). However, this gives an artificial advantage to the single-tier system, because only the Lvl 2 agents of the bilevel have the sufficient level of expertise to also be staffed in the single-tier system. For this reason, we also draw another “iso-skills” comparison, in which the two organizations have the same number of expert agents – note that a fairer comparison would be a mix of these two approaches. Our conclusions are the following:

- in the case of rescue calls (firefighters from BSPP picking up calls from lines 18 and 112), we show in the Section 2.3 that regardless of the number of staffed operators, the bilevel organization is always more efficient than the single-tier (even for the iso-staffing comparison)

when looking at the global abandonment rate of calls, or at the abandonment rate of very urgent calls. However, this better performance on abandonments is at the expense of an increased total waiting time for callers and a small increase of occupancy for agents. With the current headcounts (14 or 15 firefighters in total), we note that both single-tier and bilevel modes are equivalent and ensure almost no abandonment and no wait.

- in the case of police calls (police officers from DSPAP picking up calls from line 17), we show in the Section 2.4 that the single-tier system yields better performance than the bilevel when studied under the iso-staffing perspective and in particular if there are not enough agents. With current staffings, both architectures still give very good overall performance. This important difference with the rescue lane is explained by a different **loads ratio**. To sum it up, the LVL 1 of the police is relatively long with respect to the LVL 1 of the firebrigade, while the LVL 2 of the police is relatively short compared to the LVL 2 of the firebrigade. This is due to the fact that there are many non urgent calls in the police lane and these are given long instructions. The bilevel organization would be much more advantageous to the police treatment chain if those calls were handed over to another platform after identification – a non-emergency call center. Remark however that in an iso-skill comparison, the bilevel police call center systematically outperforms the single-tier.

We will also keep in mind the following conclusions, that apply indifferently to the two services of the PFAU and would also apply to other models adopting such layouts:

- most performance criteria of the single-tier system are highly correlated with each other, and the same goes for the metrics relative to each level of the bilevel system. Yet, the LVL 1 criteria and the LVL 2 criteria are in first approximation independent, all the more so when the system is fluid.
- the study of mixed performance criteria for the bilevel system in the Section 2 shows the major protection offered to very urgent calls, even when the LVL 2 is a bit understaffed. This is an essential robustness element of the bilevel layout, also valid when facing rises in the demand.
- the same section shows the importance to correctly size both levels and reveals for a given total headcount that there exists an optimal allocation of agents across the two levels, in order to guarantee the best behaviour of the system. This compromise may however depend upon the evaluation criterion that we choose.

PERFORMANCE OF THE PFAU AND OF ITS POSSIBLE EVOLUTIONS. – Contrary to the “proofs of concept” covered in the Section 2 – using in particular single-job simplifications, the Section 4 of Chapter 8 only simulates the real PFAU organization or some possible alternative designs for it. In particular, the interactions between the two jobs (police and firefighters) are always considered. We study the four organizations depicted in the Figure II.3, which differ from one another by the fact that some agent pools that used to be separated can be merged. Our main results are the following:

- the Section 4.2 establishes an important interest of having a collaboration between two services at the PFAU. It shows that there is an indirect support of the first level of each service brought to the whole treatment chain of the other job, by adjusting the way that the pick ups of calls from line 112 are shared. This simple mechanism diminishes the number of people needed at LVL 1 as shown by the Table 8.30.

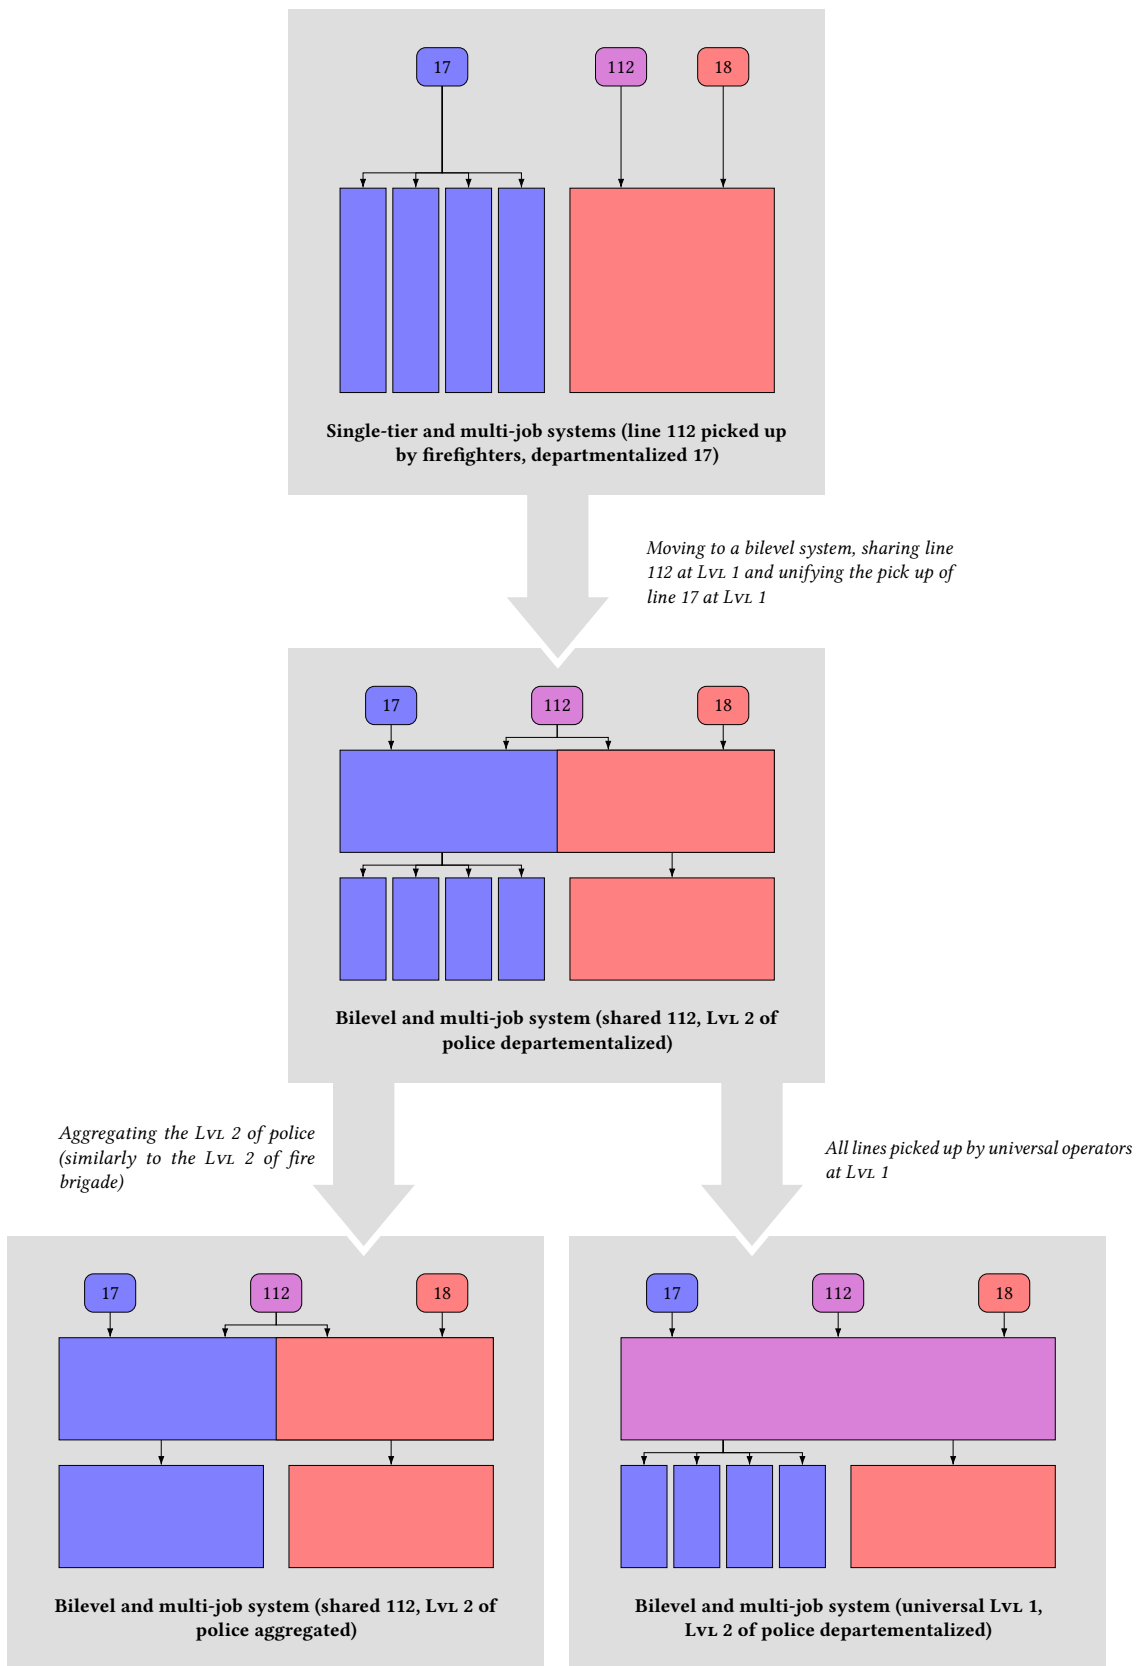


Figure II.3: Comparison of multi-job organizations

- the Section 4.3 demonstrates that aggregating the four LVL 2 of DSPAP that currently exist into a single LVL 2 (on the model of the BSPP organization) would yield a very significant gain of performance by keeping the same headcount, or alternatively would require much less operators to offer the same performance. We also show that merging several pools of first-level agents in the bilevel layout improves the robustness of the system. This would also apply to the case of several single-job and small-scale call centers willing to share first-level resource.
- we show in addition in the Section 4.4 that for the current PFAU, merging the two LVL 1 (which only share the number 112) into a single one would also require less agents (about two first-level operators) while maintaining the same performance.

We also study two situations representative of atypical arrivals of calls, to see how single-tier and bilevel layouts react to **crisis situations**:

- in the Section 5.1, we simulate the breaking out of a fire causing a massive burst of calls (peaking at 350 calls per hour) tagged as “event” by the rescue operators, but that in majority are not urgent (calls from witnesses seeing the fire from a distance). We show that in response to this event that the bilevel system gives a better protection to the very urgent calls compared to the single-tier layout, even though there are more (less urgent) abandonments with the bilevel mode. We also show that the bilevel architecture is more robust to worst-case scenarios.
- in the Section 5.2, we study how the organization behaves in case of a huge but predictable demand such as the one arising the night after Bastille Day on rescue lines (this event causes up to 650 calls per hour at the peak). We show that the bilevel system remains the best to protect the very urgent calls, but is less performant than the single-tier under the iso-staffing comparison for the global abandonment criterion. This is due to a lower loads ratio (less adapted to the bilevel) and to an increased proportion of calls needing to access the LVL 2.

We finally review how the theoretical methods agree with our simulations results. Confronting the numerical simulation approach with the analytical formulas indeed enables us to appreciate in which extent results predicted by the theory turned out to be accurate. Recall that these analytical estimates are often obtained by doing strong assumptions and approximations. Still, we can observe that:

- the theoretical throughputs of Petri nets are good estimates of the real handling speeds of the simulated organizations. This is all the more true if the number of agents are far away from the frontiers of the fluid phase of our theoretical congestion diagrams.
- the intuitive law that suggests to divide the minimum theoretical staffing levels (to reach the fluid phase) by the target occupancy of each agent type in order to derive a realistic staffing policy is well verified in practice.
- even though several ground assumptions used to establish Erlang formulas are not verified for real call centers, these analytical estimates that we recall in the Section 3.3 of Chapter 8 have a good predictive power for single-tier systems if we want to evaluate the probability that a caller waits or hangs up. They can also be cascaded to correctly predict the behaviour of a bilevel system, but this only works if one of the two level is fluid.

7

DATA ANALYSIS OF PHONE LOGS

Contents

1	Description of used databases and global counts	252
2	Throughputs and proportions of calls	255
2.1	Dependences of the throughputs	255
2.2	Arrivals process	259
2.3	Callbacks from the callers	261
3	Identification of logged agents	261
4	Service times of calls	264
4.1	Conversation time at level one	264
4.2	Conversation time at level two	268
4.3	Escorted transfer	269
4.4	Additional instruction time at level two	270
4.5	Other factors of dependence	274
5	Patience of callers	275
5.1	Kaplan–Meier estimator	275
5.2	Comparison of patience functions	277
6	Some specific days	281
6.1	New Year’s Eve of 2019 on line 17	281
6.2	Bastille Day night of 2019 on line 18	281
6.3	Fire at Bobigny on 26 June 2019 on line 18	284
6.4	Storm on 4 June 2019 on line 18	286
6.5	Notre-Dame fire on 15 April 2019 on line 18	287

The objective of this chapter is to derive the main parameters that shall be used for the mathematical modeling and the numerical simulation. In particular, we will endeavour in simulations to generate scenarios that mimic the real PFAU organization as realistically as possible. This requires a fine statistical study of the call center features, notably their dependence with respect to exogenous parameters or even their inter-dependence.

Several studies already focused on the data analysis of call centers phone logs, as recalled in the review [ALo5]. This includes the study of the arrival processes of calls (see for instance [JKo1, KW14, VVR⁺17]), the modeling of service times by usual laws (with a good fit being found for the lognormal law [Chl97]), or the modeling of people calling back the platform [AAKD08]. We shall often refer to the in-depth study of Brown et al. [BGM⁺05b], which covers the previous aspects and also deals with the modeling of callers patience.

We give in Section 1 global counts and figures characterizing the volumes of calls handled by the PFAU. The Section 2 is devoted to the study of arrival processes of calls, distinguished per emergency number, department of origin and gravity. Staffing levels and their variations are reviewed in Section 3, while we focus on the distribution of conversation times (as well as their correlations) in Section 4. We study in Section 5 the patience of callers. Using realistic patience estimates in simulation will be a significant contribution of our work in the next chapter. Finally, we present in Section 6 a few profiles of arrival throughputs and service times for some days of interest in the year 2019, when unusual or unexpected events occurred.

1 Description of used databases and global counts

In everything that follows, we build on the calls emitted and received at the PFAU during the year 2019 (more precisely from the 26 December 2018 to the 5 January 2020, in order to cover the New Year celebrations). The data come from a first aggregate made by the BSPP, using the ticket files of phone calls in the Alcatel system recorded by the Ticket Extractor software. They contain many information on the origin and the duration of the calls at the first level of instruction of the PFAU and at the second level for BSPP calls. The second level of the DSPAP uses a different information system, and we had to get the corresponding data using the *Pegase* databases (one for each department). We went through a preliminary phase of data processing and “cleaning”, and we merged these different databases. This operation was based on the (hashed) phone number of the callers to link the calls arrived at LVL 1 and handed over to the LVL 2 of DSPAP. We also tagged the accompanied transfers (the threeway conversation for Very Urgent calls generates an additional call in the databases, and we need to reattach this one to the initial call). Currently, the *Pegase* data give less information than the Ticket Extractor software, hence there are a couple of analyses that we can carry out for the LVL 2 of BSPP but not for the LVL 2 of DSPAP.

The Table 7.1 indicates how the 8,5 millions of calls in these databases are distributed, depending on their incoming or outgoing character but also depending on their entry point. Out of 6 millions of outbound calls, more than 5 millions (that is 85.8 %) arrive at the LVL 1 with the label “Emergency phone number” and the explicit mention of the dialed number (17, 18 or 112). The calls arriving at the LVL 2 from partners of the PFAU account for 6.6 % of the incoming calls, and among these about one third comes from the four SAMU of Paris area. We have reported in the table the activity of BSPP’s “coordination médicale” (7.4 % of incoming calls) but its study is not in the perimeter of our work. As far as the inbound calls are concerned, there remain 11,332 calls (0.2 % of the total) that we could not fit in any of the previous categories.

The outbound activity (2,5 millions of calls in 2019) is mainly initiated by the LVL 1 (in 70.2 % of the cases), but the vast majority of these calls are actually just data artifacts associated with inter-level transfers or interconnections between *Pegase* and LVL 1’s layer – and they bring no additional information. We shall discuss the other calls emitted by the LVL 1 in the end of Section 3. Other outgoing calls are passed from the LVL 2 (16.4 % of the total, mainly to partners, see the Section 4.2), from the “coordination médicale” of BSPP (13.3 %) or do not fit in the previous categories (0.1 % of the total).

The Table 7.2 gives a more detailed vision of the repartition of incoming calls at LVL 1 – the main material of our study. Out of the 5,1 millions of emergency calls, we have chosen to rule out up to 280,000 of them for certain analyses. Indeed, we have distinguished five categories of “non standard” calls that could alter some computations (see the description)

The first three lines of the Table 7.2 illustrate the main trends for the operated emergency lines. The number 17 is the most solicited, but features a significant fraction of abandonment (36.3 %), however three out of four of these are done during the voice prompt. We note that 27.2 % of the given up calls occurred after 30 seconds of wait. Among the line-17 calls picked up at LVL 1, about a half (46.6 %) is filtered out by the first level, and 19.8 % of the calls transferred to the second level is tagged as Very Urgent (we shall come back on the different gravity proportions in the Section 2). There are less abandonments for the calls received on line 18 (and 85.9 % of them arise on the voice prompt), and above all, more than three quarters of them (77.1 %) are handed over to the LVL 2, with 12.1 % of Very Urgent calls among the transfers. In total, there are more calls on line 112 than on line 18, but the great majority of these (91.0 %) do not go at the LVL 2. This phenomenon is well-known by the operators, who talk about a high level of “pollution” on line 112. This is due to the fact that cell phones automatically dial the number 112 if the emergency function is activated, which can often happen unintentionally (pocket calls).

	Total number of calls : 8,553,592	
	Inbound calls	Outbound calls
All levels	5,977,939	2,575,653
At level 1	5,130,011	1,806,795
At level 2	393,200	423,861
<i>(At level 2 and to SAMU)</i>	<i>126,378</i>	<i>105,510</i>
At the “coordination médicale”	443,396	342,489

Table 7.1: Global counts of calls

		Incoming at LVL 1																		
		Abandonments				Picked up at LVL 1				LVL 1 transferred to LVL 2										
Total	Total	On prompt		After prompt		LVL 1 not transferred		LVL 1 transferred to LVL 2		Total	NU	N/R	Total	VU	U	LU	NQ	EVU	N/R	Other
		Total	> 30 s	Total	> 30 s	Total	N/R	Total	N/R											
17	2,066,765	36.30	76.62	23.38	27.19	63.70	46.56	99.48	0.51	53.43	19.80	74.57	0.88	4.70	0.03	0.02	0.01			
18	1,274,347	31.79	85.89	14.11	8.36	68.21	22.37	98.93	0.73	77.11	12.14	79.92	7.23	0.02	0.54	0.15	0.52			
112	1,773,012	52.80	84.39	15.61	1.02	47.20	90.99	99.87	0.12	8.96	11.12	84.72	3.10	0.38	0.25	0.22	0.05			
Other	15,887	74.27	81.70	18.30	1.16	24.10	97.68	0.00	65.74	1.67	32.81	7.81	1.56	0.00	0.00	43.75	0.65			
Total	5,130,011	41.00	81.90	18.10	14.18	58.99	51.96	99.36	0.50	47.87	15.80	77.57	3.93	2.33	0.28	0.09	0.17			
Ruled out	281,781	5.00	47.95	52.05	15.27	94.91	12.74	77.14	16.72	86.83	9.74	79.57	6.40	3.65	0.33	0.31	0.43			
lost	104,566	0.00				99.99	0.00			99.97	6.94	76.19	8.59	7.55	0.24	0.47	0.03			
no transfer	8,601	0.79	76.47	23.53	37.50	99.10	0.00			100.00	18.89	69.62	2.72	4.60	1.34	2.73	0.00			
callback	29,208	15.16	0.38	99.62	0.93	84.84	83.93	96.05	3.93	16.03	20.90	76.99	0.38	1.38	0.05	0.30	0.04			
extra steps	25,601	4.37	9.39	90.61	99.70	95.60	29.95	21.88	67.02	68.76	27.53	61.34	6.21	2.31	1.05	1.57	1.29			
internal	131,791	6.45	77.59	22.41	3.52	93.37	5.15	76.05	2.37	94.13	10.59	83.65	5.32	0.04	0.31	0.09	0.72			

The “ruled out” calls are opposed of: “lost calls” (transferred to LVL 2 but given up before being picked up, hence making biased computation for second level durations), transferred calls for which no transfer trace is found (indicating a non standard protocol), calls interrupted and called back immediately by the LVL 1 operator (giving a biased conversation time at level one), calls for which aggregated data feature too many ticket files (meaning that the LVL 1 or LVL 2 instruction was made in two steps), and finally calls for which an origin field indicates “Internal”. There are in addition more than 16,000 calls coming from unspecified phone numbers.

Table 7.2: Detailed global counts for inbound calls. Percentages are always taken with respect to the parent total. Columns labeled “N/R” indicate unspecified gravity levels.

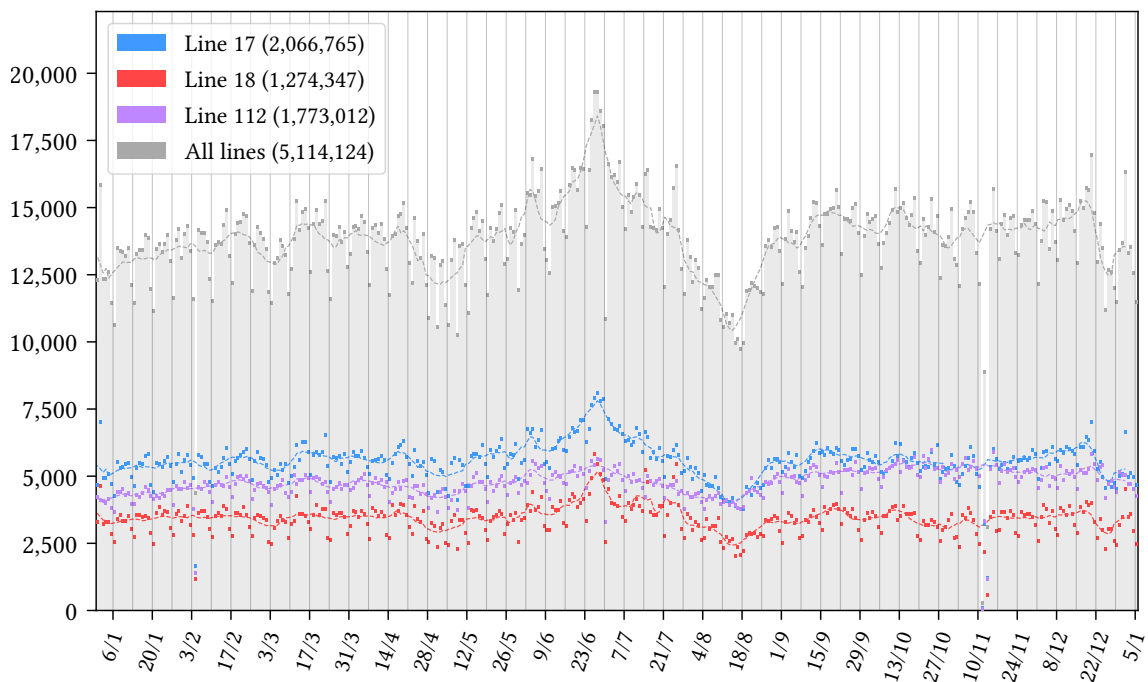


Figure 7.3: Number of 17–18–112 inbound calls per days in 2019

2 Throughputs and proportions of calls

2.1 DEPENDENCES OF THE THROUGHPUTS. — We want to identify the parameters upon which the arrival throughputs depend. It is first natural to distinguish calls based on categorical variables associated with the different treatment lanes: initial emergency line, gravity level determined at LVL 1, orientation at LVL 2. Nevertheless, the throughput of arrivals – such as other variables studied later – are prone to also depend on other factors, such as the time. This parameter is conveniently denoted by *datetime*. It includes a first component date which varies from the 1st of January to the 31st of December and takes into account an “inter-days” effect. On the other hand, the component *time* varies from midnight (12 AM) to 11:59:59 PM and models the “intra-day” aspects. We shall also study the influence of other categorical fields like the department of origin or the telephone area code (from which we can also extract the landline/cell character).

We depict in the Figure 7.3 the daily number of 17–18–112 inbound calls received at the LVL 1 of the PFAU (whether they are picked up or not) during the year 2019. The dashed curves correspond to weekly moving averages. The total number of received calls per day typically varies between 12,000 and 15,000. The peak in 2019 was reached on 26 June, with 19,364 received calls (during a heatwave). The “inter-days” effect is well visible and we can see a weekly periodicity (each vertical line marks a Sunday, often associated with fewer calls). We can also clearly identify the school holiday periods during which much less calls are received (even more visible during August).

This figures shows that the data can suffer some acquisition anomalies, that may be caused by planned technical updates or temporary equipment failures. Five days of 2019 were impacted by such phenomena (2 February, 30 June, and 11, 12, 13 November). These five days shall be systematically ruled out in our future analyses.

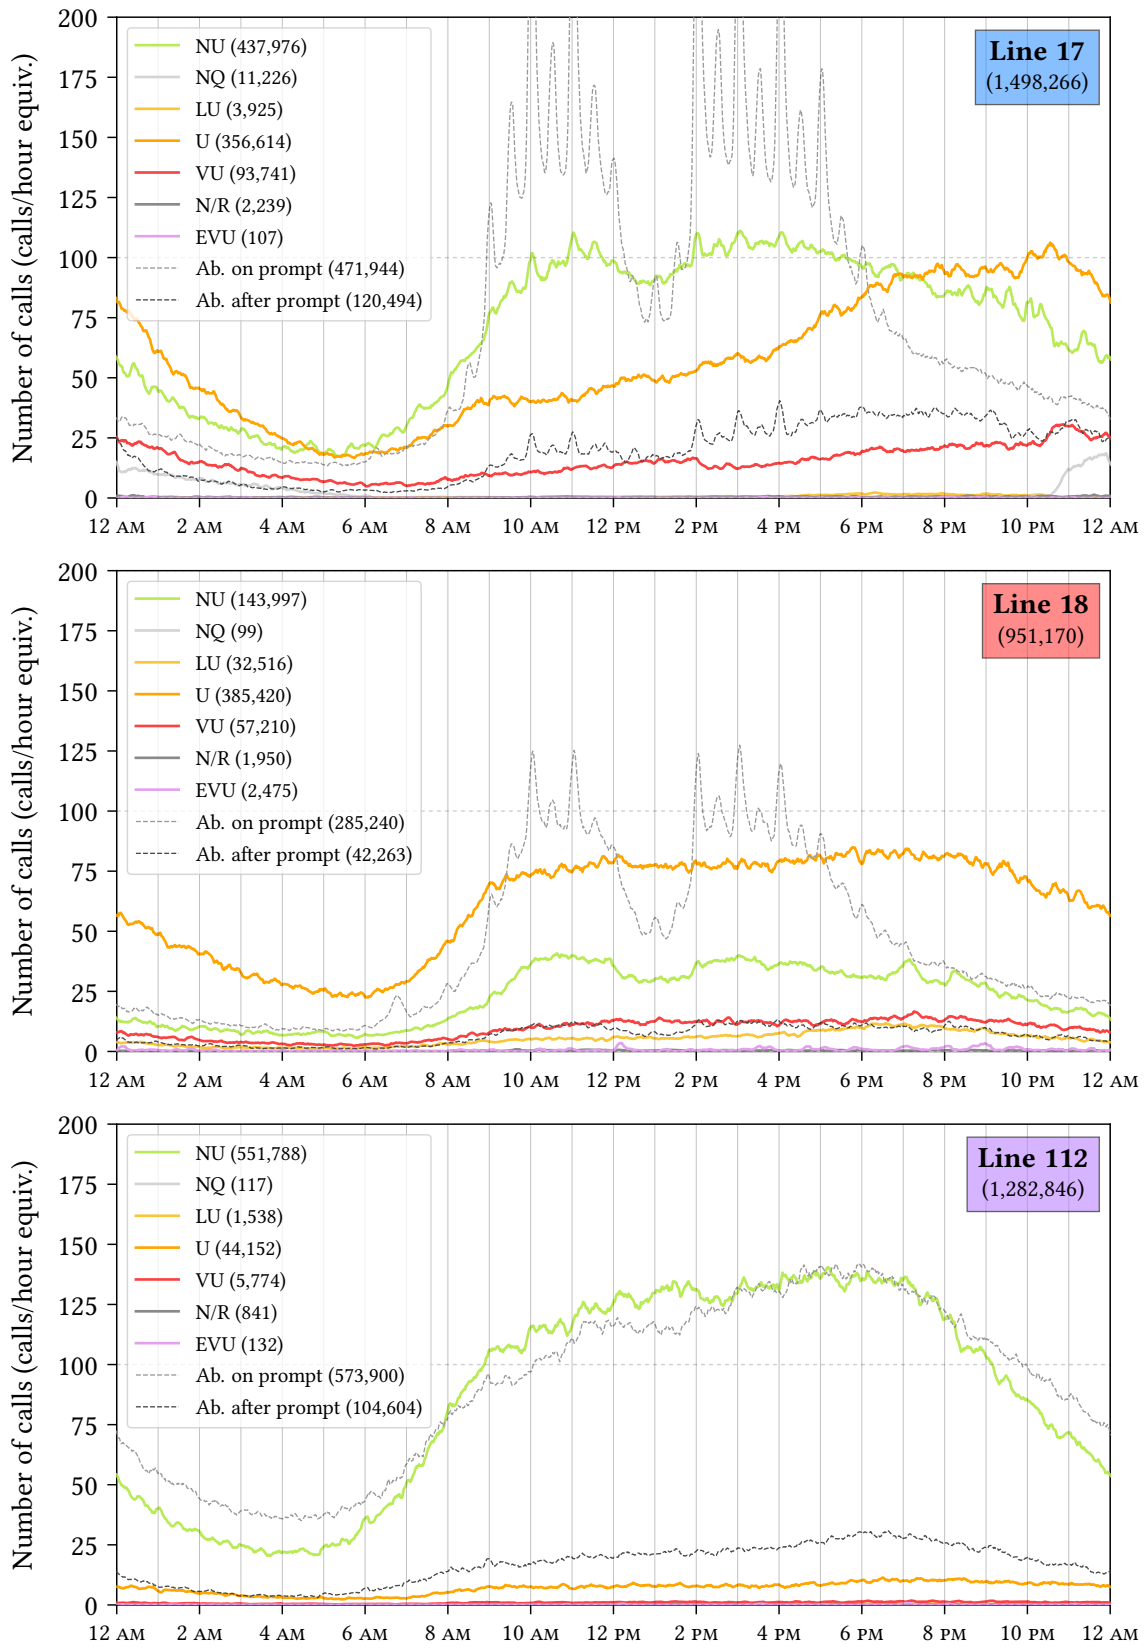


Figure 7.4: Throughputs of inbound calls per gravity, in week days, year 2019

In order to study the “intra-day” dependence of the throughputs, we compute moving averages of the number of arrivals over fixed time-windows (this gives more consistent results than mere hourly histograms). The result is always expressed in “calls per hour equivalents” (abbrev. calls/h), meaning that if 7 calls are received between 12:00 AM and 12:10 AM, we obtain a throughput of 42 calls/h at 12:05 AM (as if the load received over ten minutes was kept the same over an hour). Second, it is desirable to exhibit “typical day” profiles of arrivals, to later generate simulation scenarios. Even ruling out some known “atypical” days (like New Year’s Eve, Bastille Day etc.), we know that such a day does not exist and that micro-peaks arise each and every day due to the natural variability of arrivals. This prompts us to take an average of the daily profiles for all days of the year, in order to even out the peaks and gaps, and also alleviate the effect of atypical days – which do not contribute much in the average. On the contrary, we will give in Section 6 unaveraged profiles of arrivals corresponding to particular days, that can be scenarios of interest in simulation to model atypical events.

The Figure 7.4 gives for each emergency line (17, 18 and 112) the average throughput of received calls during week days (i.e., excluding Saturdays and Sundays) per gravity of calls, and using moving averages over ten consecutive minutes. These curves can be thought of as dynamical counterparts of the Table 7.2.

We represent nine categories of gravity. The most standard are the NU calls (for Non Urgent, filtered out by the first level) and LU/U/VU (for Less Urgent, Urgent and Very Urgent, all transferred to the second level). The category NQ corresponds to calls with No Qualification, only used by police forces. The EVU calls (Urgent Event) only arise when specific protocols to handle unusual events are activated. The class N/R gathers all the other calls whose gravity is unknown. We also distinguish the calls abandoned by the caller, that is to say when the call is ended before pick up at Lvl 1. We make a distinction between the abandonments during the voice prompt (8 to 10 seconds before giving up) or after the prompt.

The curves of Figure 7.4 call for several remarks. We first observe that there are many abandoned calls over the voice prompt during the day on line 17, and that these occur in salient peaks every exact hour (9:00 AM, 10:00 AM...) or exact half hours (9:30 AM, 10:30 AM...). These peaks mostly arise from 9:00 AM to 11:30 AM and from 2:00 PM to 5:00 PM and are less important at noon. The same phenomenon is visible for line 18 (to a lesser extent) but not for line 112, and is totally absent during the week-end (curves not depicted here). A plausible explanation is that these calls are error made from corporate landlines when external numbers or audio meeting services are trying to be reached (since this happens during work hours and every half hour). In these companies, employees often have to press the “0” key before dialing external numbers. If they do not and try to dial a number of the form 01.7x.xx.xx.xx (very common in Paris area), the call will automatically be routed towards the emergency number 17. The same phenomenon also happens for the calls to SAMU who often receive dial errors of people wanting to dial a number of the form 01.5x.xx.xx.xx. Fortunately, the voice prompt enables to get rid of these calls, because the callers will quickly figure out that they reached emergency services. All these abandonments are *invisible* to the PFAU agents. Note however that a small fraction of these make it through the voice prompt, hence the tiny periodic peaks in the curves of abandonments after prompt or non urgent calls.

As far as the non-abandoned calls are concerned, we observe that all the throughput profiles tend to be minimum around 4 AM and start to rise from 5 or 6 AM. On line 17, there is a majority of NU calls (with a little decrease during lunch hours) until 6:00 PM. The throughput of U calls continuously increases during the day and become predominant in the evening and during the night. The trend of VU calls is similar, in a lesser extent. The NQ calls (i.e., with no qualification)

also only arise during the night, after 10:30 PM. They probably correspond to a specific police procedure that we were not told about, and we will often neglect them later on.

On line 18, the average throughput of urgent calls is very stable from 9 AM to 10 PM, roughly equivalent to 80 calls/h. The VU arrivals are still approximately proportional to U arrivals, and NU calls also feature a drop between 12 PM and 2 PM. The categories LU (Less Urgent) and EVU are almost exclusively used by the BSPP, which is why the LU curve is not very visible for lines 17 and 112. The EVU curve is also negligible, since events for which this category is enabled are by essence restricted to limited time windows and are sporadic. For calls arrived by line 112, the predominant throughputs are by far those associated with abandonments after the voice prompt and non urgent calls, with no diminution during lunch hours. This bears out the idea that most of these demands are pocket calls; the callers that realize it contribute to abandonments and those who do not contribute to non urgent calls.

We can compute similar profiles during the week-ends (Saturdays and Sundays only). Doing so, we find out that the throughputs are lower during the day and higher during the night (this is particularly true for police calls, with a total of 340 calls/h on line 17 at midnight during the week-end, versus 220 calls/h on another week day). We have also checked that the proportions of different gravity categories stays identical for all days of the year 2019.

On the other hand, there are non negligible differences from one department to another, as shown in Figure 7.5. On line 17, the proportion of non urgent calls is lower in Paris (39.5 %) than in the inner suburbs (around 50 %). Likewise, there are more VU calls in Paris on line 18. The calls arrived from line 112 tend to be more transferred to Lvl 2 in Paris and Val-de-Marne (8.5 % and 7.4 % of U or VU calls) than in Hauts-de-Seine or Seine-Saint-Denis (6.5 % and 5.3 %).

For the calls that are transferred to Lvl 2 (regardless of the emergency line of arrival), we can compute the breakdown to the DSPAP or BSPP services. The calls from line 17 (resp. 18) are

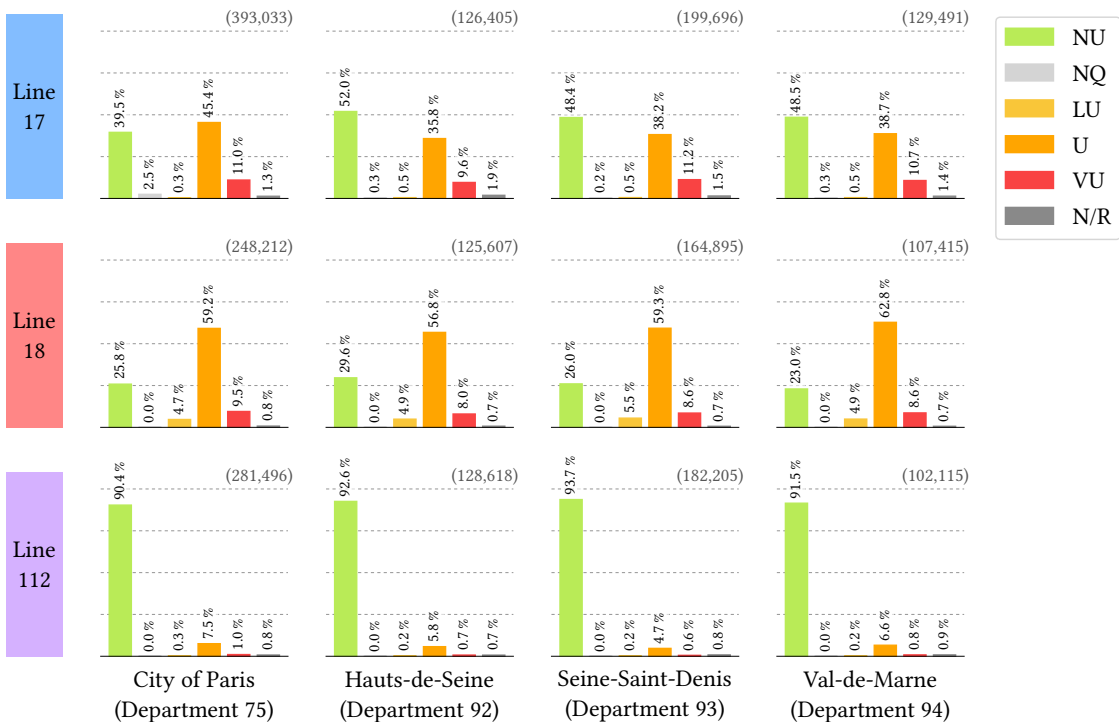


Figure 7.5: Repartition of inbound calls per gravity and department (week days, year 2019, urgent event calls and abandonments ruled out)

a priori about police (resp. rescue) matters and are expected to be handed over to the LVL 2 of DSPAP (resp. BSPP), but the PFAU organization allows to have “cross-flows”, if the caller did not contact the most appropriate number in the first place. In addition, the calls from line 112 have to be distributed to one of the two specialized second levels. For the line 17, 5.77 % of transferred calls are actually rerouted to the LVL 2 of BSPP, while 1.10 % of calls arrived by line 18 and that are transmitted for further instruction are sent to the LVL 2 of DSPAP. These proportions are stable all along the year. The calls from number 112 that are urgent or very urgent are in average sent to the LVL 2 of BSPP in 69.8 % of cases and to the LVL 2 of DSPAP in 30.2 % of cases. These proportions are subject to higher daily variations (mainly due to the low amount of transmitted calls at LVL 2 from the line 112).

We now summarize the explicit form that we retain for the throughputs. The formalism that we have used in the first part of the thesis (see in particular Chapters 4 and 5) uses a global (constant) throughput λ and proportions π_1, π_2, \dots etc. to indicate how this throughput breaks down into the different possible lanes of instruction. However, the previous observations (and in particular the variations with respect to the parameter time) prompt us to distinguish the throughputs of arrivals for each emergency line, each gravity level, each department of origin and each possible destination at LVL 2. We can always use these primitives to retrieve the different proportions of orientation in the system. We shall thus adopt the throughputs $(\lambda^{\ell,g,d,o})$, where: ℓ is the dialed emergency line and varies in $L = \{17, 18, 112\}$, g is the gravity of the call in $G = \{\text{NU}, \text{U}, \text{VU}, \text{LU}, \text{EVU}\}$, d is the department of origin of the call in the set $D = \{75, 92, 93, 94\}$ and if the call is transferred to a second level, o is the LVL 2 orientation in $O = \{\text{DSPAP}, \text{BSPP}\}$. In order to better account for the total load, we also need to consider the throughputs $\lambda_{\text{ab.}}^{\ell,d}$ of abandonments after the voice prompt (per line and department) since those should have been picked up. We can for instance recover the global throughput per line λ^ℓ via the relation

$$\lambda^\ell(\text{time}) = \sum_{d \in D} \left(\lambda_{\text{ab.}}^{\ell,d}(\text{time}) + \lambda^{g=\text{NU},\ell,d}(\text{time}) + \sum_{\substack{g \in G \\ g \neq \text{NU}}} \sum_{o \in O} \lambda^{\ell,g,d,o}(\text{time}) \right).$$

Since this throughput actually also depends on the parameter date, we convene to compute two types of scenarios. The first one will be an average profile of week day (ruling out holidays and anomalies), where the quantities $\lambda^{\ell,g,d,o}$ are the moving averages of Figure 7.4 (taking in addition the LVL 2 orientation into account). The other type of scenarios will be unaveraged and based on specific days of interest. Five such examples are given in Section 6.

2.2 ARRIVALS PROCESS. — The sole knowledge of the (deterministic) instantaneous throughput λ does not suffice to describe the arrivals of calls, since stochastic aspects are also at stake. It is quite often assumed for call centers that call arrivals follow Poisson processes, i.e., the inter-arrival times are exponentially distributed with a parameter equal to the throughput. Recall that we have briefly discussed such processes and given some theoretical elements in the Section 3.3 of Chapter 1. Poisson processes enable to have in average $\lambda \Delta t$ calls arriving in the time window $[t, t + \Delta t)$ while having “excursions”, that is to say phases where the number of received calls can differ from the average value due to the alea.

We wish to check for the PFAU that calls indeed arrive according to Poisson laws. Actually, since the throughput λ is not constant but varies according to the hour of the day (parameter time), we are going to use nonhomogeneous Poisson processes (see again our discussion in Chapter 1). To do so, comparing the whole histogram of inter-arrivals duration to an exponential law is insufficient, since it in fact arises from a mixture of several exponential distributions.

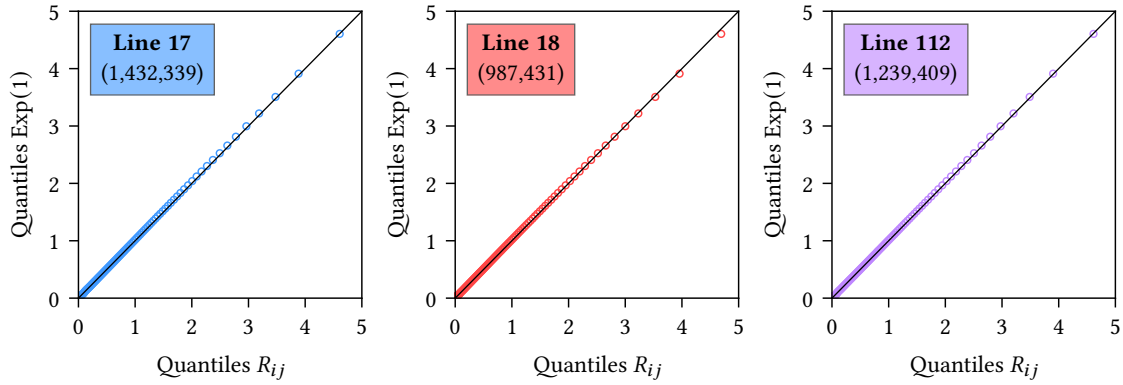


Figure 7.6: Fit of arrivals with nonhomogeneous Poisson laws

In [BGM⁺05b], Brown et al. have proposed the following statistical test to evaluate the fit of a sample to a nonhomogeneous Poisson process. We “split” a long period into short time windows of length L (here set to 10 minutes), and we rank the arrival dates of calls so that $T_{i,j}$ is the arrival time of the j -th call in the i -th interval, which contains a total of N_i calls. Under the hypothesis that the arrivals follow Poisson processes over each short time window (which is much more realistic than over the global period), then the realizations of the statistics R_{ij} defined by

$$R_{ij} := -(N_i - j + 1) \log \left(\frac{L - T_{i,j}}{L - T_{i,j-1}} \right)$$

follow an exponential law with shape parameter 1.

The Figure 7.6 gives the quantile-quantile diagrams (or Q-Q plots) between the variables (R_{ij}) and the standard law $\text{Exp}(1)$ for the three emergency lines 17, 18 and 112. It appears in a very good approximation that both variables are identically distributed (even though the usual

Line	Callbacks from the caller				
	Gravity	Number of callbacks	(% of on-prompt abandonments)	(% of after-prompt abandonments)	Median time between 1 st and 2 nd call (s)
17	NU	43,515	52.62 %	47.38 %	68
	U	44,903	60.34 %	39.66 %	41
	VU	11,759	54.47 %	45.53 %	45
18	NU	19,764	77.01 %	22.99 %	31
	U	23,522	78.45 %	21.55 %	25
	VU	5,230	83.38 %	16.62 %	18
112	NU	41,005	74.69 %	25.31 %	177
	U	3,176	61.21 %	38.79 %	35
	VU	525	55.05 %	44.95 %	29

Table 7.7: Callbacks from the callers

statistical comparison tests fail, disadvantaged by the huge size of the samples). In the rest of our study, we will then stick to the assumption that the calls arrivals follow nonhomogeneous Poisson processes with time-varying parameters given by the curves of Section 2.

2.3 CALLBACKS FROM THE CALLERS. — It is relevant to note that a non-negligible proportion of callers who abandon before pick up do **call back** themselves. Among the abandonments at LVL 1, there are 14.0 % of callbacks within the next hour on line 17, 12.5 % on line 18 and 4.9 % on line 112, accounting for about 6 % of the total picked up volume. Actually, these callbacks mostly happen within the next minute, and feature some calls that are subsequently qualified as U or VU. We observe in addition that the calls to be qualified U or VU tend to call back more quickly than those to be qualified as NU. The figures of Table 7.7 are computed on the basis of 744,508 abandonments on line 17 (among which there are 104,309 callbacks), 402,392 abandonments on line 18 (for 50,473 callbacks) and 912,207 abandonments on line 112 (for 44,969 callbacks).

For the simulation runs, it is important not to count twice these calls in our throughput profiles, otherwise we would overestimate the real demand. Hence, the throughput of after-prompt abandonments (which were “candidates” to pick up) never counts these callbacks.

3 Identification of logged agents

Since our dataset features agent tags that are unique for each operator, we can recover the staffing levels of the PFAU in real time. We can actually only figure out if an agent is on the phone or not, so over a short time window what we observe is a lot of “on/off” signals.

We convene to introduce a notion of **work shift** by considering that an agent is inactive if he/she spends more than fifteen consecutive minutes without picking up or emitting any call at LVL 1, and thirty consecutive minutes at LVL 2. This convention avoids to mark an agent as absent if no call is presented to him/her, and discards the effect of short breaks. It is however possible that an agent is present at the PFAU but not active on the phone (handling of back-office tasks). In this case, he/she will not be detected by our method. We can therefore count at any instant how many agents are active in a shift. At the first level, the data we have do not give any information on the principal job of agents (recall that police agents are asked to pick up calls from line 17 and firefighters are asked to pick up demands from line 18, while calls from line 112 are picked up by both types of operators). We convene that an agent who picks up less than 5 % of his/her calls from line 17 (resp. line 18) is a firefighter (resp. a police officer). Finally, due to the fact that the second level of police calls is managed by a different information system, we unfortunately cannot precisely recover the staffing of police agents logged at LVL 2. The level supervisors are also not counted in the following figures (even though we can detect their activity) because they do not pick up emergency calls in the same manner as standard agents.

Throughout a usual day, the number of logged agents can fluctuate quite quickly. To make up for this aspect and obtain a smoother profile, we choose to compute the average of such daily curves for multiple days. To avoid including too many heterogeneous profiles in the mean, we rule out the week-end days and the days corresponding to school holidays – this leaves us with 36 weeks of the year 2019, or equivalently 180 days. We show the result of this averaging operation in the Figure 7.8 for the LVL 1 and in the Figure 7.9 for the LVL 2 of BSPP. Since the mean is only a partial indicator (and it gives a non-integer result), we also represent the inter-quartiles interval by light-shaded areas (recall that the mean can sometimes lie outside this 50 % presence interval).

The first eye-catching phenomenon on Figure 7.8 is the periodic drop of the number of logged agents that happens at fixed hours. We guess here that this is due to the shift rotation, with

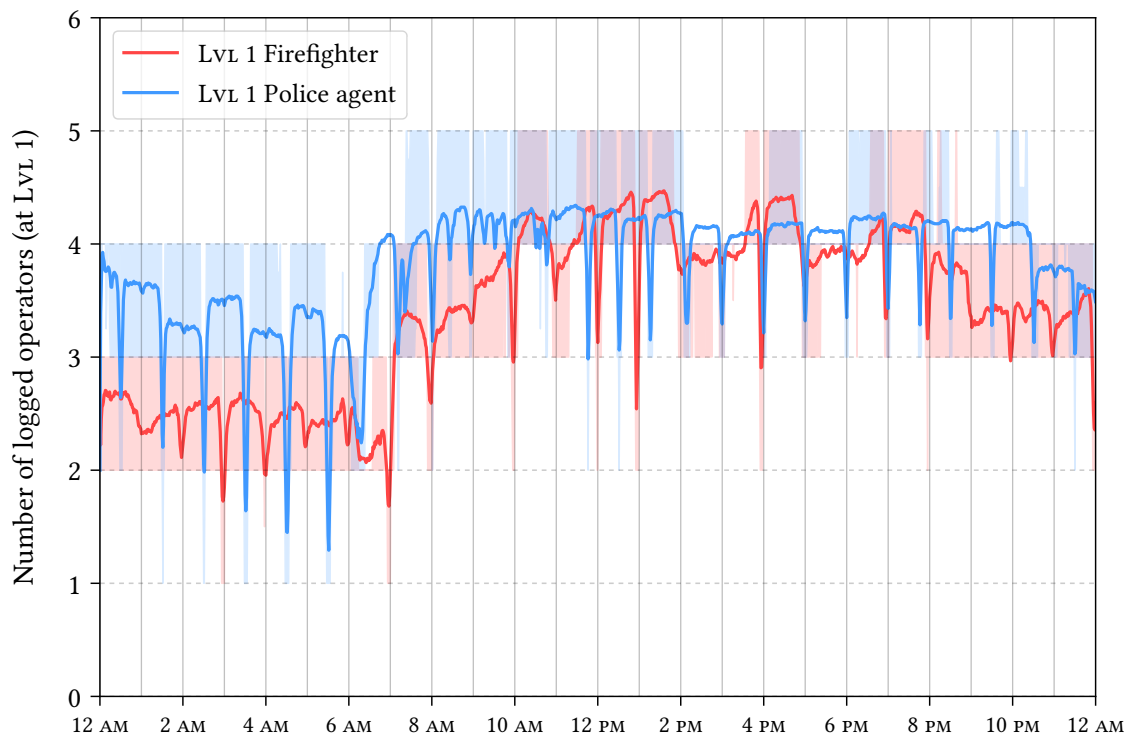


Figure 7.8: Average number of logged operators in work days (180 days) at LVL 1

agents logging out at the end of their shift before the next ones log in. Still, this is a somehow “structural” practice, since it is not skimmed by the averaging operation.

We also naturally observe a difference between night and day hours, with a typical staffing that varies from 2 to 3 firefighters during the night at the LVL 1, versus 3 to 5 during the day. Similarly, there are typically 3 to 4 police agents active in the night at LVL 1, and 4 to 5 at business hours. As far as the second level of the BSPP part of the PFAU is concerned (shown in Figure 7.9), we detect in average 4 second-level operators active in the night, and from 9 to 11 operators between 9 AM and midnight. We have also distinguished another type of agent in Figure 7.9, the so-called “field report firefighters”, who only coordinate the interventions of teams dispatched on-site, and do not pick up inbound calls from the population. We have also computed the same mean staffings during the week-end (Saturdays and Sundays), which led us to observe at the LVL 1 that there are the same number of logged firefighters but that there is typically an extra police officer present at night. Similarly, at the LVL 2 of BSPP during the week-end, there is one extra agent during the night and one less agent during the day, relatively to weekdays.

We also represent in Figure 7.10 the number of logged agents at the LVL 2 of BSPP corresponding to the night after Bastille Day (between the 14th and the 15th of July, customarily associated with more accidents or acts of urban violence). In particular, there is no averaging operation and we recover on this figure integer number of active agents. We observe that at the peak of the night, there are on this specific date up to 18 logged firefighters at LVL 2, that is to say four to five times the “standard” night staffing. This follows the increase of the demand for urgent calls (see the Figure 7.24 to come). The quickness with which the BSPP can double its staffing is here quite noteworthy. Besides the previous figures, we have also computed the histograms of agents’ shifts duration. It turns out that the typical shift of a LVL 1 police officer lasts one hour, with other characteristic duration of 45, 90, 105 or 120 minutes. This regularity is less striking for

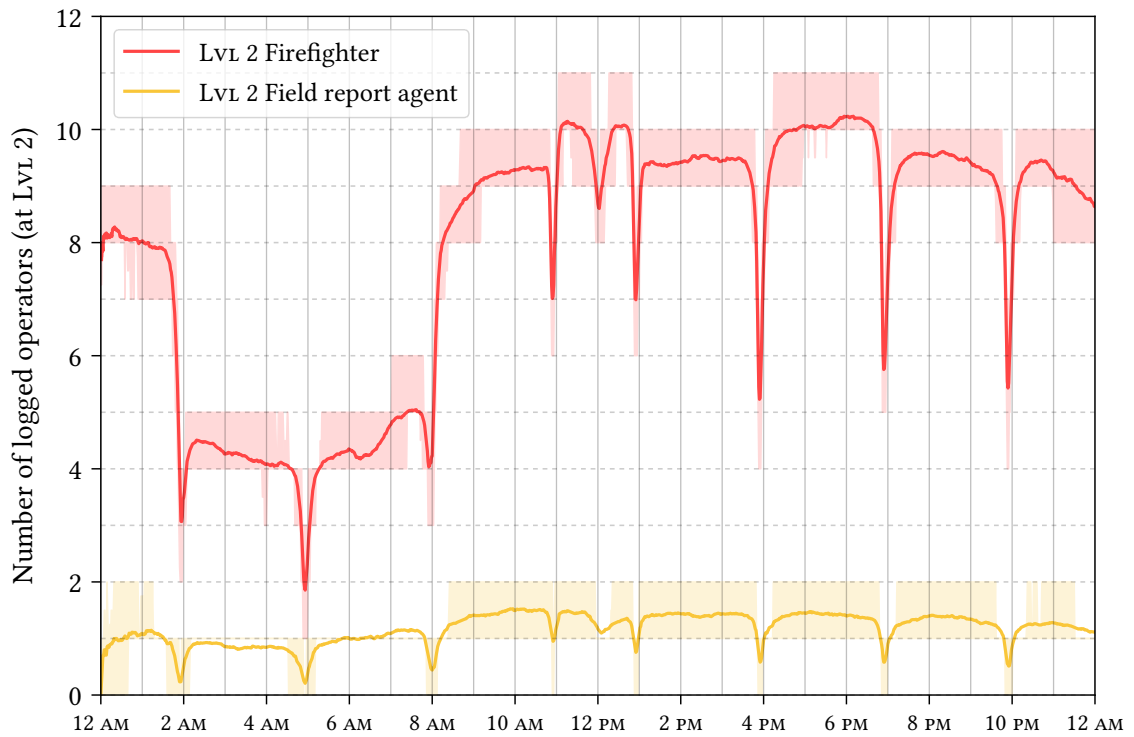


Figure 7.9: Average number of logged operators in work days (180 days) at LVL 2 of BSPP

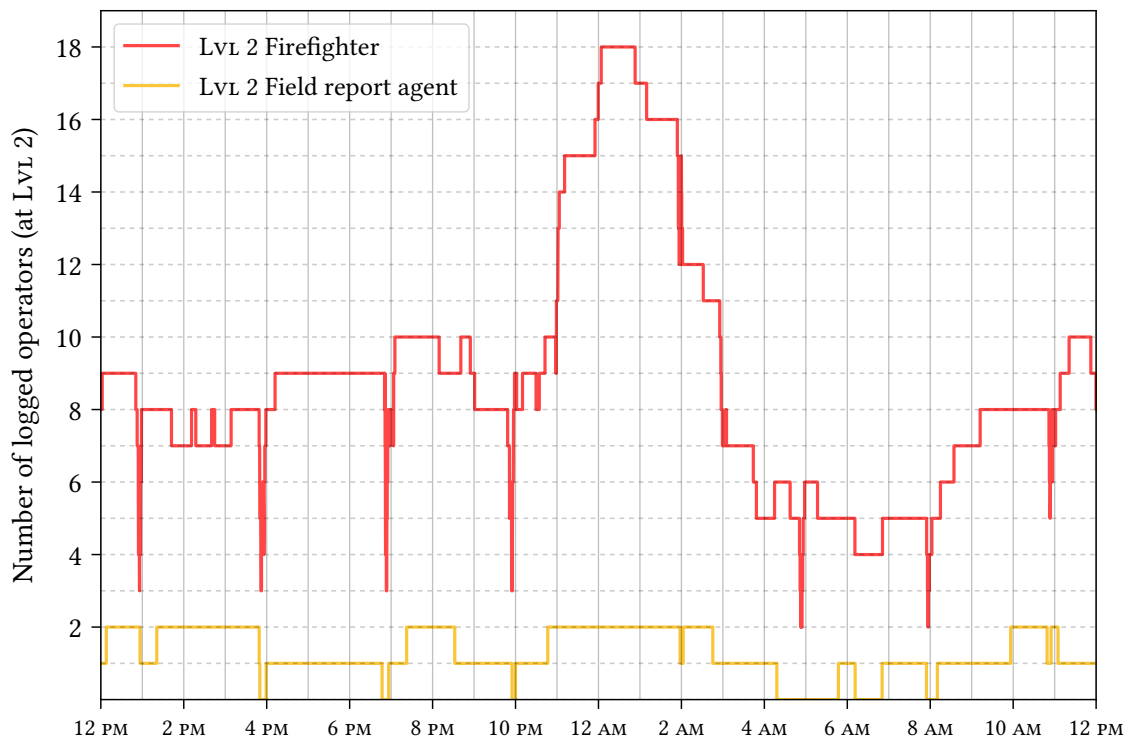


Figure 7.10: Average number of logged operators at BSPP's LVL 2, during the night from the 14th to the 15th of July 2019

firefighters, with peaks at 55 minutes and 2, 3 or 4 hours. Part of these variation can probably be explained by the fact that there are different types of work status for the agents of the PFAU.

We finish this section by precisizing that it is also possible to quantify the fact that LVL 1 operators are sometimes emitting calls. Indeed, 2.8 % of calls involving LVL 1 firefighters are dialed by them, and this goes to 5.4 % for first-level police agents. This can correspond to immediate callbacks of callers if for some reason the conversation is interrupted (the DSPAP agents have the explicit instruction to call back the people who hang up between LVL 1 and LVL 2). It is relevant to weight these volumes by their actual treatment time, which enables to obtain an “outbound workload” to be compared to the “inbound workload”. With this new look, emitted calls represent 3.8 % of the total time spent on the phone by LVL 1 DSPAP agents, and 2.9 % for the LVL 1 BSPP agents.

4 Service times of calls

Measuring the characteristic times of the various steps spent on the phone in the call center is essential for simulation. These conversation duration are dimensioning elements of the system and impact the performance of the platform – contrary to the waiting times before pick up that in turn are consequences of the performance levels. All the instruction times studied here shall be reused in the models of Chapter 8. We analyze in this section four different duration: the time τ_1 of conversation at LVL 1 between a first-level agent and the caller, the time τ_2 of conversation between the LVL 2 operator and the caller, the time τ_s of transfer (or escort time) between the LVL 1 and LVL 2 operators for Very Urgent calls (this time does not include the possible wait of the LVL 1 agent before the LVL 2 operator is available) and finally the time τ_3 of additional instruction at LVL 2 after the second-level operator hung up with the caller.

4.1 CONVERSATION TIME AT LEVEL ONE. — We depict in Figure 7.11 the histograms of the conversation time at LVL 1, split by layers corresponding to the different gravity levels. We have discarded the week-end days and naturally the days with acquisition anomalies. In addition, we have distinguished the new “NU-hangup” category, the calls labeled as non urgent by the LVL 1 operator but hung up by the caller himself/herself (which might indicate that this was an error). The calls actually marked as NU on the Figure 7.11 were hung up by the answering operator. The legend gives the mean time $\bar{\tau}$ and the median time $\tau_{1/2}$ of each layer.

The first histogram in the Figure 7.11 shows a “two-humped” profile for the conversation times at LVL 1 of calls received on the emergency line 17. The layered structure shows that the first peak, around 7 seconds, corresponds to very short calls, namely NU and NU-hangup. There are in fact plenty of the latter for short duration. The second mode, between 25 and 30 seconds of conversation, is explained in majority by U and VU calls. Those two categories have very close instruction times. Note how the distribution’s tail for NU calls is heavy – some conversations go over one and a half minute or sometimes two minutes, owing to some “advice calls”. This very spread profile tends to increase the mean and the median treatment times upward.

For the calls arrived on the emergency line 18, we retrieve very short NU-hangup calls. NU calls are treated in a quick and homogeneous manner, with a median time of 11 seconds. All the calls that are transferred to the second level (LU, U, VU and EVU) are handled with very similar characteristic times. It is interesting to remark that the average conversation time for a NU call is also similar to the one of transferred calls, although with a much lesser median. This shows that on line 18 as well there are quite long “advice calls”. On the contrary for the calls from line 112, we see that the NU layer has both a low mean and a low median time, advocating again for a large majority of “pollution calls” (with no proper conversation), versus very few advice calls.

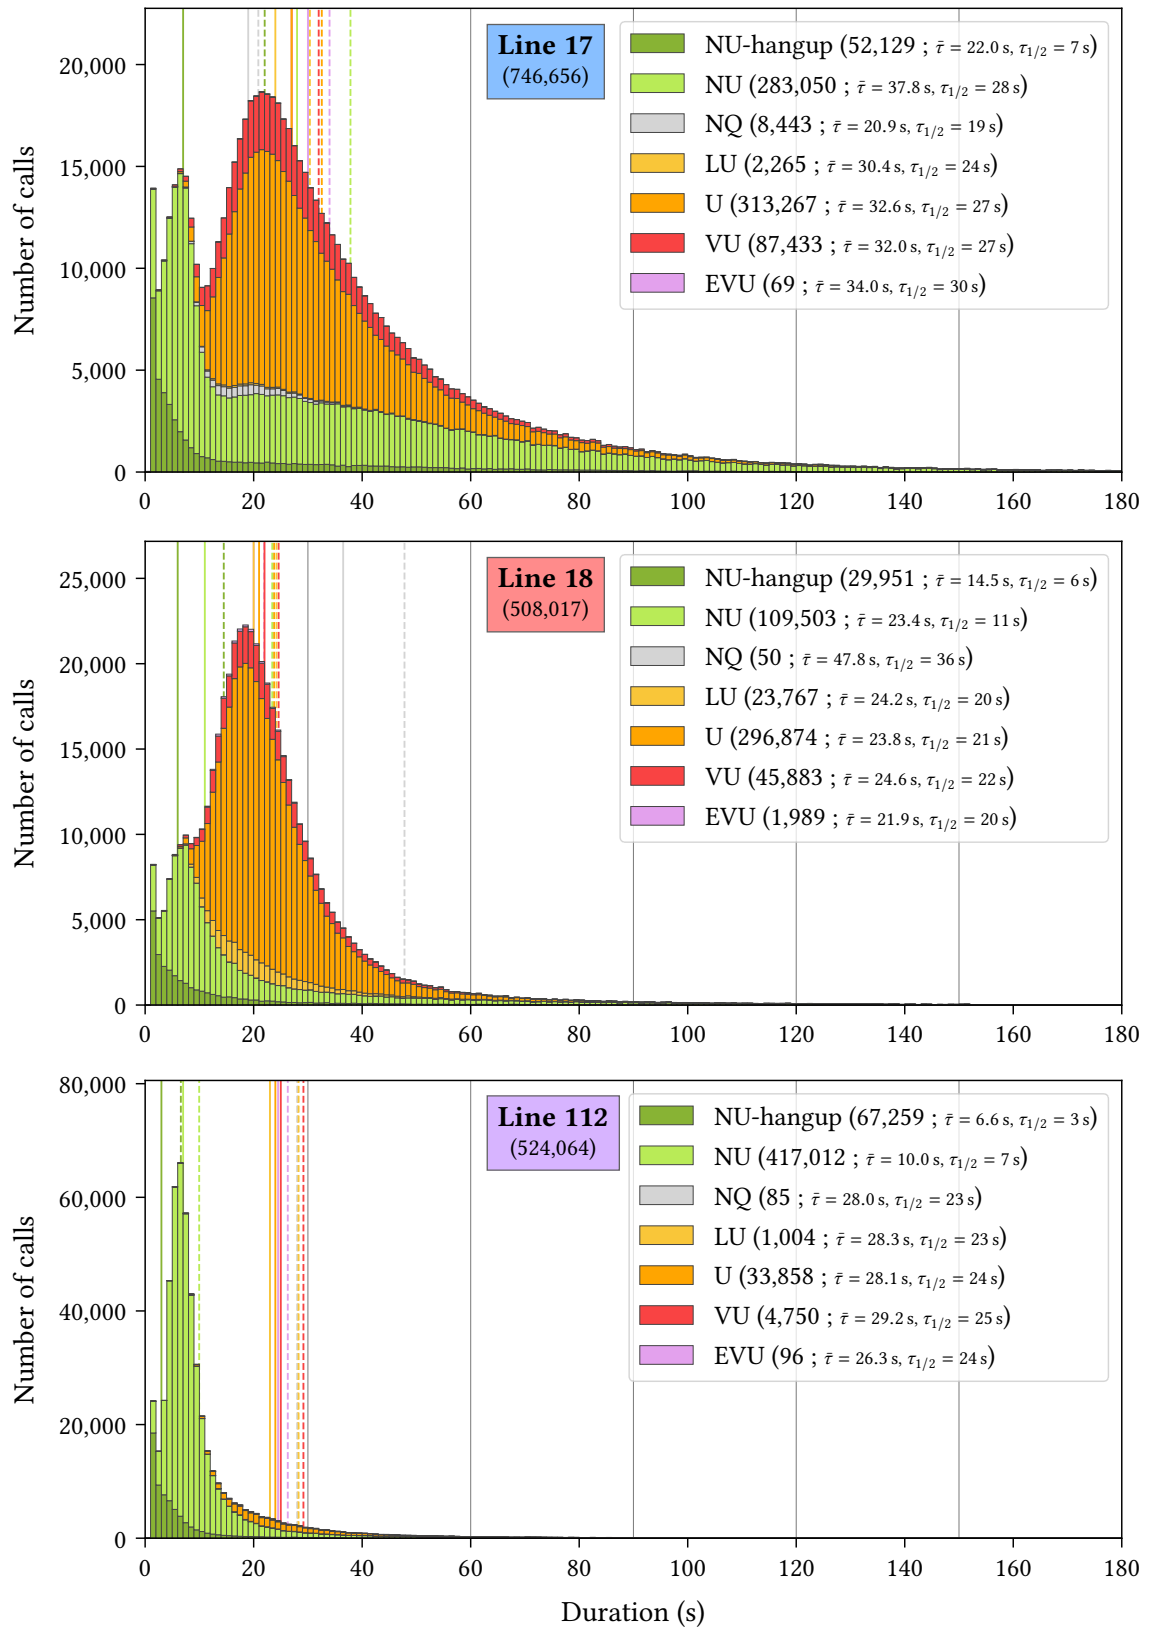


Figure 7.11: Conversation time at Lvl 1, per gravity and line (all week days of 2019)

We have already mentioned in the Section 2 the existence of “cross-flow”, i.e., some calls arrived by line 18 and transferred to the LVL 2 of DSPAP, or line 17’s calls handed over LVL 2 of BSPP. By comparing the mean conversation times at the first level with the previous distributions, it turns out that dialing the “wrong number” in the first place increases the instruction by five to ten seconds in both cases (a LVL 1-firefighter will need more time to handle a call to be transferred to LVL 2 of DSPAP, and vice versa for LVL 1-police agents having to route calls to the LVL 2 of BSPP). This longer time could be interpreted by the fact that the request is more complex (since the initial contacted number was not the most appropriate) or by the fact that first-level agents are less accustomed to picking up calls involving the other job. Nevertheless, it is important to note that being able to reroute a caller who did not reach the best emergency line to begin with in less than ten extra seconds is a great strength of the two-job layout at the PFAU. It is an advantage with respect to separated organizations that handle such transfers in less swift ways. Having universal LVL 1 operators (i.e., picking up calls from line 17, 18 and 112 indifferently) as studied in the Section 4.4 of Chapter 8 could also help lowering further these differences.

To the purpose of finding upon which parameters the conversation times may depend, we now study if the previous distributions vary from one department to another. To do so, we represent for each area and for each gravity level the cumulated distribution function F_{τ_1} of the time τ_1 . The Figure 7.12 illustrates this experiment for calls arrived on line 17. We observe that the four cumulated distribution functions are significantly different for Non Urgent calls. The curve for calls received in department 75 (city of Paris) is the one which grows the faster in the beginning, which tells us that these calls are handled more quickly than in the other areas. Then come departments 92, 93 and 94. For U calls, the four curves are much nearer to each other, and this is all the more true for VU calls. As far as calls from line 18 are concerned, we have the same findings (NU calls are more quickly treated for area 92, then 93, 75 and finally 94, while curves are the same for U or VU calls). These curves demonstrate that even though the profiles of non urgent calls vary from one zone to another, the instruction gets more and more homogeneous as the gravity increases (recall that the four areas are picked up by the *same* PFAU LVL 1 agents).

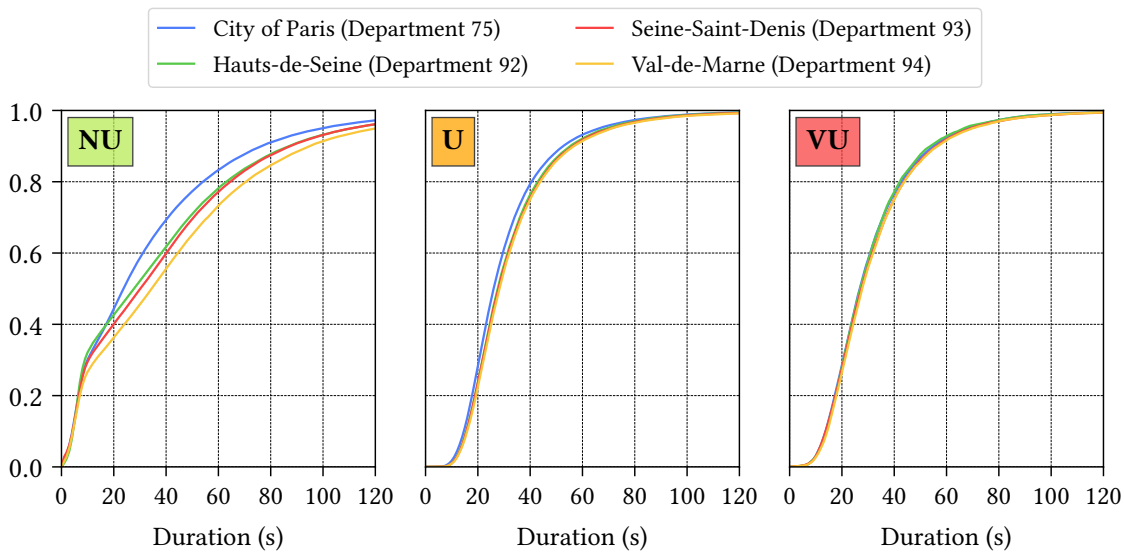


Figure 7.12: Cumulated distribution functions of the LVL 1 conversation time, on line 17 in week days and per departments

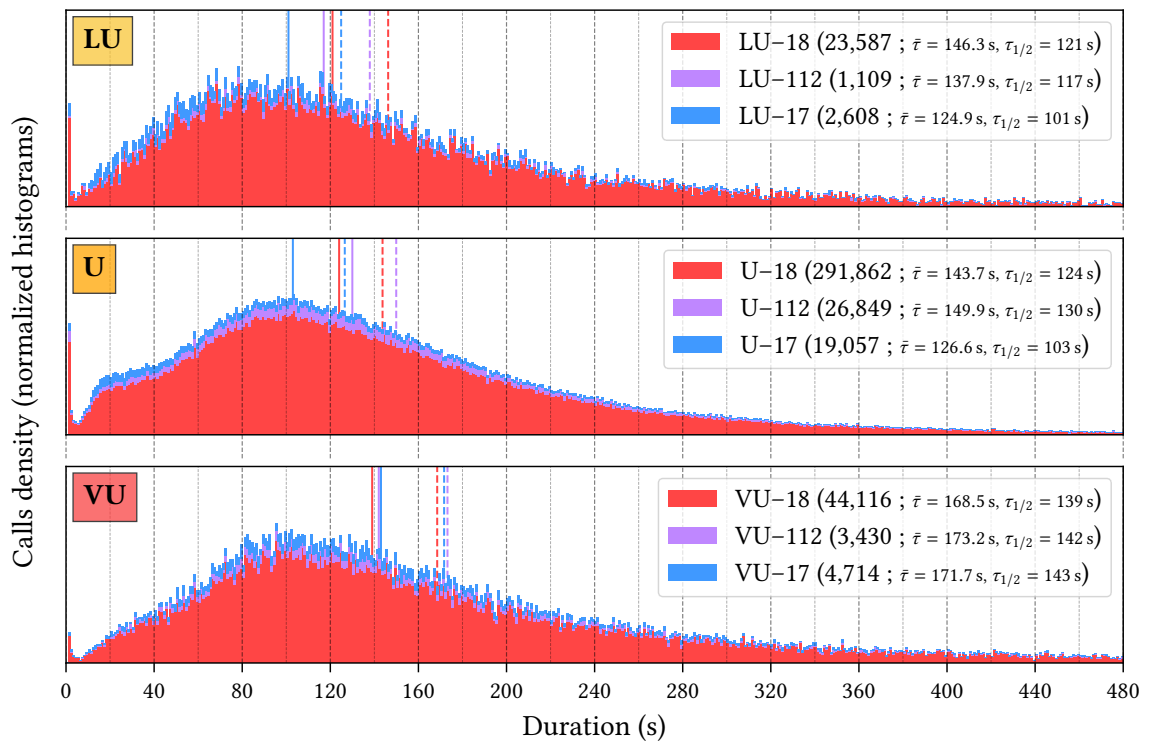


Figure 7.13: Conversation time at Lvl 2 for calls handed over to BSPP's Lvl 2, per gravity and line

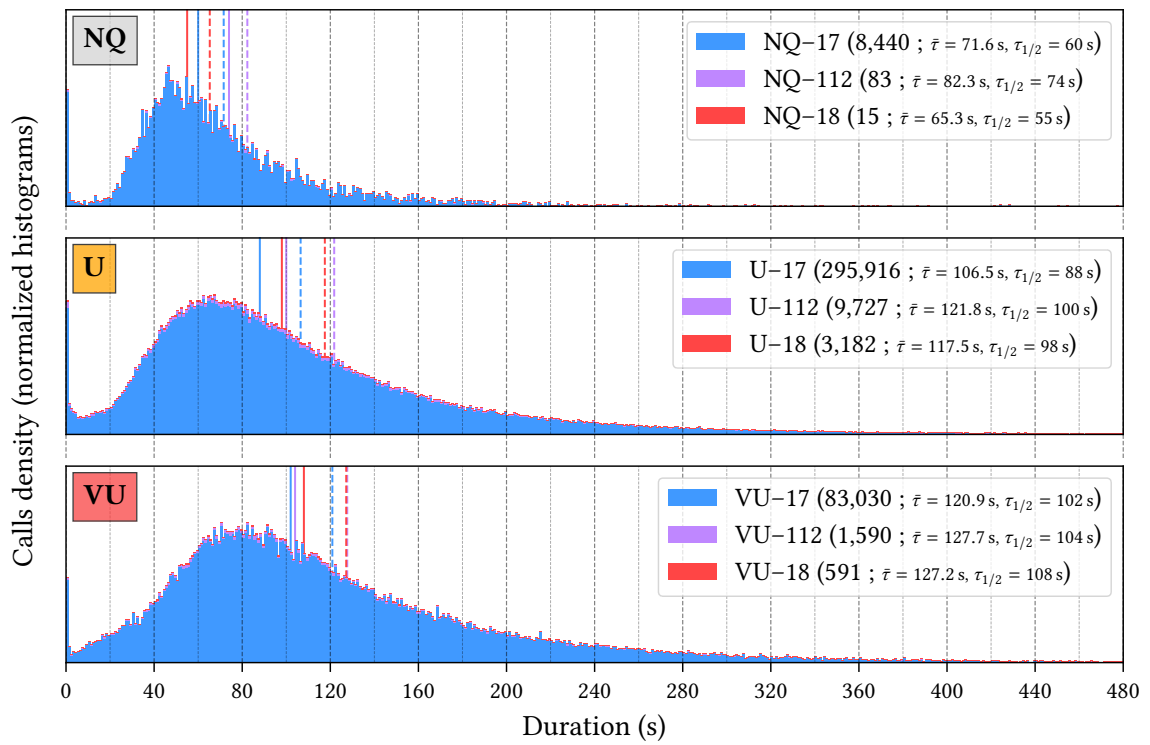


Figure 7.14: Conversation time at Lvl 2 for calls handed over to DSPAP's Lvl 2, per gravity and line

4.2 CONVERSATION TIME AT LEVEL TWO. — The Figures 7.13 and 7.14 give the histograms of conversation times at the LVL 2 of BSPP (for LU/U/VU gravities) and at the LVL 2 of DSPAP (for NQ/U/TU gravities), distinguished by initial contacted emergency number.

At the second level of BSPP, we remark thanks to the Figure 7.13 that the handling time of calls LU and U are quite equivalent. The calls arrived by line 17 are significantly quicker to treat at LVL 2 than those from line 18. On the contrary, calls from line 112 require a bit more time to be instructed (for such 112 calls, we will mostly pay attention to the U distribution and not to the LU, much less significant). The VU calls need in average 25 extra seconds of conversation compared to U calls, and feature a heavier distribution tail. It is however important to mention that in spite of these longer times, for instance to help or give instructions to the caller, this figure does not say anything on the conversation time *before* dispatching emergency response teams. The latter is in fact the only time of interest from the caller point of view. We also note that the treatment of VU calls is pretty uniform regardless of the initial emergency line used.

The conversations at the second level of DSPAP are much shorter than those at the LVL 2 of BSPP. An Urgent call arrived by the line 17 will be handled in average in 106 seconds, versus 116 or 121 seconds if the call comes from the line 112 or 18. The handling of VU calls is here also less prone to variations with respect to the initial dialed line, and is longer than U calls by approximately 15 seconds in average. Similarly to the LVL 1 of the police organization, NQ calls are treated in a rather short way at LVL 2.

We have previously computed the histograms of conversation times at LVL 1 and LVL 2, but these **marginals** do not carry information about the possible interdependence of both times. Indeed, we could imagine that a call that has been long to handle at LVL 1 would also be likely to cause longer instruction at LVL 2 as well. To sort this out, we represent in the Figure 7.15 the density of the (τ_1, τ_2) -**copula**. Recall that when the random variable $\hat{\tau}$ admits the cumulated distribution function F_{τ} , then the random variable $\hat{U} := F_{\tau}(\hat{\tau})$ follows a uniform law over $[0, 1]$. In other words, the tuple $(u_1, u_2) := (F_{\tau_1}(\tau_1), F_{\tau_2}(\tau_2))$ indicates the normalized location of the tuple (τ_1, τ_2) relatively to the marginal laws. Under the hypothesis that the conversation times τ_1 and τ_2 are independent, this tuple should admit a uniform density over $[0, 1]^2$ (see the control map on the right). On the contrary, a positive correlation should entail a bivariate density that is more concentrated along the diagonal; short (resp. high) values of τ_1 being more frequently associated with short (resp. high) values of τ_2 .

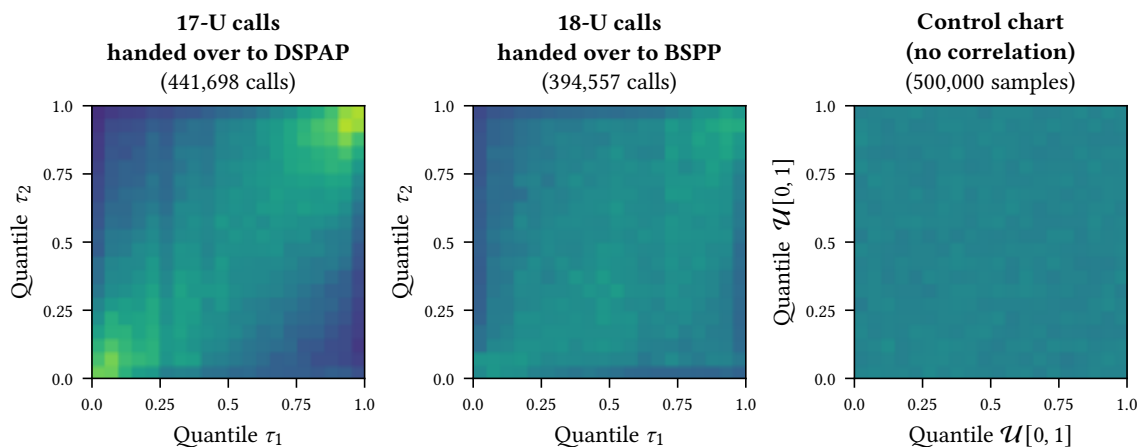


Figure 7.15: Density of (τ_1, τ_2) -copulas

The Figure 7.15 distinctly confirms that LVL 1 and LVL 2 instruction times are correlated for the calls arrived on line 17 and handed over to DSPAP's second level (the diagonal is lit up and the antidiagonal is darker). This phenomenon is less striking for calls arrived by line 18 and transferred to BSPP's LVL 2, but is still observable. This outcome tells us that for simulation purposes, it would not be realistic to generate independent values of τ_1 and τ_2 for a same call. Indeed, this would give an over-representation of tuples (τ_1, τ_2) for which τ_1 is high and τ_2 is short relatively to their respective distributions, which does not occur very often. We shall assume next that the previous copulas are **bivariate Gaussians copulas** with a correlation factor ρ . Fitting such a model to our data yields a value $\rho = 0.25$ for police calls and $\rho = 0.07$ for rescue calls (we shall apply the latter value for 112 calls as well).

4.3 ESCORTED TRANSFER. — The calls qualified as Very Urgent (VU) at the first level of the PFAU are treated differently from others, since they are escorted to the second level. The LVL 1 operator does not put the caller on hold in a waiting room, but stays with him/her on the phone until a LVL 2 agent is available to pick up. At this point, a short synchronization step takes place between the two agents (that may also involve the caller), before the LVL 1 agent is freed and the LVL 2 agent handles the call alone. The Figure 7.16 shows the synchronization times τ_s for VU calls transferred to the second levels of DSPAP and BSPP. It is given only for the second half of the year 2019 because the VU calls were not accompanied before.

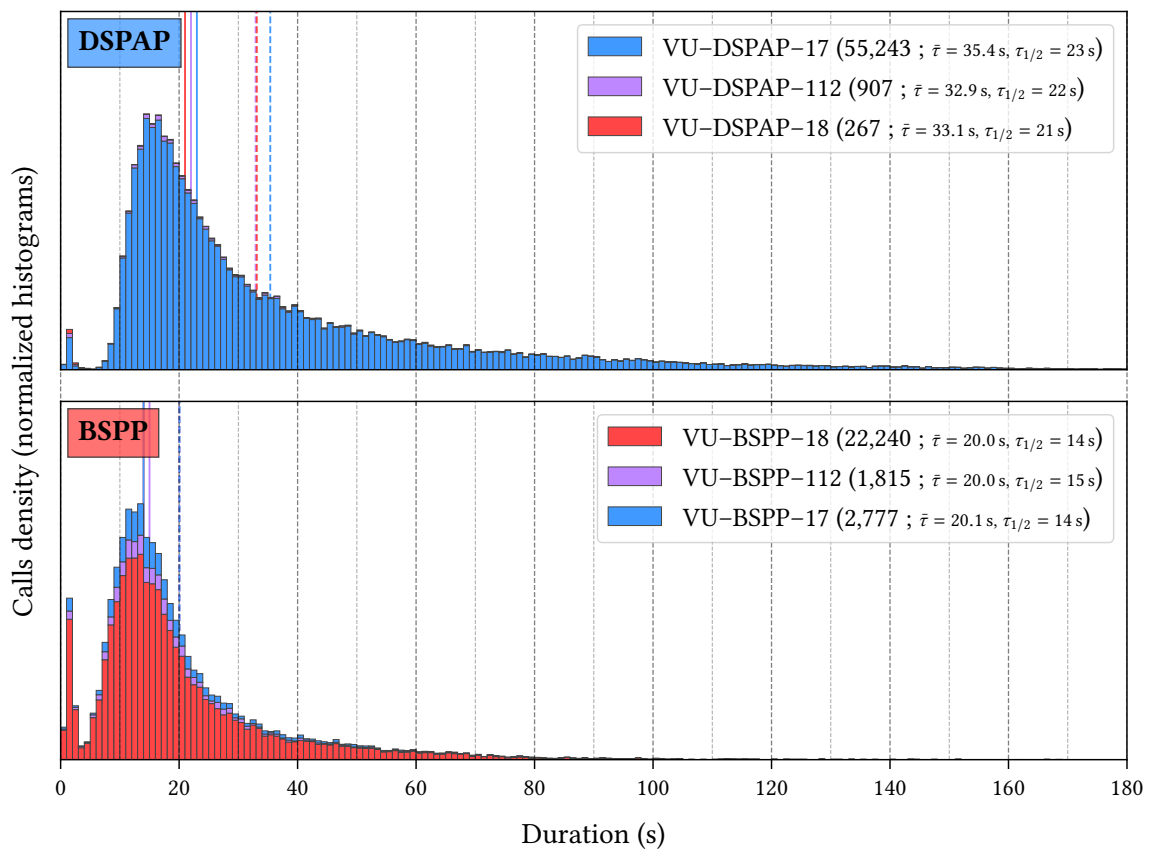


Figure 7.16: Synchronization time τ_s between LVL 1 and LVL 2 for VU calls per line (second half of 2019, week days, non standard calls ruled out)

Because the transfer calls initiated by the LVL 1 operators are internal calls, they appear in our datasets as if they were picked up from the beginning. As a result, the time lengths given in the Figure 7.16 are – unlike the other times shown before – the sum of the waiting time before the LVL 2 agent picks up and the conversation time between the two agents.

We remark that the synchronization time with a LVL 2 firefighter of BSPP is very similar regardless of the emergency number of origin, with a mean at 20 seconds and a median at 14 seconds. The calls shorter than 3 seconds probably correspond to residual transfers with no escorting or to unusual situations where this mechanism is switched off to cope with more demand. For the LVL 2 of DSPAP, the conversation time τ_3 is higher, with a 35-second mean. This phenomenon and the heavy tail of the distribution may be a consequence of more congestion periods on line 17.

4.4 ADDITIONAL INSTRUCTION TIME AT LEVEL TWO. – Once a call is hung up at the second level, it is frequent that the LVL 2 agent has to finish handling the demand of the caller, and thus remains unavailable.

In particular, the operator can make outbound calls to dispatch on-site response teams, reach the SAMU for medical support, ask a piece of advice to the BSPP’s “coordination médicale”, etc. This traceable activity in our data enable us to measure how much extra time of instruction is needed for BSPP’s requests at LVL 2 (while we do not have similar data for LVL 2 of DSPAP). In order to link (possibly multiple) outbound calls to an initial inbound call, we convene to aggregate these into “cases”. A **case** shall be made of an inbound call (17–18–112 caller) transferred at LVL 2 and of all the subsequent outbound calls passed by the LVL 2 operator having picked up the inbound request – until the next inbound call. In fact, second-level agents can also make outbound calls that are unrelated to any inbound demand. Thus, we convene that if an outbound call is passed after more than 90 seconds of telephonic inactivity, it is part of a new case. This choice is arbitrary and modifiable.

When a case gives rise to one or multiple outbound calls, we call **extra instruction time** τ_3 the duration between the moment the initial inbound call is hung up and the moment the last outbound call of the case is hung up. By convention, we let $\tau_3 = 0$ if the inbound call does not yield any outbound call. It is also possible that the outbound call(s) emitted by the LVL 2 agent are hung up before the communication ends with the caller, if the dispatch demand are made in parallel of the initial call. In this case, we also convene that $\tau_3 = 0$ since no additional time is needed on top of τ_2 to handle the case. In particular, the extra instruction time does not give indication on the possible instant of response teams dispatch.

Note that it may also happen that LVL 2 operators dispatch vehicles by digital means or using radio contact, which is an untraceable activity from a telephone point of view (the only one we have in our datasets). This is bore out by the distribution of wrap-up time (after last pick up) that is far from an “expected” exponential profile.

Out of 811,429 analyzed cases for the LVL 2 of BSPP, 37.9 % of them give rise to one or multiple outbound calls to make (cases with two or more such calls account for 4.0 % of the total). We also count 3.1 % of “only outbound” cases (communication not preceded by an inbound call during the 90-second interval). Those could correspond to a LVL 2 agent coming in support of another LVL 2 operator in case of complex cases or calls involving level supervisors.

We depict in the Figure 7.17 the histogram of extra instruction time τ_3 only when $\tau_3 > 0$ (recall that τ_3 can be zero if the case did not generate any outbound call or if those where ended

before hanging up with the caller). In order to compare only homogeneous variables, this figure is restricted only to cases leading to a single outbound call in the case. This rules out many calls (about 11,000) passed to fire stations. Those services are indeed called in 33.1 % of requests for which two outbound calls are made (probably explained by the fact that these teams have a very quick on-site response time and are therefore often called in support of another partner).

Based on a phonebook of partners collected with the BSPP, the histogram of Figure 7.17 is divided into layers corresponding to calls addressed to: physicians from the BSPP’s “coordination médicale”, to fire and rescue stations of the BSPP or stations from neighbor departments, to SAMU centers of the Paris area, to the so-called “SAMU Social” (emergency services for social distress and homeless people) and finally a last layer of either unidentified partners or who do not fit in the previous categories (transport and rail companies, electricity, gas or elevator services...).

We can observe that the internal calls made to the physicians of the BSPP are the shortest, with a mean of 50 seconds. Instructions passed to local fire and rescue stations also represent a cost of one extra minute after hang up for LVL 2 operators (if not made in parallel to the inbound call). Contacting the SAMU is a bit longer, with an average time of 98 seconds, and even longer for the SAMU Social (113 seconds).

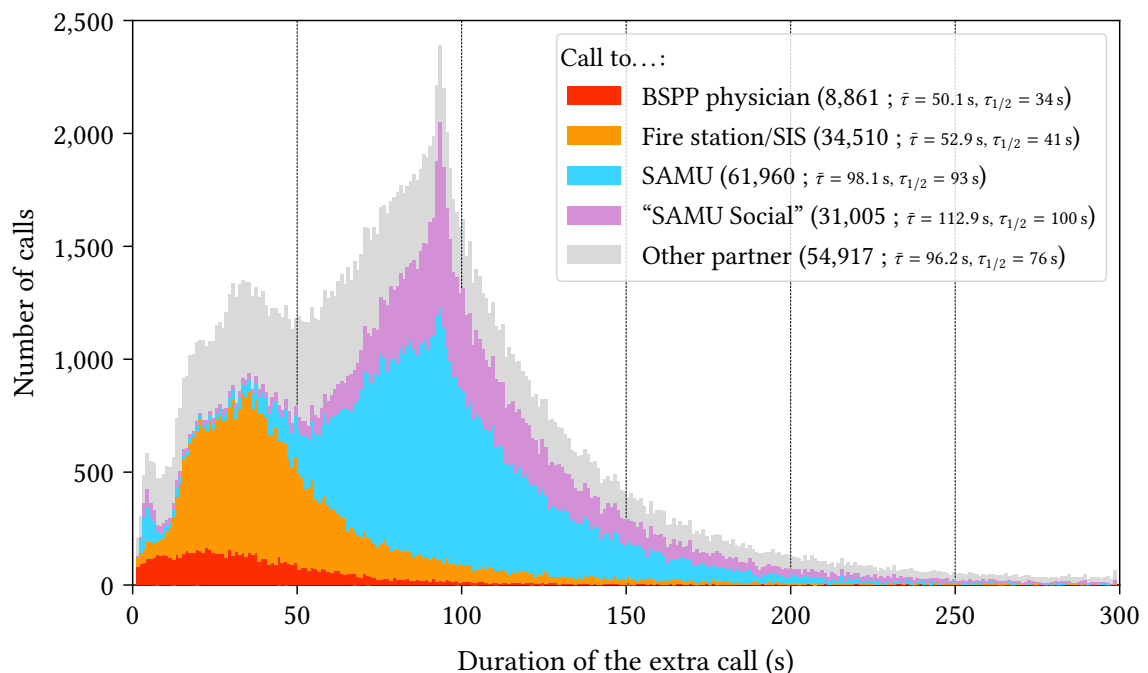


Figure 7.17: Extra instruction time of outbound calls subsequent to BSPP’s LVL 2 cases, per partner (and for cases leading to a single extra outbound call)

We now focus on the correlation between the handling times τ_2 and τ_3 . As studied in Section 4.2, we can indeed wonder whether the extra instruction time after hanging up with the caller depends upon the previous discussion time with him/her.

Preliminary experiments showed us that strong correlations were effectively at stake and we need to take them into account for simulation. However, the reasons of these dependences and the predominant role played by the value $\tau_3 = 0$ do not make the use of bivariate copulas very advantageous as before (in the Section 4.2, having $\tau_1 = 0$ or $\tau_2 = 0$ were rare events and unrepresentative of the distributions, while having τ_3 here is very frequent). We have therefore preferred an approach where the correlation between τ_2 and τ_3 is studied by discretizing the values taken by τ_2 and plotting the corresponding cumulated distribution functions F_{τ_3} . For each level of gravity, we have divided the space of values of τ_2 into **ventiles**, that is to say twenty ordered intervals that realize a partition of values of τ_2 and that all contain the same number of observations.

The result of this analysis is depicted for Urgent calls in the Figure 7.18. The first two ventiles are shown in yellow and the last two ventiles in blue. We remark that the cumulated distribution function of the first ventile starts at the value 0.88, which means that 88 % of the calls with a short conversation time τ_2 (less than 26 seconds) between the caller and the LVL 2 operator do not yield an additional instruction time. This same value goes to 82 % for the second ventile (calls with τ_2 between 26 and 43 seconds). This suggests that short calls at LVL 2 tend to be errors or need quick advice, since the short duration τ_2 does not really enable the LVL 2 operator to make outbound calls. The distribution functions of the subsequent ventiles follow this decreasing trend until the eighteenth ventile. We indeed observe that the curve of the 19th ventile goes back up and this is even more true for the last ventile (63 % of cases with no extra instruction time, versus 54 % in the 18th ventile). This phenomenon can probably be explained by the fact that for such long calls (more than 4 minutes of discussion), the LVL 2 operator remains on the line with the caller to guide or reinsure him/her. However, the possible dispatch of response teams are performed in parallel of this communication, hence not costing any extra instruction time after hanging up.

The same phenomena are visible for Very Urgent calls in the Figure 7.19. Nevertheless, the curves of each ventiles are here much closer to each other. This demonstrates that VU calls are treated with more uniformity, and suggests the existence of more standardized protocols to instruct them.

In our simulations, we shall generate the tuples (τ_1, τ_2) by resorting to bivariate Gaussian copulas as explained in Section 4.2, and then the times τ_3 from cumulated distribution functions depending upon τ_2 . This allows us to avoid unrealistic situations where τ_2 and τ_3 are both long.

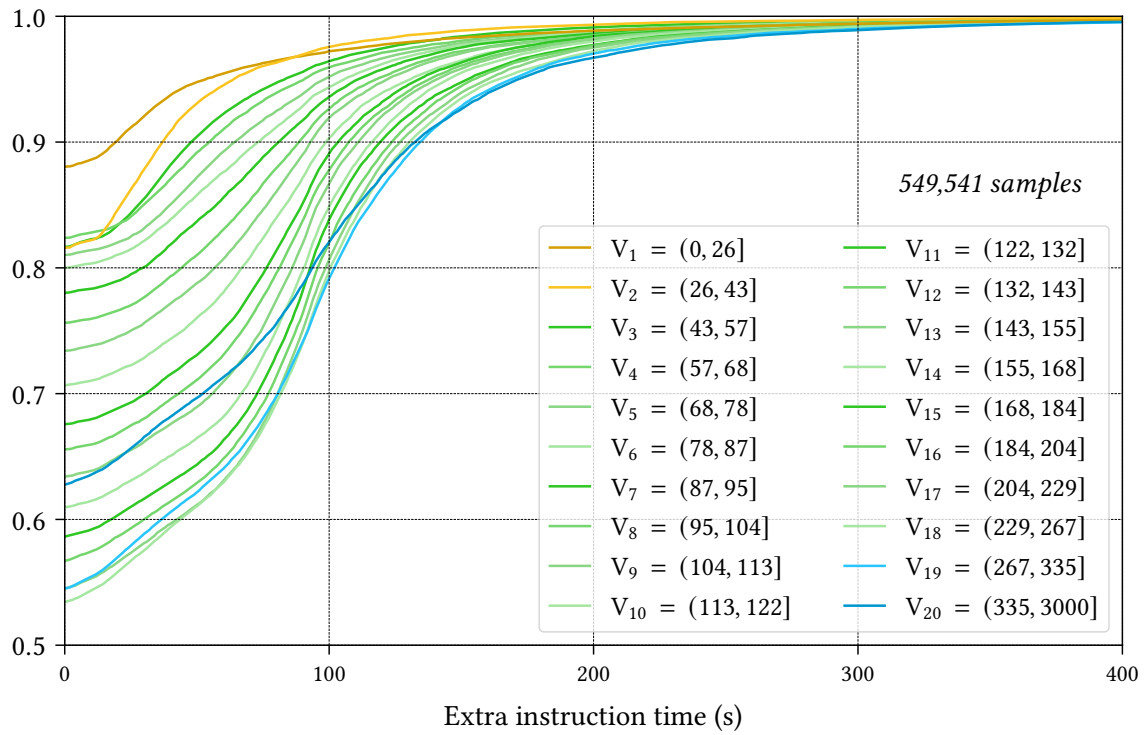


Figure 7.18: Dependence of the cumulated distribution functions of τ_3^U with respect to τ_2

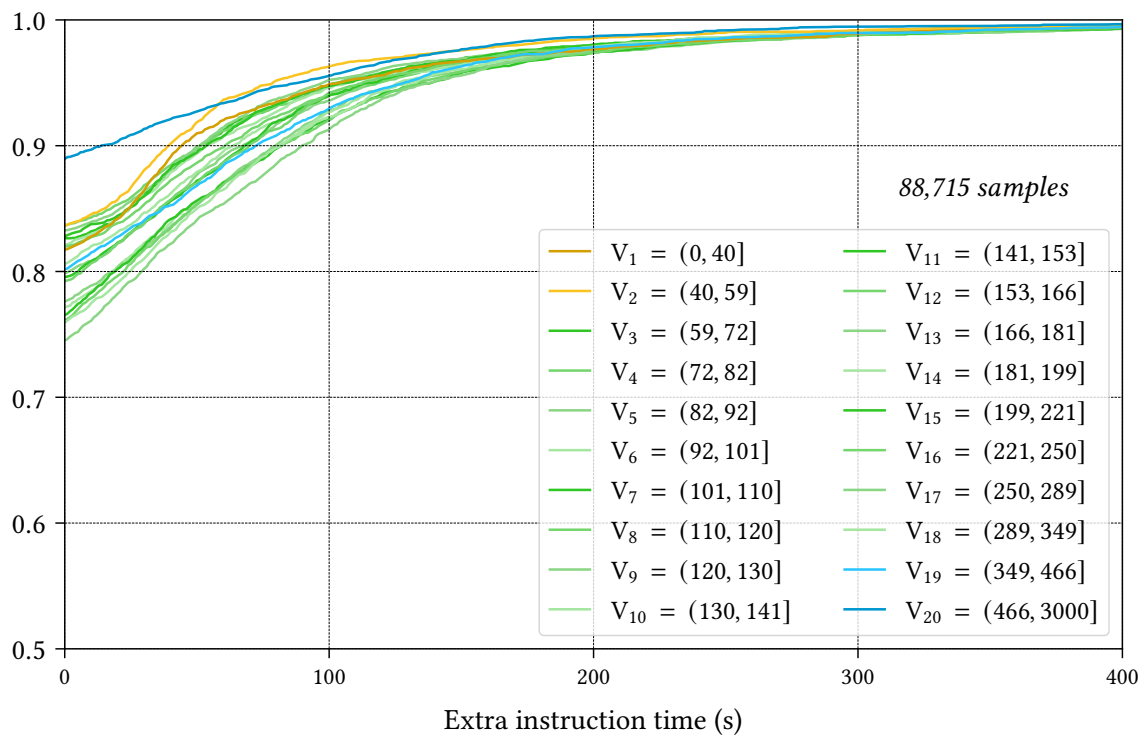


Figure 7.19: Dependence of the cumulated distribution functions of τ_3^{VU} with respect to τ_2

4.5 OTHER FACTORS OF DEPENDENCE. — We have already showed how the four instruction times τ_1 , τ_2 , τ_3 and τ_4 could differ with respect to the emergency number of entry, level of gravity, department of origin, orientation at LVL 2, and even how they were intertwined. We explore in this subsection a couple of other factors prone to influence these time lengths.

First, we may wonder whether the conversation times depend on the hour of the day. We have binned those times into 24 classes corresponding to each hour. Some characteristic times are very stable, such as the treatment time at LVL 1 of U calls arrived by line 18. On the other hand, the U calls arrived by line 17 are handled more quickly during the night (4 to 5 seconds less in average from 10 PM to 5 AM). The same is true for NU calls on both lines. The conversation time at LVL 2 of DSPAP is very steady during the day, but progressively increases between midnight and 6 AM. This could be due to more complex cases occurring during the night, or to the fact that operators tend to take more time on the phone when the demand is lower (or again to an increased level of tiredness at the end of the night).

The previous remark leads us to study whether the number of logged operators impacts the average conversation time, supposing that agents could “speed up” their instruction if they are a little understaffed. It is a difficult effect to measure, as the only knowledge of the staffing that we have is the number of active agents on the phone (see Section 3). The latter is naturally correlated to the inbound throughput of calls, therefore distinguishing between the two causes is tricky. Nevertheless, we tend to observe a positive correlation between the number of active operators on the phone and the conversation time τ_1 at LVL 1 on lines 17 and 18, regardless of the gravity level of the call. This does not overturn the hypothesis that operators indeed adapt their instruction time to the load or the staffing.

Another manner to split the calls is to use the telephone code area present in the data. French lines start with the +33 code, and we know that the device is a cell phone if the designator is followed by the number 6 or 7. Foreign users can be identified by different starting codes (+1 for the US, +44 for the UK, etc). It then turns out that among the NU requests (whether from line 17 or 18), those emitted by mobile devices have a much higher distribution tail, and a mean conversation time twice higher than calls emitted from landlines. This phenomenon does not happen for U and VU calls, which feature similar instruction times between cell phones and house devices. The calls associated with foreign callers typically have increased instruction times at LVL 1 (in average 5 extra seconds) and at LVL 2 (which is sometimes explained by the fact that LVL 2 operators can reach out to a translation platform).

In spite of the last observations, we shall choose in simulation not to take into account the above dependences. The influence of the time of the day is quite minor, and splitting too much the sample size would alter the smoothness effect obtained for distributions of large samples. In addition, it is not desirable that our simulations do integrate the correlation between instruction times and staffing levels, nor the fact that handling times can vary from one operator to another. Indeed, we prefer that the number of logged operators do not cause any implicit or feedback influence on the experiments, in order to better control its direct impact on the performance levels. Finally, distinguishing conversation times based on the origin of the call is only significant for NU calls, and this effect is anyway already accounted for in the associated cumulated distribution function. It does not seem relevant to us to extract finer corresponding throughputs.

5 Patience of callers

In order to derive realistic results by simulation, it is important to take into account the patience of the callers. By definition, it is the time that they are willing to wait for before pick up. If the waiting time exceeds the caller's patience, the latter will hang up himself/herself, and we talk about an abandoned call.

Estimating the patience of the population of callers is a difficult problem because most often, the calls are picked up before callers get impatient (which is desirable of course). Denoting by P the patience of a caller (in seconds) and by D the pick up time of this call, what we can only measure is actually the quantity

$$T := \min(P, D).$$

In most cases, T equals the pick up time D – in other words we only measure the real patience of people who have abandoned. We talk about **data censorship** (the variable P being masked by the variable D). In order to simulate the real patience of callers, we wish to recover the distribution function F_P , and it would be very crudely wrong to use in place the distribution function F_T that we can directly measure – since the majority of the data are censored by pick up. Paradoxically, the days with congestion and for which callers experience longer waiting times are very interesting to study the patience, because they provide data points closer to the real patience of people.

5.1 KAPLAN–MEIER ESTIMATOR. – A tool that is commonly used in this context is the **Kaplan–Meier estimator**, first introduced in survival analysis. In medicine and healthcare studies for instance, we often want to analyze the effectiveness of an experimental treatment by comparing several groups (with or without the treatment). Such problems regularly involve censorship of the data of interest (typically the date of death of a patient), due to the patient leaving the cohort before the end of the study. The Kaplan–Meier estimator, already applied to the analysis of call centers in [BGM⁺05b], gives an estimate of the survival function S , defined as the complementary $1 - F_P$ of the patience function F_P . In our framework, $S(t)$ will denote the proportion of callers that are willing to wait at least t seconds before abandonment, and is defined by the expression:

$$\widehat{S}(t) = \prod_{t_i \leq t} \left(1 - \frac{a_i}{n_i}\right), \quad (7.1)$$

where we denote by

- $t_1 < t_2 < \dots < t_i < \dots < t_m$ the ordered dates (values of waiting times) for which abandonments of callers arise;
- a_i the number of callers who hang up at instant t_i ;
- n_i the number of callers who have not hung up before the instant t_i (and are still waiting for pick up).

In large populations, we shall typically observe multiple abandonments every second for short times, that is to say $t_1 = 1$ s, $t_2 = 2$ s, ..., and further a few abandonments (most of time unique) occurring after long duration, for instance $t_{m-1} = 462$ s, $t_m = 603$ s. The time-length t_m is the maximal patience measured in the population.

In the absence of censorship (very theoretical hypothesis under which no call is picked up...), we have the relation $n_i = n - (a_1 + a_2 + \dots + a_{i-1})$, where n is the total number of callers. In this

case, the Kaplan–Meier estimator reduces to

$$\widehat{S}(t) = 1 - \frac{1}{n} \sum_{t_i \leq t} a_i,$$

which exactly corresponds to the empirical distribution function $1 - F_P$. Under censorship though, the definition (7.1) weights abandonments by the amount of individuals still prone to abandon.

It is important to note that the Kaplan–Meier estimator is only defined until observations are available. In theory, the patience function S tends towards 0 for very large times (no caller at all is for instance willing to wait for 24 hours on the phone). In practice, S features decreasing steps only for times at which abandonments are recorded, and the last drop reaches 0 only if $a_m > 0$ – which means that the highest time-value with observations contains abandonments. Otherwise, the patience function S remains steady and nonzero until t_m and is not defined further. We refer to the analysis of the Figure 7.20 for a discussion on this aspect.

There is a rich literature on the good behavior of the Kaplan–Meier estimator. The central question is to know whether the estimator **converges**, that is, if the computed proportion $\widehat{S}(t)$ is indeed a good estimation of $S(t)$ and what is the quality of the approximation as a function of the sample size. One of the most general results claims that when the patience levels $(P_i)_{1 \leq i \leq n}$ of callers and the pick up times $(D_i)_{1 \leq i \leq n}$ are i.i.d. variables (mutually independent) and S is continuous, then \widehat{S} converges in law towards S “as much as censorship allows it”. Indeed, denoting by τ_P the time after which the most patient of all callers gets impatient, τ_D the maximum possible pick up time, and as a consequence $\tau_T := \min(\tau_P, \tau_D)$ the maximal possible observation time, then we have according to [SW93] that

$$\sup_{t \in [0, \tau_T]} \left| \widehat{S}(t) - S(t) \right| \xrightarrow[n \rightarrow \infty]{\mathbb{P}} 0.$$

This means that with many observations, the probability that $\widehat{S}(t)$ differs from $S(t)$ tends to 0.

Using the “limit times” τ_P and τ_D is here theoretical and it is not clear if such values should be finite or infinite. Rather, it is reasonable to picture that there are finite values $\tau_P^{(\varepsilon)}$ and $\tau_D^{(\varepsilon)}$ such that $F_P(\tau_P^{(\varepsilon)}) > 1 - \varepsilon$ and $F_D(\tau_D^{(\varepsilon)}) > 1 - \varepsilon$, after which almost all calls (up to the error ε) are respectively abandoned and picked up. We expect in practice that the distribution of pick up times reaches its asymptotic value before the patience levels do, in other words $\tau_D^{(\varepsilon)} \ll \tau_P^{(\varepsilon)}$, and this prevents us from having a complete convergence of \widehat{S} . However, the fact that we can record some high pick up times allows us to recover with a great precision most of the patience functions as shown in Figure 7.20. In addition, for numerical simulations, small errors for high levels of patience will not be active nor make a difference, unless we purposefully want to study what happens in understaffed configurations. In those situations only, it shall turn useful to suggest an empirical “extension model” of the patience functions beyond τ_T .

The assumption that the random variables at stake are independent is a bit more problematic, since rigorously inexact. It is acceptable to suppose that the patience levels of all the callers are independent, but pick up times are not. Indeed, short consecutive pick up times suggest that many operators are available, which favors higher pick up times for the next callers if the agents are still busy on the phone. We consider anyway that the Kaplan–Meier estimator brings a very interesting light on the patience problem and that the very high number of observations that we have helps mitigating the possible correlations between neighbor variables.

It is possible to account for the incertitude on \widehat{S} by computing its variance. An estimator of the latter is given by the Greenwood formula, defined (with same notation as before) by:

$$\widehat{\text{Var}}(\widehat{S}(t)) := \widehat{S}(t)^2 \sum_{t_i \leq t} \frac{a_i}{n_i(n_i - a_i)}.$$

This estimator shall be all the higher as the data involved in the computation of \widehat{S} are spread, but also if the amount of data points is too small and gives less confidence in the estimated value. Several authors have shown that under the previous assumptions and if $\tau_T < \infty$, then the estimator \widehat{S} is asymptotically normal (see for instance [Gil80]), which legitimates the use of a Gaussian confidence interval. Actually, when the variance is high, it can lead to confidence bounds out of the $[0, 1]$ interval. Therefore, we will preferably use the Rothman–Wilson interval (see [ABP82]), whose bounds under a risk α are given by:

$$\text{CI}(\alpha) = \frac{K}{K + z_{1-\alpha/2}^2} \left[\widehat{S}(t) + \frac{z_{1-\alpha/2}^2}{2K} \pm z_{1-\alpha/2} \sqrt{\widehat{\text{Var}}(\widehat{S}(t)) + \frac{z_{1-\alpha/2}^2}{4K^2}} \right]$$

where $K := \widehat{S}(t)(1 - \widehat{S}(t))/\widehat{\text{Var}}(\widehat{S}(t))$ and $z_{1-\alpha/2}$ is the $(1 - \alpha/2)$ -quantile of the bilateral normal law. When the variance is low, we recover the usual result where the confidence interval is centered around the estimator and whose width is proportional to the standard deviation of the latter.

5.2 COMPARISON OF PATIENCE FUNCTIONS. — We have computed the Kaplan–Meier estimator for all the calls at LVL 1, and for all calls transferred to LVL 2. Recall that the great majority of calls received on line 17 (resp. line 18) which are transferred are handed over to LVL 2 of DSPAP (resp. BSPP). Consequently, we have compared the patience of calls of line 17 at LVL 1 with those of DSPAP’s LVL 2 calls. Similarly, patience at LVL 1 of calls from line 18 is compared to patience at LVL 2 of BSPP calls. The patience functions at LVL 1 are recovered “thanks to” the calls that were not picked up at all, while their LVL 2 counterparts are built based on the calls lost between the two levels (caller getting impatient waiting for the LVL 2 agent). Because we know the gravity levels of the transferred calls, we are able to distinguish the patience functions at LVL 2 depending on this parameter. On the other hand, the gravity of calls that were lost before LVL 1 pick up is by essence unknown. As a result, the corresponding patience curve *a priori* contains abandonments of all gravity levels. It is sound to think that most of callers who abandon (in particular for short waiting times) would have been qualified as Non Urgent if picked up at LVL 1. The analysis of callbacks from Section 2.3 shows however that even callers to-be-qualified as U or VU do hang up on their own. We have not extrapolated the (known) gravities of callers having abandoned at LVL 1 and then called back, because we considered that these situations were not representative of most abandonments. We may then only assume that an important part (impossible to quantify) of abandoned calls would have been NU. Note that it can happen that a group of people witnessing the same accident simultaneously call emergency services for a U or VU matter but all of them hang up as soon as one call within the group is handled.

The Figures 7.20 and 7.21 feature the patience functions, and each shaded area stands for 95 %-confidence intervals. We have excluded from our analysis the calls hung up during the voice prompt, reason why the LVL 1 curves do not decrease for a few early seconds.

The green curves in the Figures 7.20 and 7.21 stand for the patience functions at LVL 1. We can thus read that a little bit less than 80 % of callers dialing the number 18 are willing to wait for

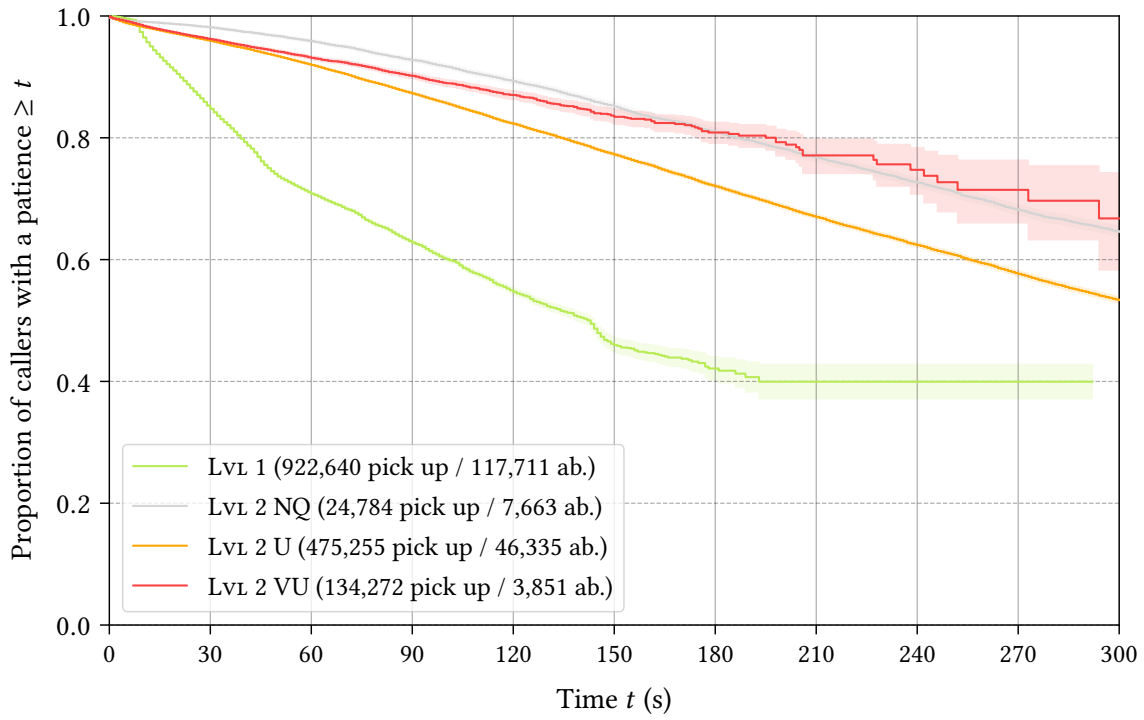


Figure 7.20: Comparison of patience functions (at LVL 1 of line 17 and LVL 2 of DSPAP)

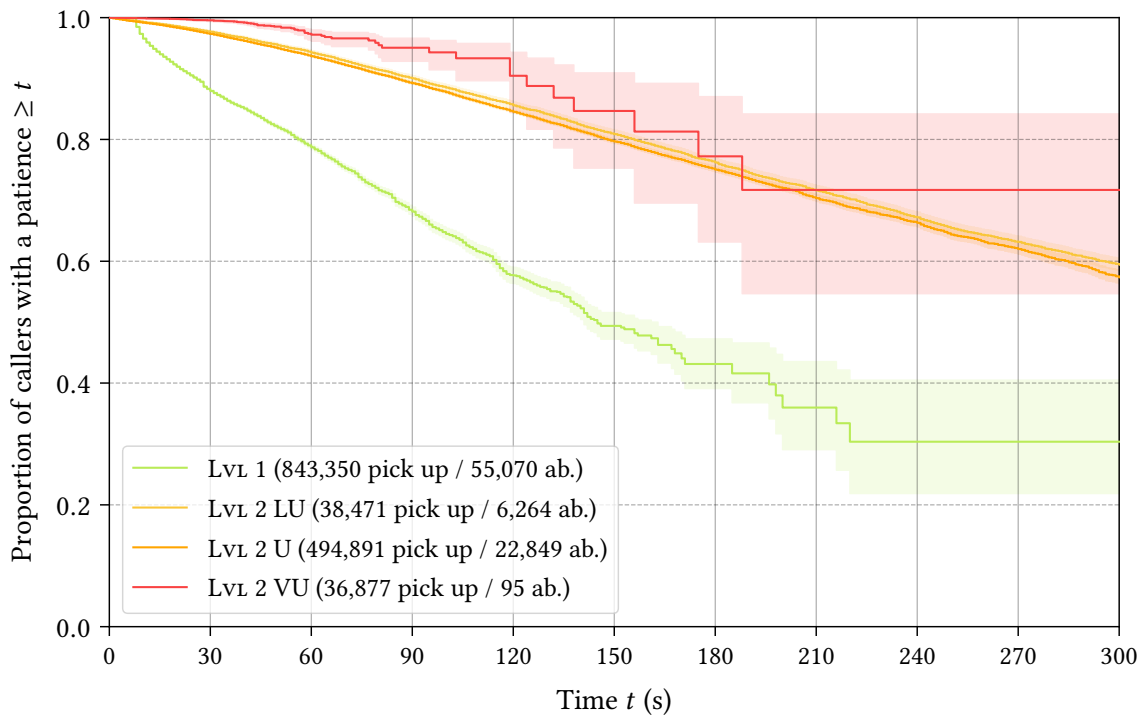


Figure 7.21: Comparison of patience functions (at LVL 1 of line 18 and LVL 2 of BSPP)

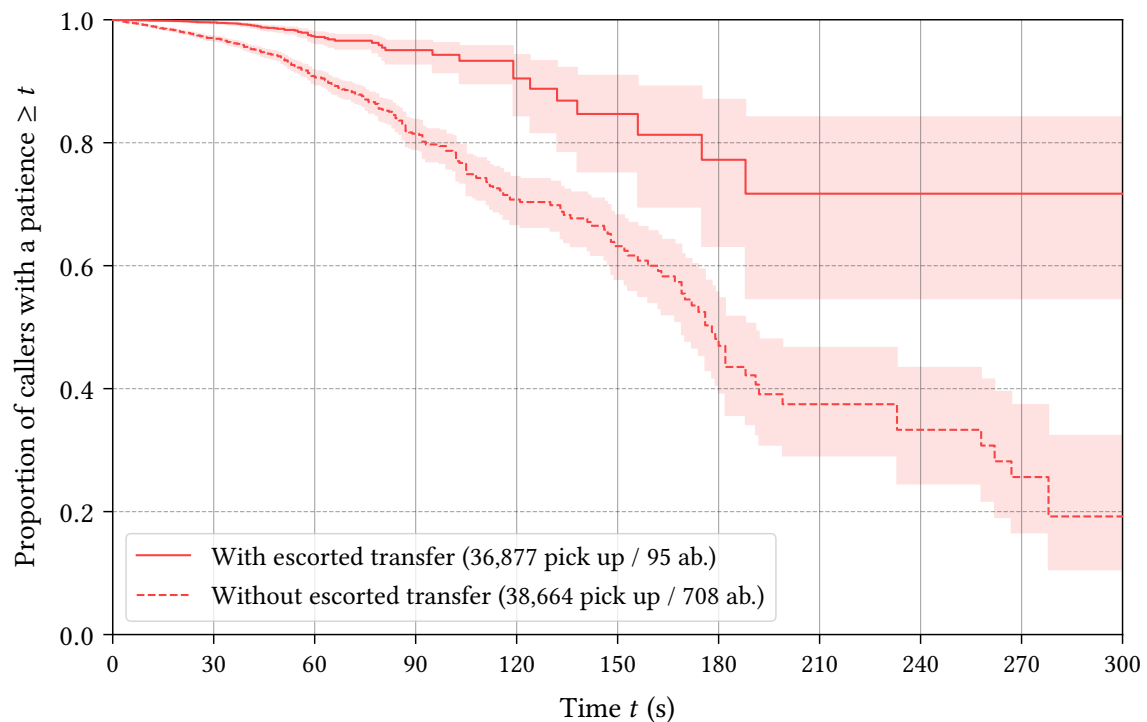


Figure 7.22: Influence of the mechanism of escorted transfer for VU calls on line 18

at least one minute, and this is about 70 % for people calling line 17. The fact that those curves remain flat after 190 seconds for line 17 and 220 seconds for line 18 means that these are the maximal abandonment times recorded at LVL 1. The representation of the curve for line 17 stops at 290 seconds because it is the longest observed pick up time in the year 2019.

Since the instruction at LVL 2 is longer than the one at LVL 1, we can witness relatively high waiting times for U calls, and we can represent the patience functions of those calls with a good level of confidence up to 10 or 15 minutes. As far as VU calls are concerned, the curves stop earlier because those calls are handled in priority and it is extremely rare to experience more than three minutes of wait. For these reasons, we chose to clip the figure between zero and five minutes of wait.

It is very interesting to note that the patience at LVL 2 is significantly higher than the patience at LVL 1. For instance on line 17, 20 % of callers have abandoned after 40 seconds of wait, while the same level of abandonments arises at 135 seconds at LVL 2 of DSPAP. We further remark that VU callers are more patient than U callers. One might argue that the patience improvement is not due to the bilevel layout but rather to the sample content. Indeed, the patience curve of LVL 1 is a mixture of NU, U and VU calls. However, this explanation does not suffice to justify the gap between the curves. Under the hypothesis that U and VU patience functions are unchanged between LVL 1 and LVL 2, we could recover from the green curves and using the proportions of gravity a patience function of NU callers at LVL 1 that turns impossible (total impatience reached in less than one minute, and then negative values). We can therefore effectively assert that first being welcomed at the platform by a quick pick up at LVL 1 enables to significantly increase the patience in view of a further instruction at the second level. We may also remark that for BSPP's LVL 2 calls that the patience of LU and U calls are very similar.

We now focus on the **influence of the escorting mechanism** for VU calls, where LVL 1 operators stay on the line with the caller until the arrival of a LVL 2 specialist. We have already mentioned in Section 4.3 that before mid-June of 2019, those calls were in fact not escorted. We have been able to leverage on the existence of both protocols in our datasets to compare the patience functions of the two situations, with or without accompaniment. This is depicted in the Figure 7.22. It proves that the escort mechanism of VU calls greatly increases the patience of callers and diminishes their tendency to hang up. More than 90 % of callers stay on the phone during two minutes when they are accompanied, versus 70 % when put on hold. The total counts also speak for themselves, both samples of the study have similar sizes, but there are **seven times less abandonments** in the escorted population.

As done in previous sections, we also review the influence of other factors on the patience levels. When distinguishing the patience functions per departments, we obtain a very high degree of similarity of LVL 1 patiences on line 18. On the other hand for line 17, there are small disparities, the callers from Paris (75) being a bit more patient than those of Val-de-Marne (94), themselves more patient than people from Seine-Saint-Denis (93) and finally Hauts-de-Seine (92). It is not easy to find an explanation for these differences, the corresponding calls could be made at different moments of the day (and the reason home/office or week/week-end may be involved). In addition, we do not know if the same voice prompts are used.

In order to study the possible dependence of patience functions upon the hour of the day, we have split the calls in four slots: morning (received from 7 AM and 1 PM), afternoon (1 PM to 7 PM), evening (7 PM to 1 AM) and night (1 AM to 7 AM). On line 17, computing the patience functions for these four windows at LVL 1 clearly shows that callers are more patient in the evening and the night (with similar profiles) than during the morning and the afternoon (close profiles as well). For instance, 40 % of callers will hang up on their own after 90 seconds during the day, versus 35 % at night. This phenomenon is much less visible for calls received on line 18, even though small differences remain (patience increased of roughly 1 % at night relatively to the day). For simulation purposes, we shall neglect the effect of the night and we will only compute the patience functions using calls received between 9 AM and 8 PM – further used regardless of the hour of the day. This is appropriate to study an average scenario during the day, and it does not underestimate abandonments if we apply these patience levels for a night scenario.

6 Some specific days

As mentioned at the end of the Section 2, we now focus on a couple of specific days in the year 2019. These could be used as scenarios that differ from the average profile, but that are still characteristic of the PFAU activity. The subsequent figures feature the throughputs of calls per gravity (NU, LU, U, VU, and EVU, as well as on-prompt and after-prompt abandonments) and the handling times (VU, U or EVU) in bold lines. The reference (average) curves of NU, U and VU calls of Figure 7.4 are given in thin lines to ease the comparison.

- 6.1 NEW YEAR'S EVE OF 2019 ON LINE 17.** – The top of the Figure 7.23 features the throughput curves for the “Nuit de la Saint-Sylvestre” (that is to say between the 31th December and the 1st January) on police number 17. Every year on this night during or after the festivities, law enforcement has to cope with public unrest or urban violence. The number of Urgent calls received on line 17 is already higher than usual in the afternoon of December the 31th, but it skyrockets from 11 PM to the end of the night – there are five times more U calls than usual. The same goes for VU calls whose average night throughput gets multiplied by six on this specific day. On the other hand, the arrivals of NU calls follow their usual behaviour. Note that we witness during the night a huge number of abandonments after the voice prompt, meaning that many calls are unanswered (and probably some of them try to call back once or several times).

The bottom figures show the cumulated distribution functions of the handling times τ_1 (still for calls received on line 17) and τ_2 (for calls handed over to LVL 2 of DSPAP). The bold lines were computed using only the calls received on New Year's Eve between 10 PM and 8 AM. We can see that both U and VU calls are handled with exactly the same characteristic times than every other day.

- 6.2 BASTILLE DAY NIGHT OF 2019 ON LINE 18.** – The night before and the night after the French national day (14th of July) also suffer a rise of civil incidents, putting all the firemen of the country in the front line. By looking at the Figure 7.24 (top), we see that until 10 PM, the throughput of NU, U and VU calls follow the usual average trend. However, from 10:30 PM, the urgent calls soar until midnight, with a maximum throughput equivalent to 500 calls per hour. In a much lesser extent, there is also a great rise of NU calls in the beginning of the night. The huge and exceptional load of U calls decreases but lasts during the entirety of the night, with a slight upturn between 4 AM and 6 AM. The normal profile is recovered in the morning of July the 15th. Note here that the throughput of abandonments remained very low (almost null), which shows that the huge demand has indeed been coped with and absorbed (recall the firebrigade staffing for this very night depicted in Figure 7.10).

The bottom part of the Figure 7.24 gives the cumulated distribution functions of the conversation times for Bastille Day night between 10 PM and 8 AM. It brings an additional key to understanding how the firefighters managed to resist the peak. Indeed, we see very clearly that the BSPP agents handled the calls from Bastille Day much quicker than usual. This holds for LVL 2 conversations but also at LVL 1 for U calls. Indeed during this night, 75 % of Urgent calls were handled in less than 1 min 30 at LVL 2, versus 3 minutes for a normal day. The VU calls are also treated 30 seconds quicker at LVL 2. The first level also goes faster, it transfers urgent calls to the LVL 2 with a 5-second gain compared to its usual pace. Although the incoming flow of calls is atypical, the previous differences illustrate that the firebrigade has a strong adaptability and dedicated protocols to handle planned events.

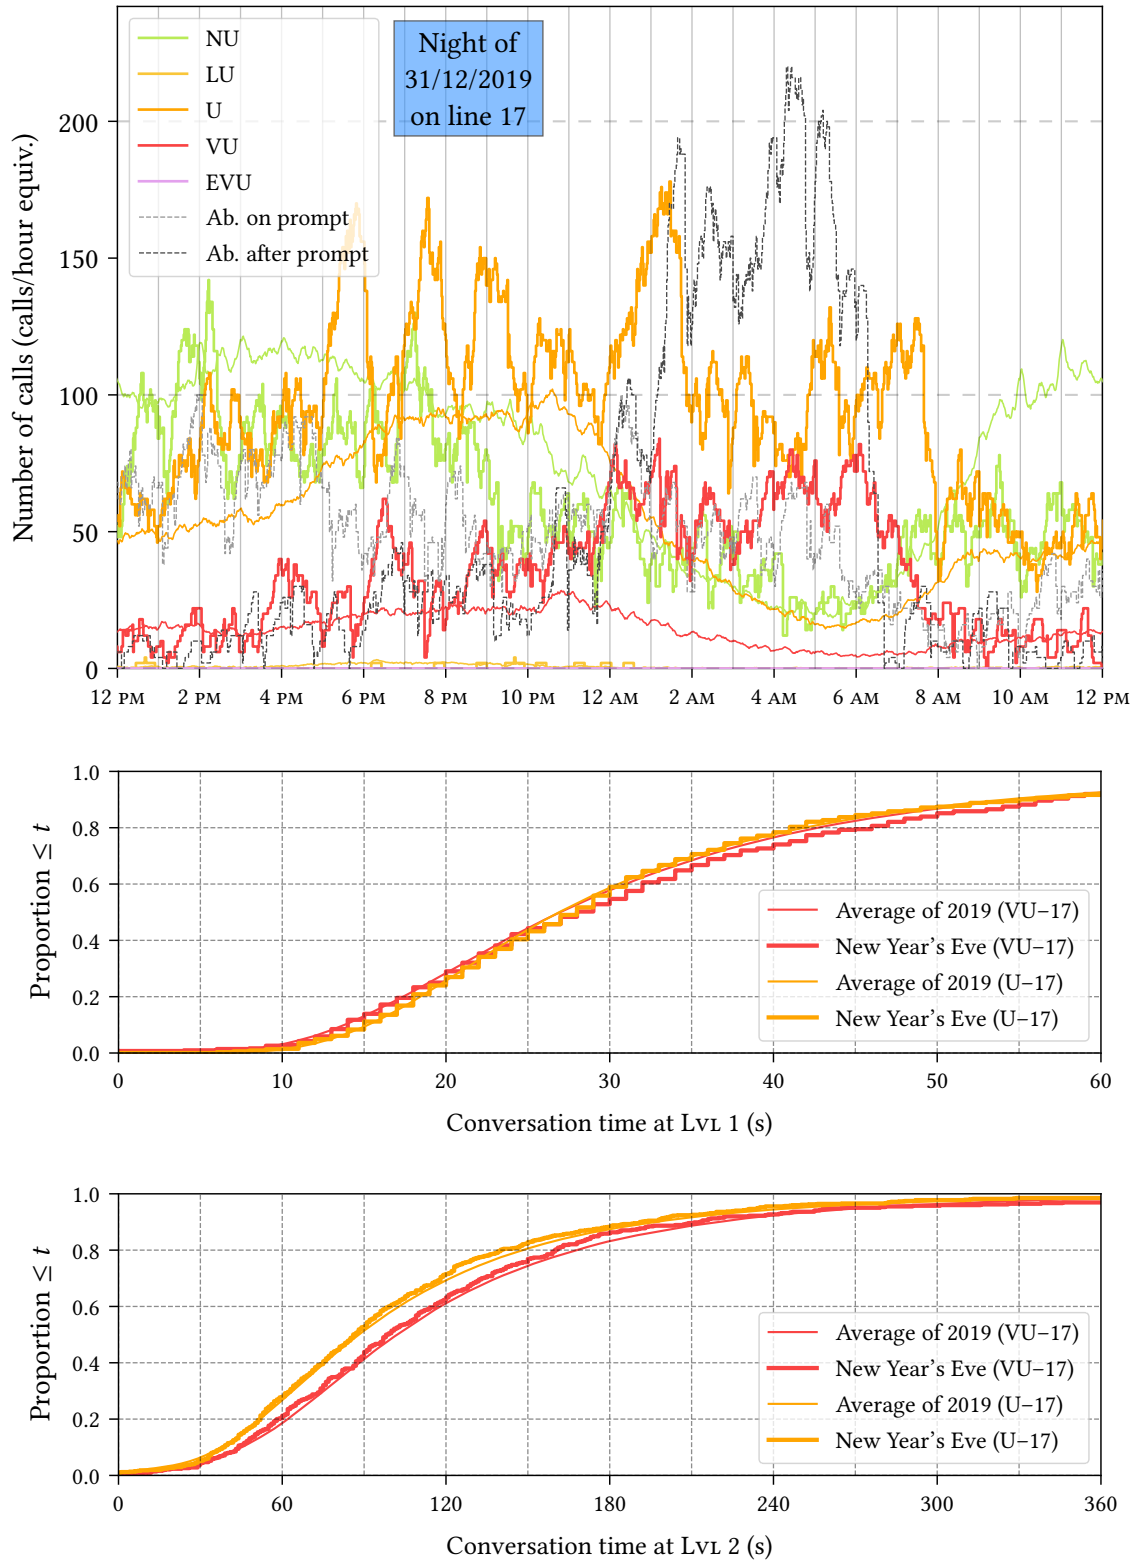


Figure 7.23: Throughputs and conversation times for the night of 31 December 2019

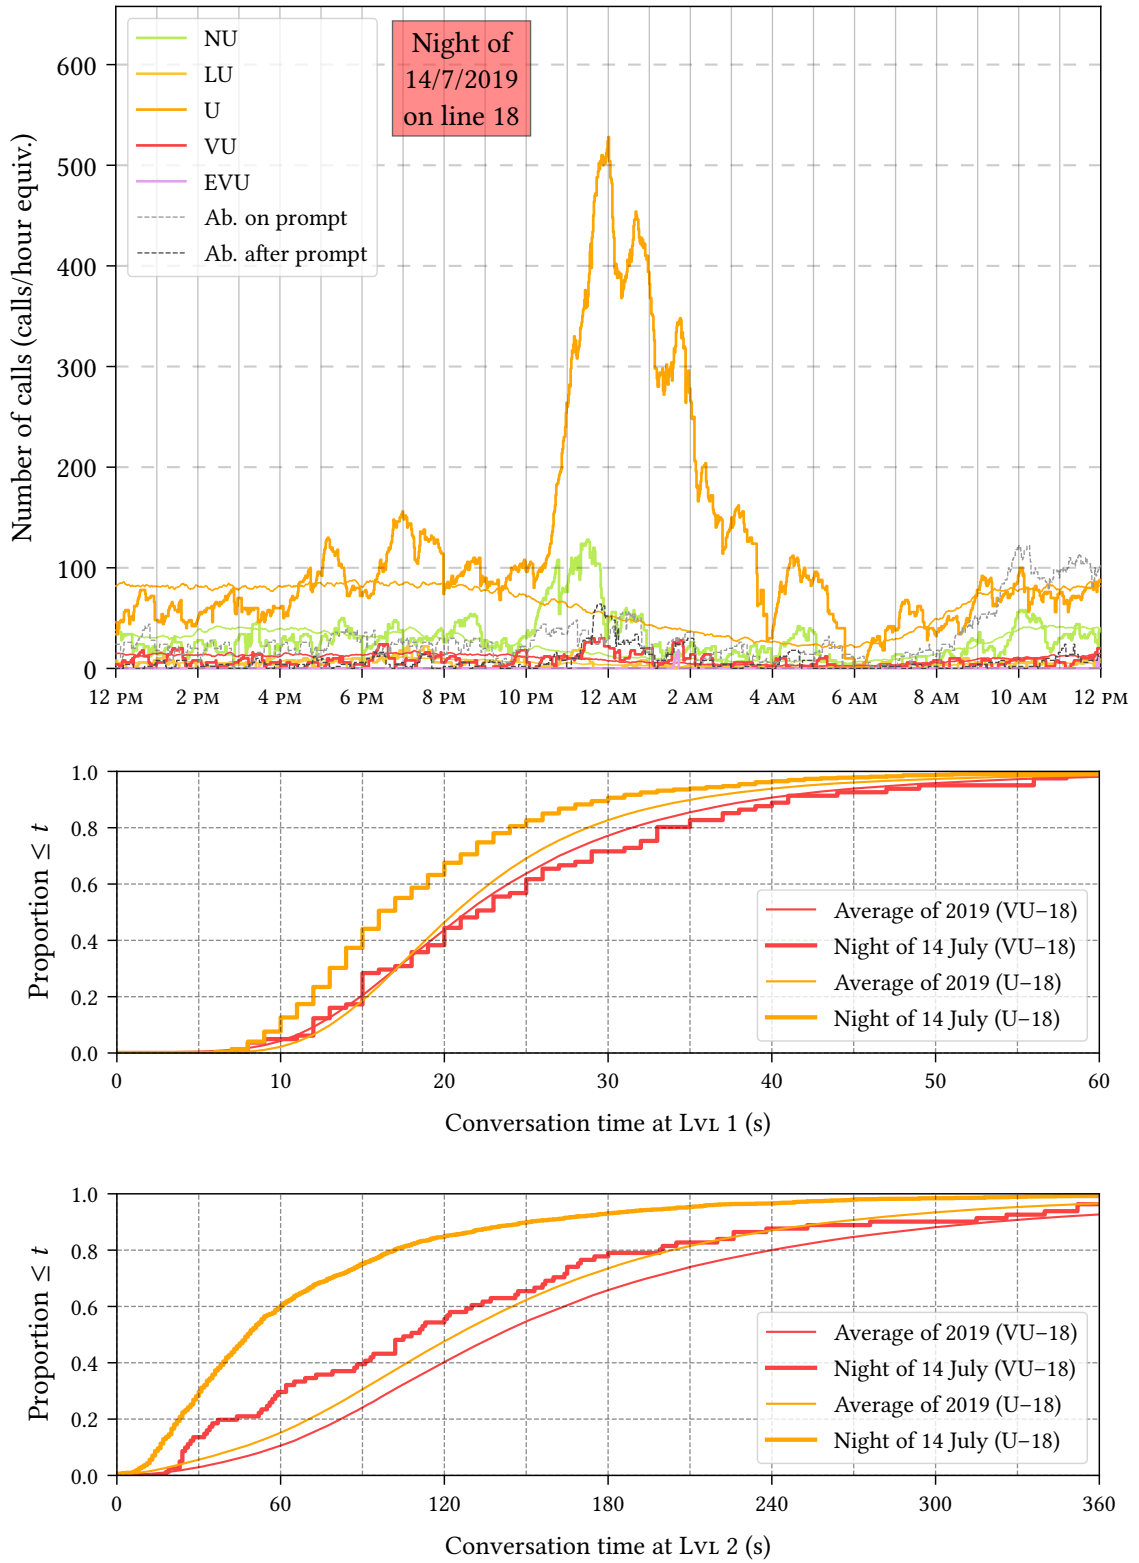


Figure 7.24: Throughputs and conversation times for the night of 14 July 2019

6.3 FIRE AT BOBIGNY ON 26 JUNE 2019 ON LINE 18. — The day of 26 June 2019 is interesting because the profiles of LU (Less Urgent) and EVU (Event Urgent) curves are atypical (see the top part of Figure 7.25). The demand on the Urgent lane is higher than usual, probably due to the heatwave experienced in France at this time. We can also suppose that this is the reason why the LU lane was activated during the afternoon. The curve of EVU calls is characteristic of a one-off, short and intense event. It corresponds to an important fire which broke out in a junkyard in Bobigny (located at one kilometer from Paris in the North-East suburb, department of Seine-Saint-Denis). It emitted a thick smoke which was visible from the Parisien ring-road ¹ – and we indeed check in the data that the received EVU calls come from both Paris and Seine-Saint-Denis.

This fire is the biggest one of the year 2019 in terms of calls labeled as EVU. The associated load peaked at an equivalent of 260 EVU calls per hour. We also remark a great number of abandonments after the voice prompt that are very probably linked to the same event (recall that this curve does not count the people who called back and were picked up, hence these calls were never handled). Hence, we can estimate that the real PFAU load associated with the Bobigny fire was equivalent to 400 calls per hour at the peak, even though most of these calls were concentrated in a 30-minute time window. In parallel of the EVU calls, we observe an important rise of VU calls, most certainly also about the fire but treated differently. The firebrigade protocol often labels calls of immediate victims needing assistance or very close witnesses as VU, while mere witness from a distance go in the EVU lane.

The bottom part of the Figure 7.25 once again shows the cumulated distribution functions of conversation times between 7:30 PM and 10 PM on 26 June. It illustrates well the singularity of EVU calls. While U and VU calls are transferred from Lvl 1 to Lvl 2 with the same typical times as a normal day, the EVU calls related to the fire are identified and handed over more quickly (about 2 to 3 seconds). Above all, at Lvl 2, those EVU calls are handled very quickly – more than 80 % conversations below one minute, versus 15 % for “classical” U calls. We note in addition that the U calls are not handled in a different way than usual, which shows that the incoming flows were well separated by the first level. On the other hand, VU calls are handled a bit quicker than usual at Lvl 2 (30 seconds less in average), since most of the calls are related to the fire.

¹ <https://www.francebleu.fr/infos/faits-divers-justice/un-incendie-dans-une-casse-automobile-a-bobigny-1561572815>

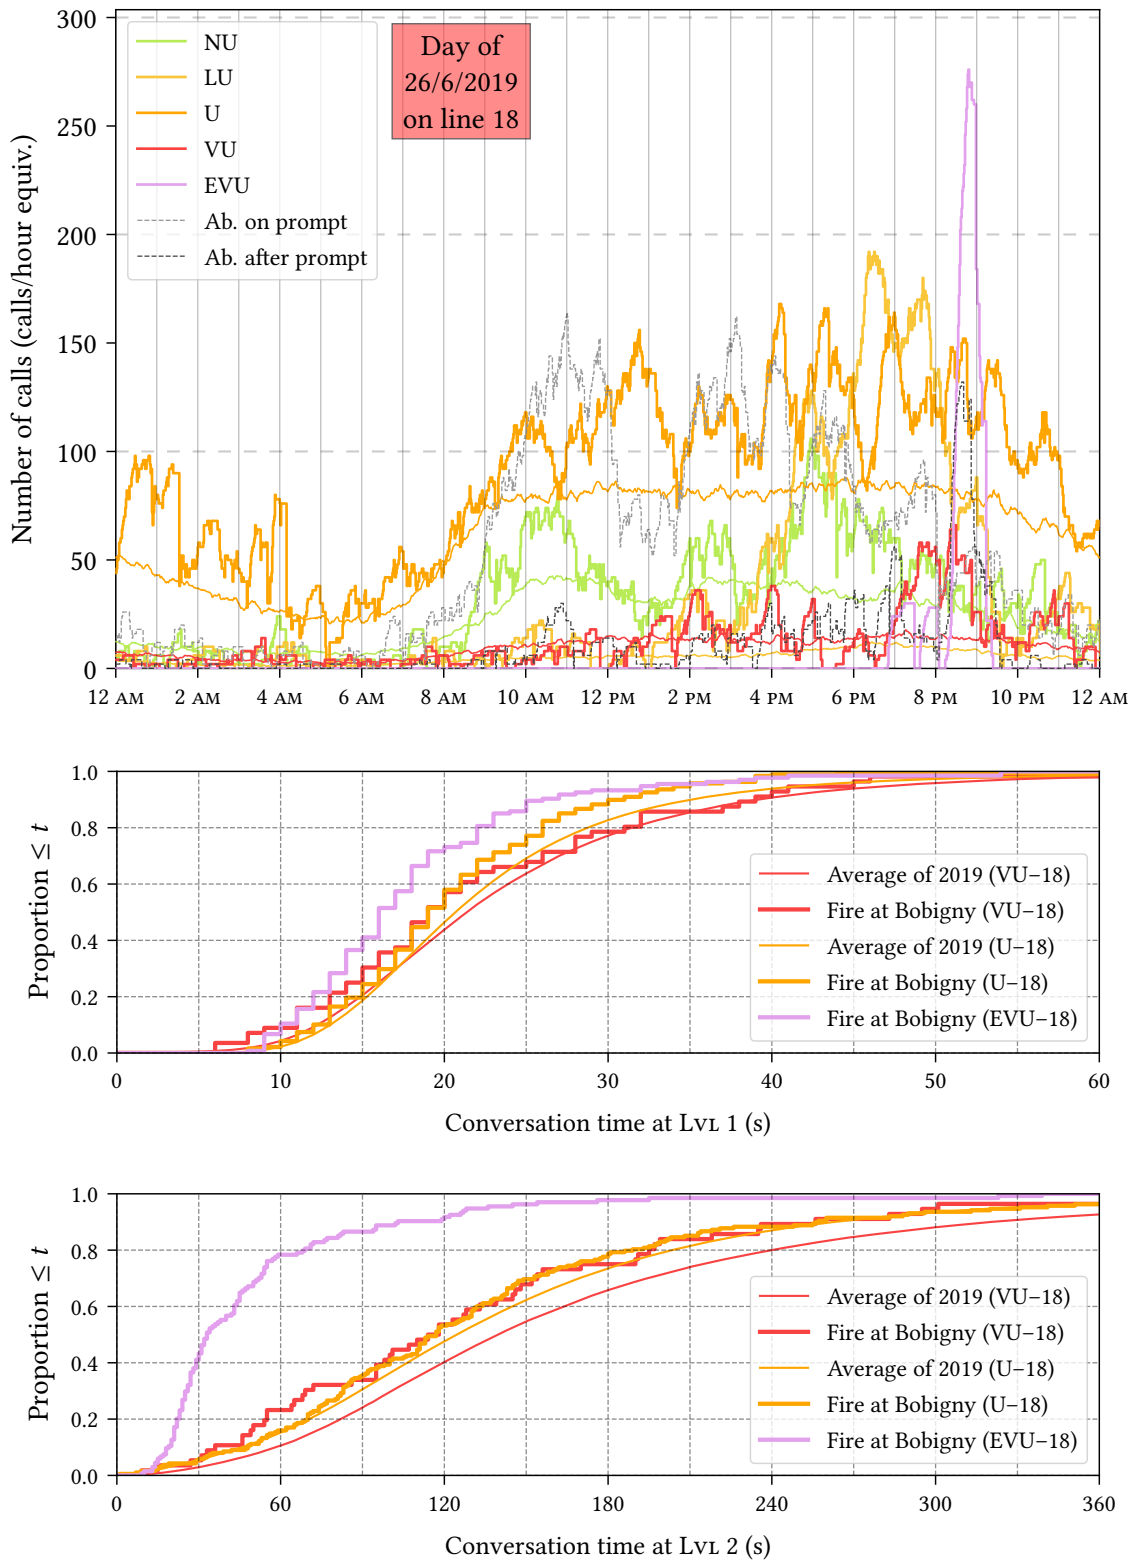


Figure 7.25: Throughputs and conversation times for the day of 26 July 2019

6.4 STORM ON 4 JUNE 2019 ON LINE 18. — The throughput curves of Figure 7.26, computed for the 4th of June, tend to oscillate around their average behavior until 4 PM. At this moment, a violent thunderstorm hit the Paris area and triggered more than 30 BSPP interventions² (falling trees, accidents, floodings...). It is interesting to note here that apart from the peak of Urgent calls just before 6 PM, the curve of Less Urgent adopted a very atypical profile with respect to its average trend. It seems that a specific instruction was given to BSPP agents, to use this LU label (usually not very much used) in order to establish a hierarchy of calls and prioritize the response. This way, the VU lane could still be reserved for life-threatening situations.

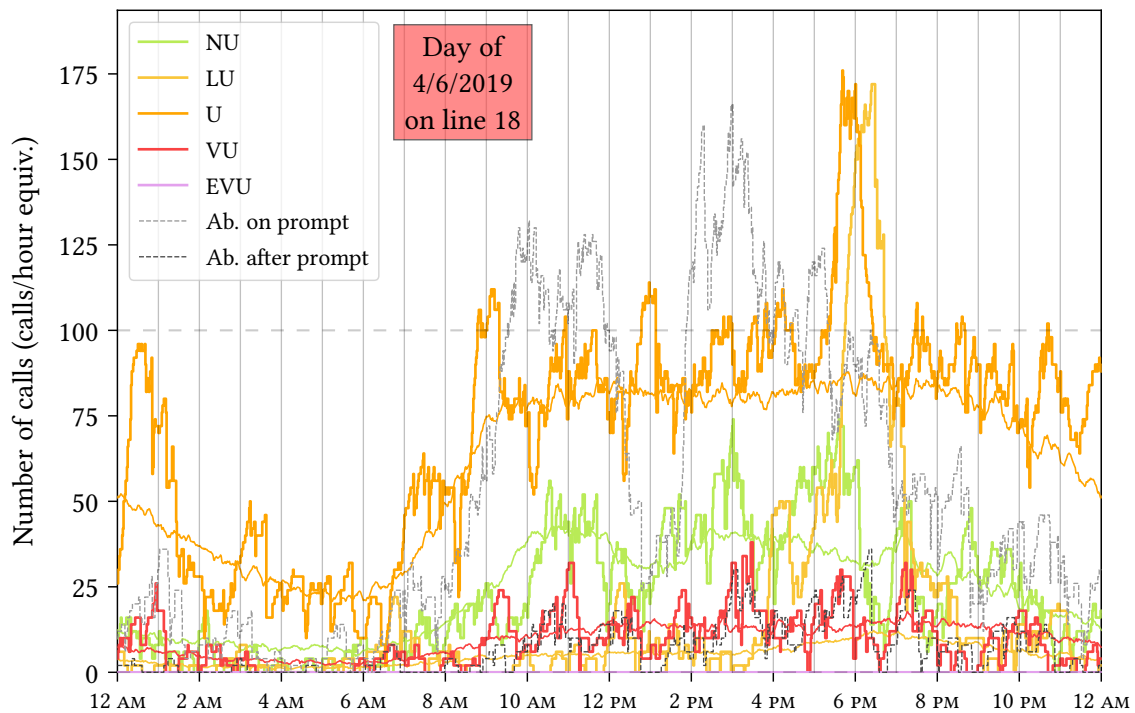


Figure 7.26: Throughputs for the day of 4 July 2019

² <https://www.leparisien.fr/hauts-de-seine-92/hauts-de-seine-le-violent-orage-fait-tomber-plusieurs-arbres-04-06-2019-8086407.php>

6.5 NOTRE-DAME FIRE ON 15 APRIL 2019 ON LINE 18. — All the eyes of the world were this day on Paris and its firefighters, after an unprecedented fire devastated the Notre-Dame de Paris cathedral from 6:20 PM. The EVU protocol was activated and received according to the Figure 7.27 an equivalent of 180 calls per hour short after 7 PM, that is to say when the smoke was visible in the city. It is interesting to note that the throughput of VU calls plummets at the same time as the EVU protocol is closed, and the NU calls skyrocket. We can guess that at this moment, the firefighters had much enough information about the fire and had already dispatched teams, hence received the instruction to not transfer the corresponding calls at LVL 2 anymore. We did not find an explanation about the peak of VU calls at 11 AM.

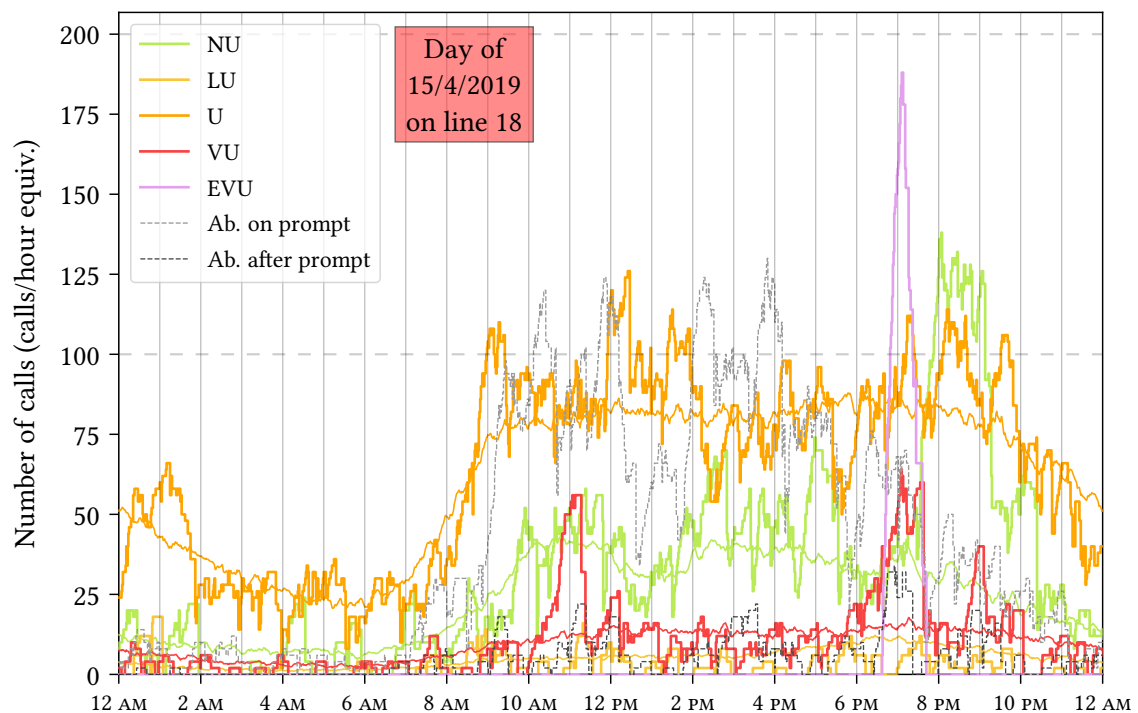


Figure 7.27: Throughputs for the day of 15 April 2019

8

PERFORMANCE ANALYSIS USING NUMERICAL SIMULATION

Contents

1	Methodology	290
1.1	Comparison of the analytical and simulation approaches	290
1.2	Interface between data analysis and performance evaluation	292
1.3	Comprehensive list of modeling assumptions	294
2	Performance analysis of a bilevel and single-job system, comparison with a single-tier layout	296
2.1	Petri net model and performance criteria for the bilevel system	296
2.2	Performance of the single-job and bilevel system	299
2.3	Comparison of the bilevel and single-tier layouts for rescue calls	303
2.4	The case of police calls and the key role played by the loads ratio	307
2.5	Discussion on the iso-staffing comparison	311
3	Validation of the analytical results by simulations	314
3.1	Throughput estimates of Petri nets used for the staffing problem	314
3.2	Influence of the occupancy rate	317
3.3	Performance measures derived from Erlang calculus	319
4	Study of the PFAU architecture and its possible evolutions	326
4.1	Petri net models for organizations with two distinct jobs	326
4.2	Study of the interactions between the two jobs of the platform	328
4.3	Effect of the system departmentalization	332
4.4	Effect of the first level unification	336
5	Behaviour of the bilevel layout under non-standard arrivals	338
5.1	Response of the bilevel layout to an unexpected peak of calls	338
5.2	Response to a planned peak, the case of rescue calls during Bastille Day	340

We study in this chapter the performance of the PFAU using simulation tools. This is a common approach in the world of call centers and for emergency purposes (see for instance [Sav69]). We refer once more to [ALo5] for a review of the simulation literature on call centers. In the more specific field of emergency call centers, we may mention the recent work of Lewis et al. [LHSCo7] on the resource allocation problem – and which gives additional references. In [vBKvdMB15], van Buuren et al. use simulation to tune the staffing of an emergency call center. See also the thesis of Petitdemange [Petzo] (carried out in partnership with three SAMU of Southern France) in which call centers are simulated using a third-party software.

Our approach differs from the above by several aspects. To begin with, we build on a very large amount of real data, as a consequence we have a precise knowledge of the distribution of the call center parameters (see our Chapter 7 for a detailed discussion). Another originality is to study bilevel call centers, in contrast with single-tier layouts that are most often simulated. Above all, we use in our study a simulation software that was specifically developed in our research team for this very application. This tailor-made solution allows us to handily take into account all the phenomena presented in the previous chapter: the intra-day variation of the call arrivals, the variability of the conversation times and the correlations between them, the patience of the callers depending on the gravity of the calls, etc. Having a simulation software of our own also enabled us to implement realistic routing rules used in the world of call centers (priority rules, first-in-first-out pick ups, etc).

The Section 1 details our methodology and gives the exhaustive list of all the assumptions we had to make. Our Section 2 deals with the simulation of a bilevel call center inspired by the PFAU, although in a single-job layout. In particular, we extensively compare the performance of the single-tier and the bilevel architectures. In the Section 3, we compare the numerical results derived by simulation with the analytical formulas computed in the Chapter 5. We study the performance of the real PFAU in the Section 4, taking into account the multi-job features. Further, we simulate several alternative layouts for the PFAU. This involves merging currently separated pools of agents. Eventually, the Section 5 is devoted to the study of the PFAU outside average scenarios, that is to say in case of unexpected peak of calls.

1 Methodology

- 1.1 COMPARISON OF THE ANALYTICAL AND SIMULATION APPROACHES.** – We come back in this chapter on the modeling of emergency call centers by timed Petri nets, extensively studied in the Part I of the thesis, in particular in the Chapters 1, 2, 4 and 5. As discussed there, those models are governed by dynamical equations from which we can derive closed forms of some performance estimates of the organizations – in particular the throughput. These analytical formulas are obtained at the expense of several approximations, mainly the relaxation of integrality constraints and the hypothesis that tokens have an infinite patience. We have also assumed that the sojourn times of the tokens were constant – this was justified by the fact that the throughput only depends on the average parameters of the net (see the Remark 4.6). In addition, we have always assumed in Chapters 4 and 5 that calls arrived according to a linear profile. This is a satisfying approximation in average but not necessarily appropriate on long time-windows.

The simulation approach aims at implementing a software that reproduces *in silico* the behaviour of an emergency call center (modeled by a timed Petri net) given a scenario of call arrivals and a choice of operator staffing. The advantage of this method is to grant a much wider modeling power and to get rid of the approximations we had to make to establish theoretical results. In particular, we may simulate the actual patience of callers. It suffices to attach to

each token a patience level μ in every “waiting place” of the net, so that if the sojourn time exceeds μ , then the token leaves the net. Likewise, we can generate instruction times that vary from one token to another, which recreates the variability of situations taking place in a same place. The throughput of call arrivals (per line, department, gravity and LVL 2 orientation) do not need to be constant anymore and can follow an arbitrary profile. Finally, the simulation approach allows to consider very general routing rules in case of conflicts between several transitions. This includes preselection and static (or token-independent) priorities introduced in the Section 2 of Chapter 1, but also dynamical (or token-dependent) priorities in which we may fire a transition by consuming an older token rather than a younger token. In particular, priorities between transitions are determined by first comparing the degree of static priority, and then using dynamical information in case of equality. The degrees of dynamical priorities are computed based on the age of the tokens, so that we always fire the oldest possible tokens in case of conflict between several of them. The Table 8.1 below summarizes the modeling frameworks offered by the two approaches.

	Analytical calculation	Numerical simulation
Tokens patience	Infinite	Arbitrary
Sojourn times	Place-dependent	Token-dependent
Arrival throughputs	Constant	Arbitrary
Conflicts resolution	Preselection and static priority	Arbitrary

Table 8.1: Modeling power of analytical and numerical approaches

An additional advantage of the numerical simulation is that it allows us to appreciate the variability of the system performance. Indeed, we can repeat multiple times a same simulation involving stochastic aspects (for instance the arrival dates of calls or their handling times) in order to derive an average performance *and* a characteristic deviation around the latter.

It is paramount to ensure that the simulation software accurately simulates the studied call center, otherwise the subsequent performance analysis would be incorrect. In other words, a simple bug or error in the implementation of the software could lead to wrong conclusions. As a consequence, we have chosen a development methodology similar to the one used for critical software. We have used the OCaml programming language [LDF⁺20], whose typing system during compilation allows one to identify many bugs in the development phase. In addition, we deployed a chain of unit and regression tests to prevent errors due to software evolution. Finally, we have designed the software so that its architecture allows an easy review of the code. This layout is a hierarchy of three components. The low-level part is the simulation kernel. It is a fairly short software (a couple hundred lines of code) whose aim is to compute the evolution of tokens in Petri nets. The middle layer implements the **routing oracle**, which performs the routing rules. This routing oracle is called by the simulation kernel to know which transitions should be fired at any instant. The high-level component corresponds to the specification of the call center that is studied. It takes in input the topology and the routing rules of a user-defined Petri net instance.

1.2 INTERFACE BETWEEN DATA ANALYSIS AND PERFORMANCE EVALUATION. — The analysis of call data presented in Chapter 7 is key to provide precise figures to our analytical and simulation methods. On one hand, our theoretical formulas indeed require the knowledge of average values of some primitive parameters, like the throughput, the conversation times and the different proportions of gravity. On the other hand, we need the complete distributions of such parameters to generate realistic simulation scenarios. In this section, we explain how we go from the data to the simulation, and finally to the computation of performance estimates.

The first step of the chain is to build simulation scenarios. As explained in Section 2 of Chapter 7, we will generate arrivals for each emergency line, each gravity level, each department of Paris area and if applicable each possible LVL 2 orientation. We assume that these arrivals follow nonhomogeneous Poisson processes (whose intensity is the throughput of received calls), which we have shown to be a very good approximation. Working with nonhomogeneous Poisson arrivals is a little trickier than classical Poisson arrivals, and we resort to the Çinlar’s method (see [Ç13, p. 99]) to get accurate interarrival times. From one simulation to another, the random numbers used to obtain the input scenarios will produce varying arrival dates, and can give different global counts of received calls on the same time-window. The throughputs used as intensity functions of the Poisson processes are user-specified. We can choose to use the average profiles computed in the Section 2 of Chapter 7 to get performance measures representative of the typical day-to-day activity of the PFAU. Alternatively, we can plug the throughputs derived from specific days as studied in the Section 6 of Chapter 7, which allows us to “replay” an event of interest.

For each call, we draw a random conversation time at LVL 1 and a patience level according to the (line-dependent, gravity-dependent, and department-dependent) histograms or cumulated distribution functions of Chapter 7. For the calls that do not have a NU gravity, we also draw a conversation time at LVL 2 (correlated to the handling time at LVL 1), another patience at LVL 2, and an extra instruction time after hang up for the LVL 2 operator (correlated to the handling time at the second level and often null). We also draw synchronization times associated with the escort step for VU calls. We insist on the fact that all these parameters do vary from one simulation run to another, but are all drawn from a common scenario.

The resulting workload obtained with the above gravity levels is not complete until we take into account the abandonments. We choose to totally ignore in simulation the calls that were abandoned on the voice prompt in real life, since they were filtered out by the system and invisible to the agents. On the contrary, we keep the throughput of after-prompt abandonments and we add it in the queue of calls to be picked up. Indeed, we consider that these calls should have been handled. However, we need to decide what are the gravity level of such calls once picked up, that was by essence unknown in the step of data analysis. We may for instance assume that 97 % of calls previously abandoned after the prompt were actually Non Urgent, and 3 % would have

ID	Arrival	Grav.	Line	Dpt.	Tau_1	Mu_1	Orient.	Tau_2	Mu_2	Tau_3	Tau_s
78	09:18:33	NU	17	93	7.0	22.0					
79	09:18:42	NU	112	75	7.0	54.0					
80	09:19:11	VU	18	92	34.0	57.0	BSPP	114.0	171.0	0.0	10.0
81	09:19:12	U	17	75	18.0	47.0	DSPAP	90.0	1726.0	0.0	
82	09:19:20	NU	112	92	18.0	110.0					
83	09:19:22	NU	112	94	8.0	10.0					
84	09:19:45	U	18	93	14.0	74.0	BSPP	207.0	433.0	73.0	

Table 8.2: Example of call arrivals and characteristics

been qualified as Urgent. Naturally, these parameters can be modified and it will be appropriate to adapt them for some scenarios – for instance during a “crisis event” where the platform is overwhelmed, most of the abandoned calls in real life were probably U, VU or EVU. We will detail these choices in the coming Section 1.3.

The Table 8.2 is an example of input data that could be given to the simulation software. The column labels τ_1 , τ_2 , τ_3 and τ_S respectively refer to the conversation time at LVL 1, the conversation time at LVL 2, the extra instruction time for the LVL 2 operator after hang up, and the synchronization time for VU calls between the two levels (escorting time). The column labels μ_1 and μ_2 refer to the patience levels of the callers at LVL 1 and LVL 2.

At the other end of the simulation chain, we need to convert the simulation results into global performance measures. Each simulation run (associated with an input scenario) will give us the waiting time of all the callers at each step, and the fact that they abandon or not. To begin with, we can compute performance estimates of interest for a single scenario/simulation run. We shall typically focus on the total percentage of abandoned calls, of calls abandoned at LVL 1 and LVL 2, or per gravity. We may also compute the average waiting times (at LVL 1 and/or at LVL 2) of all callers. Another common indicator is the **service level**, giving the proportion of calls picked up before a given threshold (typically 30 seconds, or one minute). The complete knowledge in the simulation of the activity of individual agents also enables us to compute occupancy rates per agent, and average occupancy rates per pool of agents (for instance at LVL 1 and LVL 2). We will come back more precisely on performance indicators in the Section 2.

The Table 8.3 illustrates the result we can have after aggregating performance estimates per simulations. It is derived for three different staffing scenarios, where the tuple (N_1, N_2) (number of agents at LVL 1 and LVL 2) varies in $\{(1, 5), (2, 5), (2, 6)\}$. For each one of these scenarios, three runs are performed, hence a total of nine simulations. Observe that the total number of received calls in each run is different due to the random draws involved in the Poisson processes – even though arrivals are generated using the same deterministic intensity functions. Four performance indicators are shown in the Table 8.3: the fraction of abandoned calls at LVL 1, the same at LVL 2, the service level at LVL 1 below 30 seconds, and the average wait experienced by callers at LVL 1.

In order to get a global performance of a single scenario (fixed parameters and fixed staffing), we perform a large number of times the simulation process, that is to say on input data generated from the same primitives. Then, we can take the mean values of the associated performance indicators. Recall that the variance of this mean value is inversely proportional to the number of runs. In view of the empirical standard deviations that we measured, we considered that achieving 100 simulation runs per scenario was sufficient to obtain a satisfying confidence interval on the computed performance. This number could of course be increased if necessary.

Simu_ID	Lvl1_op	Lvl2_op	Calls	%_Ab_Lvl1	%_Ab_Lvl2	%_SL30_Lvl1	Avg_wait_Lvl1
1_5_67	1	5	2082	27.71	1.29	43.65	46.74
1_5_68	1	5	2068	28.77	0.87	48.87	47.23
1_5_20	1	5	2065	29.00	1.80	44.54	48.10
2_5_95	2	5	2081	5.76	5.36	90.51	8.37
2_5_3	2	5	1977	4.95	5.07	94.89	5.65
2_5_92	2	5	2064	5.66	6.62	92.44	7.63
2_6_42	2	6	2136	4.11	3.71	93.70	5.99
2_6_33	2	6	2075	4.91	3.46	91.78	7.46
2_6_88	2	6	2076	4.52	1.61	95.20	5.62

Table 8.3: Example of aggregated performance indicators for nine simulations

1.3 COMPREHENSIVE LIST OF MODELING ASSUMPTIONS. – All the simulation results presented in this chapter must be interpreted by keeping in mind the many modeling assumptions that had to be made. Some of them have an influence on all the simulation aspects, while some others specifically impact the comparison between single-tier and bilevel layouts. The following lists detail all our assumptions.

Assumptions that impact all the simulations, regardless of the organization.

- We assumed an identical patience function before the first pick up, regardless of the gravity. As explained in the Section 5 of Chapter 7, we have made this hypothesis by lack of a canonical method that could confirm that VU and U callers are more patient than NU callers – although it comes rather intuitive. This assumption then tends to overestimate the abandonments of U and VU calls before the first pick up, independently of the studied organization.
- The callbacks are never simulated, even though we showed in the Section 2.3 of Chapter 7 that these calls account for 6 % of the total picked up volume. The reason behind this choice is that it is more complicated to attach a patience to these calls, and additional assumptions would be in order to correctly model them. On the other hand, out of consistency, the input throughput of arrivals that we consider are always net of callbacks. Therefore, they represent the real demand, not the “inflated” one which contains the people who abandon then call back.
- We have chosen not to take into account in the simulations the abandonments occurring on the voice prompt (8 to 10 seconds). Indeed, we can consider that they do not correspond to a real emergency. However, as far as the abandonments recorded after the voice prompt are concerned, we deem that they should have been picked up – this surely is arguable for calls abandoned just after the prompt, for instance after 11 or 12 seconds of wait, but a threshold need to be fixed anyway. The hitch is that by essence, we do not know what the gravity of these abandoned calls would have been. We have consequently made the following choices. Between 9 AM and 7 PM, calls arrived on line 17 and abandoned in our data are NU in 85 % of the cases and U in the remaining 15 %; on line 112, abandonments are 98 % NU and 2 % U (breaking down in equal proportions between U–DSPAP and U–BSPP); and 97 % NU and 3 % U on line 18. From 7 PM to 11 PM, these proportions become 70 % NU and 30 % U on line 17; 96 % NU, 2 % U–DSPAP and 2 % U–BSPP on line 112; and 95 % NU and 5 % U on line 18. Remark that all these choices assume that there is no VU calls among the calls abandoned on prompt. Once the gravities of the above calls have been set, we also draw in accordance patience levels and conversation times from the distributions of the calls with same gravity (and same initial line, department and Lvl 2 orientation).
- We have retained in our simulations the outbound activity of Lvl 2 operators, since it is an important part of their tasks. However, we have neglected the outbound calls emitted by the Lvl 1 agents. As mentioned at the end of Section 3 du Chapter 7, this activity represents about 3 % of their time spent on the phone – our approximation is then quite acceptable. On the other hand, we have not taken into account the calls arriving directly at Lvl 2, most of time coming from PFAU partners. The Table 7.1 tells us that these calls represent about 7 % of the total inbound calls at Lvl 1 and Lvl 2, but at least a third of these calls come from the SAMU and are picked up by a different pool of operators of the PFAU, not considered here. In fact, we can expect that the workload of Lvl 2 agents involved in the response of emergency calls of the population is underestimated by 3 % of 4 % at most.

Assumptions that specifically impact the single-tier/bilevel comparison.

- We have chosen on the figures comparing single-tier and bilevel layouts to compare *iso-staffing* performance criteria. In this view, we would for instance compare a single-tier system with 10 operators with a bilevel architecture with 3 operators at LVL 1 and 7 operators at LVL 2. It is important to insist on the fact that *iso-staffing* is very different from *iso-skill*. Indeed, the LVL 2 agents of a bilevel system have to be expert operators, much more trained than the LVL 1 agents whose tasks are more straightforward. In a single-tier system however, all the agents have to be experts (same level of competence than the LVL 2 agents of the bilevel) in order to cope with all the possible situations. In order to illustrate the influence of this methodological choice on the figures, we shall come back to it in details in the Section 2.5.
- An arguable modeling choice is that for single-tier systems, we took as the single conversation time either the LVL 1 conversation time of the bilevel system in case of NU call, or the LVL 2 conversation time of the bilevel otherwise (U or VU call); up to a random extra time between 1 and 6 seconds. This boils down to assuming that the filtering step at LVL 1 and the calls pre-qualification does not save time for the LVL 2 instruction – which remains to establish in practice. It is also an optimistic convention for the single-tier architecture, since we have reasons to believe that it is precisely thanks to the two-tier layout that the handling times of the bilevel system are so “short” and appropriate to the calls – owing to the more standardized protocols made possible by this system. The operators confirmed to us that before adopting the bilevel system, the conversation times of the single-tier system were significantly longer.
- We have not been able to measure the workload induced by the extra instruction after hanging up the calls at LVL 2 for police tasks, since the corresponding information were not in the dataset we were given. In our simulations, we assumed that the police calls had post-LVL 2 instructions identical to those of the rescue/firebrigade lane, in terms of handling times and proportions of calls indeed entailing an extra instruction time. In reality, it is plausible that this workload is underestimated and that police protocols are a bit longer than rescue ones. Extending the extra instruction times after LVL 2 would make our results more favorable to the bilevel architecture, by increasing the load ratio (see Section 2.4).
- We have shown in the Section 5 of Chapter 7 that for U and VU calls, the inter-level patience was greatly increased with respect to the patience before LVL 1 (see Figure 7.20). However, we have not computed a proper correlation between those two levels due to the inherent censorship of the datasets. Hence, we have drawn for simulation purposes independent patience levels μ_1 before LVL 1 and μ_2 before LVL 2. However, combinations with a high value of μ_1 and a small value of μ_2 would be unrealistic – and unfair for the bilevel system. Thus, we have chosen to avoid these by slightly tweaking the generated patiences in our scenarios. For calls such that $\mu_2 \leq \mu_1$ and $\mu_2 \leq 30$ s, we decided to update the LVL 2 patience to $\mu_2 := \max(\frac{1}{2} \times \mu_1, 30 \text{ s})$. This adjustment only affects 3 % of the total of U and VU calls, and hardly changes the patience distributions – while preventing unrealistic patience tuples. More generally, the fact that inter-level abandonments seen in our datasets come from “relevant demand” is questionable. This is all the more true for abandonments occurring after a short wait – and these calls are responsible of the possible low values of μ_2 drawn in simulations. Indeed, it can happen that several witnesses of a same event initiate multiple emergency calls in parallel, and that they hang up once one of them is transferred to the LVL 2, or as soon as a response team arrives. There could also be technical failures or errors. Thus, the above corrections seem perfectly appropriate to improve the realism of simulations.

2 Performance analysis of a bilevel and single-job system, comparison with a single-tier layout

As discussed in the [Foreword](#) and illustrated with the [Figure II.2](#), the aim of this section and of the following one is to play the role of a “scientific demonstrator” of our two complementary approaches to compare different emergency call center layouts. In particular, we focus here on one of the most important aspects of the transformation brought to the treatment of emergency calls by the PFAU, that is to say going from a single-tier to a bilevel architecture. To avoid overcomplicating things, we have chosen to first study an idealized [single-job](#) platform, picking up only one type of calls (either police or rescue). We will derive very rich conclusions from this simplified framework already, and the study of [multi-job](#) systems will be conducted in the [Section 4](#).

We recall that the following simulations are based on a large-scale organization (covering four densely populated French departments), with high throughputs of calls and significant staffing levels. Simulating smaller platforms may lead to different results in some extent.

2.1 PETRI NET MODEL AND PERFORMANCE CRITERIA FOR THE BILEVEL SYSTEM. — We start by representing in [Figure 8.4](#) the timed Petri net associated with a bilevel system, according to the description given in the [Foreword](#) (however for a single-job version). This Petri net needs to be compared with the simpler net of [Figure 5.2](#) in the [Chapter 5](#).

The blue arcs symbolize the circuit of LVL 1 operators (there are N_1 of them), while the circuit of LVL 2 agents (there are N_2 of them) is formed by orange arcs. Recall that we have an inbound flow with intensity λ (possibly non constant) coming from the upmost transition. All the inputs and exits of the system are represented by the purple double-arrows (\rightleftarrows). A caller can leave the network before his/her call is picked up by a LVL 1 operator if not handled before the patience time μ_1 , by taking the striken-arc (\dashrightarrow). Otherwise, the conversation with the LVL 1 agent happens in a (not gravity-dependent) place with duration τ_1 (caller-dependent).

Once this first conversation is over, the token can be sent to three different lanes depending on the gravity level, either NU (Non Urgent), U (Urgent) or VU (Very Urgent). For the sake of simplicity, we have not depicted the LU (Less Urgent) lane, merged here with the U category. If the call is NU, the caller leaves the system since “filtered out by the first level” and the LVL 1 operator comes back in his availability pool – ready to pick up new calls. In case the call is U, then the caller is put on hold in an inter-level waiting room, and the LVL 1 operator is still released. The caller may get impatient and leave the platform on its own if not picked up before the time μ_2^U . If the call is VU, then the LVL 1 operator stays on the line with the caller in a pre-LVL 2 waiting place (after the first-level instruction has ended). The caller may still get impatient and hang up after the time μ_2^{VU} , which triggers the release of the LVL 1 agent.

The LVL 2 operators must pick up the VU calls in priority (if there are any), which is symbolized by the double-tip arrow ($\rightarrow\rightarrow$). The firing of the corresponding transition starts the synchronization step (or escort mechanism, or threeway conversation) involving the caller, the LVL 1 agent and the LVL 2 agent. After a conversation that lasts τ_5 , the LVL 1 operator is released and the caller enters a conversation place with duration τ_2 . Once the consultation is over, the caller leaves the platform and the LVL 2 agent spends a time τ_3 (possibly null) in the extra instruction place if necessary (see the [Section 4.4](#) of [Chapter 7](#)). Finally, LVL 2 operators come back in their availability pool.

The next key step is to select the performance criteria that we will want to compute at the end of the simulations. It is relevant to focus on a limited number of metrics that are of practical interest. In this view, we recall the three criteria highlighted by the “mission MARCUS”

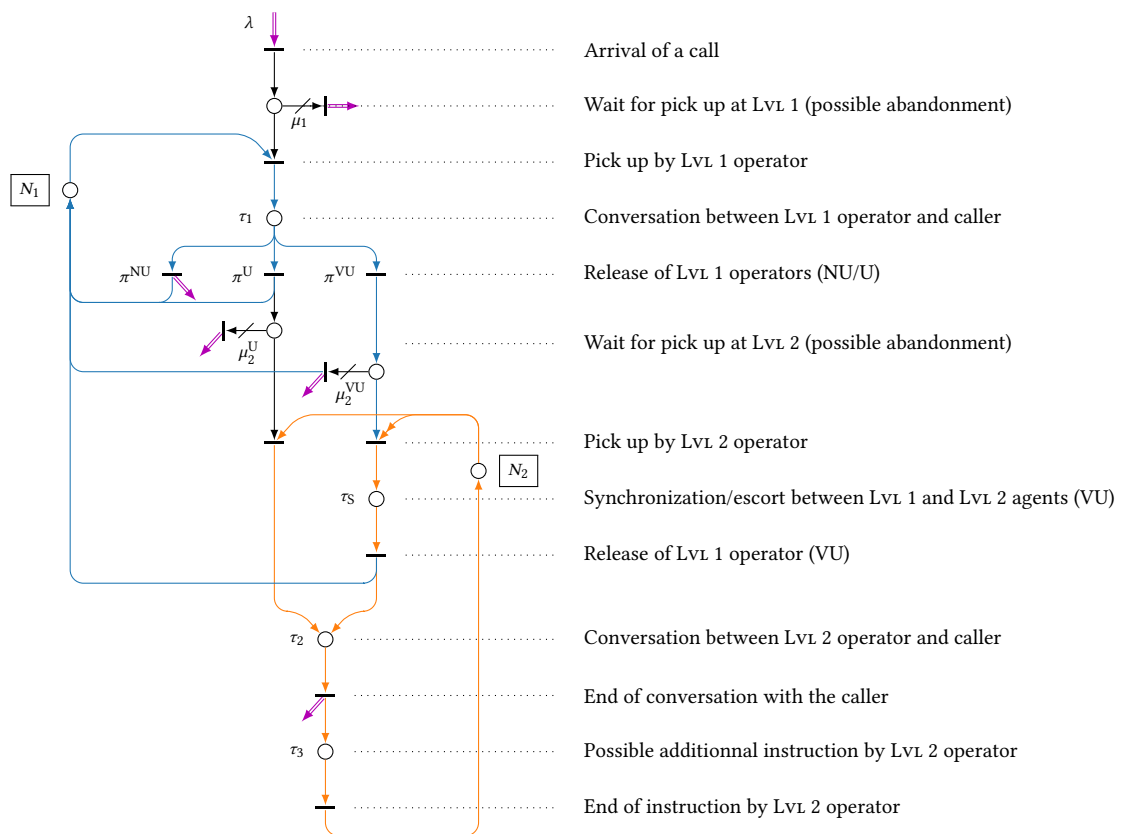


Figure 8.4: Petri net model of a PFAU-type single-job and bilevel call center

(see [MP19]), commissioned by the French *ministère de l'Intérieur* and *ministère des Solidarités et de la Santé* (respectively Home Affairs and Health Ministry):

- the **admission rate**, which is the proportion of picked up calls relatively to the number of received calls. It is the complementary of the **abandonment rate**.
- the **service level**, proportion of calls picked up before a given time threshold (among the picked up calls). The voice prompt is included in this threshold – this way, it is a “customer-oriented” criterion. A close metric is the **average waiting time** of picked up calls.
- the **occupancy factor** (Occ) of agents, given by the proportion of their time spent on the phone.

Owing to its complex topology, the bilevel layout gives us an important amount of performance criteria to compute, based on the previous list. Indeed, we can evaluate these indicators (abandonment rate, average wait, occupancy, etc) depending on the gravity of calls – since they do not follow the same path – but also at LVL 1 and LVL 2, since these are independent in some extent.

Actually, because we want to end up with a system where both levels are sufficiently staffed, it will be relevant to focus on “mixed” criteria, that combine the performance of both levels. This way, we may measure the *global* abandonment rate of calls in the network, since it is useless to have a low abandonment rate at some level but a high one at the other level. Similarly, we may compute the *cumulated* waiting time throughout the complete instruction chain, including not only the wait before the first pick up for NU calls but also the inter-level wait for U and VU calls. As far as the service levels are concerned, we choose contrariwise to only count the proportion of calls picked up quickly at LVL 1, which does correspond to an admission rate. We consider that the swiftness of pick up at LVL 2 is a bit less critical, because life-threatening situations shall already be identified at LVL 1 and transferred to the VU lane – to be handled in priority. In addition, we have seen in the Section 5 of Chapter 7 that the patience of callers (and thus their proneness not to hang up) was greatly increased after the first pick up.

Remark that most performance metrics that we can bring up to study the organization are correlated. In other words, a good performance on one criterion is often associated with a correct sizing of the agents pool, and subsequently to good performance on other criteria. This phenomenon is depicted in the Figure 8.5 which gives the correlation matrix of seventeen indicators. The color intensity indicates the strength of a correlation between two variables; blue (resp. red) hues being associated with positive (resp. negative) correlations. We directly see a “block structure” of this matrix, showing that the variables corresponding to LVL 1 performance feature high correlations. The same goes for the performance of LVL 2, but these two groups of variables are reasonably independent with each other. These results would a priori prompt us to only keep a very limited number of independent criteria, to recover the other ones based on the correlation relations – that are not necessarily linear. However, we accept to have a certain degree of redundancy to work with the performance indicators that are customarily used in operations management. In addition, some correlations are weaker when distinguished for VU calls, which justifies to look at those metrics separately from others.

To finish, we recall that the analyzes that follow are only **quantitative** measures of the organizations’ performance. They do not say anything about the possible **qualitative** aspects (quality of the instruction as perceived by the caller, work conditions of the operators, and so on).

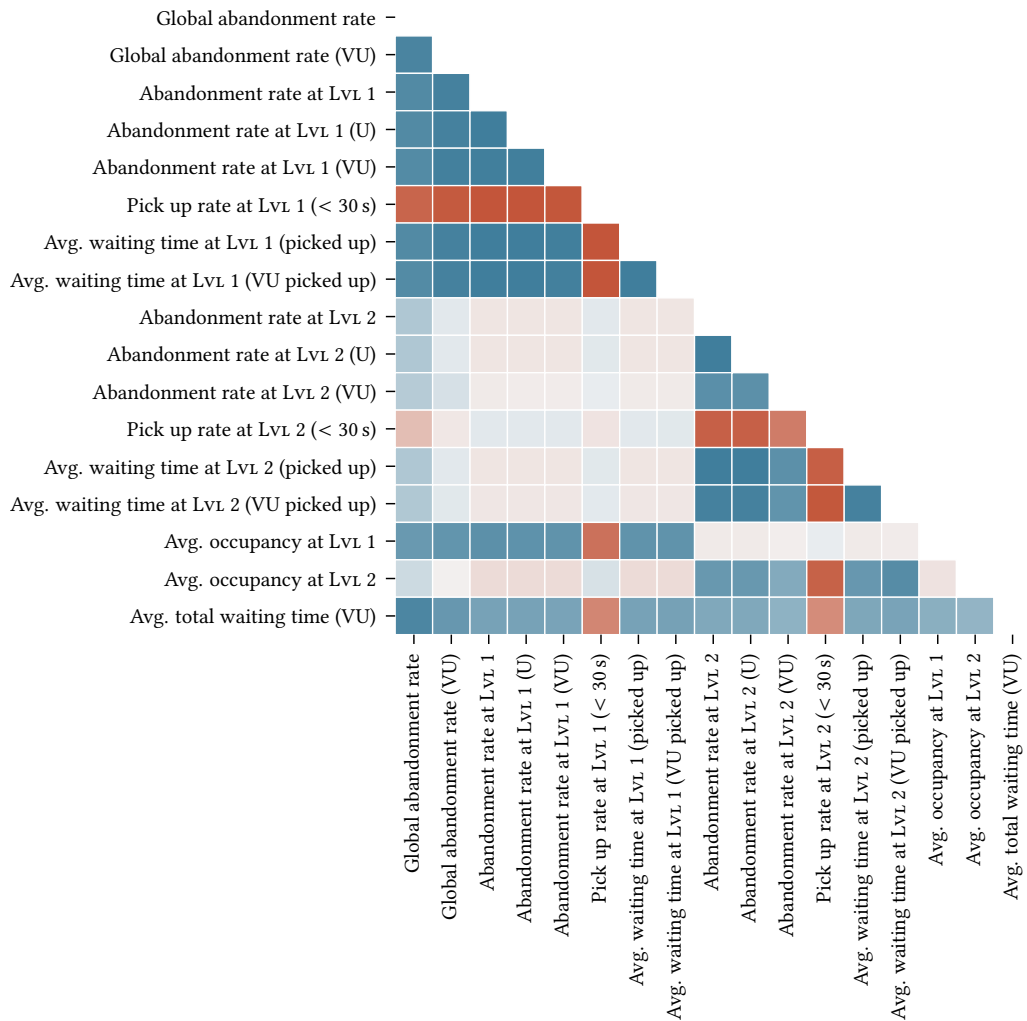


Figure 8.5: Correlation matrix of the bilevel system

2.2 PERFORMANCE OF THE SINGLE-JOB AND BILEVEL SYSTEM. — We depict in Figure 8.6 performance charts associated with six different criteria. Since there are two parameters to be tuned (the number N_1 of operators at Lvl 1 and the number N_2 of operators at Lvl 2), it is convenient to adopt a representation using heatmaps.

All these experimental results are obtained from simulation campaigns that follow the average arrivals scenario in a week day for emergency lines 18 and 112 between 9 AM and 5 PM. We naturally use the corresponding patience levels and conversation times.

Let us explain how to read the maps of Figure 8.6. We can for instance see on the first map that for a bilevel system with three agents at Lvl 1 and a single operator at Lvl 2, we experience a 33 % abandonment rate among all the received calls. This rate would fall to 5 % with $N_1 = 4$ and $N_2 = 5$. The cells outlined in black indicate the best performance with a fixed value of $N_1 + N_2$. As a result, if we have a total of seven operators, then we have to choose the combination $N_1 = 3$ and $N_2 = 4$ to minimize the global abandonment rate (equal to 11 %), as other allocations give higher values. When the criterion is already optimal (null abandonment rate or zero waiting time), we frame the whole optimal area. We have also drawn in dashes the line with equation $N_2 = rN_1$,

where r is the critical ratio found by the analytical approach and defined by the equation (5.6) – recall that beyond this ratio, the bilevel system is fluid. We refer to our Section 3.1 of Chapter 5 for the detailed computation of this ratio.

Focusing on the different abandonment rates, we observe that we reach less than 1 % of global abandonment rate as soon as we have 11 operators ($N_1 = 4$ and $N_2 = 7$). However, if we are to only focus on the abandonments of VU calls, this performance is achieved only with 9 agents ($N_1 = 4$ and $N_2 = 5$). We retrieve here the key interest of the priority routing, offering an increased protection to VU calls. As for the cumulated wait (before LVL 1 and between both levels), the effect of the priority is once again very clear. For a same staffing choice $N_1 = 3$ and $N_2 = 5$, a U call will in average wait for 100 seconds, versus an average of 30 seconds for a VU call. With two extra second-level agents, these waiting times respectively fall to 21 seconds and 11 seconds.

At this stage, it is worth mentioning that different criteria can lead to different optimal allocations of resources between the LVL 1 and the LVL 2. For instance with 9 agents, the lowest (global) abandonment rate is achieved for $N_1 = 3$ and $N_2 = 6$, while the lowest average wait of urgent calls is reached when $N_1 = 2$ and $N_2 = 7$. This phenomenon is very common in multi-criteria optimization, and may result in increasing the staffing for all criteria to be satisfied – even though each individual criterion could be satisfied with less agents but using different breakdowns. We point out that the optimal repartitions of agents between both levels for the total average wait are in good approximation close from the critical ratio found by the theory (see again Section 3.1 of Chapter 5).

It is always relevant to check the values of our indicators in the limit cases. If $N_1 = 1$, then we always observe at least 37 % of abandonment for all calls, even though N_2 is very high. This rate falls down to 31 % for VU calls, which may seem counter-intuitive since all the abandonments take place at LVL 1 and having in mind the assumption that the patience of callers at LVL 1 do not depend upon their gravity (discussed in the Section 5 of Chapter 7). The heatmaps of waiting times also show that the average wait before hanging up is the same for U and VU gravities. The previous difference is actually explained by the fact that what we observe is a mixture of calls from lines 18 and 112. However, the latter have a patience lower than the former by approximately 10 seconds, and the weight of U and VU calls are much more important on line 18 than on line 112; hence fewer VU abandonments than global abandonments. Now, let us take $N_2 = 1$ and a high value of N_1 . We observe in average 29 % of global abandonment but 38 % of VU abandonment. The reason behind this new difference is that Non Urgent calls (accounting for a significant part of the flow) are all handled, and make the global abandonment rate decrease. This shows the importance to distinguish several performance measures on U or VU calls, to avoid being “polluted” by the weight of NU calls.

All the preceding numerical evaluations must be compared with the occupancy rate of each level, explaining why these values may seem too low for some agents who know the platform well. This way, even though a configuration with $N_1 = 4$ and $N_2 = 7$ enable us to have less than 1 % of abandonments in total, this puts the LVL 2 agents under an effective occupancy of 73 %. Such a value is too high in practice and is not allowed. The maximum occupancy of each level, contrary to global abandonment or total waiting time, are not mixed criteria. This explains why in first approximation, the corresponding heatmaps are invariant along one of the two dimensions. We can actually observe a slight coupling effect. For instance on the LVL 2 occupancy map, increasing the staffing of LVL 1 from $N_1 = 1$ to $N_1 = 2$ intensifies the occupancy rate... because much more calls now arrive at LVL 2 from the LVL 1. Each occupancy factor is then not totally independent from the other level (for small values of N_2 , we similarly observe than changing N_2 can alter the occupancy of the LVL 1 due to the escorting mechanism).

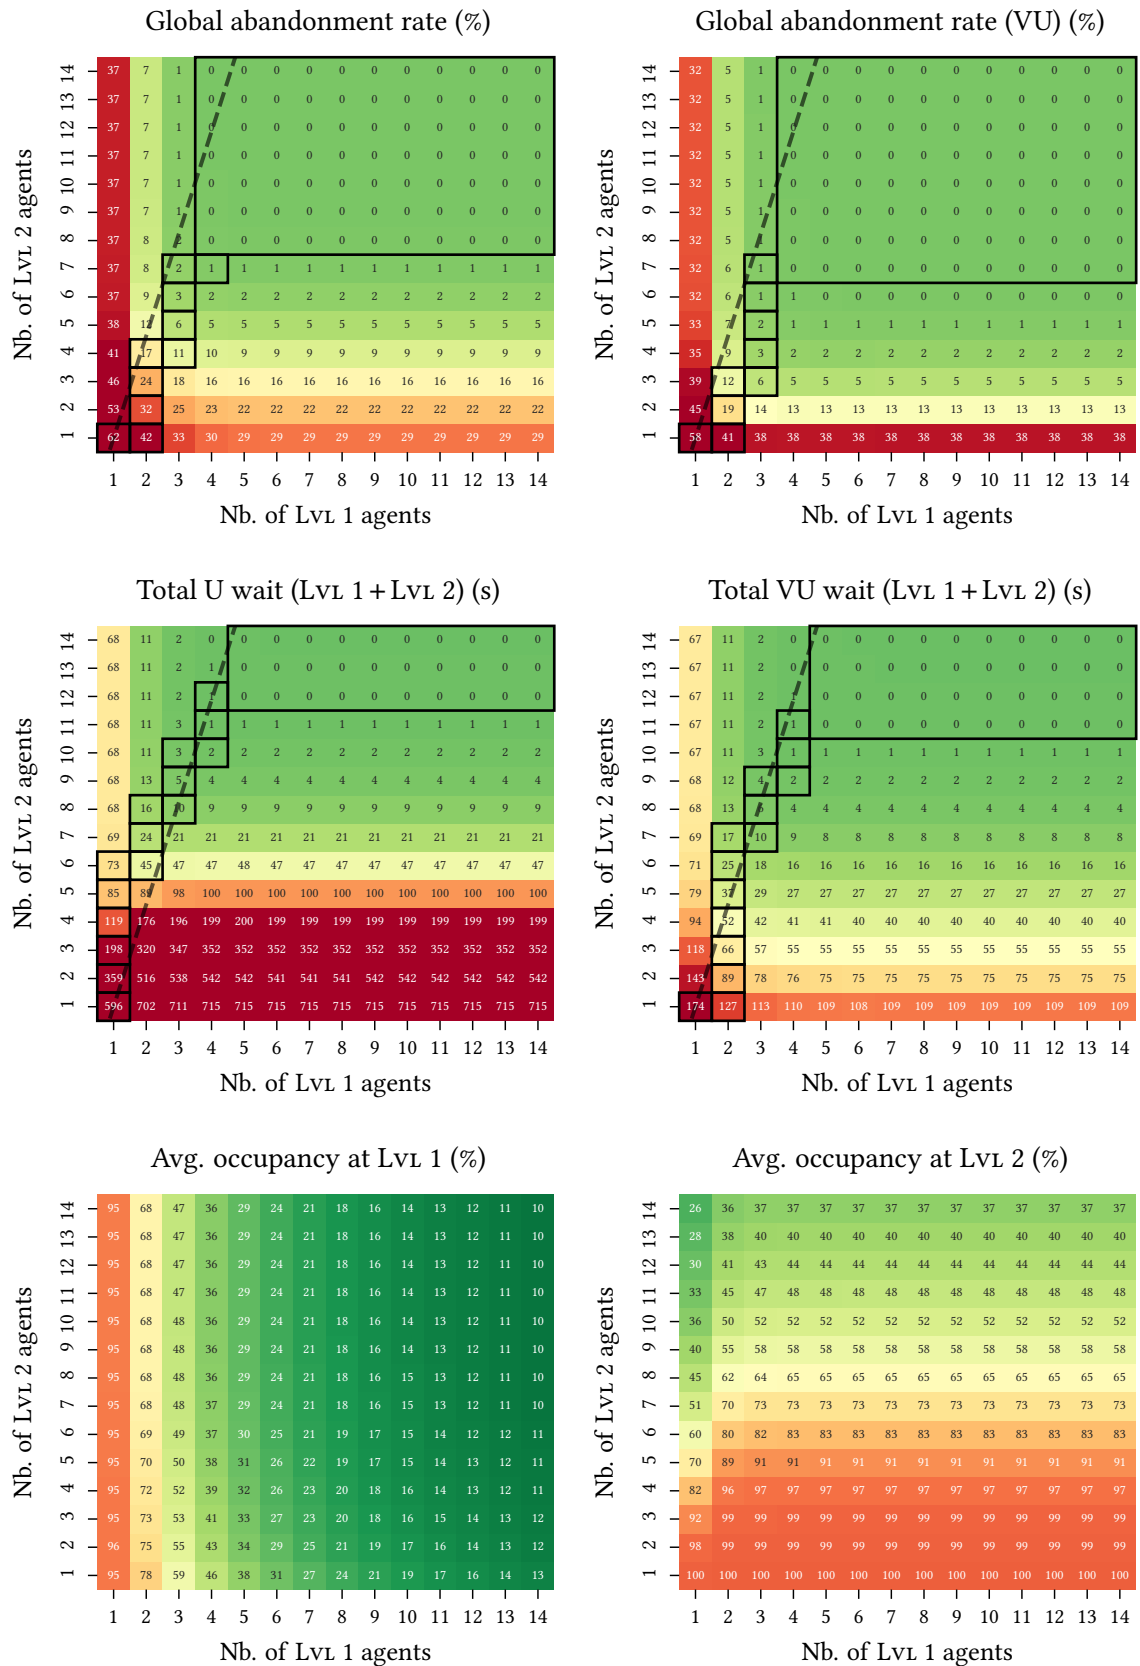


Figure 8.6: Performance maps of the single-job bilevel system

We finish this section by giving in Figure 8.7 a complementary look on the heatmaps of the Figure 8.6. We depict the values of three performance criteria for a fixed sum $N_1 + N_2$. This corresponds to an antidiagonal cross-section of the bidimensional heatmaps. In the Figure 8.7, we chose to fix the total staffing to 14 agents, which is the typical firefighters headcount of the PFAU to address calls from lines 18 and 112. Along the x -axis, we make the component N_1 vary out of this total of 14. Even though we only have measure points for integer values, we represent an interpolation based on Hermite cubic splines (such polynomials preserve monotonicity) to make the reading easier. Because our results come from an averaging operation realized on the performance criteria

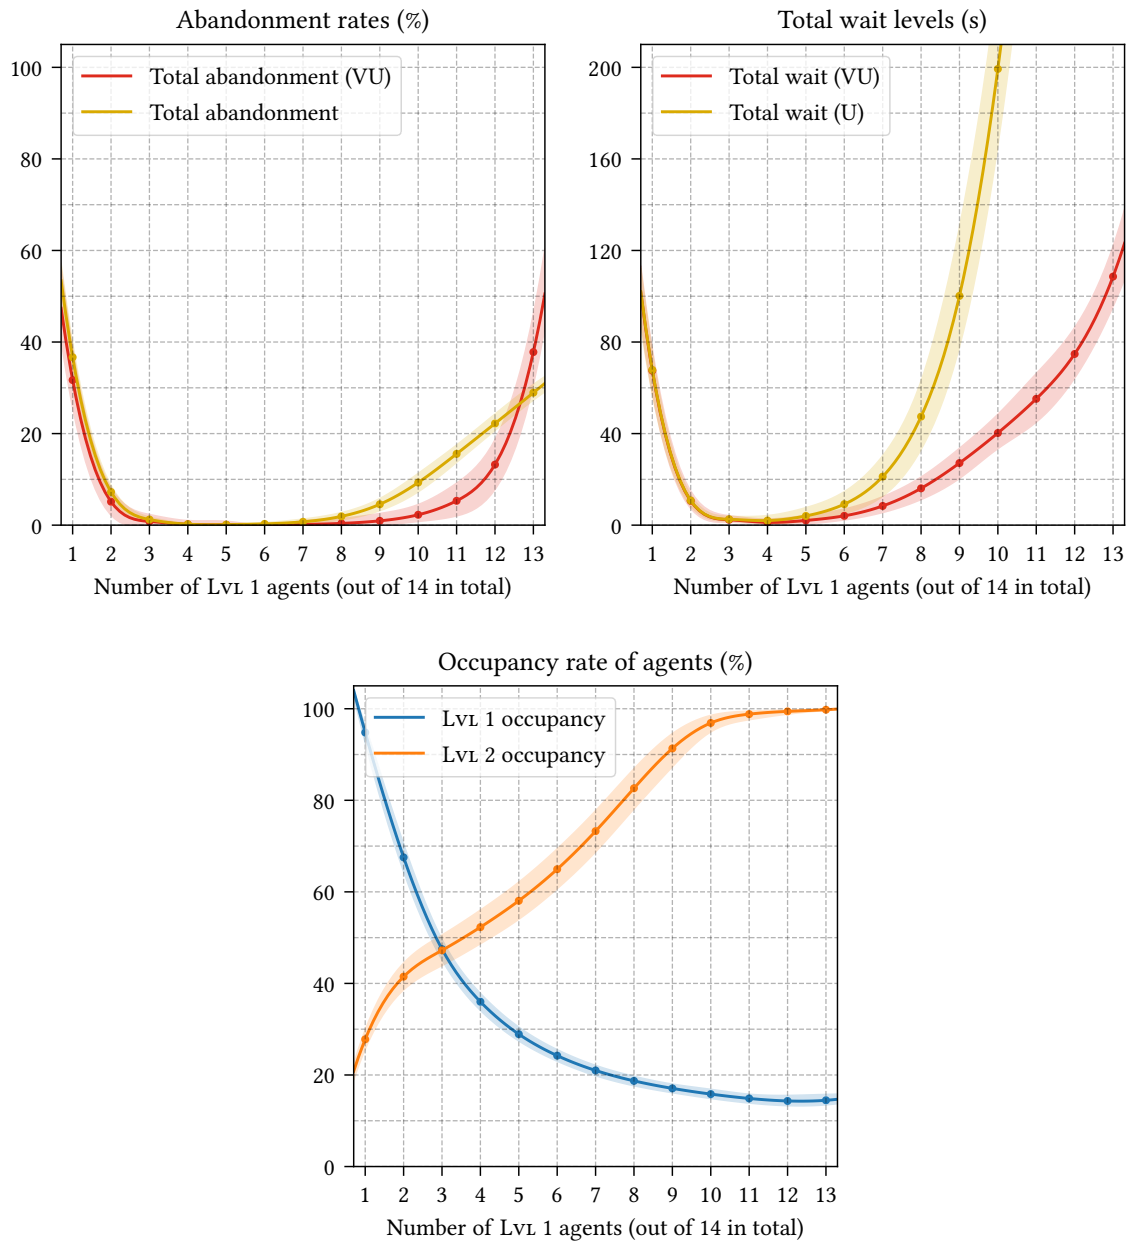


Figure 8.7: Effect of the breakdown of the staffing across the two levels

of 100 individual simulations, we also represent the inter-quartile range (half of simulations fall in this interval) to help appreciating the variability of the indicators.

The first two charts of the Figure 8.7 are about the global abandonment rate and the total mean waiting time, that is to say two mixed criteria. We thus verify that putting all the resource either at LVL 1 or at LVL 2 gives bad performance. Rather, a smart trade-off must be determined. We can observe here that the configuration $N_1 = 4$ and $N_2 = 10$ (this is the typical day-to-day staffing) leads to very good performance. We also retrieve in this figure the impact of the priority mechanism, which much less undermines the wait or abandonment of VU calls compared to the other ones, and even though the LVL 2 is a bit understaffed. Finally, we recall for information purposes the occupancy curves of both levels, which rather arise as mono-level metrics and hence feature opposite optima. Nevertheless, we remark that the configuration $N_1 = 4$ and $N_2 = 10$ offers acceptable occupancy values for the two levels of the platform.

2.3 COMPARISON OF THE BILEVEL AND SINGLE-TIER LAYOUTS FOR RESCUE CALLS. — On the model of the previous section, we briefly introduce in the Figure 8.8 a Petri net model of a single-job and single-tier emergency call center. We invite the reader to compare this model with the one of Figure 2.4 that we discussed in the Example 2.6 in Chapter 2. Here, we allow abandonments of callers before pick up but also consider the place of additional instruction after the caller left the system. The evolution in the net goes as follows. The caller tokens arrive with a throughput λ via the uppermost transition. The token waits in the first place and can only cross the pick up transition once an operator token is available in the place symbolizing the pool of agents. The integer N stands for the total number of such tokens present in this place at the beginning of the execution. If the wait of the caller token exceeds the patience level μ , the caller hangs up and the token leaves the net. Otherwise, the caller and operator tokens are consumed by the firing of the pick up transition, to produce a token in the conversation place. This token fires the next transition after the time τ is elapsed, moment at which the caller leaves the net. The operator may be held during the time τ' to finish the instruction of the case, before going back in the agent pool.

For the simulations of this single-tier system, we shall assign to the tokens a conversation time τ related to the conversation times τ_1 and τ_2 that the tokens would have had in the bilevel system. If the call gravity is NU (stopped at LVL 1 in the bilevel layout), we set $\tau := \tau_1 + \xi$, and

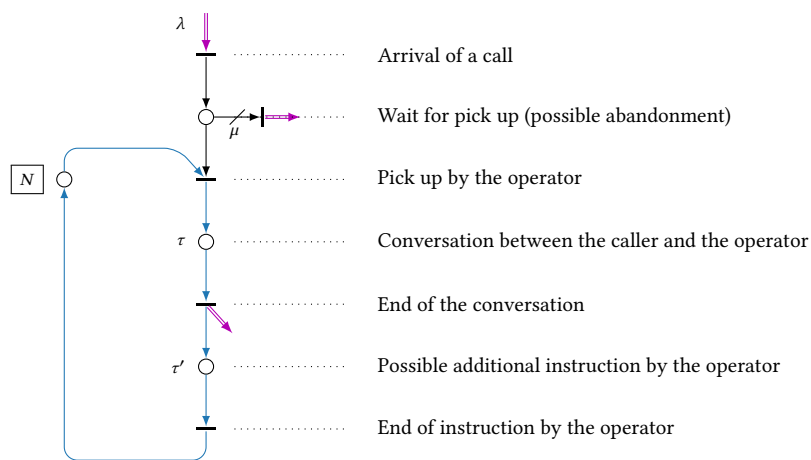


Figure 8.8: Petri net model of a single-tier and single-job call center

otherwise (the call went through LVL 1 and LVL 2) we choose $\tau := \tau_2 + \xi$. The duration ξ is a random variable between 0 and 6 seconds to model the fact that operators in the single-tier system have to cope with an increased need of adaptation from one call to another – on the contrary, conversations in the bilevel system are more homogeneous. The extra instruction time after hang up is identical in both systems, with $\tau' := \tau_3$.

We now compare the two architectures, using the previously introduced metrics. Computing the criteria of the single-tier organization is less subtle than for the bilevel system due to the reduced complexity of the system. We still retrieve strong correlations between all the indicators though (abandonment rate, mean waiting time, service levels, occupancy...).

The Figure 8.9 depicts seven performance criteria that are comparable between both layouts, as a function of the total number of agents – that is to say the number N for the single-tier system and the number $N_1 + N_2$ in the bilevel system. For the curves of the bilevel system, we have chosen to represent the values of the indicators according to an allocation of N_1 and N_2 which optimizes the following **mixed criterion**

$$\begin{aligned} C(N_1, N_2) = & 20 \times \max(\text{Abandonment_rate_VU}(N_1, N_2), 1\%) \\ & + 15 \times \max(\text{Global_abandonment_rate}(N_1, N_2), 1\%) \\ & + 1 \times \min(\text{Total_wait_U/VU}(N_1, N_2), 5\text{ s}) \end{aligned}$$

This weighted and conditional criterion aims at first minimizing the total abandonment rate of VU calls, and then for smaller values of the latter to minimize the global abandonment, and finally the total wait of U and VU calls when there are not many abandonments. The choice of this criterion is not universal and other weights and thresholds may lead to slightly different optimums. In addition we need to ensure that this criterion is decreasing relatively to N_1 and N_2 . This way, a curve associated with the bilevel system in the Figure 8.9 is *not* obtained by using the optimal breakdown of the corresponding criterion (for such representations, see the black paths of the heatmaps of the Figure 8.6). On the contrary, all the performance measures of the bilevel system are directly comparable from one chart to another. We will present at this end of the section an alternative global optimization scheme, once performance thresholds are chosen.

Looking at the two first charts on the abandonment rates (either global or restricted to VU calls) reveals that the bilevel system outperforms the single-tier system when the call centers are understaffed (and for the “iso-staffing” comparison). This is due to the increased patience of callers between the two levels and to the pick up with priority of VU calls in the bilevel layout. The understaffed area let us foresee the impact on the abandonment rates of a “crisis event” during which the incoming throughput would suddenly rise. However, under a standard flow of arrivals and with a number of agents greater than 12 (which is the case in practice), the single-tier and the bilevel systems are both excellent and guarantee abandonment rates below 1%. For such staffing levels and still for the iso-staffing comparison, the single-tier system actually gets in front of the bilevel by a tiny margin for the global abandonment rate, but not for the VU abandonment rate.

For the total waiting times of U and VU picked up calls, the single-tier architecture performs systematically better. By design, the U calls (majority of the demand) have to wait longer than VU calls when there is a lack of agents. However, even for the VU calls, the average total wait is longer in the bilevel system than in the single-tier one. This is due to the extra rigidity of the treatment chain that comes with the split of the staff into two separated pools. In particular, the callers experience before the LVL 2 a wait equivalent to the one before being picked up by a single-tier call center ran with N_2 agents (instead of N) and having to handle all the U and VU

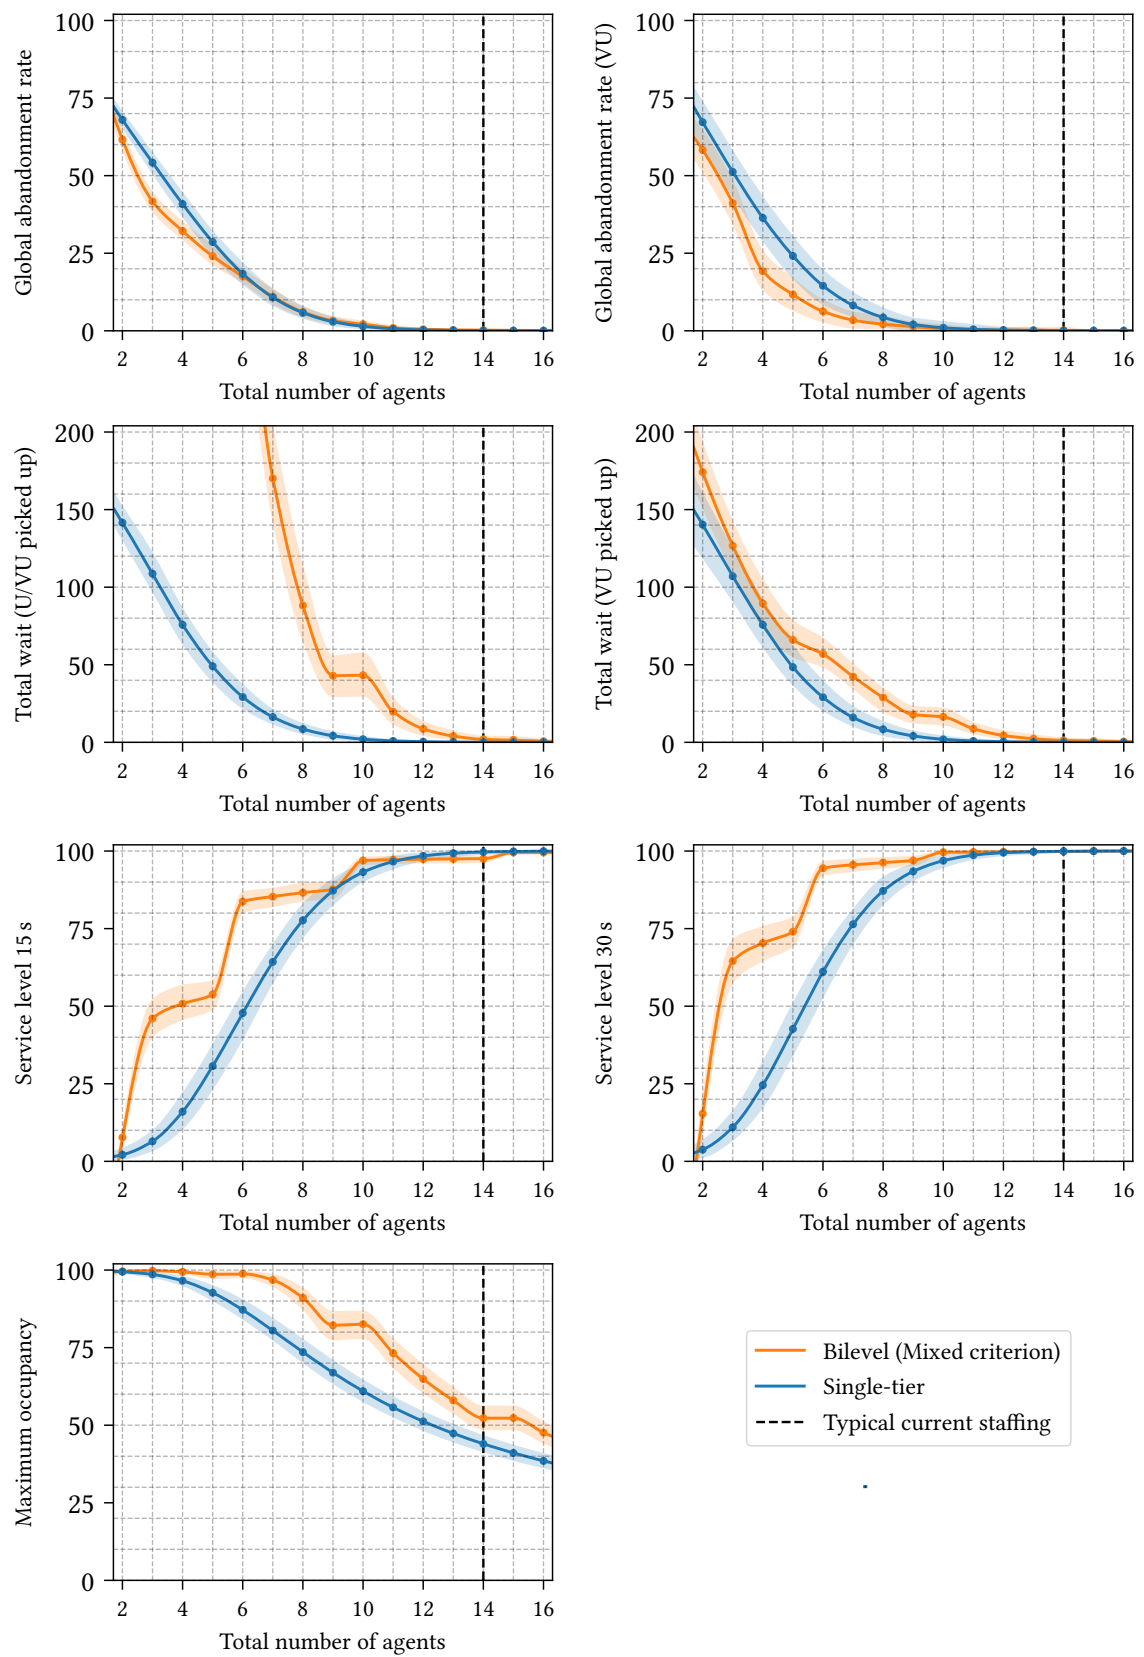


Figure 8.9: Comparison of both systems (single-tier and bilevel) on multiple criteria (for **rescue**-type single-job organizations, nominal flows, from 9 AM to 5 PM (≈ 300 calls/h))

calls. Still, we note that for more than 13 or 14 agents (current staffing of the PFAU), the total wait in the bilevel system is equivalent to the wait in the single-tier system – that is to say almost zero.

As far as the pick up times (that we can evaluate through the service levels) are concerned, the bilevel system does much better than the single-tier when the platforms are understaffed. Note that while the SL-30 (service level in less than 30 seconds) is higher in the bilevel layout regardless of the staffing, this is not always true for the SL-15. Indeed, between 9 and 12 agents, the SL-15 of the bilevel center is slightly inferior to the SL-15 of the single-tier center. This is due to our choice of mixed criterion C and our goal to first attain a low rate of VU abandonment – so that we consolidate the LVL 2 as soon as the LVL 1 is fluid. We can distinctly see *plateaus* on the SL-15 curve, they are characteristic of bivariate optimization. Each *plateau* corresponds to a constant number of LVL 1 agents, and each jump corresponds to a new LVL 1 agents being added to the pool – this highly impacts the service levels.

Finally, we depict the maximum average occupancy of the bilevel system (that is to say the largest occupancy factor between the LVL 1 pool and the LVL 2 pool) to compare it to the occupancy rate in the single-tier layout. It turns out for a same staffing (i.e., still under the iso-staffing comparison) that the bilevel system results in an increased occupancy of agents, by 5 % to 10 %. For instance with 14 agents, we observe a 45 % average occupancy in the single-tier center, versus 52 % for a well-tuned bilevel center. Note that qualitative studies may show that the latter situation is preferable for the operators, exposed to more homogeneous tasks to achieve.

We conclude this section about the cross analysis of single-tier and bilevel layouts by providing minimum staffing levels that need to be chosen in practice, in order to reach performance thresholds for each criteria. We use again the goals determined by the MARCUS report [MP19] which recommends: an admission rate of 99.95 % of calls (that is to say less than one abandonment per 2000 calls, to be compared with the average of 2400 rescue calls received by the PFAU between 9 AM and 5 PM in a typical week day), a service level in 15 seconds greater than 90 % (from the caller point of view, hence a pick up in less than 6 seconds if the voice prompts lasts for 9 seconds), a service level in 30 seconds greater than 99 %, and finally an occupancy rate not exceeding 60 % at both levels.

Based on our simulations, we compute for each performance criterion the minimum number of staff to reach the corresponding threshold. In the case of the single-tier system, it suffices to minimize N (total number of agents) under the constraint that performance thresholds are attained.

Criterion	Target	Single-tier minimum staff	Bilevel minimum staff
Admission rate	$\geq 99.95 \%$	14	15 (5+10)
Service level – 15 s	$\geq 90 \%$	10	5 (4+1)
Service level – 30 s	$\geq 99 \%$	12	6 (5+1)
Maximum occupancy	$\leq 60 \%$	11	12 (3+9)
All criteria		14	15 (5+10)

Table 8.10: Single-job staffing table, typical rescue arrivals from 9 AM to 5 PM (≈ 300 calls/h)

For the bilevel system, we follow a lexicographic minimization scheme by first minimizing the total number $N_1 + N_2$ of agents. Then, if several combinations (N_1, N_2) with identical sum satisfy the thresholds, we choose the one with the minimum N_2 component. Indeed, we deem that LVL 2 agents are more qualified resource, hence more expensive. To finish, we compute the minimum staffing level which allows to reach not only each individual performance goal, but to reach all of them together.

With the above choices of thresholds, these optimization computations give us the Table 8.10. Recall that this table is associated with simulations dealing only with rescue calls (line 18, and a fraction of line 112) according to standard and average weekday throughputs from 9 AM to 5 PM.

Consistently with the previous remarks, we observe by reading the Table 8.10 that the bilevel organization requires more agents than its single-tier counterpart to reach the same performance in terms of admission rate (15 versus 14) and occupancy (12 versus 11). As it was already mentioned, the service levels are rather mono-level criteria (governed by the number of LVL 1 operators), hence the optimal staffing in the bilevel layout only offers a partial view of the global performance of the corresponding call center. We will keep in mind that for these specific choices of thresholds, the 99.5 % target of admission rate is the most ambitious and agent-consuming goal. In fact, its importance is so predominant that optimizing all criteria is here equivalent to maximizing this single one. In order to reach all the objectives under standard and rescue arrival rates, we will thus recommend that the single-tier call center uses 14 operators, while the bilevel system should under the same conditions opt for 5 agents at LVL 1 and 10 agents at LVL 2.

2.4 THE CASE OF POLICE CALLS AND THE KEY ROLE PLAYED BY THE LOADS RATIO. – We have just observed that in the (idealized) “rescue single-job” setup – where firefighters have to pick up calls from lines 18 and 112 – that the bilevel architecture was equivalent or better than the single-tier one relatively to global and VU abandonments, and for an iso-staffing comparison. In this new section, we adapt the previous case study on “police single-job” systems. We consider idealized call centers gathering the four departments of the Paris area, whether they adopt the single-tier or bilevel layout (contrary to the real-life PFAU call center whose second level is in fact departmentalized). We are going to see that this new study unveils major differences between the two jobs, that we can mainly explain by asymmetrical loads.

Let us formally define the last notion. The **load** (sometimes also referred to as the offered load) associated with a pool of agents is the product of the throughput of calls that they have to pick up by the average handling time of these requests. The result of this operation is dimensionless but is customarily expressed in *Erlangs* (symbol E). The loads are characteristic and dimensioning parameters for call centers, reason why these terms are found in analytical computations of Petri nets and Erlang formulas as well (see the Sections 3.1 and 3.3 to come).

The loads give a (partial) quantitative indication of the number of agents required for each role. In the case of infinite patiences and in affine stationary regimes, they correspond to the absolute minimum of operators needed to reach a fluid behaviour. Actually, to cope with the variability of the arrivals and to guarantee satisfying occupancy levels, it is often appropriate to take at least twice as much agents than the load. However, this remains a relevant quantity to compare the size of the different staffs in the network.

In the current bilevel PFAU system, it is natural to compute separately the loads of each level. At LVL 1, we will distinguish the loads per emergency line, in order to represent the workloads respectively induced by line 17, line 18 and line 112. At LVL 2, we will in contrast distinguish the BSPP load from the DSPAP load. We may divide further these loads in order to distinguish the gravity levels or the departments of origin.

LVL 1	17	112	18
	2.11 (42/15/26/17)	0.50	1.02
LVL 2	DSPAP		BSPP
	2.93 (51/12/22/14)		5.60

Table 8.11: Repartition of loads (standard PFAU arrivals from 1 PM to 5 PM)

The Table 8.11 gives the loads per level and dialed line or orientation at LVL 2, under standard arrivals (that is to say average throughputs and handling times in weekdays of the year 2019) between 1 PM and 5 PM. In the cells associated with line 17 or DSPAP lane, the figures in brackets refer to the breakdown (in percentages) of the load across the four departments, in the order 75/92/93/94.

A very important fact must be noted. The load of calls from line 17 at LVL 1 is almost twice the one of calls from line 18 (2.11 E versus 1.02 E), but this goes the other way around at LVL 2 when comparing DSPAP and BSPP loads (2.93 E versus 5.60 E). This “**morphological inversion**” have several explanations. First, as seen in the Section 2 of Chapter 7, there are a bit more calls on line 17 than on line 18 (in average 206 calls/h versus 147 calls/h on this particular time slot). But more importantly, the fraction of calls not transferred to the LVL 2 is much higher for police calls (65 % of NU gravity) than rescue calls (32 % of NU gravity). Furthermore, as seen in the Section 4 of Chapter 7 and more specifically in the Figure 7.11, the calls from line 17 picked up by police agents at LVL 1 are handled with an average time longer than for calls from line 18 picked up by firefighters. This is all the more true for NU calls, with a higher median time than the rest of calls from line 17, contrary to calls from line 18 whose NU part is handled quicker than the rest of calls. Note that we could not compare the load of the extra instruction times (see Section 4.4 of Chapter 7) because we could not measure these times for the police. It could be that such protocols are longer for DSPAP than for BSPP.

We can summarize this asymmetry by taking the **load ratio** of LVL 2 over LVL 1. This yields a ratio of about 4 for rescue calls and close to 1 for police calls (by allocating 112 loads on rescue calls, as it was the case before the multi-job PFAU came into service). This strong imbalance is absolutely key to explain the performance of the bilevel system. Indeed, the architecture in two different levels behaves all the better as the LVL 1 is quick and fluid relatively to the LVL 2 (hence associated with a high ratio). Remark also that we only benefit from the patience increase of the bilevel system (see Section 5 of Chapter 7) after having passed the LVL 1. But above all, the priority mechanism for Very Urgent calls between the two levels deeply relies on the fact that all calls are picked up at LVL 1. Hence, we shall expect much better performance of the bilevel layout when the loads ratio (LVL 2)/(LVL 1) is high.

In fact, we have seen that the throughput of calls tend to significantly increase on line 17 during the evening, and at the same time the demand on line 18 decreases. It is relevant to recompute the loads on the time-window from 7 PM to 11 PM, presented in the Table 8.12.

In the evening, we verify that the load at LVL 1 for calls from line 17 increases by 14 % relatively to the load in the afternoon, but the LVL 2 load for DSPAP rises even more, by a 65 % factor. This is a bit explained by the fact that there are more calls, but mainly because the proportion of U and

LVL 1	17	112	18
	2.40 (43/15/25/17)	0.43	0.93
LVL 2	DSPAP		BSPP
	4.83 (50/13/23/14)		5.62

Table 8.12: Repartition of loads (standard PFAU arrivals from 7 PM to 11 PM)

VU calls goes from 35 % to 54 %. However, even with this different load structure, the loads ratio of the police lane “only” reaches 2, versus approximately 4 for the rescue lane. This shows that even on this time-slot with more demand, the bilevel system fits more with the morphology of the rescue job than the police one.

Having introduced the concept and typical values of the loads, we now come back to the single-tier / bilevel comparison, on the case of police calls. To this purpose, we simulate two organizations that are fictitious in the sense that we aggregate the four departments of Paris area. Recall once again that the current police part of the PFAU (contrary to the rescue part) has a single LVL 1 for police calls but four separated LVL 2 for DSPAP (semi-departmentalized system). To keep things simple in this section, we forget the different lanes per departments and we study single-tier and bilevel centers which have to pick up *all* the calls from line 17 of the whole Paris area, as well as the small part of 112 calls that are in fact police matters. These two systems that we deal with are topologically equivalent to those introduced for rescue calls, depicted in the Figures 8.4 and 8.8 and studied in the Section 2.3. The only difference will be that we consider different volumes and handling times (police parameters between 7 PM and 11 PM, and no longer rescue parameters between 9 AM and 5 PM), which precisely modifies the loads ratio.

The Figure 8.13 gives the global abandonment rate for police calls and the specific abandonment rate of VU police calls of our two fictitious systems, as a function of the total number of operators (the allocation of agents between LVL 1 and LVL 2 for the bilevel center builds on the same mixed criterion as before). We can here observe that the aggregated single-tier system always outperforms the bilevel layout under the iso-staffing view. Yet, we would have expected that VU calls are more protected by the bilevel system, thanks to the priority mechanism, and knowing that minimizing VU abandonments is the predominant part of our target (mixed) criterion.

It is very important to note as of now that this superiority of the single-tier system for the iso-staffing comparison vanishes for a sufficient number of staff – greater than 14 police agents, which is ensured by the current staff of the PFAU. Beyond this point, both the single-tier and the bilevel centers yield extremely few abandonments. However, looking at the understaffed regimes (here between 9 and 13 agents) remains of interest, since it lets us foresee the degraded phases that could arise if an important peak of calls arrived (see in particular the application in the Section 5.1 to come). More importantly, the fact that the single-tier system appears to be better than the bilevel layout regarding global abandonments and VU abandonments seems to contradict the result we did read in the Figure 8.9. Recall that the latter was about rescue calls (lines 18 and 112) between 9 AM and 5 PM, and that the bilevel architecture for the iso-staffing comparison systematically gave less abandonments than the single-tier.

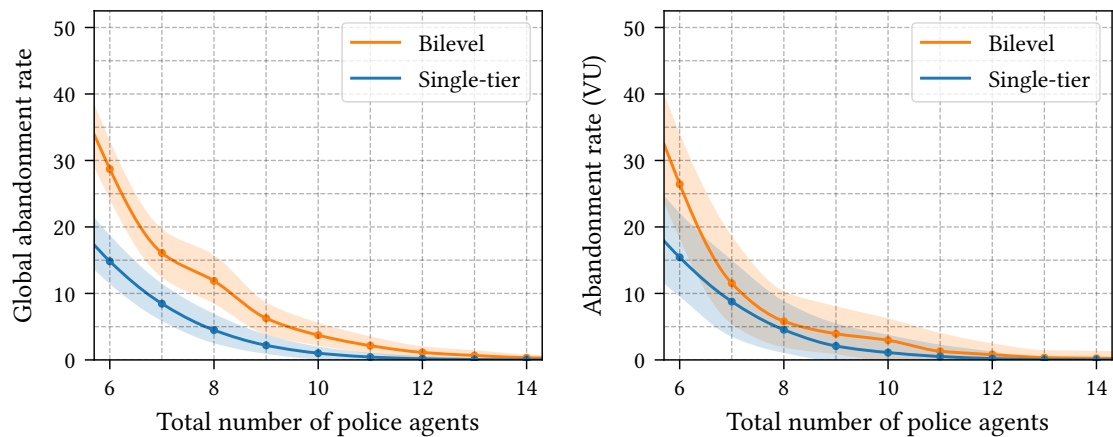


Figure 8.13: Comparison of both systems (single-tier and bilevel) on abandonment rates (for single-job and aggregated **police** organizations, from 7 PM to 11 PM, i.e., ≈ 230 calls/h)

To explain this seemingly paradox and evaluate the interest of the (single-job) bilevel system, it is necessary to recall the involved loads. Based on the previous tables, we obtain that the loads of the rescue part of the PFAU during the day (base scenario of the Figure 8.9) were 1.02 E for the LVL 1–18, 0.50 E for the LVL 1–112 and 5.60 E for LVL 2–BSPP, hence a ratio of $R \approx 4$. Conversely for the police part of the PFAU during the evening (base scenario of the Figure 8.13) the loads were 2.40 E for LVL 1–17 versus 4.83 E at LVL 2–DSPAP, hence a ratio of $R \approx 2$.

The Figure 8.14 is derived from the two first charts of the Figure 8.9, where we simulate single-job organizations picking up calls from lines 18 and 112 between 1 PM and 5 PM. For the sake of readability, we no longer represent two performance curves (one the single-tier and one for the bilevel), but rather a single one for the *difference* Bilevel – SingleTier (abb. BL – ST). This way, a +10 % value (resp. –10 %) means that the bilevel architecture achieves 10 more points (resp. less points) of abandonment than its single-tier counterpart. We will keep in mind that a curve on the “positive” side indicates a better performance for the single-tier layout, while a curve with “negative” values means that the bilevel performs better.

The blue curves in the Figure 8.14 correspond to the real loads ratio (still LVL 2 over LVL 1) of $R \approx 4$ for rescue calls. They are drawn from simulations that are identical to those used in the Figure 8.9, hence consistent with the day-to-day reality of this flow. However, the orange and green curves are fictitious. Indeed, we have purposefully reduced the conversation times at the LVL 2 and increased those of LVL 1, in order to obtain a loads ratio divided by two ($R \approx 2$, in orange) or three ($R \approx 1.3$, in green) with respect to the reality. The iso-staffing comparison of the single-tier and bilevel systems then shows that even though the bilevel system was always better than the single-tier for abandonments when $R \approx 4$ (the blue curves are negative), we have the opposite trend when R diminishes. Indeed, for global abandonments, the orange and green curves become positive. For abandonments of VU calls, when $R \approx 2$, the single-tier and the bilevel centers are approximately equivalent, with a small advantage to the single-tier system for small numbers of staff (we retrieve here figures similar to the Figure 8.13). When $R \approx 1.3$, the single-tier center always outperforms the bilevel.

Although the case of the police is not *exactly* comparable with the case of rescue (the loads ratio is not the only parameter that governs the global performance of the system, we may also consider the patiences, some second-order effects of the proportions of gravities...), the analysis of the Figure 8.14 very clearly demonstrates the systematic **advantage of the bilevel system** (with

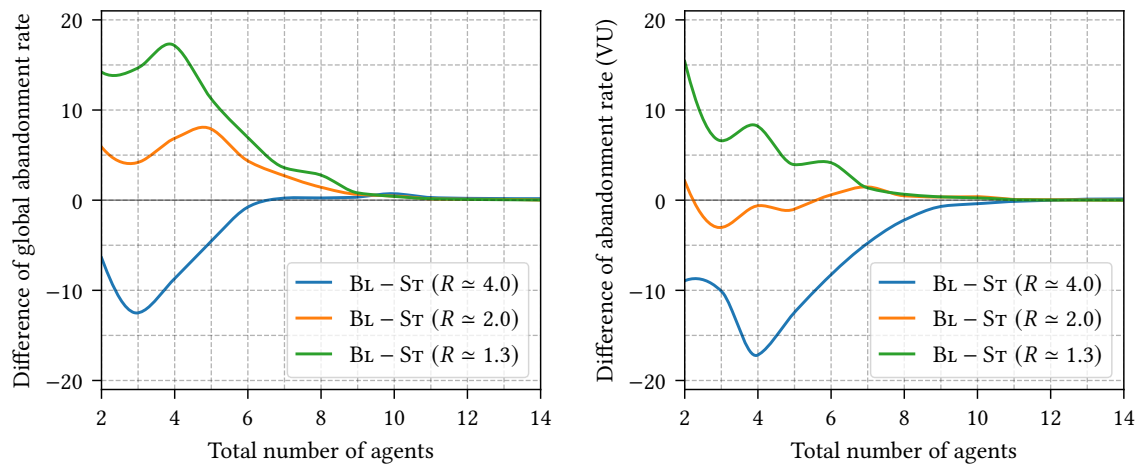


Figure 8.14: Influence of the (LVL 2)/(LVL 1) loads ratio R , for **rescue** calls as in Figure 8.9

respect to the single-tier and under the iso-staffing comparison) **when the loads ratio LVL 2 over LVL 1 is high** (at least 3, or 4). In this configuration, we can indeed afford a “small LVL 1” and a “large LVL 2”, and we benefit from the bilevel characteristics (more uniform tasks, priority of VU calls, bigger inter-level patience...). Under a lower ratio (close to 2), these advantages do not compensate the performance loss caused by the division of the staff in two pools of the same size compared to a single-tier layout with a big pool of agents. This holds even for the VU calls.

It is extremely important to insist on the fact that this analysis and its conclusion follow from our choice, recalled in Section 1.3, to compare single-tier and bilevel organizations under the **iso-staffing** look. A much different perception of the bilevel performance is obtained by comparing the systems under an “iso-expert” view. We discuss this aspect in the Section 2.5 that follows.

To sum it up, we can conclude that unlike the rescue calls, the bilevel layout is for the current police organization less interesting than the single-tier one if we compare the total number of agents. This is due to a quite different morphology: longer handling times at LVL 1, shorter at LVL 2, more NU calls... In particular, the advantages of the bilevel layout are more visible when the calls are filtered quickly by the LVL 1, and this condition is not met if too many Non Urgent calls slow down this first level. On the other hand, in order to benefit from the qualitative strengths of the bilevel system without degrading the quantitative performance, modifications of the organizations could be implemented. For instance, a recommendation could be that the NU calls at LVL 1 should be eliminated quicker, and shorter advice to callers should be given. To still meet the demand of the population, these calls should be quickly transferred to another dedicated platform, no more about emergency calls. This would result in increasing a lot the ratio R , which as shown before is a key factor contributing to the efficiency of the bilevel system; this would therefore allow to downsize a bit the LVL 1 to put more resource at the LVL 2.

2.5 DISCUSSION ON THE ISO-STAFFING COMPARISON. — As mentioned in the Section 1.3, we have made the choice in this chapter to make **iso-staffing** comparisons of single-tier and bilevel systems. In other words, we have always studied a single-tier with N agents on one hand and a bilevel system with $N_1 + N_2 = N$ on the other hand. This appears as a right choice from the mathematical point of view, in particular if one wants to compare the evolution equations at stake and the associated dynamics. From an operations point of view however, this is very arguable. Indeed, iso-staffing does not at all mean **iso-skill**. In the bilevel system, the LVL 2 agents (who as a

reminder only handle U and VU calls) are highly trained experts of emergency calls, they know how to handle complex situations and can dispatch a whole range of response means. On the other hand, LVL 1 operators have the mission to orientate and prioritize the calls, which involves performing more standardized tasks and requires a less advanced training. A single-tier layout cannot rely on such operators and must be staffed only with expert agents, equivalent to the LVL 2 resources of the bilevel system.

We can thus observe how much the staffing computation of the Table 8.10 is eloquent for the bilevel system. To cope with the rescue demand, the latter requires 5 agents at LVL 1 and 10 agents at LVL 2. This surely is one extra agent with respect to the 14 people needed by the single-tier layout, but the change of organization actually saves four expert agents! In order to highlight this phenomenon, we could find reasonable to compare single-tier and bilevel centers under an “expert iso-staffing” look; that is to say N agents for the single-tier and N_2 agents at LVL 2 for the bilevel. This representation has the drawback of hiding the number of staff of the first level and its role. It also raises the question of the choice of the staffing level at LVL 1 in the comparison. Indeed, it would seem wrong to compare a single-tier system with $N = 6$ agents, with a bilevel system with $N_1 = 10$ first-level agents and $N_2 = 6$ second-level operators.

Still, to illustrate the influence of this methodological choice in our comparisons, we depict both types of curves in the Figure 8.15. The shown criteria are the global and VU abandonment rates, and we consider either rescue or police organizations (once again aggregated in the case of the police). The figures associated with the rescue job (resp. the police job) are exactly those previously shown in the Figure 8.9 (resp. Figure 8.13), so that all the charts of the “left column” were already shown before. These charts follow the “iso-staffing” convention, for which the x -axis corresponds to the total number of agents in the call center. However, the charts of the “right column” compare the systems under the “iso-skill” convention. Hence, the x -axis gives the total number N of agents for the single-tier curve but only the number N_2 of expert agents for the bilevel curve. Visually, this operation looks like a shift to the left for the curves associated with the bilevel system (with a magnitude given by the number of staff at LVL 1). For instance, we can check on the two first charts (global abandonment rate for rescue calls) that the measure point of coordinates (10 agents, 3 %) has moved to the position (6 experts, 3 %). This is because this point was obtained for the combination of four LVL 1 operators and six LVL 2 operators. From one column to the other, the performance curves of the single-tier system do not change.

In particular, we can appreciate in the Figure 8.15 how much influence the drawing convention has on the performance comparison of our two systems. When reasoning in terms of identical number of expert agents (iso-skills), the bilevel layout appears to be very much superior to the single-tier layout. This holds for the rescue calls but also for the police. In a way, the fact that the bilevel rescue organization outperformed the single-tier one under the iso-staffing comparison was almost a surprise due to the split of the agents pool into two smaller groups. Managers of emergency call centers may on the contrary find very acceptable to increase the total staffing by a couple of agents (but less expert ones) to benefit from the advantages of the bilevel organization.

To wrap this discussion up, we have to acknowledge that the iso-staffing comparison is rather unfair to the bilevel system because it does not differentiate LVL 1 and LVL 2 operators, but on the other hand the iso-skill comparison is a bit too generous since it omits the LVL 1 operators. A fairer approach would probably be to compare the performance of a single-tier system with N agents with the one of a bilevel under a “hybrid” number $\alpha N_1 + N_2$ of staff, where $0 < \alpha < 1$. This puts the difficulty on determining a relevant value of α , and this requires the expertise of managers and operators of the call center. On our end, we may only indicate that any “fair comparison” lies somewhere between the two limit cases illustrated by the two columns of the Figure 8.15.

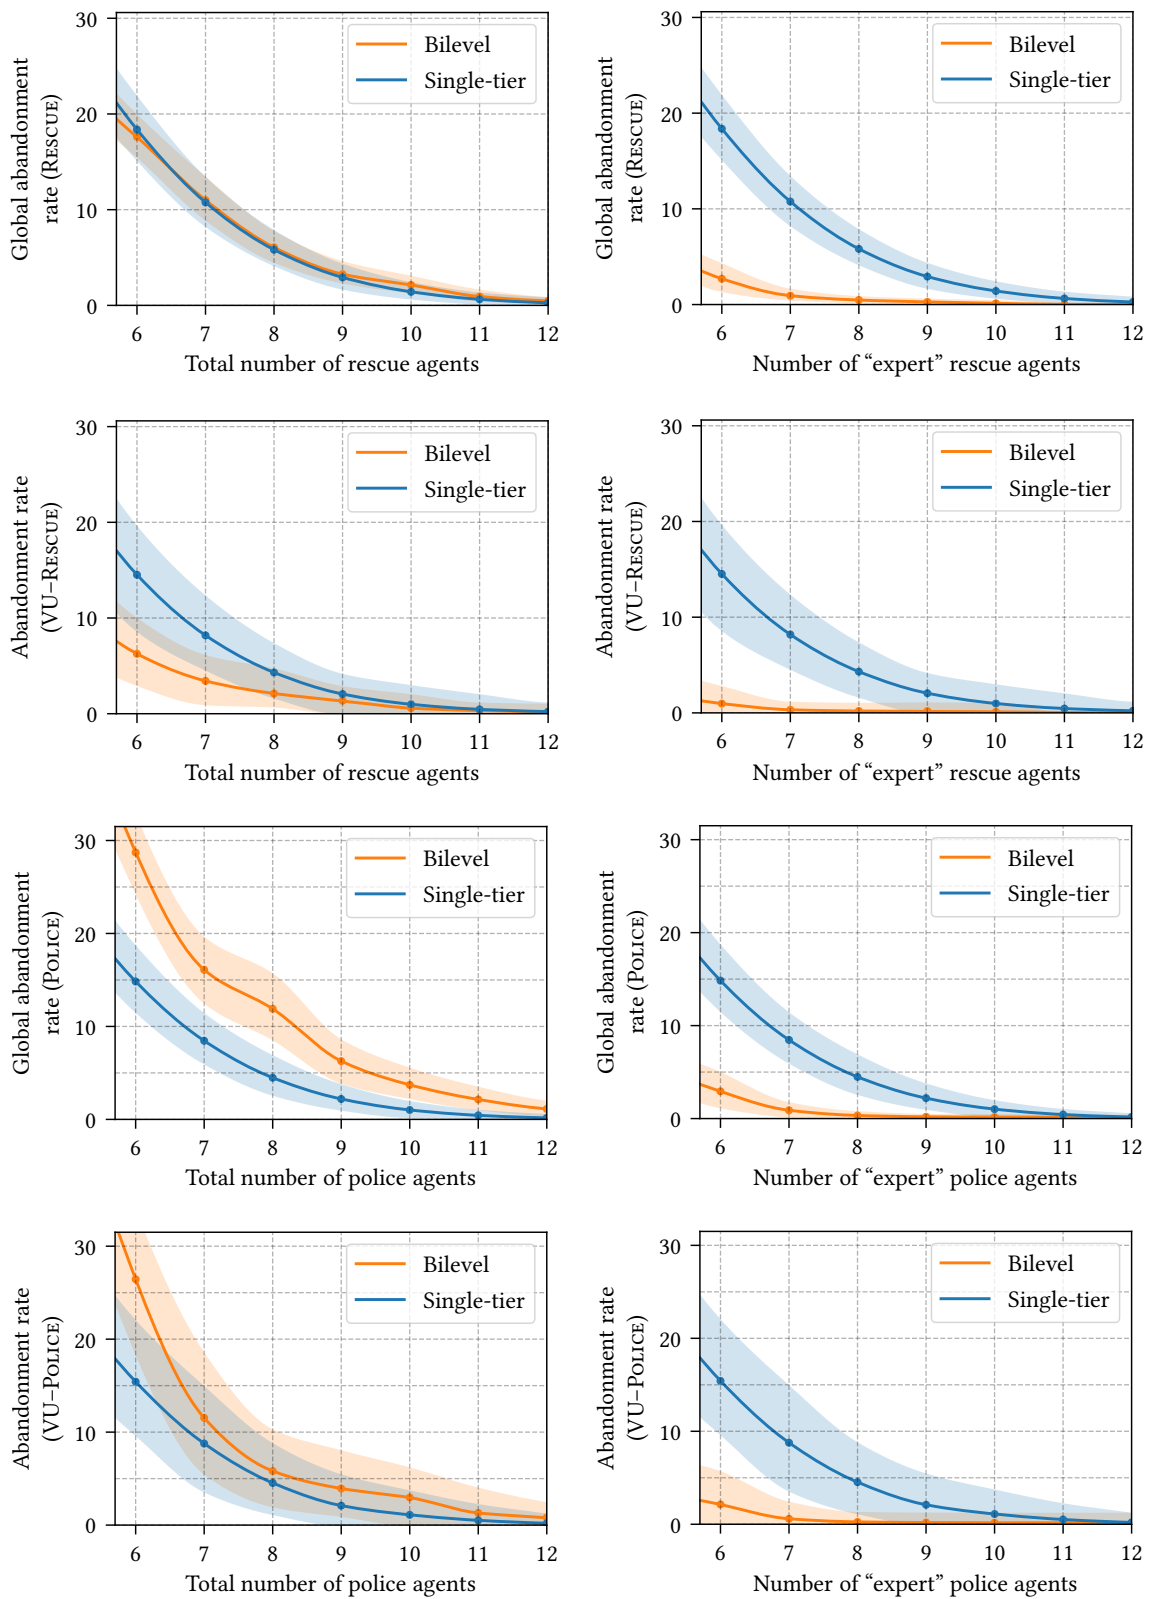


Figure 8.15: Influence of taking into account LVL 1 agents for bilevel/single-tier comparison (**rescue** and **aggregated police** organizations, from Figure 8.9 and 8.13)

3 Validation of the analytical results by simulations

In this section, we compute using our simulation approach different performance estimates that can also be predicted by theoretical techniques. In particular, we reuse several formulas introduced in the Section 2 of the Chapter 5. We shall see that the results of our two complementary methods (theory and simulation) do agree, which is a way of “cross-validating” them.

We also devote a specific subsection to the so-called Erlang calculus, which designates a couple of tools and formulas often used in the performance evaluation of call centers (see again our general [Introduction](#)).

3.1 THROUGHPUT ESTIMATES OF PETRI NETS USED FOR THE STAFFING PROBLEM. – In what follows, we build on the analytical results obtained in the Section 2 of the Chapter 5. For the single-tier system, whose Petri net is given in Figure 8.8 and assuming that callers have an infinite patience ($\mu = \infty$), we can show that the throughput ρ^{ST} of picked up calls by the N operators of the organization is given by

$$\rho^{\text{ST}} = \min\left(\lambda, \frac{N}{\tau + \tau'}\right).$$

This simple equation tells us that if the headcount N is large enough, only the arrivals flow does limit the pick up rate and hence $\rho^{\text{ST}} = \lambda$. Otherwise if N is too small, the pick up throughput is equal to the “handling speed” of agents, that is to say their number divided by their average total instruction time, as a result $\rho^{\text{ST}} = N/(\tau + \tau')$. For all calls to be picked up, the previous equation imposes to choose a number of agents N larger than $N_{\text{ST}}^* := \lambda(\tau + \tau')$.

We have depicted in the Figure 8.16 the comparison of the experimental throughput (computed with our simulations) with the throughput predicted by the theory. The simulated Petri net is the single-tier system with the real arrivals and conversation times derived from the data analysis of the PFAU in the Chapter 7 – assuming that only calls from lines 18 and 112 are received to correspond to a single-job treatment chain. For each value of N between 1 and 15, we have realized 100 simulations on the 1 PM – 5 PM time window with the average scenarios. A first batch of simulations is performed assuming that the patience of the callers is infinite (left-hand chart), and another batch uses the real patience levels from our data analysis (right-hand chart).

The solid line indicates the mean value of pick up throughput as measured over the 100 simulations, while the light-shaded area around gives the 90 % confidence interval (the five best and five worst simulations being ruled out). The black profile corresponds to the theoretical throughput ρ^{ST} ; the horizontal asymptote has y -value λ (the average intensity of arrivals over the whole time window) and the breakpoint occurs at the x -value $N = N_{\text{ST}}^*$. We do observe that in the infinite patience case, the simulation results perfectly agree with the analytical predictions, for all values of N . However, as we may have expected, the theoretical profile overestimates the actual throughput (measured in simulation) when the real patience is taken into account. The maximal difference is observed close from the breakpoint, when $N = N_{\text{ST}}^*$. This can be explained as follows. Recall that our theoretical model establishes that taking $N > N_{\text{ST}}^*$ is sufficient to pick up all the calls in a stationary regime and under the infinite patience assumption. In realistic scenarios though, the fact that some callers hang up themselves and the variability of the arrivals are prone to cause more abandonments, and hence a pick up rate far from the limit λ . It thus

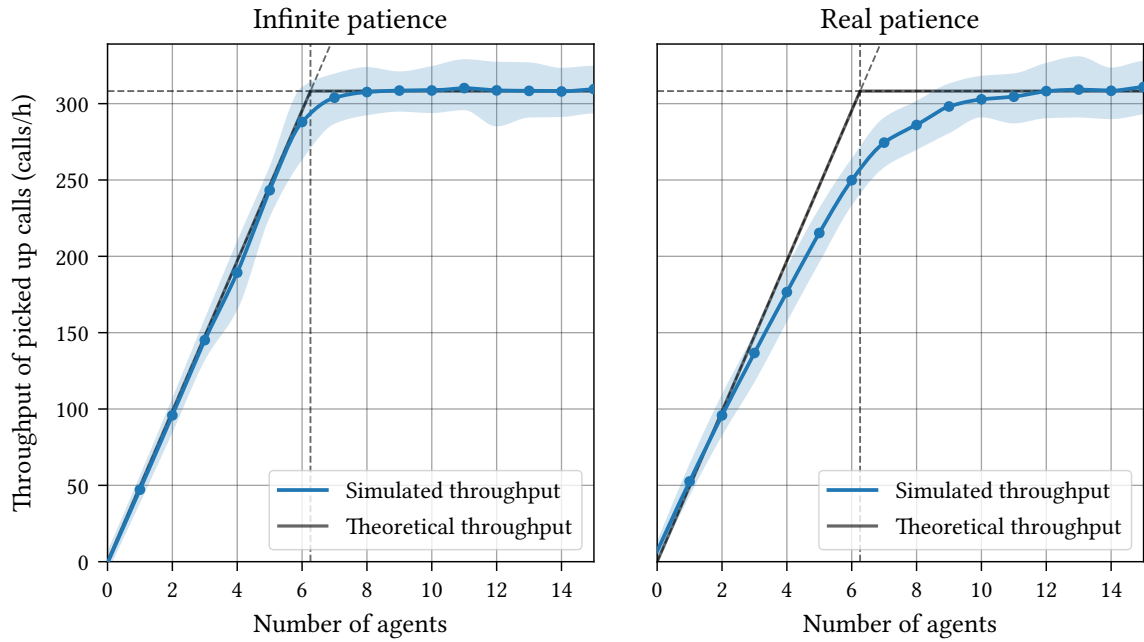


Figure 8.16: Throughput in the single-job and single-tier system

comes necessary to choose a value of N significantly greater than N_{ST}^* (a little less than twice this value) for these phenomena to vanish and retrieve a 100 % pick up rate.

As for the bilevel layout, we build on the Petri net of the Figure 8.4. As shown in the Proposition 5.2 in the Chapter 5 (under the infinite patience assumption, i.e., $\mu_1 = \mu_2^U = \mu_2^{VU} = \infty$), we can compute the analytical expressions of the two throughputs ρ_1^{BL} and ρ_2^{BL} . They respectively give the pick up throughput of calls at LVL 1 and LVL 2. Recall that the LVL 2 throughput ρ_2^{BL} is the sum of pick up throughputs of U and VU calls at the second level. In particular, the priority mechanism guarantees that even in case of a small lack of LVL 2 operators, the VU calls shall still be protected while the treatment pace of U calls will be degraded. Since the Petri net of the Figure 8.4 takes into account the extra instruction time τ_3 – contrary to the Petri net of the Figure 5.2 in the Chapter 5, we present below formulas from the Proposition 5.2 adapted accordingly.

We still use the two critical ratios r and r' introduced in the equation (5.6), where π^U and π^{VU} refer to the proportions of calls qualified as U and VU in the total flow with intensity λ :

$$r' = \frac{\pi^{VU}(\tau_2 + \tau_3 + \tau_S)}{\tau_1 + \pi^{VU}\tau_S} \quad \text{and} \quad r = \frac{(\pi^U + \pi^{VU})(\tau_2 + \tau_3) + \pi^{VU}\tau_S}{\tau_1 + \pi^{VU}\tau_S}.$$

Recall that if the quotient N_2/N_1 (number of operators at LVL 2 divided by the number of operators at LVL 1) is below r' , then there are not enough second-level agents to handle the mere VU calls transferred by the LVL 1. If $r' \leq N_2/N_1 \leq r$, there are enough second-level agents to handle the VU flow handed over by the LVL 1 but not the entirety of U calls (inferior priority). If finally $N_2/N_1 \geq r$, then the whole flow at LVL 2 transferred by the LVL 1 can be handled by second-level operators (even though the first level does not pick up all the incoming flow). The critical ratio r is the same as the one used in the Section 2.1 and in particular in Figure 8.6.

We obtain the following formulas for our two throughputs of interest

$$\rho_1^{\text{Bl}} = \min \left(\lambda, \frac{N_1}{\tau_1 + \pi^{\text{VU}} \tau_S}, \frac{N_2}{\pi^{\text{VU}} (\tau_2 + \tau_S + \tau_3)} \right)$$

at the first level, and

$$\rho_2^{\text{Bl}} = \begin{cases} (\pi^{\text{VU}} + \pi^{\text{U}}) \rho_1^{\text{Bl}} & \text{si } N_2 \geq r \min(N_1, N_1^*) \\ \frac{N_2 - \pi^{\text{VU}} \rho_1^{\text{Bl}} \tau_S}{\tau_2 + \tau_3} & \text{si } r' \min(N_1, N_1^*) \leq N_2 \leq r \min(N_1, N_1^*) \\ \frac{N_2}{\tau_2 + \tau_S + \tau_3} & \text{si } N_2 \leq r' \min(N_1, N_1^*) \end{cases}$$

at the second level, where $N_1^* := \lambda(\tau_1 + \pi^{\text{VU}} \tau_S)$.

It is interesting to note that contrary to the throughput ρ^{St} of the single-tier system, the throughput ρ_1^{Bl} is here not only constrained by the intensity λ of call arrivals and the handling speed of first-level agents, but also by the handling speed of second-level operators. Indeed, if the number of agents at LVL 2 becomes too low, not all Very Urgent calls can be handled by the LVL 2. Hence, they shall pile up at LVL 1 and hold the first-level operators who have to wait before escorting VU calls. The mathematical expression is therefore consistent with the intuition and covers all possible congestion phases, even the ones not observed in practice. Note once again the ‘‘cascading effect’’ since the expression of the throughput of picked up calls at LVL 2 depends upon its counterpart at LVL 1.

Similarly to the single-tier architecture, the above equations enable us to determine minimum staffing levels $N_{\text{Bl}-1}^*$ and $N_{\text{Bl}-2}^*$ at both levels, in order to reach a fluid phase where all calls are picked up with no delay:

$$N_{\text{Bl}-1}^* = \lambda(\tau_1 + \pi^{\text{VU}} \tau_S) \quad \text{and} \quad N_{\text{Bl}-2}^* = \lambda((\pi^{\text{U}} + \pi^{\text{VU}})(\tau_2 + \tau_3) + \pi^{\text{VU}} \tau_S)$$

We have depicted in the Figure 8.17 the experimental throughputs ρ_1^{Bl} and ρ_2^{Bl} computed in simulation. The colored dots correspond to the real measurements (still an average over 100 simulations for each value of N_1 and N_2). The colored surfaces are convenient continuous extensions given to help visualizing the global behaviour. The black dashed lines outline the theoretical profile of each throughput, predicted by the analytical formulas. The simulation points in both charts were obtained by using the real patience levels of callers (see again Section 5 of Chapter 7), while the theoretical polyhedral complexes assume infinite patiences. Similarly to the single-tier case, we observe that the simulated throughput corresponds very well with the predictions of the theory when the number of agents puts us far enough from the phases limits. This is in particular the case when N_1 and N_2 are high. The errors are bigger close from the phase transitions. We represent the difference between the analytical and the simulated models by dotted vertical lines.

From the theoretical point of view, it is interesting to focus on the throughput ρ_1^{Bl} and observe that the measurement error is important in the (small) phase where the number of agents at LVL 2 is limiting. Typically for $N_2 = 0$, the analytical model predicts a zero pick up throughput at LVL 1 in stationary regimes. That is because under the infinite patience assumption, all first-level agents become ‘‘paralyzed’’ by VU calls to be transferred at LVL 2 but with no operator to actually take these transfers. In reality when abandonments can occur, a well-staffed LVL 1 will still be

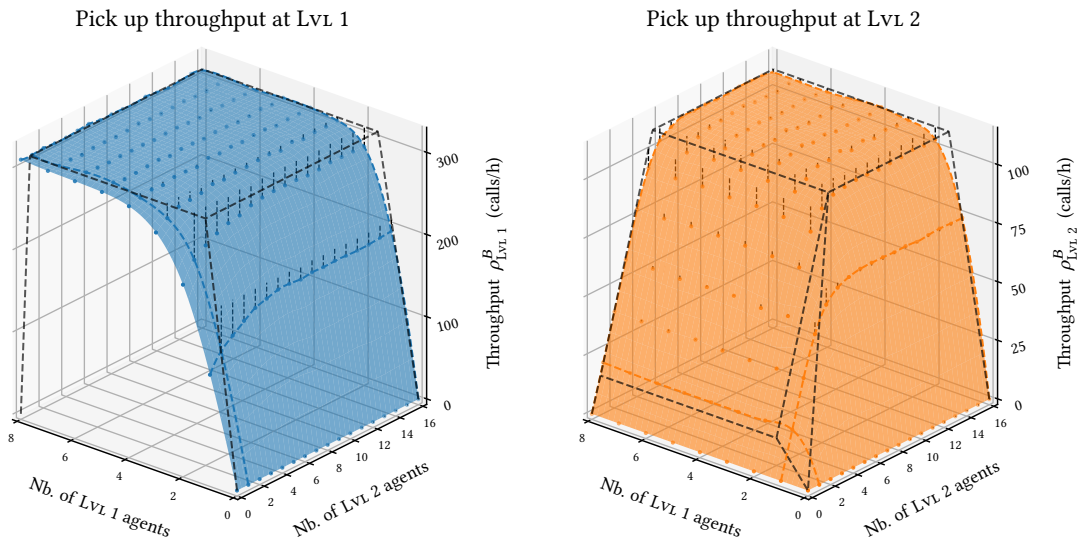


Figure 8.17: Throughput of the (single-job) bilevel system

able to pick up all arriving calls even though there are no second-level agents, because VU callers (a minority of the total volume) will unfortunately get impatient and hang up. This illustrates the fact that the analytical evaluation of the throughput under infinite patience is not always an upper envelope of the throughput under finite patience, in spite of the visual feeling that we get when looking quickly at the Figure 8.17. Still, we will keep in mind that both modelizations agree with each other in the other phases.

3.2 INFLUENCE OF THE OCCUPANCY RATE. — The previous analytical computations do not take into account the **occupancy factor** (hereafter abb. Occ) of the call center agents, defined as the proportion of their time spent on the phone. As introduced and explained in Section 6 of our Chapter 4, the Petri net models behave as if all the agents worked “non-stop”, with no pause or wrap-up time as long as there are pending calls. This is the reason why in the previous charts of this section, the computed values of N^* could seem very low (typically 7 agents for a single-tier system, way below the real staffings).

In order to reach a reasonable occupancy in practice, it is in order to divide the minimal computed number of agents by the target occupancy, hence taking an effective staffing N such that $N \geq N^*/\text{Occ}$. For instance, to reach a 50 % occupancy, we will choose N such that $N \geq 2N^*$. This is consistent with the intuition that for agents to be twice less busy on the phone, we need to double the staffing. Occupancy-wise, we expect two regimes depending upon the staffing. If the number of logged agents is less than N^* , the workload will be too high for these agents who will then have to work at full rate, and thus $\text{Occ} = 1$. Otherwise, that is to say if N is bigger than N^* , we expect a hyperbolic decrease of the occupancy under the form $\text{Occ} = N^*/N$, hence the full expression

$$\text{Occ} = \begin{cases} 1 & \text{if } N \leq N^* \\ \frac{N^*}{N} & \text{if } N \geq N^* \end{cases} \quad (8.1)$$

The law (8.1) is purely empirical. As a consequence, we have chosen to challenge it with experiments using our simulation software. The Figure 8.18 represents the product $N \times \text{Occ}$, easier to visualize than the occupancy factor itself. Indeed, we expect if (8.1) holds that $N \times \text{Occ}$ coincides with N for $N \leq N^*$ and with N^* for $N \geq N^*$, or equivalently that

$$N \times \text{Occ} = \min(N, N^*).$$

The use of the experimental occupancy in these charts is also an excellent manner to compare the analytical value of N^* (using theoretical formulas and the call center average parameters) and the asymptotic value of $N \times \text{Occ}$ that is fully obtained through experiments.

The left-hand chart corresponds to simulations performed for a single-tier system (average throughputs of the year 2019 in week days, rescue calls between 9 AM and 5 PM). We see that the experimental value of $N \times \text{Occ}$ stays very close from the profile predicted by the theory, apart from a gap around the breakpoint when $N = N_{\text{ST}}^*$. For higher values of N , we can observe that the hyperbolic decrease of the occupancy factor as a function of N is an excellent approximation. In addition, the horizontal asymptote coincides very precisely with the theoretical staffing, hence also validating this value. The second chart shows equivalent simulations for the bilevel system. We have depicted a first curve giving the product $N_1 \times \text{Occ}_1$ and another one giving $N_2 \times \text{Occ}_2$, since each level has its own headcount and thus its own occupancy. However, we have represented the product $N_1 \times \text{Occ}_1$ associated with the LVL 1 when the LVL 2 was well-staffed (16 agents). Similarly, we have shown $N_2 \times \text{Occ}_2$ when the first level is well-staffed (16 agents). This was done for the curves not to be “polluted” by cross-slowdown effects. Here as well, our theoretical predictions do well agree with the experimental measurements, all the more so when N_1 and N_2 become greater than the limit values $N_{\text{BL-1}}^*$ and $N_{\text{BL-2}}^*$. From now on, we will always admit that the empirical relation (8.1) is valid.

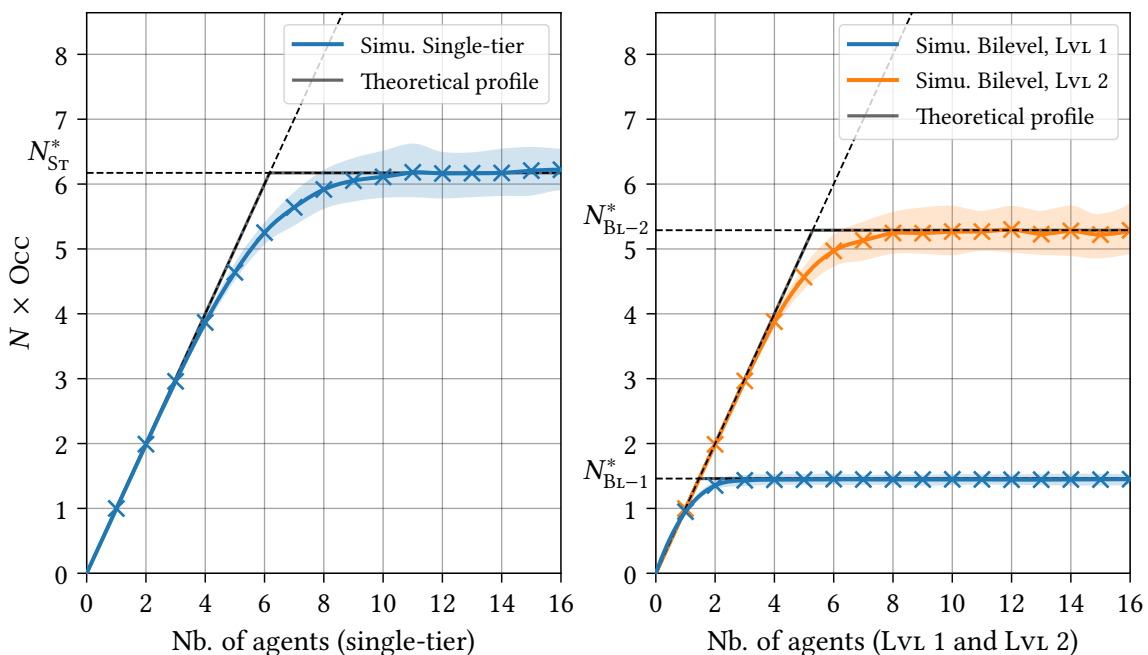


Figure 8.18: Product $N \times \text{Occ}$, theory vs. simulation

3.3 PERFORMANCE MEASURES DERIVED FROM ERLANG CALCULUS. – The Erlang calculus refers to a range of results obtained in the field of **queueing theory**. As discussed in the **Introduction**, queues are a complementary approach to Petri nets and quite popular in the call center domain. One advantage of Erlang calculus is to provide analytical formulas for several probabilistic performance criteria: the probability that a caller waits before pick up, the average waiting time, the abandonment probability, etc. Unfortunately, these results are more often obtained under strong mathematical assumptions for the computations to be feasible. For instance, the conversation times are most of times assumed to be exponentially distributed (which is not at all the case in practice). In addition, Erlang formulas typically assume simple call centers topologies. Handling more complex centers or cascading several layers is challenging in queueing theory, in particular under general distributions of call centers parameters – while the modularity of the Petri net approach comes convenient for such aspects.

Nevertheless, Erlang formulas are frequently used by managers or operators of call centers. Consequently, it remains interesting to compare their predictions with the real simulated performance. This way we can estimate in which extent Erlang formulas are valid and if we can extrapolate them for more complex systems, in order to have performance evaluation heuristics.

There are customarily three main models and Erlang formulas found in the literature: Erlang-B, Erlang-C and Erlang-A (ranked by increasing “complexity”). They all apply to single-tier architectures with a finite number of operators.

The **Erlang-B** model assumes call arrivals following a homogeneous Poisson process with parameter λ , conversation times exponentially distributed with mean τ , and infinite patience of callers. There are N operators (or servers) but also N “service lanes”, meaning that a call is blocked and rejected if all servers are occupied when it arrives. Using Kendall notation, the chain is of type $M/M/N/N$. In particular, a caller never waits, because either an operator is available when the call arrives and it is immediately picked up, or all servers are busy and the call is blocked and rejected.

The Erlang-B formula, first featured in [Erl17], gives the blockage probability of a call (denoting by B the random binary variable with value 1 if the call is blocked)

$$\mathbb{E}(B) = \mathbb{P}(B = 1) = \frac{(\lambda\tau)^N}{N!} \frac{1}{\sum_{k=0}^N \frac{(\lambda\tau)^k}{k!}} := E_B(N, \lambda\tau) .$$

We retrieve here the term $\lambda\tau$, sometimes referred to as the **load** (or the “traffic”). It is dimensionless but often expressed in *Erlangs*. This load is a key parameter in the sizing of call centers, already featured in our studies of Petri net models.

In the **Erlang-C** model, also featured in the article [Erl17], the chain is of type $M/M/N$. Compared to the Erlang-B model, there is an infinite amount of lanes and therefore no blockage. The callers will thus wait for operators to become available but are still equipped with an infinite patience – note that in real call centers there actually are a finite number of waiting lanes but that is much higher than the number of operators, so this is a relevant approximation. The Erlang-C

formula gives the probability that a caller experiments a wait after arriving and before being picked up (the nonnegative random variable W_C corresponds to the waiting time) whenever $N > \lambda\tau$

$$\mathbb{P}(W_C > 0) = \frac{\frac{(\lambda\tau)^N}{(N-1)!(N-\lambda\tau)}}{\sum_{k=0}^{N-1} \frac{(\lambda\tau)^k}{k!} + \frac{(\lambda\tau)^N}{(N-1)!(N-\lambda\tau)}} := E_C(N, \lambda\tau) = \frac{N}{\frac{N-\lambda\tau}{E_B(N, \lambda\tau)} + \lambda\tau}.$$

If $N \leq \lambda\tau$, there are not enough servers to handle the entirety of the load. As a result, calls (with infinite patience) will pile up and $\mathbb{P}(W_C > 0) = 1$. Actually, when the call center is sufficiently staffed (with $N > \lambda\tau$), Erlang calculus also provides the distribution function of the waiting time W_C before pick up via the relation $\mathbb{P}(W_C > t) = E_C(N, \lambda\tau)e^{-(N-\lambda\tau)t/\tau}$. We obtain in particular the average waiting time before pick up

$$\mathbb{E}(W_C) = \frac{E_C(N, \lambda\tau)\tau}{N - \lambda\tau}.$$

The average number of calls waiting for pick up (in other words the length L_C of the queue) can also be computed, and verifies Little's law, that is to say $\mathbb{E}(L_C) = \lambda\mathbb{E}(W_C)$. If $N \leq \lambda\tau$ though, we shall have $\mathbb{E}(W_C) = \infty$ and $\mathbb{E}(L_C) = \infty$.

Finally, in the **Erlang-A** model (that is in fact due to Palm in his article [Pal57]), the chain is of type $M/M/N + M$. Compared with Erlang-C, there is this time a patience with finite expectation, that is exponentially distributed with a mean μ . Among the three models, Erlang-A is the most representative of the reality – although it is also the least famous, maybe due to its advanced mathematical complexity. The Erlang-A formulas allow one to characterize the waiting time or the proportion of abandonments by using the following function A (built from the incomplete gamma function)

$$A(x, y) := xe^y y^{-x} \int_0^y t^{x-1} e^{-t} dt.$$

We obtain the probability that a caller waits (counterpart of the Erlang-C formula but this time valid for any value of N) by

$$\mathbb{P}(W_A > 0) = \frac{A\left(\frac{N\mu}{\tau}, \lambda\mu\right) E_B(N, \lambda\tau)}{1 + \left(A\left(\frac{N\mu}{\tau}, \lambda\mu\right) - 1\right) E_B(N, \lambda\tau)}.$$

Analytical computations based on the Erlang-A assumptions also allow to quantify the proportion of abandoned calls $\mathbb{E}(Ab)$ (the random variable Ab equals 1 if the caller hangs up, i.e., when the service time exceeds his/her patience, and 0 otherwise). Indeed, the probability of abandonment conditioned by a nonzero wait is known:

$$\mathbb{E}(\text{Ab} \mid W_A > 0) = \frac{1}{\frac{\lambda\tau}{N} A\left(\frac{N\mu}{\tau}, \lambda\mu\right)} + 1 - \frac{N}{\lambda\tau}.$$

Furthermore, the averages of the waiting time W_A or of the queue length L_A can also be determined by the equilibrium equation $\mathbb{E}(L_A) = \lambda\mu\mathbb{E}(\text{Ab})$ and Little's law $\mathbb{E}(L_A) = \lambda\mathbb{E}(W_A)$.

To finish, even though the probabilities to wait or to give up are quantities of interest for call center managers, the indicator that is the most looked at probably remains the service level (or grade of service), giving the proportion of calls picked up below a certain time threshold. For high performance objectives (low pick up times) and significantly inferior than the average patience, it is convenient to approximate the service level by the cumulated distribution function of W_A – waiting time of all calls, whether they are picked up or abandoned. The latter is given by

$$\mathbb{P}(W_A \leq t) = 1 - \mathbb{P}(W_A > 0) e^{-\left(1 + \frac{N\mu}{\tau}\right) \frac{t}{\mu}} e^{\lambda\mu(1 - e^{-t/\mu})} \frac{A\left(\frac{N\mu}{\tau}, \lambda\mu e^{-t/\mu}\right)}{A\left(\frac{N\mu}{\tau}, \lambda\mu\right)}.$$

Let us now study the adequacy of Erlang models with our simulations. We focus on Erlang-C and Erlang-A models, which both assume arrivals of calls that follow a homogeneous Poisson process, a virtually infinite waiting queue (no limit on the number of lanes), exponentially distributed service times and a constant number of servers to pick up the calls. They only differ by the fact that Erlang-A assumes a finite patience of callers (also exponentially distributed) and enables them to hang up on their own, while the Erlang-C model is built on the infinite patience hypothesis.

We first deal with a single-job and single-tier system, whose Petri net was depicted in Figure 8.8. We show in Figure 8.19 three performance criteria of interest predicted by the Erlang calculus and as functions of the number of operators. We systematically distinguish the case of infinite patience (it corresponds to the Erlang-C model, on the left charts) and of real patience (Erlang-A model, represented in the right charts). All the theoretical profiles are obtained from the above formulas, and we use as numerical values for the parameters the average values of the PFAU. We compare these theoretical curves to the experimental results obtained as before by simulations under standard arrivals from 1 PM to 5 PM for rescue calls. The conversation times and the patiences of callers used in our simulations are the real ones of the PFAU, measured and presented in the Chapter 7, hence are not exponentially distributed (hypothesis made in Erlang calculus).

The two first charts give the probability that a caller has to wait for a positive duration after he/she arrives and before pick up. Whether we look at the real or infinite patiences, we observe an excellent fit between the simulated values and those predicted by the probabilistic computation. The absolute error is always below 2.6 %, except at the breakpoint in the Erlang-C model where we find a 4.9 % error between the experiment and the theory. We check for the limit cases (either many or very few operators) that the Erlang-C and Erlang-A profiles give the same estimates, but they are very different inbetween. For instance, the Erlang-C model predicts 100 % of positive wait for 6 operators, versus 65 % in the Erlang-A framework due to the fact that many callers will abandon.

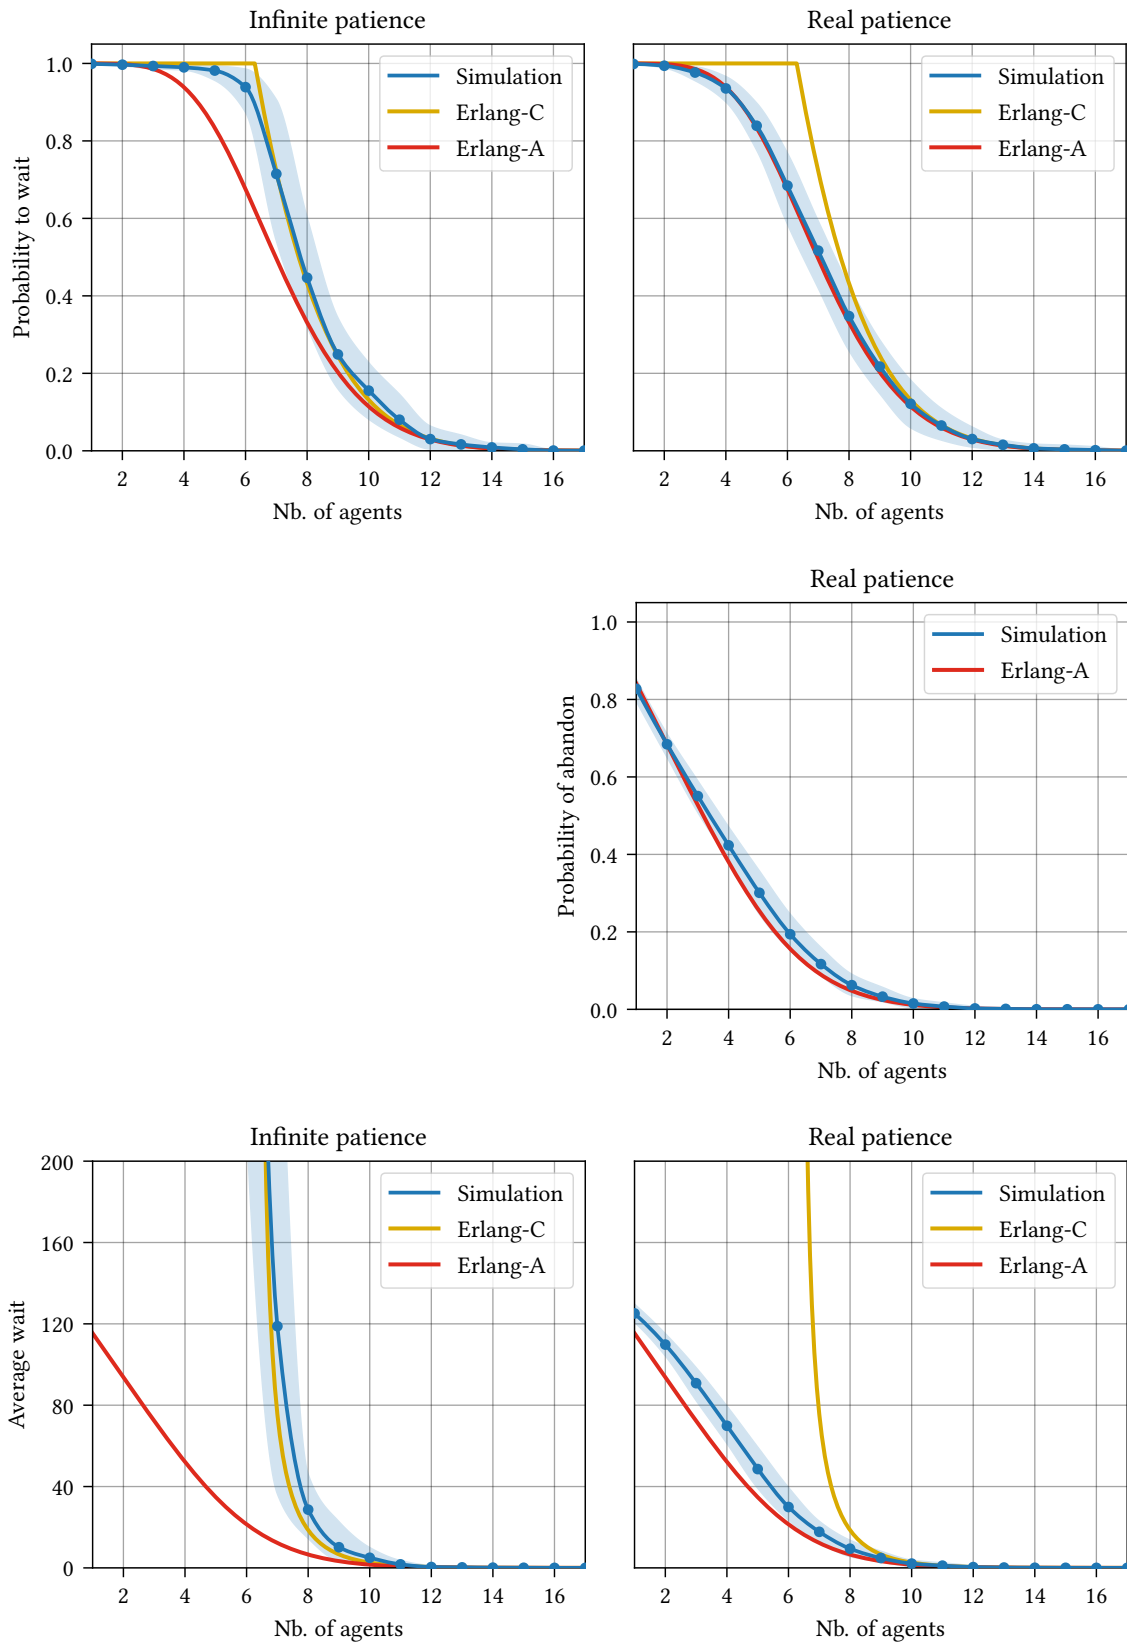


Figure 8.19: Probability to wait and to abandon, average wait, comparison with Erlang models (single-tier and single-job system)

The second row is devoted to the abandonment probability, a computation that is not applicable to the infinite patience case, hence only one chart is shown in the finite patience case. We can observe that the theoretical and the experimental profiles still quite agree, in spite of a maximal absolute error of 4 % near the critical staffing $N = \lambda\tau$.

Finally, we show in the last row the average waiting time before pick up (or abandonment). Contrary to appearances (due to the strong slopes), the Erlang-C model does not provide a good prediction of this value when N is too low. Close from the critical staffing, we have indeed a relative error of about 50 %. The prediction of the Erlang-A model is also wrong by 15 or 20 seconds in average for small values of N , even though the global look of the curve is acceptable.

This phenomenon actually comes from the fact that Erlang models and their assumptions tend to overestimate the distribution function of the wait for the real system, and therefore they predict an average wait lower than the one measured by the experiments. We have depicted in the Figure 8.20 the cumulated distribution functions $t \mapsto \mathbb{P}(W \leq t)$ for different values of N (on the right-hand chart for the real patience and on the left-hand chart for infinite patience; recall that for the latter this function is not defined if $N < \lambda\tau$). For a given value of N (and a same color), the solid curves give the profiles unveiled by our simulations, while the dashed lines give the CDF predicted by the theory. We indeed observe that the real distribution functions are below those of Erlang models, hence higher average waits in our experiments than in theory. In the infinite patience case (thus for the Erlang-C model), we may note that the cumulated distribution functions that we obtain remain roughly exponential functions, but with a different shape parameter than the one predicted by the theory. We reckon that this difference comes from the fact that the real distribution of the service times are not exponential, which impacts the global distribution of the waiting time and changes its shape. A rich literature characterizing such deviations exists but still involves approximation methods. We may however remark that for high values of N , the theoretical predicted wait is close from the one measured in practice – in this case only short waiting times are recorded.

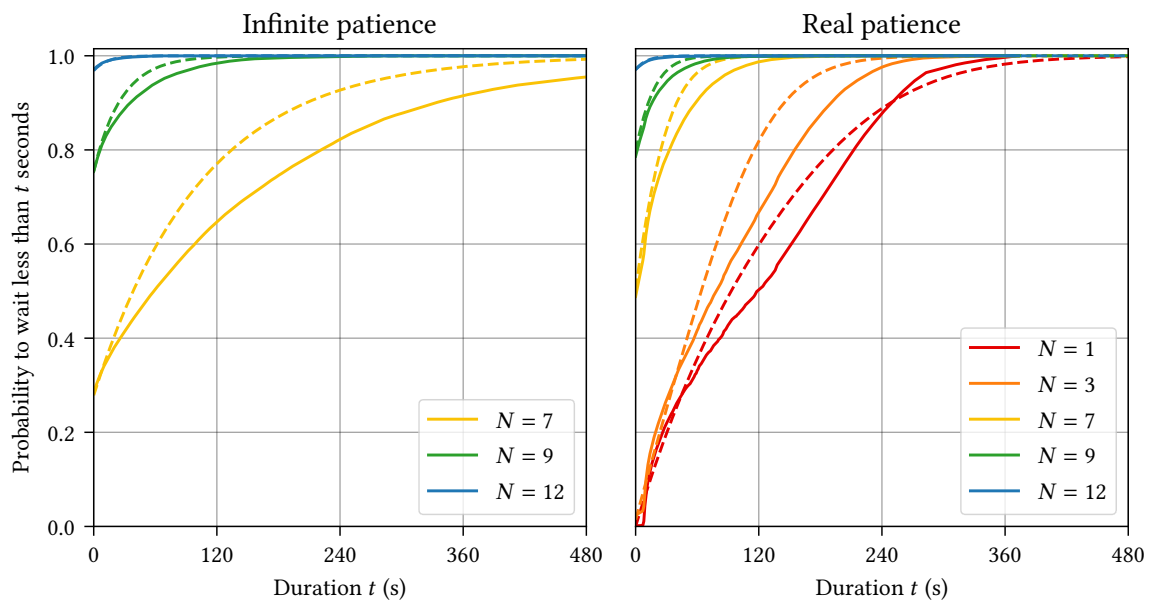


Figure 8.20: Cumulated distribution functions of the waiting time vs. Erlang models (single-job and single-tier system)

To finish, we explore with the Figure 8.21 in which extent the Erlang formulas can also be applied to a bilevel architecture, even though they are fundamentally derived from a single-tier topology. Our intuition is the following; in a well-staffed bilevel system where both levels are fluid, the treatment chain of the calls behaves like two independent and cascaded single-tier systems, and the output flow of the first level is a Poisson process (just like its input flow).

We let aside the case of the infinite patience, that is interesting from the theoretical point of view but in any case is superseded by the Erlang-A model. We want to see how the latter fits with our performance measures for the bilevel system at both levels, using the real averaged PFAU parameters in the analytical formulas.

In the two first charts of the Figure 8.21, we have shown the probability that a caller who has just arrived (either at LVL 1 or at LVL 2 after having passed the first level) has to wait before pick up. The waiting probability at LVL 1 is given as a function of the number of first-level agents, and the probability to wait at LVL 2 as a function of second-level operators. However, these values also depend a bit on the staffing number of the other level, all the more so when this other level is not sufficiently staffed. To illustrate this effect, we therefore always represent the waiting probability of a level when the *other* level is either very badly staffed ($N_1 = 1$ or $N_2 = 1$) or very well staffed ($N_1 = 8$ or $N_2 = 16$). It is interesting to observe that at LVL 1, the probability to wait is higher when the LVL 2 is under-staffed (the agents of LVL 1 are more busy because they are held to transfer their VU calls), but this is the opposite at LVL 2; a caller will wait longer if the LVL 1 is well-staffed since there are more tasks transferred to the LVL 2, thus less availability of second-level operators. This is consistent with the intuition and also with the remarks made in the Section 2.1 after the heatmaps. We have also depicted with the solid red line the probability to wait such as predicted by the Erlang-A model, even though it comes from a single-tier topology and uses unrealistic distributions of service times and patiences. We observe that this simple model does very well predict the probability to wait at LVL 1 when the LVL 2 is well-staffed, and the probability to wait at LVL 2 when the LVL 1 is also well-staffed (about 4 % of maximal absolute error, observed for small values of N_2). As expected, the theoretical result does not hold when one of the other level is under-staffed (more than 15 % of absolute error at LVL 1 and 50 % at LVL 2).

As far as the abandonment probability before pick up is concerned, we do the same comparison in the second row of the Figure 8.21 and draw identical conclusions. When the other level is well-staffed, the Erlang-A model correctly predicts the abandonment fraction at one level as a function of the corresponding staffing, using only the mean conversation time and the mean patience at this level. However, for the average waiting time (still taken for all calls, that is to say either before pick up or before abandonment), the Erlang-A formulas have a decent predictive power at LVL 1 but a bad one at LVL 2 – even when the LVL 1 is well-staffed. We have for instance a theoretical result that is about 60 seconds lower than the real wait measured in simulation for 4 second-level agents. Just like the single-tier system, we explain this difference by the distribution of the conversation time at LVL 2, whose shape features a high variance (that is much less the case at LVL 1). The Erlang-A model comes back in agreement with the simulation measurements for high values of N_2 (and close from those used in practice), but this mostly involves low values and does not particularly strengthens our faith in this theoretical model.

To summarize, we may therefore conclude that despite its simple assumptions, the Erlang-A model provides precise tools to quantify the probability to wait and to abandon (the latter being one of our fundamental performance criteria) in a single-tier system. This also holds in a well-staffed bilevel system. On the other hand, it does not give a good evaluation of the average waiting time (that is a criterion related to the service levels). This constitutes an additional argument to resort to numerical simulation.

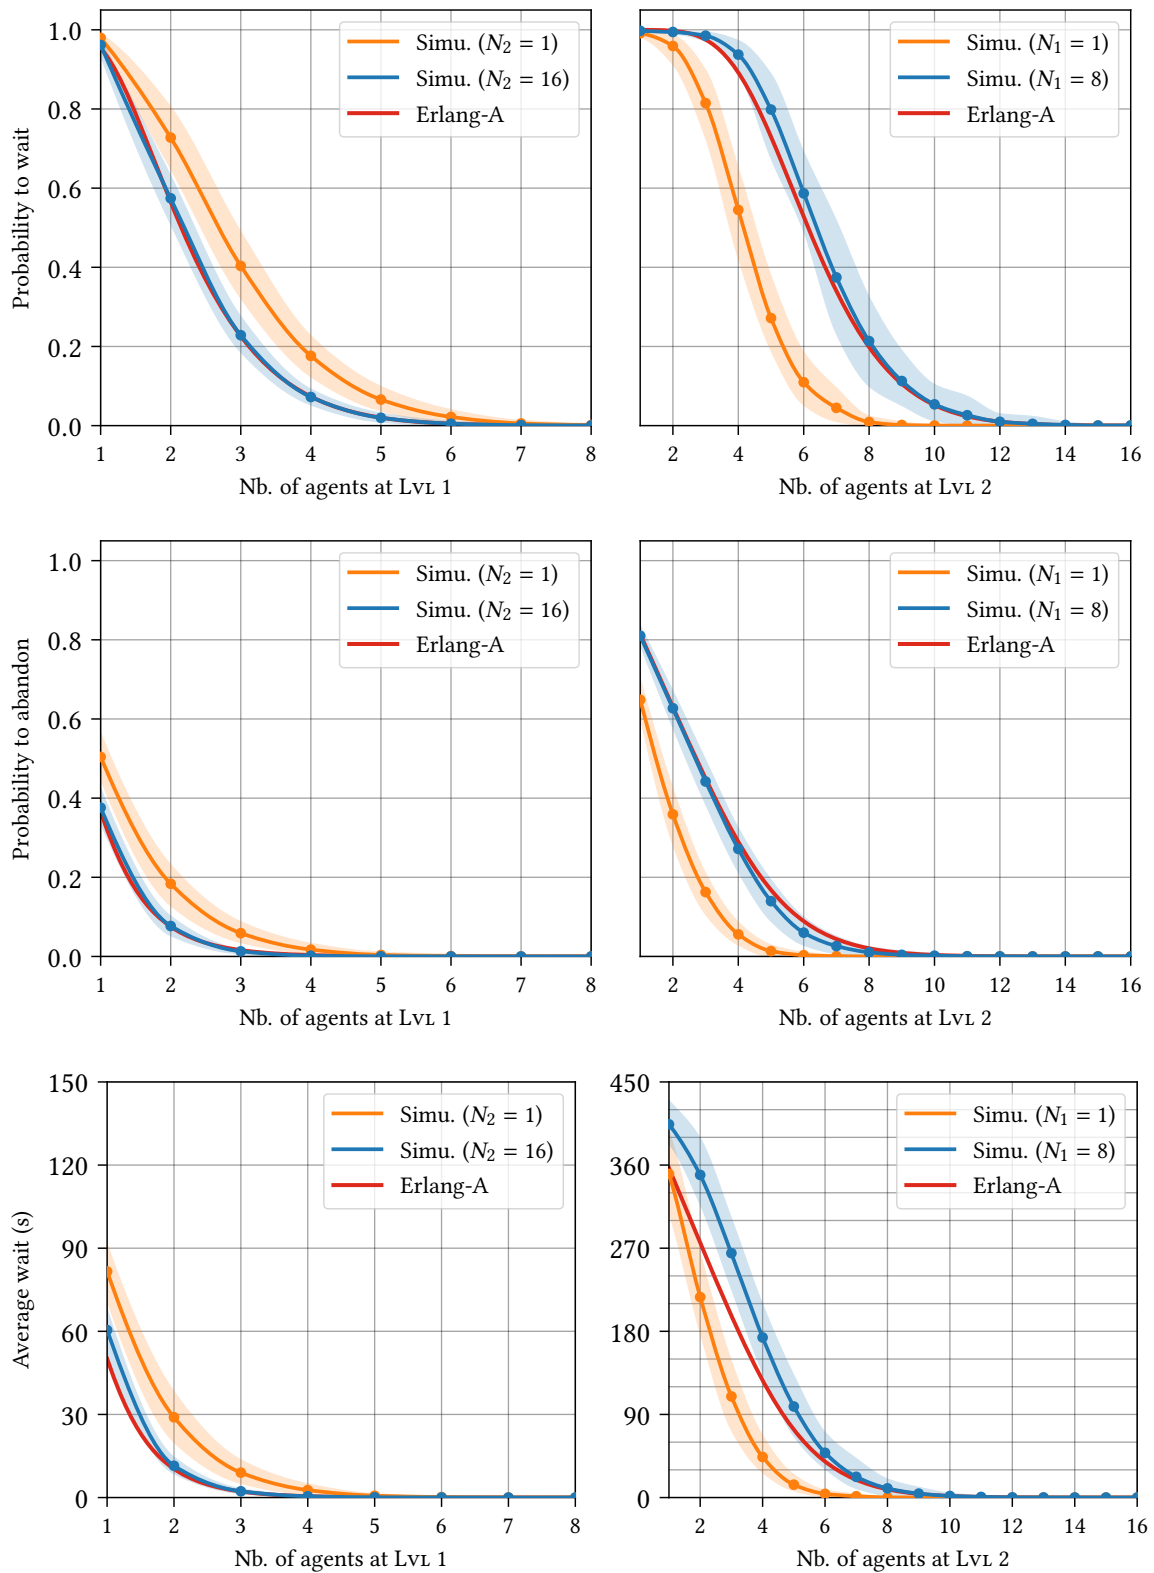


Figure 8.21: Probability to wait and to abandon, average wait, comparison with Erlang models (bilevel and single-job system)

4 Study of the PFAU architecture and its possible evolutions

In this section, we focus on models that imitate as much as possible the real behaviour of the PFAU, contrary to the Section 2 that sometimes considered fictitious organizations (single-job assumptions) to present proofs of concept. In particular, the results of the current section are always obtained by simulating **multi-job** organizations, involving both police and rescue tasks. Recall however that we still deal with large-scale emergency call centers in terms of demand and staffing (four highly populated French departments).

4.1 PETRI NET MODELS FOR ORGANIZATIONS WITH TWO DISTINCT JOBS. – We start by introducing the Petri net models that shall be used to simulate the multi-job layouts, extending the Petri net of Figure 8.4 of the previous section where detailed evolution rules are recalled. We depict in Figure 8.22 a Petri net very accurate with respect to the real PFAU, except for the fact that the second level of DSPAP is here aggregated, and it is departmentalized in practice. The subcomponents of the net for the two lanes are thus symmetrical. We have also implemented a Petri net where the LVL 2 of DSPAP is effectively departmentalized but it is not shown here for the sake of readability. This modification of the organization (either departmentalized or aggregated) will be at the heart of the Section 4.3.

Just like the single-job case, the calls from line 18 enter a waiting place, where only a first-level BSPP agent can pick them up. If the wait is too long, the callers can hang up and leave the system. The same goes for callers who dialed the number 17, involving first-level DSPAP officers. However, for people having dialed the emergency line 112, they can either be picked up by a BSPP or DSPAP agent at LVL 1, depending on the respective availability of both services. From the operators point of view, there is a fair treatment between calls from 17 or 18 on one hand and calls from 112 on the other hand. Indeed, the first call that arrives will be the first one picked up (regardless of the emergency line). In case a firefighter and a police agent are both available and a call from 112 arrives, it is the operator who has been idle for the longest time that will have to pick up the call – this mechanism is symbolized by the 🕒 icon.

As in the single-job case, and as described in the Section 2 of the [Foreword](#), a call picked up at LVL 1 is subject to different routings depending upon its nature and its gravity degree. The Non Urgent calls have a conversation with the first-level agents before leaving the system and making the operators available again. The Urgent or Very Urgent calls are transferred to a LVL 2 agent from the most appropriate service, with the escorting mechanism for VU calls. In the current multi-job setting however, we explicitly represent the possibility that a call is “rerouted” between the two levels. A caller who dialed the number 18 (a priori to reach the firebrigade) can as a result be put in communication with the second level of DSPAP (a priori for police matters). This is done by crossing the U-POL transition after discussing with a first-level firefighter from BSPP – shown by the arrow marked with a (★) symbol.

The operators of LVL 2 still serve in priority the VU calls, as indicated by the double-tip arcs. Between two fireable transitions with equal priority level, an operator token will always be used to fire the one where a caller token has been waiting for the longest time to access the second-level. After a possible escorting conversation (after which the LVL 1 operator is released), the token enters the final conversation place. When this exchange is over, the caller leaves the system and the second-level operator can perform additional post-instruction tasks before coming back available again.

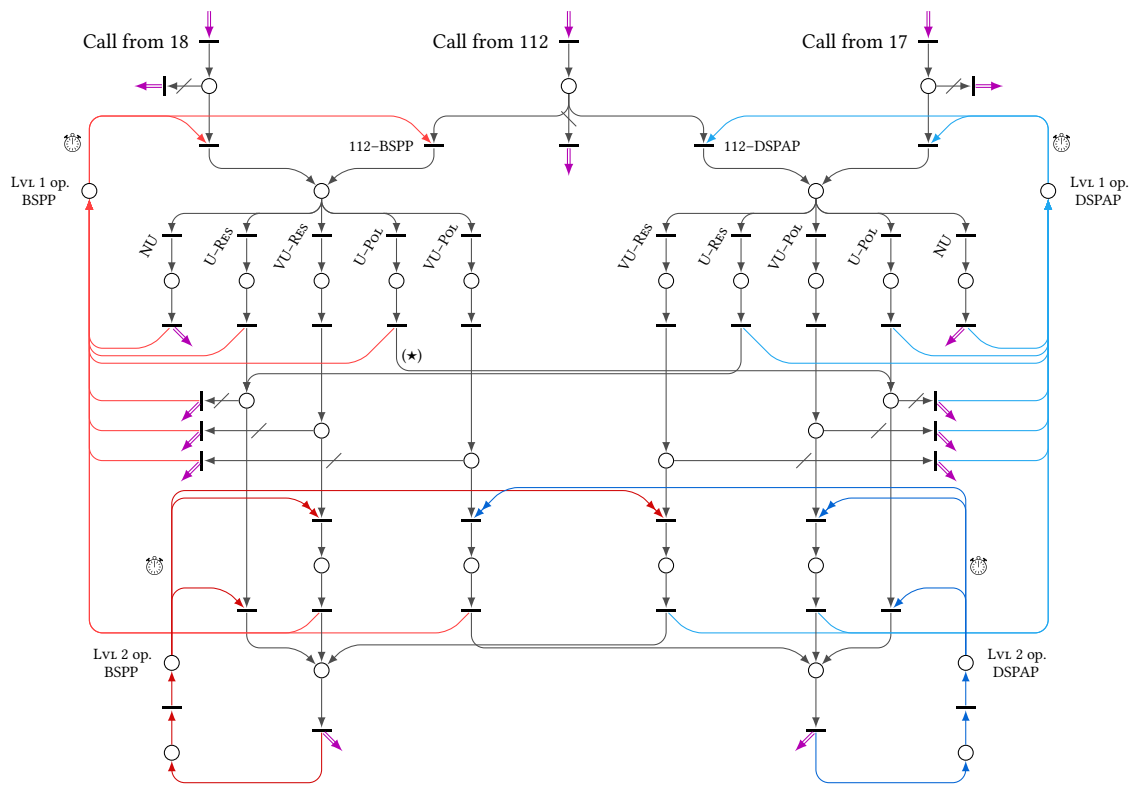


Figure 8.22: Petri net of a simplified PFAU where the handling of police calls at Lvl 2 is not departmentalized

Modeling the departmentalization of the handling of police calls at LVL 2 requires to reproduce the whole chain of treatment of the bottom-right corner of the Figure 8.22 four times, into four independent blocks each using a different pool of operators. This also involves bringing up more arcs between the LVL 1 and the four LVL 2, to correctly send the calls in terms of their department of origin. The (rare) transfers from line 18 to the second-level of DSPAP also need to be duplicated.

4.2 STUDY OF THE INTERACTIONS BETWEEN THE TWO JOBS OF THE PLATFORM. — One of the main features of the current PFAU is to bring together two different services under the same roof. This way, the police and the rescue services (respectively staffed with DSPAP and BSPP agents) are able to perform “reroutings” very efficiently. In addition, the pick up of the number 112 is shared by the first levels of each service. We present below a couple of analyzes illustrating the coupling and the interactions of the two jobs.

The Figure 8.23 allows us to appreciate the effect of the **shared treatment of the number 112**. It gives the proportion of calls from 112 picked up by the first level of DSPAP, as a function of both LVL 1 staffings (we naturally obtain the proportion of calls picked up by BSPP agents as the complementary to 100). This heatmap is computed when the second levels of the two services are fluid (that is to say 4 police officers at the police LVL 2 in department 75, 3 agents in each police LVL 2 of departments 92, 93 and 94, and 10 agents at the LVL 2 of BSPP).

This map is consistent with the intuition, confirming that the first level of one job takes over when the other one is not enough staffed. For instance, for a choice of $N_1^{\text{POL}} = 4$ and $N_1^{\text{RES}} = 2$, then 63 % of calls from 112 are picked up by police agents, versus 7 % with the converse repartition. We observe that the computed values are not symmetrical, and with for instance 4 operators at each first level, the firefighters pick up about 70 % of calls from 112. This is a somewhat expected outcome since the load (see Section 2.4) that weighs on the police agents at LVL 1 is twice the one received by firebrigade agents (1.84 E versus 0.94 E). Hence the “two thirds / one third” breakdown for the remaining 112 load (0.44 E).

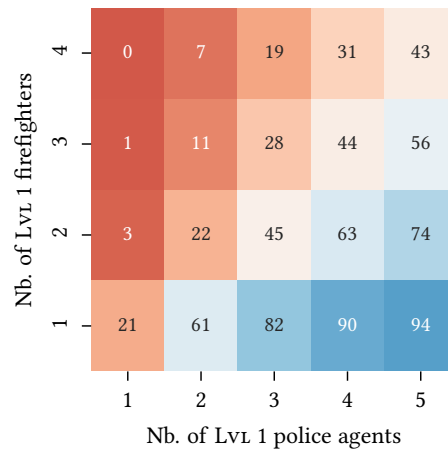


Figure 8.23: Proportion of line 112 picked up by the LVL 1 of DSPAP (current PFAU organization, standard flows, from 1 PM to 5 PM, $N_2^{\text{POL}} = 13$ and $N_2^{\text{RES}} = 10$)

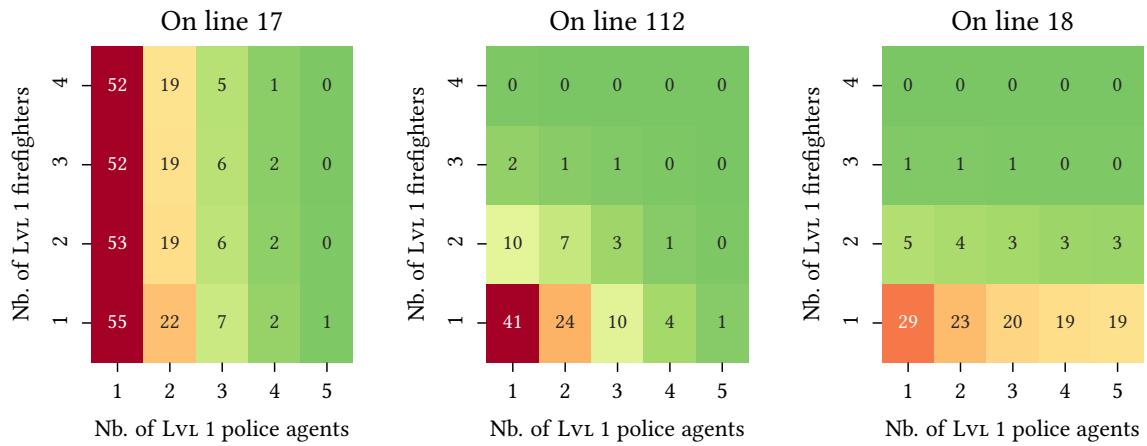


Figure 8.24: Abandonment rates at LVL 1 (current PFAU organization, standard flows, from 1 PM to 5 PM, $N_2^{POL} = 13$ and $N_2^{RES} = 10$)

The Figure 8.24 gives the abandonment rates at LVL 1 for each line, still as a function of the two first-level staffings and for a fluid cross-section of the second levels. It provides complementary information to the Figure 8.23.

On the first chart of the Figure 8.24, giving the abandonments at LVL 1 from line 17, we can check as a first approximation that the obtained figures are invariant by addition of LVL 1 firefighters (vertical stripes). Similarly, the chart for calls arrived by line 18 is nearly invariant by addition of first-level police agents (horizontal stripes). These phenomena are expected since only police agents pick up the line 17 and only firefighters pick up calls from line 18. Yet, we note slight supporting effects, with 3 % less abandonments on line 17 if $N_1^{POL} = 2$ and N_1^{RES} goes from 1 to 2. The Figure 8.23 confirms that this drop comes from the fact that this new first-level firefighter alleviates the load of first-level police agents on the 112 calls – the proportion of 112 calls picked up by the latter going from 61 % to 22 %. In the middle chart (relative to 112 calls), we see this time a “diagonal” pattern since increasing either one of the two LVL 1 staffings makes the corresponding abandonment rate decrease. Once again, the asymmetrical figures observed between lines 17 and 18 (for instance abandonments rate of 55 % and 29 % when $N_1^{POL} = N_1^{RES} = 1$) are a direct consequence of the different loads on both lines (two to one ratio). The same goes for the decreases of abandonment rates on line 112 when we go from a (1, 1) to a (1, 2) or (2, 1) configuration.

We now want to get a global view of the effect of the different staffing numbers of the PFAU – there are currently seven different pools of operators – on its performance. To this end, we depict in the Figure 8.25 a correlation matrix between several performance criteria and the headcounts. The retained metrics are: the abandonment rate, the average total wait (LVL 1 plus LVL 2) for U and VU calls, and the occupancy factors of agents; as the ones already introduced to study the single-job systems in Section 2.

The color intensities indicate the importance of the correlations (in the sense of Spearman), blue hues (resp. red hues) being associated with positive (resp. negative) correlations.

In the two first rows of the Figure 8.25, we observe (consistently with the intuition) that the global abandonment rate of Urgent rescue calls and the average wait of U/VU rescue calls essentially depend on the rescue staffings N_1^{RES} and N_2^{RES} . However, they are also weakly correlated

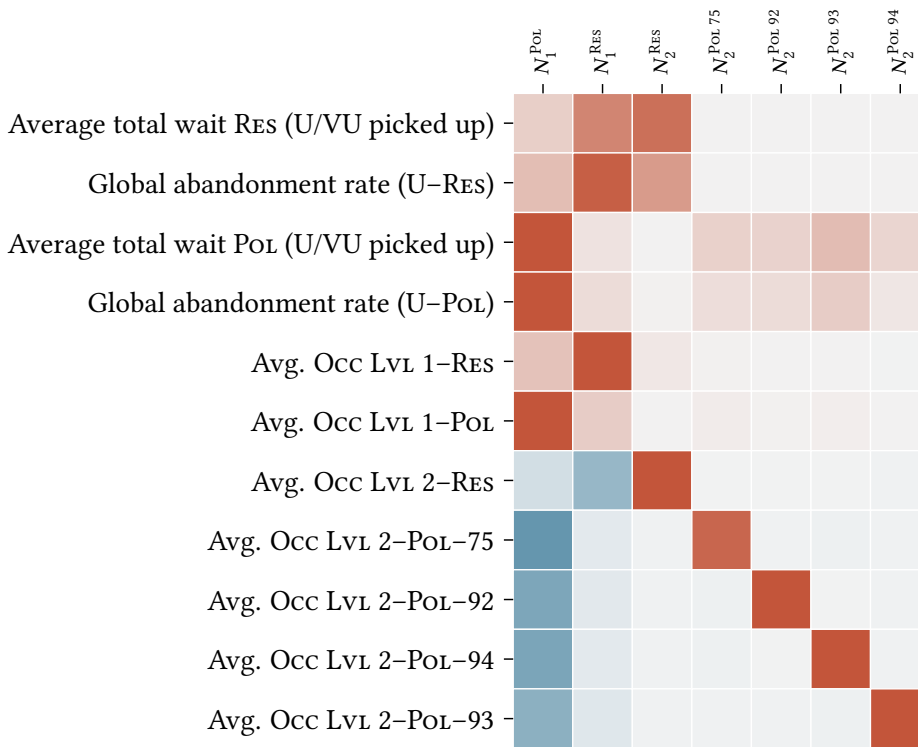
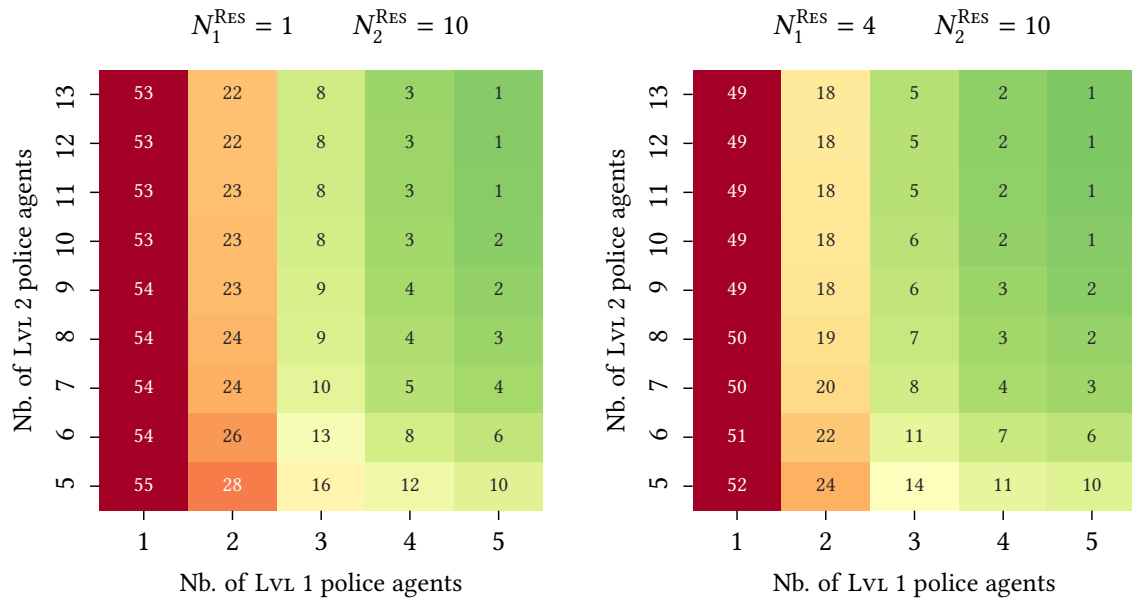


Figure 8.25: Dependence of performance criteria with respect to staffings (current PFAU organization)

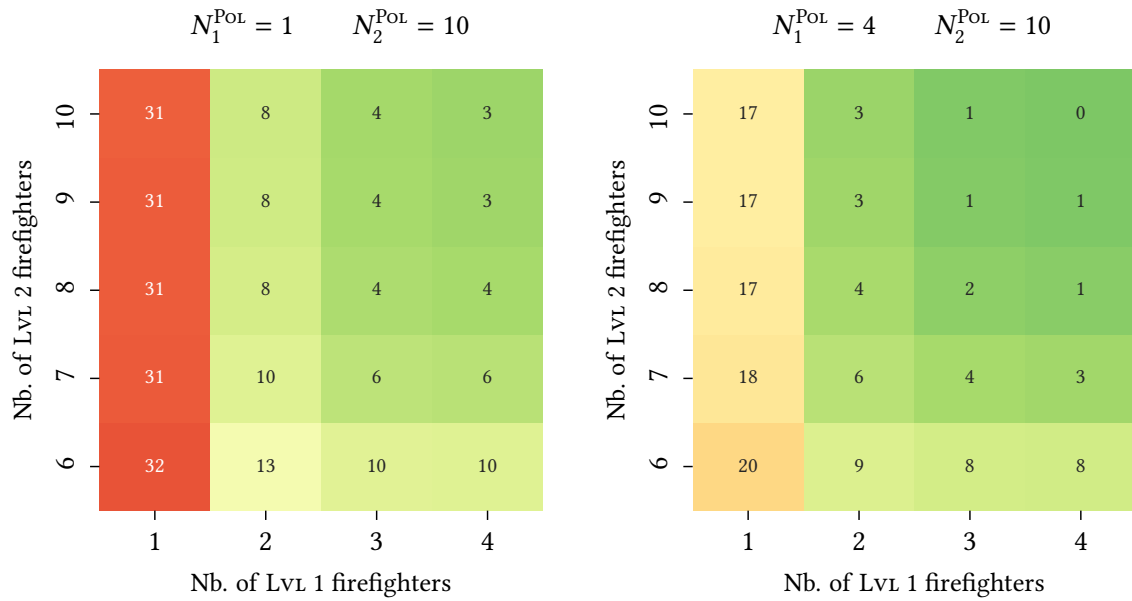
with N_1^{POL} , since the police agents at LVL 1 can indirectly help the firefighters by taking the 112 calls off their load, as seen before. However, these rescue-related criteria are in very good approximation independent of the LVL 2 staffings of police services. We find a symmetrical phenomenon for the police jobs. Their average wait and abandonment rate strongly depend upon N_1^{POL} and on the staffing numbers at police second levels (the correlations are less strong due to the departmentalization). There is a weak but not negligible dependence on N_1^{RES} , but not on N_2^{RES} .

We have also indicated the occupancy factors of each pool of agents. We naturally check that they are all highly correlated and decreasing with respect to the staffing level of the corresponding pool. We note that the average occupancy of the LVL 1 of rescue services decreases a bit when increasing the police staffing at the first level – and vice versa. At the second level, we remark as in the single-job case that the occupancy of LVL 2 firefighters increase if the LVL 1 staffings (either rescue or police) increase, because more calls are transferred. The occupancies of police agents at LVL 2 similarly increase when we raise the staffing of the first police level, but not much when raising the number of LVL 1 firefighters. We may explain this asymmetry by the departmentalization of the police LVL 2 (which “dilutes” the correlations) but also by the fact that the rescue LVL 1 transfers much less calls to the police LVL 2 from line 18 than police LVL 1 transfers calls to the rescue LVL 2 from line 17 (as seen in Chapter 7).

We finish this subsection on the interactions between the two jobs hosted at the PFAU by showing how the LVL 1 staffing of one job impacts the global abandonment rate (that is to say calls abandoned either at LVL 1 or LVL 2) of the other job, or in other words that a **mutual support** exists. For instance, the top part of the Figure 8.26 shows two performance maps relative to the global



(a) Abandonment rate for U-POL calls, effect of support by LVL 1 RES



(b) Abandonment rate for U-RES calls, effect of support by LVL 1 POL

Figure 8.26

abandonments rate of Urgent Police calls as functions of the police staffing levels at LVL 1 and LVL 2. These two maps only differ by the N_1^{RES} quantity, i.e., the staffing of LVL 1 firefighters. It is set to a single agent on the left map, and to four agents on the right map, while the staffing N_2^{RES} remains set to 10. We can see that this evolution from 1 to 4 first-level firefighters makes the police global abandonment rate drop from 53 % to 49 % if there are many second-level police agents and a single first-level police operator; and from 22 % to 18 % if there are two police agents at LVL 1. This reinforcement has almost no effect if the police LVL 1 is already well-staffed. As far

as the rescue job is concerned, if we look at the abandonment rate of Urgent rescue calls, going from 1 to 4 police agents at LVL 1 has a much higher effect, going from 31 % to 17 % (that is to say the half) when we have only one firefighter agent at LVL 1 and many at LVL 2. Even for three or four rescue agents at LVL 1, increasing the staffing of the police first level has a noticeable contribution (abandonment drop of two or three points).

Once again, this very different supporting effect between rescue and police jobs that we measure on small headcounts (leftmost columns of each map) is explained by their respective loads. With $N_1^{\text{RES}} = 1$, the single firefighter at LVL 1 is almost unable to cope with the load from line 18 that is of 0.94 E, and can even less handle the total 18 + 112 load of 1.38 E. As a result it causes many abandonments on line 18 (which contains many U calls). With many police agents at LVL 1, the police officers can handle the calls from 112, which leaves 0.94 E to handle for the unique LVL 1 firefighter. This agent will have to work with a nearly full occupancy and many abandonments will still occur, but much less than before. Now if we deal with the police job, the first-level load from line 17 equals 1.84 E, as a result with or without the support of firefighters, a single police agent at LVL 1 will not be capable of handling the entirety of the flow, and we expect to see as many abandonments on the first column ($N_1^{\text{POL}} = 1$) of both maps. When taking $N_1^{\text{POL}} = 2$, we start measuring an abandonment rate that is characteristic of a limit regime, and finally when $N_1^{\text{POL}} = 3$ we can observe a controlled rate.

4.3 EFFECT OF THE SYSTEM DEPARTMENTALIZATION. – In this section, we aim at quantifying the impact on performance of aggregating several departmentalized organizations into a single one. Intuitively, unifying small systems by keeping a constant global staffing is always advantageous because of a smaller risk a saturation due to the scale-up. Indeed under a constant flow, four organizations with two agents each have more chances to meet saturation than a single one with eight agents.

Let us first describe the models studied in what follows. We choose to only focus on the police part of the PFAU, since the DSPAP currently operates with a semi-departmentalized layout. As a reminder, the pick up and the instruction of calls from line 17 (and a part of calls from 112) for the whole city of Paris and its inner ring (four French departments in total) are done by an **aggregated first level**, common to all departments. However, if a call is qualified as U or VU by this first level, it is transferred to a **departmentalized second level**, where the operators only deal with calls from a single area (department of origin of the call).

A question of interest is to know how the performance of this organization would evolve if the four current LVL 2 were merged into one, by keeping the same total staffing. Since the main focus of our work remains the comparison of single-tier and bilevel layouts, we are going to simulate on identical scenarios the four models presented in the Table 8.27. The current organization is referred in this table as “Departmentalized bilevel”, which does not correspond to one bilevel system by department but indeed to the half-aggregated-half-departmentalized layout described above. Note that the two aggregated models correspond to the ones already studied in the Section 2.4.

For the simulations, the study of the single-tier systems amounts to running single-job (police) Petri nets – four independent nets in parallel in the departmentalized case, a single one for the aggregated case. We choose for the study of bilevel layouts to launch simulations on Petri nets derived from the current multi-job PFAU (the real departmentalized one and a fictitious platform where the police LVL 2 is aggregated). However we set the firebrigade staffings to 5 agents at LVL 1 and 11 agents at LVL 2 (consistently with the minimal staffing computed in Section 2.3);

Single-tier departmentalized	Four independent one-level systems
Single-tier aggregated	A single one-level system
Bilevel departmentalized	One common LVL 1, four independent LVL 2
Bilevel aggregated	One common LVL 1, one common LVL 2

Table 8.27: Four models of DSPAP organizations depending on the number of levels and the departmental structure

in what follows we only let the police headcounts free to study their effect. In the four models, the inbound flow given as input in our scenarios always contains the entirety of the calls from line 17, and the calls from line 112 that are Urgent or Very Urgent and are police matters. In the bilevel scenarios only, the first-level police agents can also pick up the rest of the calls from line 112 (since it is shared with the firefighters of LVL 1). This constitutes an extra load for them in bilevel layouts that we will quantify in order to draw a fair comparison between the different organizations.

We now properly tackle the **multi-criteria** comparison of the four models gathered in the Table 8.27. For each one of them, the Figure 8.28 gives as a function of the total police staffing the following criteria: the global abandonment rate of all police calls, the abandonment rate of only VU police calls, the average total wait of U and VU calls (sum of pre-LVL 1 wait and inter-level wait) and the maximum occupancy factor of all pools of police agents (whether they are staffed LVL 1 or LVL 2). For departmentalized or bilevel organizations, the computed value of a criterion for a given total staffing is always obtained by breaking down this total headcount into the different pools of agents, optimizing a mixed criterion. As in the Section 2.3, this mixed criterion first aims at minimizing abandonments of VU calls, then the global abandonment rate, and then the waiting time. Note that this criterion is common to all the charts. In this way, and even though the optimal repartition of 20 agents across 4 independent single-tier call centers is maybe not the same to achieve a minimal abandonment or a minimal total wait, all the curves for a same architecture follow in the Figure 8.28 the same “descent path” in terms of agents allocation. This is the correct way to compare all the systems, but changing too much the definition of the mixed criterion also ends up changing the comparison.

The first chart of the Figure 8.28 shows that the two departmentalized systems (whether they have one or two levels) are equivalent in the iso-staffing comparison and for the global abandonment rate criterion, beyond 13 or 14 operators in total. For 13 agents for instance, four “parallel” single-tier call centers (with 5 agents in Paris (75), 2 in Hauts-de-Seine (92) and 3 in Seine-Saint-Denis (93) and Val-de-Marne (94)) achieve as many abandonments (between 3 % and 4 %) across the four departments than the current bilevel PFAU with 5 police agents at LVL 1, 4 at the LVL 2 of DSPAP-75, 1 at the LVL 2 of DSPAP-92 and DSPAP-94 and 2 at the LVL 2 of DSPAP-93. Above all, we note the **much better performance of the aggregated systems**, where we have merged the teams of all departments. Still with 13 agents, the aggregated bilevel system suffers only 0.5 % of global abandonment rate (putting 6 police agents at the LVL 1 and

7 at the LVL 2), and the aggregated single-tier system reaches 0.05 %. Looking at the multi-job systems, we can verify (not depicted in the figure) that unifying the four LVL 2 of police services also has an influence on the rescue chain. Indeed with 13 police agents in total, the LVL 1 of DSPAP in the departmentalized bilevel system picks up 24 % of the calls from 112, versus 34 % for the aggregated bilevel.

This very strong effect of aggregating the departments is confirmed in the second chart, featuring the abandonment rate of VU calls. We clearly distinguish two groups of curves: those associated with departmentalized modes, and those associated with aggregated modes. Here, the aggregated bilevel has a very similar performance than the aggregated single-tier, the difference with the previous chart being that VU calls are handled in priority at the LVL 2 of the bilevel layout. Beyond a total of 15 operators, the abandonment rates for VU calls in both aggregated systems fall below 0.05 % (one loss every 2,000 calls). For the departmentalized architectures, we may also point out another strength of the bilevel system (that was already visible for global abandonment). Indeed, even though the mean performance are almost the same, we observe by looking at the 90 % confidence intervals that the abandonment rates for the bilevel system have a much lower standard deviation than for the single-tier. Note how the green envelope is largely included in the blue envelope for typical staffings between 12 and 19 agents in total. This proves that thanks to its aggregated first level, the bilevel architecture is more robust, and more resilient to “worst-case scenarios”. Consequently, aggregating *only* the first levels of several departments thanks to the bilevel layout is already very advantageous.

As far as the total wait of U and VU calls is concerned, we may note the significant delay of the departmentalized bilevel layout compared to other models (at least 15 extra seconds for less than 14 agents) due to more divided pools and small integer staffing constraints – however we need to keep in mind that these profiles have been obtained by optimizing a mixed criterion that favours having fewer abandonments over having short waiting times. Interestingly, the aggregated bilevel layout is quicker than four independent single-tier systems when we take more than 14 operators, even though it is 5 or 10 seconds longer below this threshold. For sufficient staffings, the aggregated bilevel lags behind the aggregated single-tier by less than 2 seconds, so it is actually very competitive.

Finally, as it was already highlighted for the single-job and rescue layout, we remark that the bilevel organizations (for the iso-staffing comparison key) make their agents a little busier on the phone, with occupancies 10 points higher than their single-tier counterparts. Nevertheless, very satisfying levels of occupancy are reached as soon as we take a total of 15 or 16 agents.

The previous analysis brings us to the computation of minimal staffing for the police pools. The Table 8.29 summarizes the minimum number of agents that we need to have in order to satisfy all the criteria selected by the MARCUS report (see again the Section 2.1). These headcounts are taken relatively to the entirety of police calls between 7 PM and 11 PM (the moment of the day with most police activity), regardless of the model. When there are several pools of agents, the total staffing is always given by first specifying the number of first-level agents (if the system is bilevel) and then the 75/92/93/94 counts (if the system is departmentalized). As in the equivalent table that was presented for single-job call centers handling only rescue calls (see the Section 2.3), each individual performance criterion may involve a different optimal breakdown of the staffing. As a result, we need to read the “All criteria” row to obtain a global dimensioning policy. The very high target of admission rate (99.95 %) is still the most challenging to reach and drives the final optimization scheme. For this reason, the last row coincides with the first one.

In order to reach less than 0.05 % of abandonments (which will also guarantee that all the other criteria are satisfied), the numerical simulations recommend that four single-tier systems

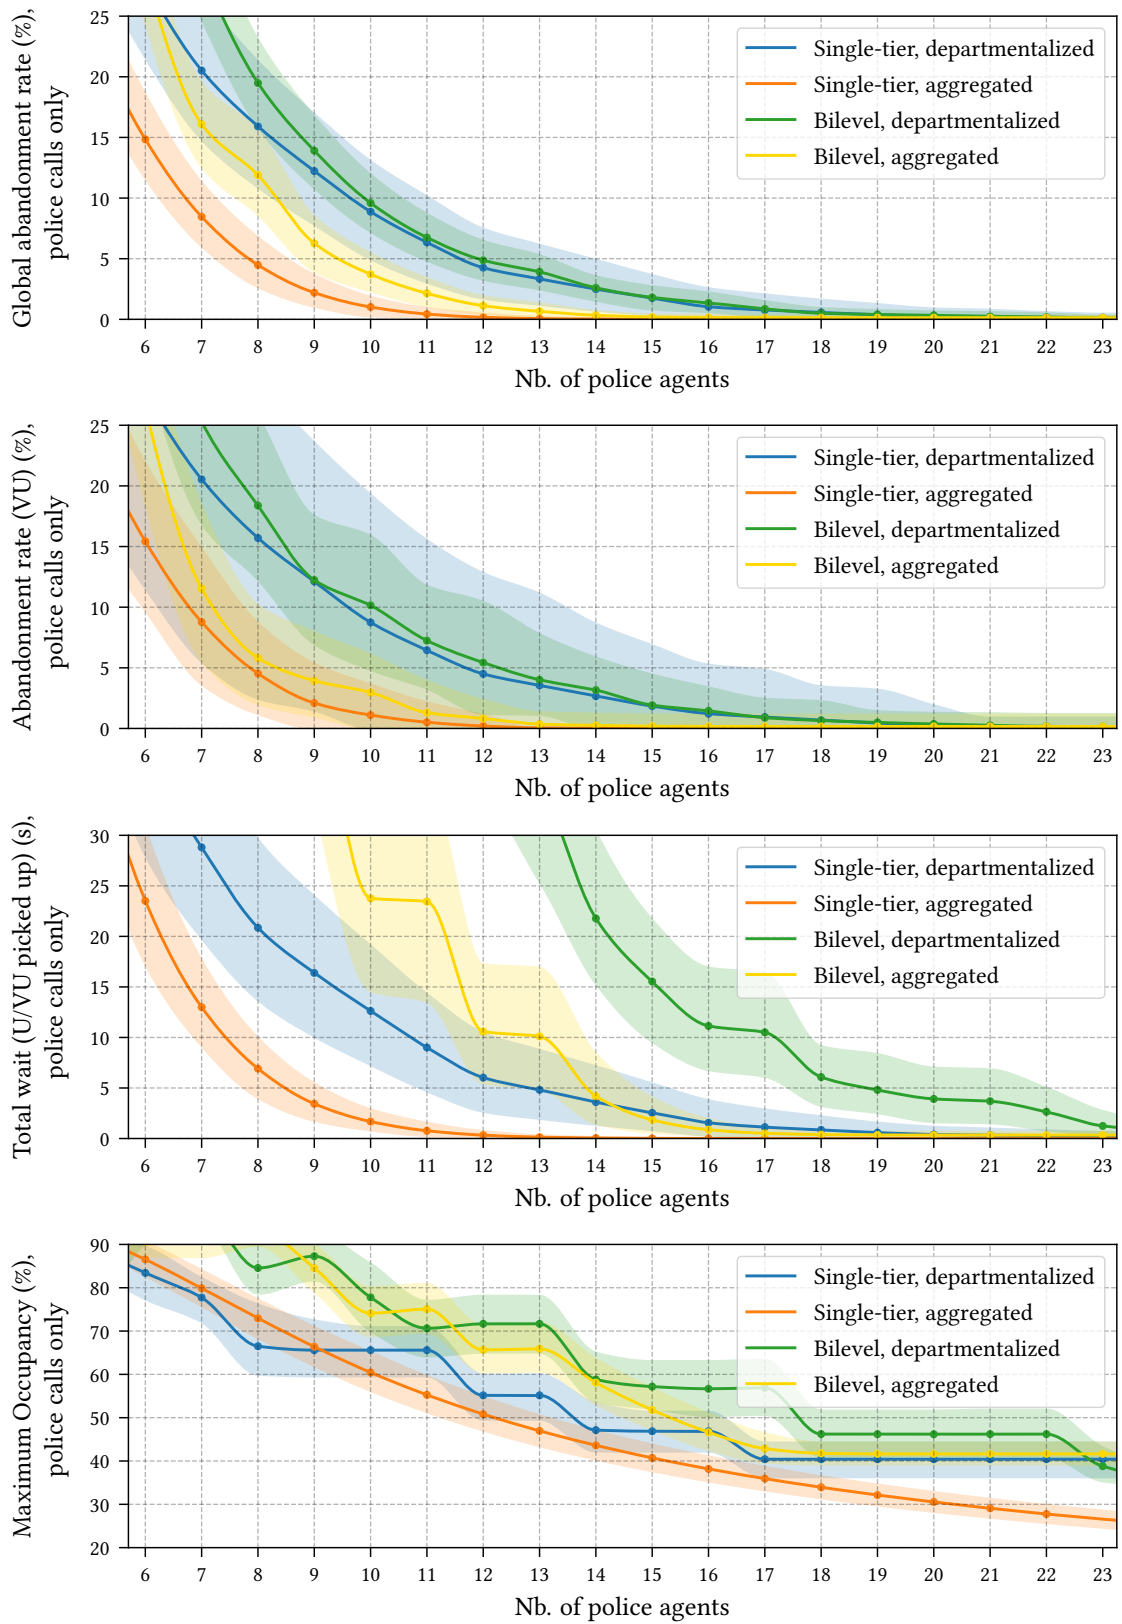


Figure 8.28: Comparison of four organizations on the police flows: single-tier vs. bilevel and departmentalized vs. aggregated (standard flows 75–92–93–94, from 7 PM to 11 PM)

Criterion	Target	Minimal staffing			
		Departmentalized		Aggregated	
		Single-tier ($\times 4$)	Bilevel	Single-tier	Bilevel
Admission rate	$\geq 99.95\%$	24 (8+4+6+6)	25 (8+6+3+5+3)	13	17 (7+10)
Service level – 15 s	$\geq 90\%$	13 (6+2+3+2)	10 (6+1+1+1+1)	10	9 (5+4)
Service level – 30 s	$\geq 99\%$	19 (7+3+5+4)	11 (6+2+1+1+1)	11	10 (6+4)
Maximum occupancy	$\leq 60\%$	12 (5+2+3+2)	13 (4+4+1+2+2)	10	12 (4+8)
All criteria		24 (8+4+6+6)	25 (8+6+3+5+3)	13	17 (7+10)

Table 8.29: Police staffing table (with well-staffed rescue services), typical police arrivals from 7 PM to 11 PM, (≈ 230 calls/h)

should be staffed with 8 agents in Dpt. 75, 4 in Dpt. 92, 6 in Dpts. 93 and 94 – making a total of 24 (expert) agents. For the current PFAU system (aggregated at LVL 1 but not at LVL 2), we recommend having 25 agents in total, with 8 police officers at LVL 1, 6 of them at LVL 2 of Dpt. 75, 5 at LVL 2 of Dpt. 93 and 3 at the second levels of Dpts. 92 and 94. As seen before, the aggregated systems spare many resource, as a unique single-tier system requires 13 (expert) agents, versus 17 for an aggregated bilevel – 7 first-level operators, 10 expert second-level officers.

We recall that for bilevel systems, the staffings featured in the Table 8.29 to satisfy the Service Levels are derived by single-criterion optimization schemes. However, these metrics rely almost only on the number of first-level agents, hence we have a very good Service Level below 15 seconds in a departmentalized bilevel with a total of 10 agents (6 at LVL 1 and a single officer at each LVL 2), but this does not yield good global abandonment rates. We have chosen to let these service level indicators in particular for their interest in the single-tier systems.

4.4 EFFECT OF THE FIRST LEVEL UNIFICATION. – In a similar view than the one discussed at the beginning of the Section 4.3 on the aggregation of the departments, we expect that two separate LVL 1 (one for the police, one for the firebrigade) and with two distinct agent pools should be less fluid than a single “universal” LVL 1 – in which all operators can pick up all the calls. In this subsection, our goal is to quantify how many operators could be spared by such a unification.

We here deal with a unified LVL 1 that keeps the current three emergency lines, their throughputs, their patience levels and conversation times. This is a little bit different than a “generalized 112” with all emergency lines merged and probably slightly modified primitives. This choice is the most realistic in the short-run and also the most handy to perform based on the data we have. We are thus going to simulate a bilevel system with a unified LVL 1 and assume that the universal first-level operators handle a call from line 17 (resp. from line 18) the same way as a current first-level police agent (resp. firefighter).

In the light of our previous experiments on the determination of minimum staffings, we choose to look for the minimum number of agents that allows us to reach a global abandonment

	Minimum staffing (bilevel layout, departmentalized police)		
	One LVL 1 per job (shared 112)	One universal LVL 1	
Standard arrivals 1 PM – 5 PM	$N_1^{\text{POL}} = 7$	$N_1^{\text{RES}} = 4$	$N_1^{\text{UNIV}} = 9$
	$N_2^{\text{POL}} = (5, 2, 3, 2)$	$N_2^{\text{RES}} = 11$	$N_2^{\text{POL}} = (5, 2, 3, 2)$ $N_2^{\text{RES}} = 11$
Standard arrivals 7 PM – 11 PM	$N_1^{\text{POL}} = 8$	$N_1^{\text{RES}} = 4$	$N_1^{\text{UNIV}} = 10$
	$N_2^{\text{POL}} = (7, 3, 4, 3)$	$N_2^{\text{RES}} = 11$	$N_2^{\text{POL}} = (6, 3, 4, 3)$ $N_2^{\text{RES}} = 11$

Table 8.30: Staffing table of the current PFAU versus a PFAU with unified LVL 1

rate that is less than 0.05 % (recall that this is the most restrictive criterion of the MARCUS report). When several allocations put us below this target, we always retain those which also minimize the number of second-level operators because they are more skilled and costly resource.

The Table 8.30 gives the staffing levels determined by our method, for the current model (one LVL 1 per job) and for the fictitious model where the LVL 1 is unified. In both cases, we keep the police lane departmentalized at LVL 2. We also make distinct computations for the afternoon time-window 1 PM to 5 PM, and for the evening slot 7 PM to 11 PM. In the afternoon, we observe that the unification of the first level allows us to obtain identical performance (in terms of global abandonment rate) with two less operators, going from 7 police agents and 4 firefighters at LVL 1 to 9 universal first-level agents. The headcounts are unchanged at LVL 2. The same difference arises in the evening, with 10 universal operators needed at LVL 1 versus 8 + 4 when splitting this level into two pools. In addition, our optimization scheme shows that it is also possible to take off a second-level police officer from the DSPAP of department 75 without exceeding the threshold of 0.05 % abandonment rate. What actually happens is that there are two configurations letting us reach the performance target: either keeping 7 second-level agents at DSPAP-75 and taking 9 first-level universal operators, or taking 10 universal agents at LVL 1 and 6 at the LVL 2 of DSPAP-75 (in the first case, the very few abandonments recorded will mostly occur at LVL 1, and in the second case at LVL 2). Even though it is eloquent for our analysis to claim that we can save three agents by unifying the LVL 1, we still prefer to minimize the number of “expert” agents at LVL 2. Naturally, it is also possible to unify the two first levels and keep the current total staffing. This offers a very significant gain of performance to the unified layout.

It is interesting to note that contrary to the optimization schemes performed before, the Table 8.30 is obtained by letting *all* the headcounts free in the optimum search and under all 17–18–112 arrivals. Recall that the Table 8.10 was obtained for a single-job organization handling 18–112 calls and the Table 8.29 used a multi-job Petri net but with already fixed rescue staffs (5 first-level agents and 10 second-level agents). This difference explains why we find here that only 4 firefighters are sufficient at LVL 1 in the global PFAU, due to the help of the 7 first-level police agents who support them picking up the line 112 and alleviate their load.

We can consequently conclude that the first column of the Table 8.30 gives the total **staffing level recommended for the current PFAU** under typical arrivals. The second column demonstrates that unifying the first level always allow to “save” two agents at LVL 1, and sometimes an extra police officer at LVL 2.

5 Behaviour of the bilevel layout under non-standard arrivals

So far, we have only focused on simulating scenarios where the arrivals of emergency calls followed average and “typical” throughputs. We now deal with the performance of the single-tier and bilevel layouts in “crisis” situations, where by definition we move away from the standard scenario. We consider either planned or unplanned non-standard events. Indeed, keep in mind that some peaks of calls can be predicted in advance, such as those of New Year’s Eve or Bastille Day in France. In these cases, the staffing are raised in consequence to cope with the unusual demand, and the new staffing policy can be thought of as a shift of equilibrium. For unpredicted events though (fire, explosion, attack...), massive call arrivals will hit the day-to-day staffing of the platform, which precisely corresponds to an out-of-equilibrium situation.

5.1 RESPONSE OF THE BILEVEL LAYOUT TO AN UNEXPECTED PEAK OF CALLS. — We choose for this analysis to study the impact of a “rescue-type” event, because the current BSPP chain of treatment is already well adapted to the bilevel architecture (see our previous discussions in Section 2.3 or Section 4.3). We place ourselves in the 1 PM to 5 PM time-window that we simulated before. We consider the usual average throughputs computed from all week days of the year 2019, but we add an (unusual) extra flow extracted from the fire that broke out in a junkyard in Bobigny on the 26 June – it is presented in details in the Section 6 of Chapter 7. This fire triggered many emergency calls that were tagged as EVU (Event Urgent) by the first-level of the rescue lane, mainly around 8 PM when the smoke was visible from Paris’ outer ring. Many extra VU calls were also received. We have taken from this particular day of 26 June the EVU calls and the extra VU calls (net of the usual daily VU throughput) to add them on top of the average throughputs between 1 PM and 5 PM. This gives us a realistic scenario of an important fire occurring in the afternoon of a “normal” day.

The Figure 8.31 shows the throughputs of calls per gravity that we use for this crisis simulation, computed over five rolling minutes. The reader may note that the peak of the EVU curve reaches a maximum of 350 calls/h, while the one of the Figure 7.25 on which it is based only peaked at 270 calls/h. This comes from the fact that we took into account in our simulation the calls that were abandoned after the voice prompt during the real fire. Those demands peaked at 120 calls/h and we chose to break them down into 90 % of calls which would have been labeled as EVU, and 10 % that would have been VU (recall that our throughput curves never count the callbacks). The virtual event that is depicted in the Figure 8.31 is simulated between 1 PM and 5 PM, but actually starts at 1:30 PM. From 2:15 PM, we can see that the throughput of VU calls exceeds the one of NU calls, and reaches almost 75 calls/h just after 3 PM. This is almost four times the usual VU demand. The VU calls that may experience the longest wait are those arriving after 3:30 PM, because at this moment the EVU solicitation skyrockets.

To evaluate how these crisis profiles are absorbed by our call center models, we give them as input scenarios to a single-tier and single-job (rescue) center staffed with 14 agents, and to the current (bilevel) PFAU organization where there are 5 firefighters at LVL 1 and 10 firefighters at LVL 2. This is consistent with the recommendations of the Table 8.10 (and in this case the police lane uses 8 first-level agents, and 6+3+5+3 police officers in the four departmentalized LVL 2, following the recommendations of the Table 8.29). In particular, these two concurrent models are staffed so that they achieve less than 0.05 % abandonment rate on the usual demand. In these two models, the EVU calls are handled quicker than the U and VU calls, using conversation times of the Figure 7.25. In average, we counted in our simulations 888 calls on line 18 per run (on

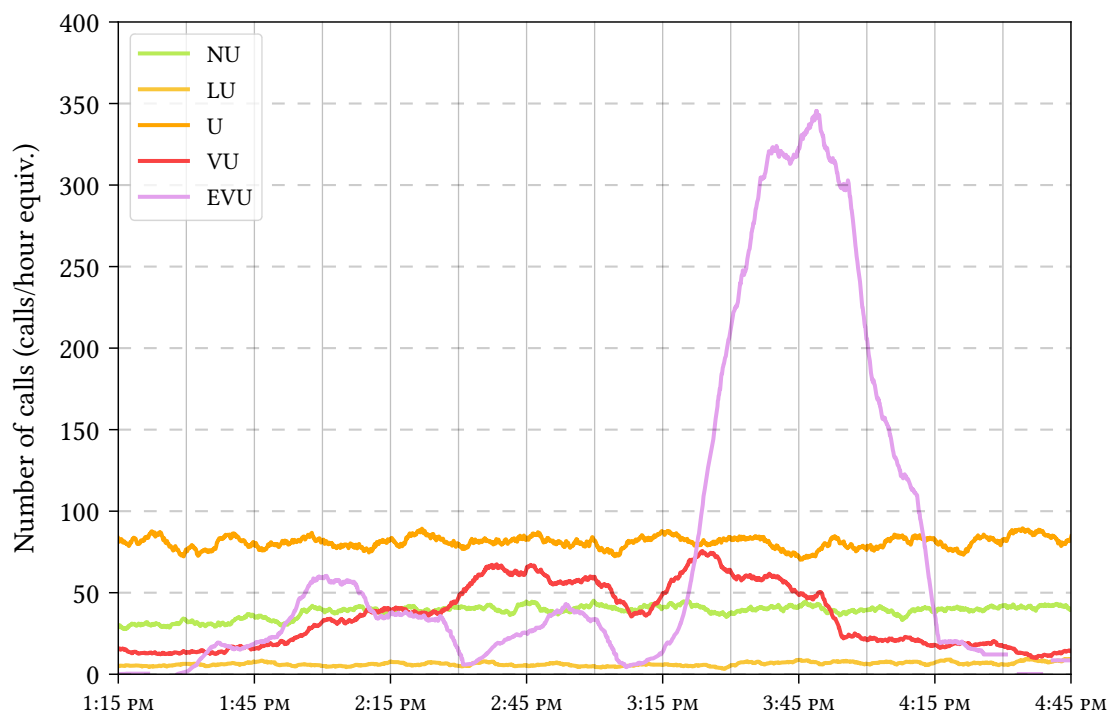


Figure 8.31: Throughputs used to simulate a massive peak of calls on line 18 (based on a fire in Bobigny)

Criterion	Average performance		Worst-case performance	
	Single-tier	Bilevel	Single-tier	Bilevel
Abandonment rate (global)	1.64 %	5.40 %	3.24 %	9.55 %
Abandonment rate (VU calls)	0.95 %	0.28 %	2.42 %	1.35 %

Table 8.32: Performance of the single-tier and bilevel layouts facing an expected peak of calls such as Bobigny fire

the 4-hour time-window) with an average of 135 VU calls. In standard scenarios – that is to say without the EVU and extra VU calls, there are in average 592 calls in total, and 61 VU calls.

The Table 8.32 gathers the abandonment rates computed by our simulations on the two architectures, either for all calls or specified only to Very Urgent demands. We distinguish the average performance (mean abandonment rate over 100 simulation runs) and a “worst-case” performance. The latter is computed by taking the 95th centile of the empirical distribution, which is in practice the minimum abandonment rate recorded across the five worst simulations (recall indeed that each simulation scenario is a bit different due to random draws). Using these metrics, we can observe that when facing this unexpected crisis and without increasing the usual staffing, the single-tier system suffers 1.64 % of global abandonment, versus 5.40 % for the bilevel – this is three times less. However, the single-tier has a 0.95 % abandonment rate for VU calls, versus 0.28 % for the bilevel – this is three times more. Digging a bit in the results of these experiments, we find out that for the bilevel system, these VU abandonments almost exclusively arise at Lvl 1. When they are picked up by the Lvl 1, the priority mechanism fulfills its role and prevents callers from hanging up. On the contrary, the 5.40 % of global abandonments in the bilevel case are mainly due to U or EVU calls waiting for a second-level operator. The first-level filter has therefore played its role, enabling the rescue agents to handle in priority the very urgent and life-threatening cases.

In the worst-case scenarios, that sure are rare but still exist, we observe the same qualitative phenomenon. The global abandonment rate is about three times higher in the bilevel system compared with the single-tier, but the VU calls are still better protected by the bilevel, with twice less abandonments for this specific gravity degree. This illustrates that the bilevel architecture is more robust to undesirable events because the confidence interval of the VU abandonment rate is more narrow.

5.2 RESPONSE TO A PLANNED PEAK, THE CASE OF RESCUE CALLS DURING BASTILLE DAY.

– Contrary to the previous subsection, we here review the case of an event for which the massive arrival of calls is predictable. Every year, the night of Bastille Day (French National Day) between the 14 and the 15 July receives much more rescue calls than usual. We have already detailed in the Section 6.2 of Chapter 7 the throughput of calls on this particular night in the year 2019. We are going to reuse these profiles in simulation to evaluate the response of our call center models. We also explained before that the handling times of the calls were significantly shorter this particular night on the rescue lane; this is the proof that firefighters resorted to much more standardized procedures, adapted for Bastille Day calls.

In terms of loads, this is an important element to consider. Even though we noted that the throughput intensity of Urgent calls was (at the peak of the night) seven or eight times higher than a usual night, the corresponding instruction times were typically divided by two. As a result, the total load weighing on agents for rescue calls is “only” multiplied by four (see again the Section 2.4 for our introduction on loads).

For our simulations of “Bastille night”, we then naturally use the above throughputs and service times. We have also chosen to modify the breakdown of calls abandoned after the voice prompt (actually there were very few of them this night), to allocate 80 % of them on the U gravity, 15 % of NU calls and finally 5 % of VU calls. The simulated time-window goes from 10 PM (14 July evening) to 4 AM (15 July morning), which covers almost the entirety of the peak.

The Figure 8.33 gives the performance of single-tier and bilevel systems facing this particular event in terms of global and VU abandonment rates of calls received on line 18. It uses the same models than Figures 8.9 and 8.13. For the bilevel system, the optimal repartition of a given total staffing across the two levels is still obtained by optimizing a mixed criterion (first minimizing VU abandonments, then all abandonments, then the waiting time). We check in particular that this mixed criterion allows to achieve much better VU abandonment rates for the bilevel layout compared to a single-tier layout, and using the iso-staffing comparison. The two architectures become equivalent beyond 21 or 22 operators in total (at this point, the optimal breakdown in the bilevel is to take 7 agents at LVL 1 and 14 or 15 agents at LVL 2). Yet, we note that the single-tier

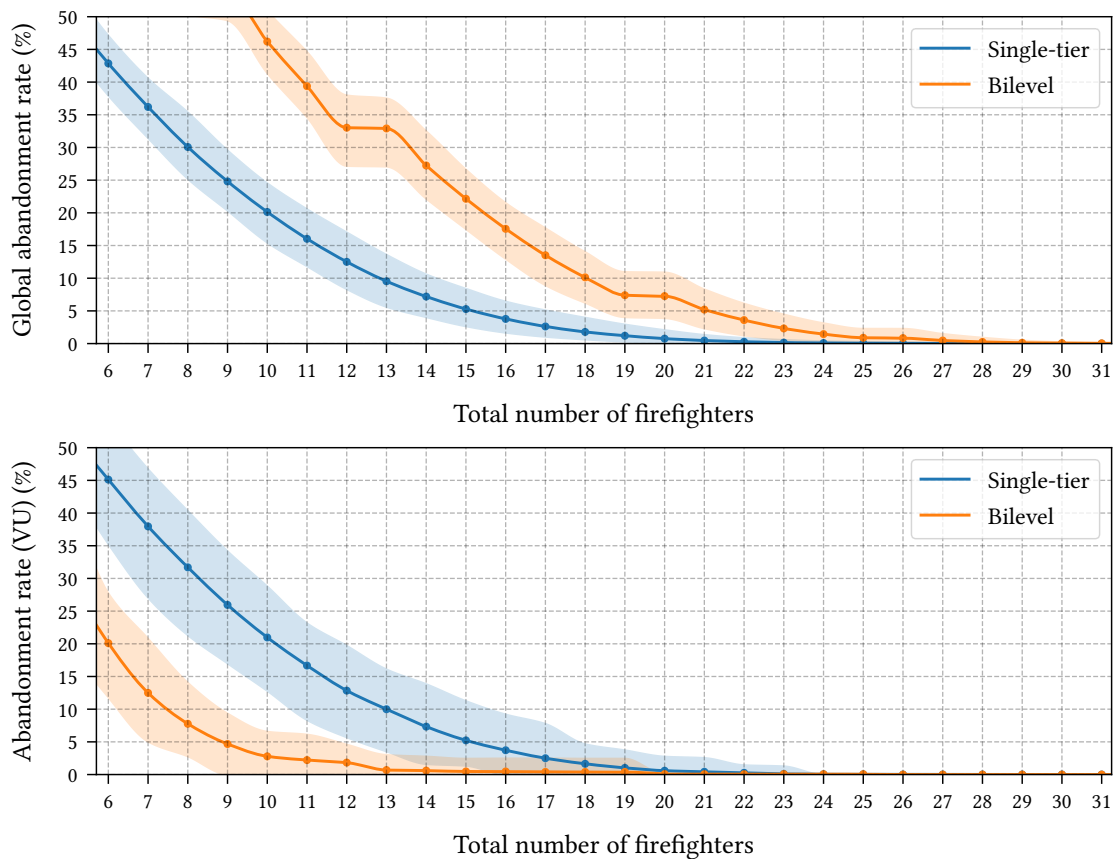


Figure 8.33: Comparison of the single-tier and bilevel systems in the night of Bastille Day 2019 (flows of lines 18 and 112, between 10 PM and 4 AM, i.e., 630 calls/h)

system is always better as far as the global abandonment rate is concerned (for the iso-staffing comparison). This orange curve could be improved a lot by increasing the proportion of agents at LVL 2, but this would alter a bit the abandonment rate of VU calls, and the latter has priority in our mixed criterion.

In view of the conclusion of the Section 2.4 which establishes that the loads ratio $(LVL\ 2)/(LVL\ 1)$ is a key indicator characterizing the good behaviour of the bilevel layout, it is relevant here to recompute this ratio during Bastille night. From 10 PM to 4 AM during this exceptional event, we have a typical load of 3.00 E at LVL 1 via the number 18, of 0.47 E at LVL 1 via the number 112, and of 10.49 E at LVL 2 of BSPP. This yields a ratio of roughly $R \simeq 3$, significantly lower than the standard ratio of the rescue lane that is typically equal to 4.

This decrease of the loads ratio partly explains why the global abandonment rate in the bilevel system is not that good compared to the single-tier. However, the role of the proportion of calls transferred at LVL 2 is also at stake. Indeed, in the standard regime (average day-to-day throughputs), only 34 % of calls from lines 18 and 112 are qualified U or VU and are handed over to LVL 2 (this fraction is very much lowered down by the calls from 112, see again the Table 7.2). During Bastille night though, this proportion goes to 65 % due to the massive arrival of U calls on line 18. A bilevel organization that is tuned to not lose any VU call (typically giving the performance of the Figure 8.33) will yet tend to put an important number of agents at LVL 1 for all calls to be quickly picked up and then leverage on the priority mechanism at LVL 2 to prevent VU abandonments. Now that two thirds of calls require an instruction at LVL 2 (versus only one third usually), this allocation of agents do not let enough resource to quickly handle all U calls at LVL 2, that are in important amount and do not have priority, hence higher abandonments. We easily check by digging in the simulation results that the global abandonment rate in the bilevel system shown in the first chart of the Figure 8.33 is almost exclusively due to U calls being hung up between the two levels – and not to calls being abandoned before LVL 1 and their initial qualification. We retrieve here once again the structural choice to protect Very Urgent calls before all the others, even if it means loosing more non-prioritized calls.

CONCLUSION AND PERSPECTIVES

As a conclusion of this work, we summarize the main contributions of the thesis in both the theoretical and applied fields. Then, we discuss a few problems that arose in this journey and were left unsolved.

Contributions rundown

It is always valuable in general to build bridges between apparently distinct pieces of mathematics, as it offers new perspectives, and sometimes new results on the initial fields. In this sense, the correspondence theorems between continuous-relaxed Petri nets and semi-Markov decision processes that we presented in Chapter 4 hold a singular position in our work. These results consolidated a lot the observations first made by Cohen, Gaubert and Quadrat in [CGQ95], by rigorously covering the general case of Petri nets with inputs and with underlying multichain structure.

We recall below our main contributions, all related in different ways to the piecewise affine dynamics of Petri nets or SMDPs.

MONOTONIC PETRI NET DYNAMICS. – We have extensively used the previous correspondences to describe the dynamical properties of monotonic Petri nets. In the continuous-relaxed setting, we have precisely characterized the throughput vector (Corollaries 4.4 and 4.5) and we proved the convergence of the trajectories towards the associated constant-rate function up to a bounded term (Corollary 4.2), regardless of the initial conditions. This was done under quite weak assumptions on the holding times of the net, following the efforts made in the Chapter 3 to allow most sojourn times of SMDPs to be null. We obtained a complete characterization of the congestion phases of the systems in terms of their resources, represented by a polyhedral complex. We showed that this tool could be used in practice for staffing recommendations (Section 6 of Chapter 4).

We have been able to present advances of interest on the study of transient regimes of Petri nets and correspondingly on the convergence of the value iteration algorithm for SMDPs in stochastic shortest path configuration. We have indeed established for the latter the convergence of the value towards a constant cost vector when the horizon approaches infinity (Theorem 3.41). Once again, this property was demonstrated regardless of the initial conditions (or terminal costs). We have characterized the ultimate contraction rate of the dynamics in this case (Theorem 4.12), which allowed us further to exhibit settings where the convergence occurred in finite time (Theorem 4.13). This property is key for organizations modeled by Petri nets, and we probably would not have been able to uncover it without the correspondence results. We have computed

explicit upper bounds on convergence times for a class of hierarchical Petri nets (Theorem 4.19), which is another performance estimate of great practical importance. In the emergency call center application, this was used to bound the absorption time of a peak of calls. We point out once again that many of these results are also new in the simplified framework of Markov decision processes, that uses a logical time instead of the physical time.

The above results also have consequences on the study of discrete Petri nets, in which tokens and firings remain integer quantities. We have quantified the deviation between the trajectories of these nets and their continuous-relaxed counterpart and showed that it was bounded for similar inputs and preselection routings (Theorem 2.19). This provided a bound between the throughput vectors of discrete and continuous-relaxed nets (Remark 4.3). We established that the approximation of discrete Petri nets by continuous-relaxed nets was asymptotically precise when scaling the number of tokens navigating in the net, even in the presence of many nondeterministic phenomena (Theorem 2.23). This result is an important contribution for the comparison of both types of models, and legitimizes the use of the above tools to study the discrete framework as well.

SYSTEMATIC ANALYSIS OF CONGESTION PHASES. — Our Chapter 5 was devoted to the presentation and illustration of a methodology to study the congestion phases of systems described by piecewise affine dynamics, typically arising from Petri nets. Using the tools of [ABG15], we have computed such polyhedral complexes even for non-monotonic models. We found that counter-intuitive and somehow paradoxical behaviours could arise in degraded phases due to the priority rules. We emphasized in this chapter the systematic and modular characters of this technique, still leading to the computation of minimum staffing recommendations or layout comparisons.

We have applied this methodology to several models of emergency call centers. As far as performance are concerned, we provided in each case the minimum number of agents to reach a congestion-free regime. Architecture-wise, we detailed many elementary mechanisms that it is possible to build with priority rules: allocation of different flows with distinct priority, reservoir system, asymmetrical or symmetrical assistance of resources... we showed that the comparison of multiple alternatives involved the computation of piecewise affine hypersurfaces.

COMPARISON WITH SIMULATED DISCRETE-EVENT SYSTEMS. — The Part II of the thesis allowed us to validate our theoretical approach based on continuous-relaxed Petri nets, as we showed that our throughput estimates were satisfying approximation of the real processing rates for simulated call centers. We retrieved by these numerical methods the congestion phases predicted by the theory.

In view of applications to real-life emergency call centers, this case study turned out to be very fruitful. We have been able to demonstrate multiple advantages of a well-tuned two-tier layout: better protection given to the very urgent calls, improvement of the patience levels, better resistance to peak of calls, diminution of the minimum number of skilled resources, etc. We exhibited the appropriate conditions to respect for a two-tier organization to be more efficient than a single-tier one, by showing that the instruction at the first level of a two-tier architecture should remain quick compared to the second one. We also measured the resource savings obtained by merging different pools of agents, as well as enabling synergies between two separate groups of operators (for example to pick up the calls from number 112).

PROPERTIES OF POSYNOMIAL SYSTEMS. — Our focus in Chapter 6 on simplified systems that arise from fixed-points equations associated with Petri net dynamics allowed us to show new results on classical posynomial systems and tropical posynomial systems. We established in both cases that the general associated decision problems were NP-hard, but we identified tractable classes of instances. These involved purely geometric conditions on the relative positions of the posynomial exponents: the pointed character of the configuration of exponents and the existence of a colorful vector. We showed under these conditions that tropical posynomial systems were equivalent to Markov decision processes and that they could be solved by linear programming. In comparison, we proved under the same assumptions that classical posynomial systems could be solved by geometric programming. We gave necessary conditions for the set of colorful vectors to be nonempty, and conversely, we stated sufficient nonvacuity conditions under the form of a conjecture, that we proved for three or less bodies.

Open problems and perspectives

Let us now come back on several questions that remain open at this stage, and that could be addressed in future work.

MORE GENERAL SEMI-MARKOV FRAMEWORKS. — We discussed at the end of the Section 3.2 in Chapter 2 and also in the Section 2.8 of Chapter 3 that our techniques developed for Petri nets and semi-Markov decision processes with constant holding and sojourn times could carry over to a more general framework with stochastic times. However, this was mainly restricted to cases where the distributions of these times did not contain zero in their adherence. Although this is a rather mild requirement in practice, it remains a problem of interest to determine the minimal assumptions on the distributions of sojourn times that allow one to prove that the value function of SMDPs is still defined, and verifies well-posed dynamic programming equations. We explained in the Section 2.8 of Chapter 3 that addressing this issue involves dealing with switched renewal equations, or Volterra integral equations of the second kind. We reckon that several standard results and proof techniques from this field could be applied to our setting.

Another related question, and that could sound natural in the SMDP community, would be to extend our results to SMDP configurations with non finite state and action spaces. It was indeed an appropriate setting to study the problems arising from Petri net models, but many other fields of performance evaluation lead to SMDP model on infinitely countable or compact structures. As recalled in the Chapter 3, the transition probabilities at stake may also depend upon the time. By analogy with the rest of the literature, we already know what would be the candidate dynamic programming equations in these more expressive settings. Yet, the challenge is to adapt the treatment of the non-Zeno behaviours and the convergence results established for the value iteration algorithm, which made use at some points of the finite structure of the state-action space.

Following the “strong law of large numbers” for Petri nets established in Theorem 2.23, it would also be desirable to derive large-deviation results or Chernoff bounds to better understand the difference between discrete and continuous-relaxed dynamics of Petri nets when stochastic events are involved. This echoes risk-sensitive or quantile-based criteria in the world of Markov and semi-Markov decision processes, a field that is still very active nowadays.

TRANSIENCE BOUNDS OF MONOTONIC DYNAMICS. — Our study of the transient regimes and input catch-up of Petri nets (or equivalently the convergence of the value iteration for SMDPs

in SSP configuration) in Chapter 4 led us to the computation of explicit transience bounds. We explained why the obtained expressions featured reasonable terms and ratios, however, they were obtained after some simplifications. As a consequence, we reckon that there is room for improvement on the result stated in Theorem 4.19. The subsequent Proposition 4.21 features a refined bound but only for the example (EMS-A). We illustrated there that a symbolic propagation scheme on affine stationary regimes (reminiscent of the policy iteration algorithm) may be worth-studying. We think that the field of network calculus could also bring appropriate techniques to tackle this question.

More generally, it is of interest to better understand the second-order properties of Petri nets or SMDP trajectories. Our study of the input catch-up may be thought of as the easiest problem in this very hard class! Similar results in the field of Markov decision processes teach us that the periodic character of SMDP classes plays an essential role at this scale, which suggests further that arithmetic conditions need to be taken into account. Working in the semi-Markov setting may in some cases bring a regularizing effect (at least when the time distributions are absolutely continuous) that cannot occur in the usual Markov case, where all the sojourn times are integer.

PROPERTIES OF NON-MONOTONIC DYNAMICS. — Although the core of our theoretical contributions on Petri net dynamics was brought to monotonic models, we have repeatedly recalled in Chapters 1 and 5 the difficulties raised by the introduction of priority rules – intrinsically non-monotonic. We have exhibited with the Petri net of Figure 1.14 a priority layout yielding ill-posed dynamics, partly due to our wish to allow null holding times for some places of the net. It would be an interesting problem to find minimal well-posedness conditions, under which we could guarantee that a unique trajectory arises from the dynamical equations associated with non-monotonic models.

Yet, the most important problem that remains unsolved for this class of nets is the question of the convergence of the growth-rate of trajectories towards the throughput computed by solving min-plus lexicographic systems on germs of affine functions. Indeed, even though these analytical expressions are very often reasonable and are good throughput candidates, it was shown in [ABG15] that trajectories could fail to converge towards these expected limits, owing to some arithmetic phenomena.

There are still open problems about the solutions of the lexicographic systems on germs themselves. We may indeed have expected that the counter functions associated with transitions having the highest degrees of priority remained nondecreasing functions of the resources. Yet, the paradox established in Section 3 of Chapter 5 proves that this is not true. We think that this negative result can be due to a certain layout of priorities, where high priority transitions are indirectly served by low priority transitions. It is possible that ruling out such patterns could allow one to show that some properties of monotonic models actually carry over the non-monotonic framework.

COMPUTATION OF THE THROUGHPUT COMPLEX. — In the Chapter 5, we have presented the congestion phases of several models of non-monotonic Petri nets, with up to ten full-dimensional polyhedral cells in \mathbb{R}^3 . These regions were determined by hand, and it is a problem of practical interest to determine the computational cost of obtaining all these phases in general, starting from a min-plus lexicographic system on germs of affine functions. Recall that such systems can feature conditional equations, some terms being switched on or off depending on the values of other throughputs. We have shown in addition that multiple policies could determine phases of the complex that were not full-dimensional. One would want to design an

algorithm to find the full-dimensional phases of such systems while avoiding the enumeration of all the policies, since there are exponentially many of them. Of course, one would need to use symbolic techniques to obtain the analytical expressions of the throughput in each cell and the formal inequalities that delineate the validity regions. Further, it would be desirable to perform computations on the resulting polyhedral complexes to compare the congestion phases of different models. More generally, these problems are tied to the computation of so-called tropical varieties [JSV17] and the computation of mixed subdivisions of polytopes [DLRS10].

MIXING POLYHEDRAL COMPLEXES WITH PROBABILISTIC ESTIMATES. — One disadvantage of the approach based on continuous-relaxed Petri nets is that it does not provide probabilistic estimates such as the average waiting time of tokens, probability for a token to wait, etc. This is inherent to the fact that we focus on the long-run behaviour of fluid models. Indeed, for a given place of the net in such a stationary regime, either no delay is experienced by the tokens of the place and this place remains empty, either an infinite amount of tokens pile up inside it and the delay goes to infinity. In particular this approach is blind to the notion of patience. As we discussed it in the Section 3 of Chapter 8, these effects typically cause our polyhedral formulas to be overapproximations of the real throughput. If there are enough resources, the difference is negligible, but on the contrary, a significant error can be made in the neighborhood of some frontiers between the different phases.

The throughput functions that are depicted in Figures 8.16 and 8.17 suggest that a corrective term could be added to our polyhedral formulas to better represent the real throughputs. We reckon that this error term could be computed by models and heuristics of queueing theory (that do not relax integrality constraints). This is in fact reminiscent of the concept of queueing networks, in which different blocks are cascaded and product-form solutions are computed.

PROBLEMS DERIVED FROM POSYNOMIAL SYSTEMS. — We finally point out that several problems related to our study of posynomial systems in Chapter 6 remain open. To begin with, we only studied posynomial systems on \mathbb{R}^n – or $(\mathbb{R}_{\max})^n$, while those featured in the Chapters 4 and 5 were defined on tuples. It would therefore be interesting to see if some of our techniques carry over to this more general case. In addition, we did not study the posynomial systems for which a colorful vector does not exist. In fact, the notion of colorful vector was introduced as a sufficient condition for the complementarity slackness equations in the proof of Theorem 6.7 to yield a solution of the initial posynomial system. We do not know if these sufficient conditions are necessary and there might be milder hypotheses under which we can still resort to linear programming techniques to solve tropical posynomial systems.

On the colorful interior subproblem, the Conjecture 6.23 remains open for more than four convex bodies. We believe that it could be proved by induction on the dimension of the ambient space, although the real challenge remains to go from geometric properties of sets given as intersection of convex bodies to more analytical notions involving the existence of a Perron vector for a given nonnegative matrix (actually, one needs to show that a so-called Z-matrix is an M-matrix, see [FP62]).

From the computational point of view, another question of interest is to determine the complexity of finding a colorful vector whenever the colorful interior is nonempty. We know that provided we have a strong separation oracle (see [GLS12]), this problem can be solved in oracle-polynomial time. However, separating a point from the frontier of the colorful interior is a nontrivial problem. The complexity status of computing the tangent hyperplanes to $n - 1$ convex

bodies in dimension n is itself unknown (the simpler constructive proof of their existence builds on Brouwer's fixed-point theorem [LvHK96]).

It is also worth-mentioning that the analysis of the colorful interior is of interest under different notions of convexity. Indeed, when the n convex bodies at stake are in the *tropical semifield* \mathbb{R}_{\max}^n and are *tropically convex*, we still obtain that there are n tropical hyperplanes (each one being tangent to $n - 1$ convex sets) delineating the colorful interior. The latter is in this case the interior of a tropical simplex. Remarkably, this colorful interior is the locus of apices of tropical hyperplanes that separate the initial convex bodies, in the sense that each one of these n bodies lies in exactly one of the n sectors of the hyperplane. As a result, the colorful interior problem solves the “Tropical Support Vector Machine” question in tropical machine learning [GJo8, ABG20a].

SUMMARY OF OUR WORK DURING THE COVID-19 PANDEMIC

The Covid-19 virus began to spread in France in January 2020. Like many other countries, the number of contaminations soared during the following weeks, until national lockdown came into effect on 17 March. The diffusion of the disease resulted in a rapid surge of emergency calls passed to the *Service d'Aide Médicale Urgente* (SAMU) via the national emergency number 15 (we refer once again to our presentation of French emergency services in the [Introduction](#)). On 13 March, 26,000 emergency calls arrived in the four SAMU of the *Assistance Publique – Hôpitaux de Paris* (AP-HP), four times the usual demand. This same day, we were contacted by AP-HP to provide assistance to the SAMU teams of Paris area in forecasting their staffing and logistics needs. To fulfill these missions, we spent the first month of lockdown at the *SAMU de Paris*. This appendix describes the work that we have done during this period and some subsequent projects that were carried out to predict the evolution of the pandemic.

1 Adaptation of the organizations

The SAMU centers of Paris area began to modify their response procedures around the 15 February, when the demand started to rise significantly. We refer to the feedback article presented by Télion, Marx, Dautreppe and Carli [[TMDC20](#)] for the case of SAMU 75 (city of Paris), see also [[LGV⁺20](#)] for the case of SAMU 93. During our time spent with the different teams, we had the opportunity to witness the huge logistic adaptation of the SAMU and AP-HP, who quickly set up tens of additional computer and phone workstations. Many reinforcement were also called upon. As far as call handling procedures are concerned, several modifications were brought to the usual system (of which a simplified version is described in the [Section 1](#) of [Chapter 4](#)), including:

- the **duplication** of the call handling circuit, in order to keep the regular inbound flow of urgent calls on one hand and the Covid-19 calls on the other hand. Indeed, the usual flow of calls arriving at the SAMU is well-known to be polymorphic and to require skilled agents or physicians to address it. Instead, the Covid-19 calls were of a much more homogeneous nature, and the less severe requests (accounting for a large volume of calls) could be addressed by less trained resources, such as medical students who volunteered to help the SAMU.
- the roll-out of an **interactive voice response** (IVR) in the SAMU 75, SAMU 92 and SAMU 93, consisting in a short message asking the patient to type the number 0 on his/her phone

if he/she was calling for Covid symptoms without respiratory difficulties. In this case, the patient was sent to the lane of Covid calls; otherwise he/she remained in the regular lane. Reorientation in the other lane was possible at all subsequent steps of the instruction in case the responding agents detected the initial choice was wrong.

- the recourse to a “**hang up and call back**” procedure, that was used for mild Covid calls in order to call back the patient either by the same agent or from a different workstation. The rationale behind this process was to free some inbound lines that are in high but limited amount, and spread the workload on outbound lines as well. This ensured that more inbound calls could be handled at the same time.
- the sending of **text messages** (SMS) on the cell phones of patients who presented very mild Covid-19 symptoms or no symptoms at all. These text contained instructions and allowed the patients to reply if they wanted to be called back by the SAMU (which happened in one case out of three).

We refer to [TMDC20] for the full details of these procedures in the case of SAMU 75. The SAMU 94 did not implement an IVR to perform the triage of the calls between the regular and Covid-19 lanes. Instead, they set up a first level of pick up by a dedicated pool of agents, who then very quickly dispatched the calls.

2 Contributions to the dimensioning of response systems

- 2.1 INNOVATIONS MODELING AND STAFFING COMPUTATION.** — The core part of our work with the SAMU has been to estimate the staffing needs of the different call centers, depending upon each organization topology and parameters. To this purpose, we reused the methodology detailed in this thesis, namely the representation of a real organization by means of a Petri net (Chapter 1), the approximation of its behaviour by its continuous-relaxed counterpart (Chapter 2), and the computation of its congestion-free phase (Chapters 4 and 5).

We actually enriched a bit our approach by implementing interval arithmetic in our estimates, due to the variability of the exogenous parameters and the need to include the uncertainty intervals on the epidemic growth rates computed by epidemiologists. We delivered analytical formulas and numerical applications in different scenarios designed in constant collaboration with the physicians of the SAMU. We illustrate the Covid-19 response models that we have studied in the two Petri nets below. The net depicted in Figure A.1 represents the organization of the SAMU 75 (city of Paris). The right part is similar to the Petri net of Figure 4.2, with a first instruction by a Medical Regulation Assistant (MRA) and a second possible instruction by a General Practitioner (GP) or an emergency physician. The MRA circuit is still depicted in blue, the GP circuit in purple and the emergency physician circuit in red. Consistently with the above description, the whole call handling circuit has been duplicated after the initial IVR. If the patient indicated that his/her call was about mild Covid-19 symptoms, he/she did not go in the regular lane but in the Covid lane, depicted in green. The calls of this room were picked up by another pool of MRAs, the so-called “Covid MRAs”. They could hand over the calls directly to the emergency physicians if the patient underestimated the severity of his/her symptoms. Otherwise, they could transfer the call to a dedicated room with different levels of priority. This room, not fully depicted but represented in orange, was staffed with reinforcement GPs and medical students to address the non severe Covid cases.

The net depicted in Figure A.2 represents the organization of the SAMU 94 (department of Val-de-Marne). As explained in the previous section, no IVR was implemented in this department, which explains why the design looks quite different from the one of Figure A.1. This model was built on the ground model (EMS-B) of Figure 5.6, that was studied in the Section 3 of Chapter 5. In the Figure A.2, we use double arrows to stand for already classified flows of calls, so that we may treat some of their calls with priority relatively to others.

During the Covid-19 pandemic, the SAMU 94 used a first level of instruction (depicted in yellow in Figure A.2) to quickly qualify the request of the patient. If it was either a regular call or a Covid-19 call with severe symptoms, the call went in a regular lane (such as the other SAMU). Otherwise, it was sent in a different lane dedicated to mild Covid calls. For the latter, a longer response was given by a level-two MRA and the “hang up call back” scheme was used (represented with the boxed cross). Afterwards, the patient was called back by a GP.

The calls sent in the regular lane were further handled by usual MRAs, possibly in a priority order specified by the first-answering MRAs. The patients could then be transferred to usual GPs, or to two pools of emergency physicians, whether the call was about Covid or not. The two pools of emergency physicians used the principle of monitored reservoir studied in the Section 3 of Chapter 5. A “shortcut link” was also planned (in dashed yellow) if the first-responding agent detected a vital threat and an immediate need to reach the emergency physician and bypass the intermediary layers. Note that even when not explicitly depicted, this procedure always exists.

Note that the Petri nets of Figures A.1 and A.2 were both simplified representations of the real systems in the sense that the role of agents not directly involved in the instruction of inbound calls were not depicted. This includes the MRAs who actually coordinate the fleet of ambulances or Mobile Intensive Care Units.

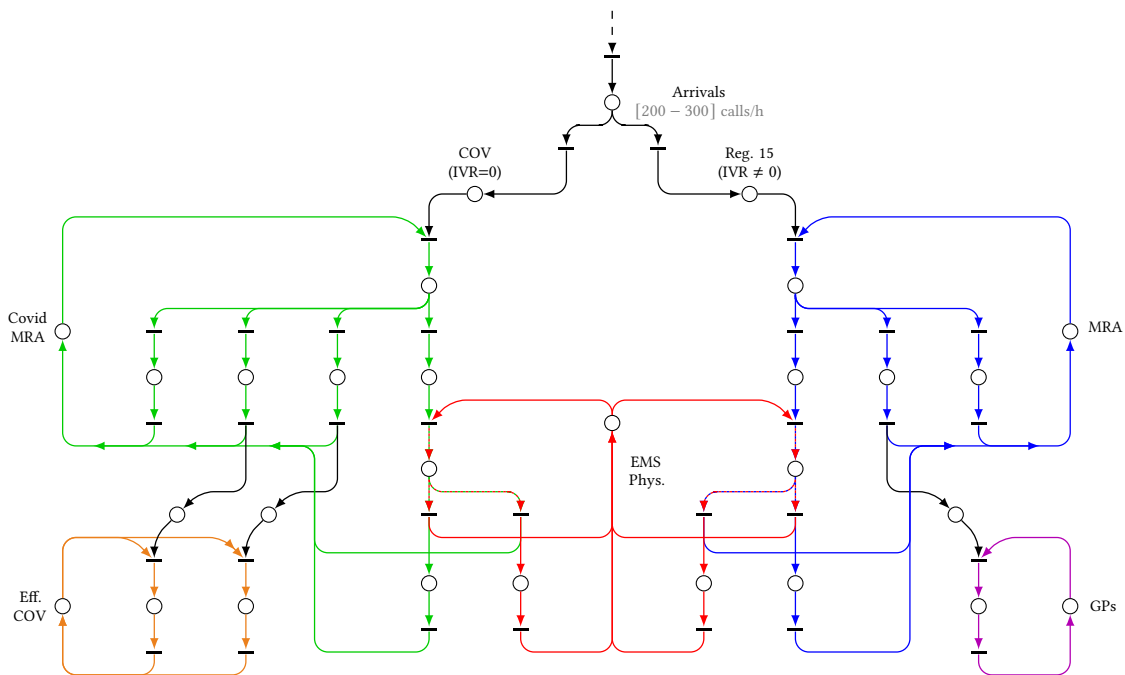


Figure A.1: Simplified Petri net of the SAMU 75 organization in the middle of the Covid-19 crisis

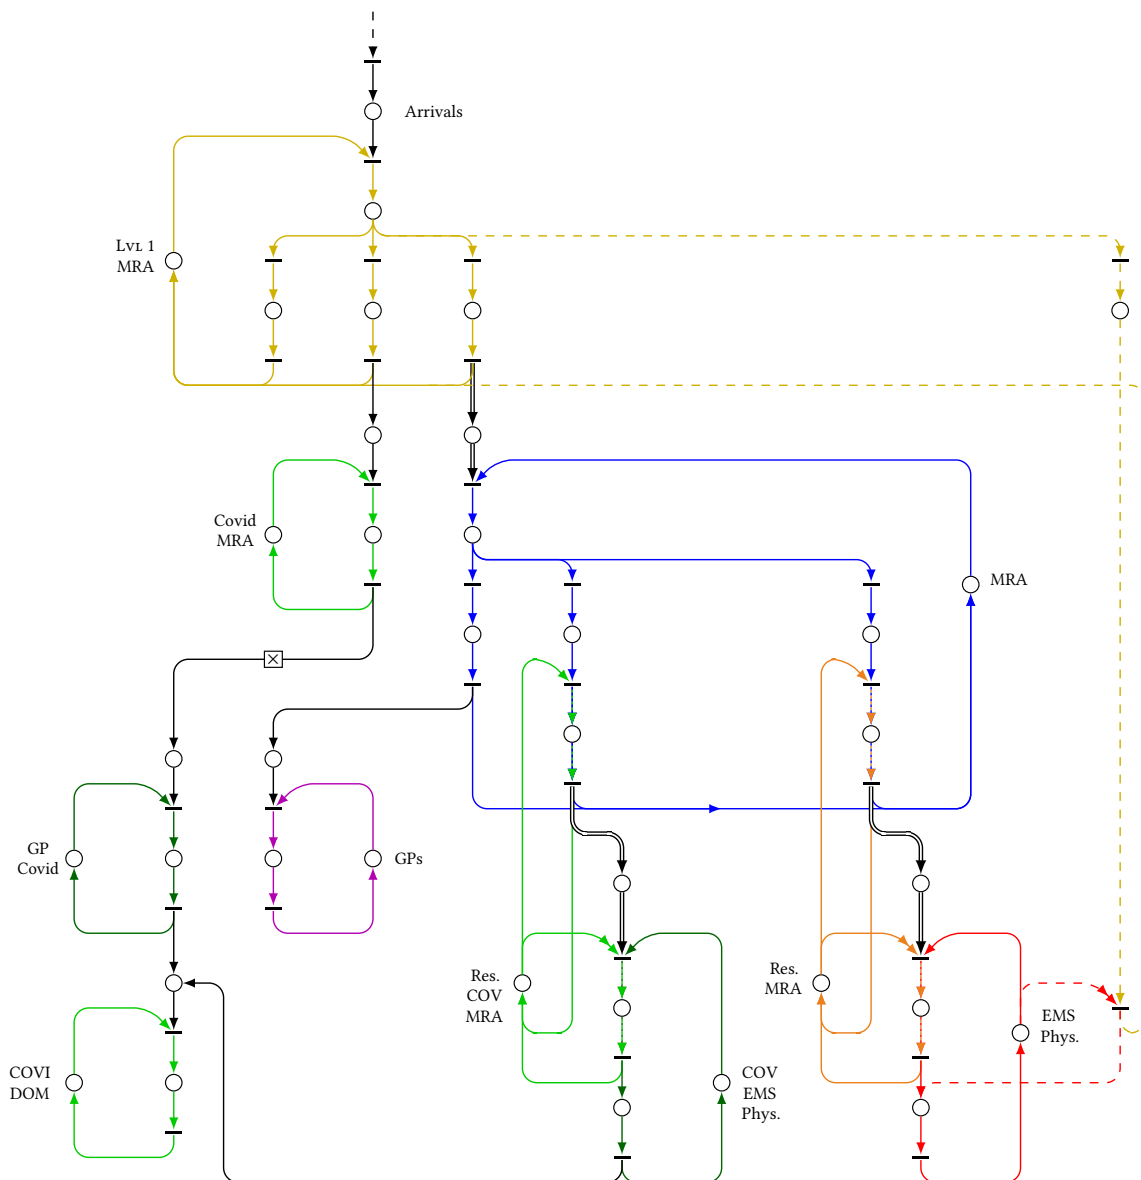


Figure A.2: Simplified Petri net of the SAMU 94 organization in the middle of the Covid-19 crisis

2.2 DATA ANALYSIS AND CONGESTION DIAGNOSIS. — Another early part of our work has been from March 13 to understand the magnitude of the inbound flow of calls and to provide day-to-day reporting to the AP-HP. We reused the data analysis tools that we developed in 2019 for our collaboration with the four SAMU of AP-HP (and who led to the reports [ABG20b] delivered in January 2020). We also performed data fusion with other sources of information on the calls, in particular those made available by the telephone provider Orange.

It allowed us to identify some congestion links upstream of the SAMU centers or uneven load balancing between the different technical equipments on 13 March. These elements helped the logistic teams to strengthen the calls routing facilities and were used to assess that the consolidation operations were successful.

In the following days, we used our data analysis codes to measure the characteristics of the new procedures implemented by the SAMU (service times and experienced wait in the Covid lane, overall proportion of Covid-related calls and degrees of severity...). These parameters were used in our modified models to compute staffing recommendations, as presented in the previous section.

3 Contributions to the forecast of the epidemic

We summarize in this section a subsequent work that we have realized on the forecast of the Covid-19 spread in Paris area and the analysis of the different phases of evolution, using the SAMU data. Indeed, the number of calls passed to the SAMU showed to be among the earliest indicators of the Covid dynamics [Cov20], with up to ten to fourteen days of advance on critical care admission. This work led to the publication of the article [GAA⁺20]. This contribution was featured in the article [Cov20], which discusses similar aspects on a broader basis and is intended to a medical audience.

- 3.1 CLASSIFICATION OF CALLS.** — The first step of the work was about data categorization. Using the information filled by the SAMU agents and decisions taken for each call, we classified together with the SAMU physicians the Covid-19 calls in three categories: the calls resulting in the dispatch of a Mobile Intensive Care Unit, those resulting in the dispatch of an ambulance and those resulting in no dispatch decision. If $C(t)$ denotes the number of contaminations on day t , that we cannot measure directly, we postulate that each one of the three previous epidemic observables Y is a fraction of the daily number of contaminations shifted by a certain aggravation delay, i.e., $Y_i(t) = \pi_i C(t - \tau_i)$ with i in $\{1, 2, 3\}$. This form is justified by the mathematical analysis which follows.
- 3.2 A MULTI-COMPARTMENT TRANSPORT PDE MODEL.** — To analyze the evolution of these observables and take into account the presence of delays in the contamination process, we use an SEIR partial differential equation model, inspired by the model of Kermack and McKendric [KM27]. We assume that the epidemic is in an early stage phase, so that the infected population is much smaller than the susceptible population. This allows us to linearize the epidemic dynamics, and to introduce an evolution operator that is order-preserving (such as those considered in the Section 2 of Chapter 2). Results from Perron–Frobenius theory (see for instance [Per06]) entail that provided no changement in sanitary measures alter the epidemic evolution, the number of newly contaminated individuals grows exponentially. More generally, this approach applies to the quasi stationary regimes.
- 3.3 PIECEWISE AFFINE APPROXIMATION OF THE LOGARITHM OF OBSERVABLES.** — When the sanitary policies change (beginning or end of lockdown, mandatory masks outdoors, etc.), this modifies the growth or decay rate of the epidemic. As a result, the epidemic observables can be modeled by piecewise affine functions (with as many linear pieces as there are phases of sanitary measures). Determining the evolution rates on the different phases amounts to perform a best piecewise linear approximation. In order to compute a robust approximation, we minimize ℓ_1 -norm criteria, which is made difficult by the fact that the objective function is non-smooth and non-convex.

We derived from this methodology epidemic alarms based on doubling times, designed to trigger alarms when the probability of undesirable events (like the resumption of the epidemic) exceeded a chosen threshold. We refer again to [GAA⁺20] for the complete methodology.

4 Study of the Interactive Voice Response of SAMU 92

We had the opportunity to work in collaboration with the SAMU 92 (Hauts-de-Seine) on their IVR data. This automatic voice server offered the caller to either press the “0” key for Coronavirus Syndrome with no respiratory difficulties; or stay on line for any other reason. The main question was to determine if the implementation of such a system was efficient and safe. To this purpose, we first collected the phone logs of the SAMU 92 from March the 19th to April the 26th, giving the IVR choice among: “Yes” if the patient chose to press the “0” key after listening to the introductory message, “No” if the patient did not press any key and listened to the whole message, or “May be” if the patient did not press any key but the call was picked up quickly (so the patient *may have wanted* to press the “0” key). We merged these data with medical regulation files to attach additional information on the patient such as the age, gender, profile (caller her/himself or third-party) and also the decision made by the physician.

The Table A.3 below compares the groups IVR Yes (2,846 calls) and IVR No (12,111 calls). We observe that the IVR allowed almost 15 % of these calls to be directed to a specialized provider where they waited to be processed by staff trained within a few days to deal exclusively with COVID cases. This led to decrease the number of calls handled by the conventional dispatchers and gave them more time to handle more severe emergency calls. Moreover, because only 0.1 % of “IVR Yes” calls needed an Advanced Life Support ambulance, we can assume that the use of IVR is safe (very few severe cases within this flow).

The above material is drawn from a joint work with DR Margot Cassuto, DR Jeremie Boutet, DR Guillaume Douge, DR Anna Ozguler, DR Michel Baer and DR Thomas Loeb from SAMU 92, *Assistance Publique – Hôpitaux de Paris*. A conference communication on this topic was accepted to the European Emergency Medicine Congress 2022, another one was submitted to the Resuscitation Science Symposium 2022, in addition to an article being prepared for a medical journal.

	IVR Yes N = 2,846	IVR No N = 12,111	P
Mean age	41.3 years old (SD 23.1)	44.7 years old (SD 23.5)	< 10 ⁻³
Gender			< 10 ⁻³
– Male	971 (34.1 %)	4,970 (41.0 %)	
– Female	1,284 (45.1 %)	6,381 (52.7 %)	
– Not collected	591 (20.8 %)	760 (6.3 %)	
Caller (patient her/himself)	2,340 (82.3 %)	6,937 (57.3 %)	< 10 ⁻³
Decision			< 10 ⁻³
– ALS* ambulance	2 (0.1 %)	141 (1.1 %)	
– BLS** ambulance	181 (6.3 %)	2,659 (22.0 %)	
– No transport	2,655 (93.6 %)	9,309 (76.9 %)	

*ALS: Advanced Life Support, **BLS: Basic Life Support

Table A.3: Comparison of patients’ age, gender, caller profile and decision taken at dispatch center among IVR-Yes and IVR-No

INDEX

A

affine regimes (invariant halfline), [122](#), [149](#), [165](#), [192](#)

asymptotic expansion results, [73](#), [81](#), [125](#), [129](#), [149](#), [163](#)

C

call center architectures

aggregated and departementalized systems, [332](#)

bilevel, [194](#), [242](#)

escorted transfer, [136](#), [195](#), [202](#), [280](#), [296](#)

monitored reservoir, [201](#)

multi-job layouts, [326](#)

mutual support of agents, [215](#)

single-tier, [57](#), [242](#)

call center features

agents shifts, [261](#)

arrivals throughputs, [255](#)

load, [158](#), [307](#)

occupancy factor, [185](#), [211](#), [300](#), [317](#)

patience, [138](#), [195](#), [275](#), [290](#)

service times, [264](#)

conflict, [24](#)

conflict-free, [61](#), [140](#)

congestion diagram, [155](#), [199](#), [204](#), [209](#), [213](#), [217](#)

control theory, [1](#), [62](#)

conversion rate, [53](#)

correlation of parameters in call centers, [268](#), [272](#), [274](#), [298](#), [329](#)

correspondence theorem, [140](#), [142](#), [148](#), [161](#)

D

discrete-event simulation, [2](#), [242](#), [290](#)

E

Erlang formulas, [4](#), [319](#), [320](#)

event graphs, [22](#), [134](#)

evolution operator, [64](#), [119](#), [142](#)

evolution semigroup, [65](#), [117](#), [142](#), [164](#)

F

fixed-point property, [104](#), [111](#), [127](#), [222](#)

French emergency call centers, [5](#)

PFAU, [6](#), [194](#), [241](#), [251](#), [290](#)

SAMU, [5](#), [136](#), [185](#), [201](#)

G

geometric convergence, [100](#), [169](#)

H

hierarchical systems, [173](#)

I

induction principle, [39](#), [110](#), [143](#)

invariant

P-invariant, [51](#)

stoichiometric invariant, [66](#), [72](#), [142](#)

T-invariant, [21](#), [22](#), [68](#)

K

Kaplan–Meier estimator, [275](#)

L

law of large numbers, 77, 81
lexicographic systems, 123, 150, 192
linear programming, 159, 223

M

Markov decision process, 92, 226
see also semi-Markov decision process

N

network calculus, 3, 181
nonexpansive mappings, 119, 125
nonlinear Perron–Frobenius theory, 104, 168, 171

P

peak of calls, 168, 281, 338
performance evaluation, 1
periodic phenomena, 65, 89, 125, 159
Petri net, 2, 18
 continuous-relaxed Petri net, 49
 controlled Petri net, 23
 earliest behaviour, 32, 54
 scaled Petri net, 69, 149
 simultaneously fireable transitions, 29, 52
 throughput, 148, 149, 153, 314
 timed Petri net, 31
piecewise affine mappings, 7, 62, 105, 122, 154, 220
Poisson process, 43
 nonhomogeneous Poisson process, 43, 74, 259, 292

Q

queues or queueing networks, 2, 69, 158, 185, 319

R

routing rules for Petri nets, 24
 dynamical allocation, 291
 preselection, 24

priority routing, 26, 190
proportional preselection routing, 51
race policy, 41

S

semi-Markov decision process, 92, 226
 average-cost, 123, 148
 discounted models, 92, 97, 107, 140, 226
 policy, 95
 stopping costs, 93, 98, 115, 141
 strategy, 95
 terminal costs, 93, 98, 106
 value function, 97
 value iteration algorithm, 108, 166
simulation, 290
stochastic shortest path, 127, 161
stochastic times, 41, 74, 120, 144, 154, 291

W

weighted sup-norm contraction, 105, 169

Z

Zeno
 non-Zeno Petri net, 34, 143
 Zeno behaviour for Petri nets, 33
 Zeno behaviour for SMDPs, 100

BIBLIOGRAPHY

- [AABG20] Marianne Akian, Xavier Allamigeon, Marin Boyet, and Stéphane Gaubert. A convex programming approach to solve posynomial systems. In *International Congress on Mathematical Software*, pages 241–250. Springer, 2020.
- [AAC⁺12] Bernardetta Addis, Roberto Aringhieri, Giuliana Carello, Andrea Grosso, and Francesco Maffioli. Workforce management based on forecasted demand. In *Advanced decision making methods applied to health care*, pages 1–11. Springer, 2012.
- [AAKD08] M. Salah Aguir, O. Zeynep Akşin, Fikri Karaesmen, and Yves Dallery. On the interaction between retrials and sizing of call centers. *European Journal of Operational Research*, 191(2):398–408, 2008.
- [AB06] Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Springer Science & Business Media, 2006.
- [AB09] Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.
- [AB17] Venkatachalam Anantharam and Vivek S. Borkar. A variational formula for risk-sensitive reward. *SIAM J. Control Optim.*, 55(2):961–988, 2017. arXiv:1501.00676.
- [ABCG21] Xavier Allamigeon, Marin Boyet, Baptiste Colin, and Stéphane Gaubert. Optimisation de la performance de centres de traitement d’appels d’urgence. Technical report, Inria & CMAP, École polytechnique, 2021.
- [ABFG⁺93] Aristotle Arapostathis, Vivek S. Borkar, Emmanuel Fernández-Gaucherand, Mrinal K. Ghosh, and Steven I. Marcus. Discrete-time controlled Markov processes with average cost criterion: A survey. *SIAM Journal on Control and Optimization*, 31(2):282–344, 1993.
- [ABG15] Xavier Allamigeon, Vianney Bœuf, and Stéphane Gaubert. Performance evaluation of an emergency call center: tropical polynomial systems applied to timed Petri nets. In *International Conference on Formal Modeling and Analysis of Timed Systems*, pages 10–26. Springer, 2015.
- [ABG17] Xavier Allamigeon, Vianney Bœuf, and Stéphane Gaubert. Stationary solutions of discrete and continuous Petri nets with priorities. *Performance Evaluation*, 113:1–12, 2017.
- [ABG20a] Xavier Allamigeon, Marin Boyet, and Stéphane Gaubert. Computing tropical SVM. Unpublished manuscript, 2020.

- [ABG20b] Xavier Allamigeon, Marin Boyet, and Stéphane Gaubert. Études des données de téléphonie (février 2019), modélisation et premières indications de dimensionnement. Technical report, Inria & CMAP, École polytechnique, 2020.
- [ABG21] Xavier Allamigeon, Marin Boyet, and Stéphane Gaubert. Piecewise affine dynamical models of Petri nets – application to emergency call centers. *Fundamenta Informaticae*, 183(3-4):169–201, 2021.
- [ABG22] Xavier Allamigeon, Marin Boyet, and Stéphane Gaubert. Computing transience bounds of emergency call centers: a hierarchical timed Petri net approach. *arXiv preprint arXiv:2202.02752*, to appear in the *proceedings of Petri Nets 2022 conference*, 2022.
- [ABP82] James R. Anderson, Leslie Bernstein, and Malcolm C. Pike. Approximate confidence intervals for probabilities of survival and quantiles in life-table analysis. *Biometrics*, 38(2):407–416, 1982.
- [AD94] Hassane Alla and René David. Petri nets for modeling of dynamic systems: A survey. *Automatica*, 30(2):175–202, 1994.
- [AD98] Hassane Alla and René David. Continuous and hybrid Petri nets. *Journal of Circuits, Systems, and Computers*, 8(01):159–188, 1998.
- [AEH04] Július Atlason, Marina A. Epelman, and Shane G. Henderson. Call center staffing with simulation and cutting plane methods. *Annals of Operations Research*, 127(1):333–358, 2004.
- [AFG⁺14] Xavier Allamigeon, Uli Fahrenberg, Stéphane Gaubert, Ricardo D. Katz, and Axel Legay. Tropical Fourier–Motzkin elimination, with an application to real-time verification. *International Journal of Algebra and Computation*, 24(05):569–607, 2014.
- [AG03] Marianne Akian and Stéphane Gaubert. Spectral theorem for convex monotone homogeneous maps, and ergodic control. *Nonlinear Analysis. Theory, Methods & Applications*, 52(2):637–679, 2003.
- [AG13] Marianne Akian and Stéphane Gaubert. Policy iteration for perfect information stochastic mean payoff games with bounded first return times is strongly polynomial. *arXiv preprint arXiv:1310.4953*, 2013.
- [AGGCG19] Marianne Akian, Stéphane Gaubert, Julien Grand-Clément, and Jérémie Guillaud. The operator approach to entropy games. *Theory of Computing Systems*, 63(5):1089–1130, 2019.
- [AGN11] Marianne Akian, Stéphane Gaubert, and Roger Nussbaum. A Collatz–Wielandt characterization of the spectral radius of order-preserving homogeneous maps on cones. *arXiv preprint arXiv:1112.5968*, 2011.
- [AGQS19] Marianne Akian, Stéphane Gaubert, Zheng Qu, and Omar Saadi. Solving ergodic Markov decision processes and perfect information zero-sum stochastic games by variance reduced deflated value iteration. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 5963–5970. IEEE, 2019.

- [AL04] Frédéric Adnet and Frédéric Lapostolle. International EMS systems: France. *Resuscitation*, 63(1):7–9, 2004.
- [AL05] Athanassios N. Avramidis and Pierre L’Ecuyer. Modeling and simulation of call centers. In *Proceedings of the Winter Simulation Conference, 2005.*, pages 9–pp. IEEE, 2005.
- [Alt99] Eitan Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.
- [AMBC87] Marco Ajmone Marsan, Gianfranco Balbo, and Gianni Conte. *Performance models of multiprocessor systems*. Mit Press, 1987.
- [AMBC⁺98] Marco Ajmone Marsan, Gianfranco Balbo, Gianni Conte, Susanna Donatelli, and Giuliana Franceschinis. Modelling with generalized stochastic Petri nets. *ACM SIGMETRICS performance evaluation review*, 26(2):2, 1998.
- [AMCB84] Marco Ajmone Marsan, Gianni Conte, and Gianfranco Balbo. A class of generalized stochastic Petri nets for the performance evaluation of multiprocessor systems. *ACM Transactions on Computer Systems (TOCS)*, 2(2):93–122, 1984.
- [Ban05] Jerry Banks. *Discrete event system simulation*. Pearson Education India, 2005.
- [BC63] Richard Bellman and Kenneth L. Cooke. *Differential–Difference Equations*. Academic Press New York, 1963.
- [BCMP75] Forest Baskett, K. Mani Chandy, Richard R. Muntz, and Fernando G. Palacios. Open, closed, and mixed networks of queues with different classes of customers. *Journal of the ACM (JACM)*, 22(2):248–260, 1975.
- [BCOQ92] François Baccelli, Guy Cohen, Geert J. Olsder, and Jean-Pierre Quadrat. *Synchronization and Linearity*. Wiley, 1992.
- [BD01] Maury Bramson and Jim G. Dai. Heavy traffic limits for some queueing networks. *Annals of Applied Probability*, pages 49–90, 2001.
- [Bel57] Richard Bellman. A Markovian decision process. *Journal of mathematics and mechanics*, 6(5):679–684, 1957.
- [Ber11] Dimitri P. Bertsekas. *Dynamic programming and optimal control* 3rd edition, volume II. Belmont, MA: Athena Scientific, 2011.
- [Ber12a] Patrick Bernard. The Lax–Oleinik semi-group: a Hamiltonian point of view. *Proceedings of the Royal Society of Edinburgh Section A: Mathematics*, 142(6):1131–1177, 2012.
- [Ber12b] Dimitri P. Bertsekas. *Dynamic programming and optimal control: Volume I*, volume 1. Athena scientific, 2012.
- [Ber12c] Dimitri P. Bertsekas. Weighted sup-norm contractions in dynamic programming: A review and some new applications. *Dept. Elect. Eng. Comput. Sci., Massachusetts Inst. Technol., Cambridge, MA, USA, Tech. Rep. LIDS-P-2884*, 2012.

- [BFG96] François Baccelli, Serguei Foss, and Bruno Gaujal. Free-choice Petri nets - an algebraic approach. *IEEE Transactions on Automatic Control*, 41(12):1751–1778, 1996.
- [BG01] Anne Bouillard and Bruno Gaujal. Coupling time of a (max,plus) matrix. In *Proceedings of the Workshop on Max-Plus Algebras, a satellite event of the first IFAC Symposium on System, Structure and Control (Praha, 2001)*. Elsevier, 2001.
- [BGF94] François Baccelli, Bruno Gaujal, and Serguei Foss. Structural, temporal and stochastic properties of unbounded free-choice Petri nets, 1994.
- [BGM00] Fabio Balduzzi, Alessandro Giua, and Giuseppe Menga. First-order hybrid Petri nets: a model for optimization and control. *IEEE transactions on robotics and automation*, 16(4):382–399, 2000.
- [BGM05a] Anne Bouillard, Bruno Gaujal, and Jean Mairesse. Throughput in stochastic free-choice nets under various policies. In *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 2131–2136. IEEE, 2005.
- [BGM⁺05b] Lawrence Brown, Noah Gans, Avishai Mandelbaum, Anat Sakov, Haipeng Shen, Sergey Zeltyn, and Linda Zhao. Statistical analysis of a telephone call center: A queueing-science perspective. *Journal of the American statistical association*, 100(469):36–50, 2005.
- [BHJ08] Imre Bárány, Alfredo Hubard, and Jesús Jerónimo. Slicing convex sets and measures by a hyperplane. *Discrete & Computational Geometry*, 39(1-3):67–75, 2008.
- [BL91] François Baccelli and Zhen Liu. Comparison properties of stochastic decision free Petri nets, 1991.
- [BL05] Eric Buist and Pierre L’Ecuyer. A Java library for simulating contact centers. In *Proceedings of the Winter Simulation Conference, 2005.*, pages 10–pp. IEEE, 2005.
- [BL19] Laurent Barret and Yves Lambert. Mission d’audit des 8 SAMU d’Île-de-France. Technical report, Agence Régionale de Santé d’Île-de-France, 2019.
- [BM98] François Baccelli and Jean Mairesse. Ergodic theorems for stochastic operators and discrete event networks, 1998.
- [Bœu17] Vianney Bœuf. *Dynamics of a two-level system with priorities and application to an emergency call center*. PhD thesis, École polytechnique, 2017.
- [Bon07] Blai Bonet. On the speed of convergence of value iteration on stochastic shortest-path problems. *Mathematics of Operations Research*, 32(2):365–373, 2007.
- [Bou14] Anne Bouillard. *Algorithms and efficiency of Network calculus*. PhD thesis, Ecole Normale Supérieure (Paris), 2014.
- [Bow00] Fred D.J. Bowden. A brief survey and synthesis of the roles of time in Petri nets. *Mathematical and Computer Modelling*, 31(10-12):55–68, 2000.
- [BP79] Abraham Berman and Roger J. Plemmons. *Nonnegative matrices in the mathematical sciences*. Academic Press, 1979.

- [Bra98] Maury Bramson. State space collapse with application to heavy traffic limits for multiclass queueing networks. *Queueing Systems*, 30(1):89–140, 1998.
- [Bru17] Hermann Brunner. *Volterra integral equations: an introduction to theory and applications*, volume 30. Cambridge University Press, 2017.
- [BS78] Dimitri P. Bertsekas and Steven E. Shreve. *Stochastic Optimal Control: The Discrete Time Case*. Mathematics in Science and Engineering 139. Elsevier, Academic Press, 1 edition, 1978.
- [BT91] Dimitri P. Bertsekas and John N. Tsitsiklis. An analysis of stochastic shortest path problems. *Mathematics of Operations Research*, 16(3):580–595, 1991.
- [Buc99] Peter Buchholz. Structured analysis approaches for large Markov chains. *Applied Numerical Mathematics*, 31(4):375–404, 1999.
- [Bur56] Paul J. Burke. The output of a queueing system. *Operations research*, 4(6):699–704, 1956.
- [BV04] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [BY13] Dimitri P. Bertsekas and Huizhen Yu. Stochastic shortest path problems under weak conditions. *Lab. for Information and Decision Systems Report LIDS-P-2909, MIT*, 2013.
- [CBFN13] Bernadette Charron-Bost, Matthias Függer, and Thomas Nowak. Transience bounds for distributed algorithms. In Víctor Braberman and Laurent Fribourg, editors, *Formal Modeling and Analysis of Timed Systems*, pages 77–90, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [CC88] Jacques Carlier and Philippe Chrétienne. *Problèmes d’ordonnancement: modélisation, complexité, algorithmes*. Masson, 1988.
- [Ç13] Erhan Çinlar. *Introduction to stochastic processes*. Courier Corporation, 2013.
- [CCS89] Javier Campos, Giovanni Chiola, and Manuel Silva Suárez. Properties and steady-state performance bounds for Petri nets with unique repetitive firing count vector. In *PNPM*, pages 210–220, 1989.
- [CDF91] Giovanni Chiola, Susanna Donatelli, and Guiliana Franceschinis. Priorities, inhibitor arcs and concurrency in P/T nets. In *Proc. of ICATPN*, volume 91, pages 182–205, 1991.
- [CEHMR12] Yudith Cardinale, Joyce El Haddad, Maude Manouvrier, and Marta Rukoz. Web service composition based on Petri nets: Review and contribution. In *International workshop on resource discovery*, pages 83–122. Springer, 2012.
- [CGP⁺94] Sylvain Cappell, Jacob E. Goodman, János Pach, Richard Pollack, and Micha Sharir. Common tangents and common transversals. *Advances in Mathematics*, 106(2):198 – 215, 1994.

- [CGQ95] Guy Cohen, Stéphane Gaubert, and Jean-Pierre Quadrat. Asymptotic throughput of continuous timed Petri nets. In *Proceedings of the 34th Conference on Decision and Control*, New Orleans, December 1995.
- [CGQ98] Guy Cohen, Stéphane Gaubert, and Jean-Pierre Quadrat. Algebraic system analysis of timed Petri nets. In Jeremy Gunawardena, editor, *Idempotency*, Publications of the Isaac Newton Institute, pages 145–170. Cambridge University Press, 1998.
- [Cha00] Cheng-Shang Chang. *Performance guarantees in communication networks*. Springer Science & Business Media, 2000.
- [CHB13] Bertrand Cottenceau, Laurent Hardouin, and Jean-Louis Boimond. Modeling and control of weight-balanced timed event graphs in dioids. *IEEE Transactions on Automatic Control*, 59(5):1219–1231, 2013.
- [Chl97] Edward Chlebus. Empirical validation of call holding time distribution in cellular communications systems. In *Proceedings of the 15th International Teletraffic Congress*, pages 1179–1188, 1997.
- [CL08] Christos G. Cassandras and Stéphane Lafortune. *Introduction to discrete event systems*. Springer, 2008.
- [Cov20] APHP-Universities-INRIA-INSERM Group Covid-19. Early indicators of intensive care unit bed requirement during the Covid-19 epidemic: A retrospective study in Ile-de-France region, France. *PLoS one*, 15(11):e0241406, 2020.
- [Cru91] René L. Cruz. A calculus for network delay. I. Network elements in isolation. *IEEE Transactions on information theory*, 37(1):114–131, 1991.
- [CS16] Venkat Chandrasekaran and Parikshit Shah. Relative entropy relaxations for signomial optimization. *SIAM Journal on Optimization*, 26(2):1147–1173, 2016.
- [CT80] Michael G. Crandall and Luc Tartar. Some relations between non expansive and order preserving maps. *Proceedings of the AMS*, 78(3):385–390, 1980.
- [CT05] Dongyan Chen and Kishor S. Trivedi. Optimization for condition-based maintenance with semi-Markov decision process. *Reliability engineering & system safety*, 90(1):25–29, 2005.
- [CY01] Hong Chen and David D. Yao. *Fundamentals of queueing networks: Performance, asymptotics, and optimization*, volume 4. Springer, 2001.
- [DC64] John S. De Cani. A dynamic programming algorithm for embedded Markov chains when the planning horizon is at infinity. *Management Science*, 10(4):716–733, 1964.
- [DE95] Jorg Desel and Javier Esparza. *Free choice Petri nets*. Cambridge university press, 1995.
- [DF68] Eric V. Denardo and Benett L. Fox. Multichain Markov Renewal Programs. *SIAM J.Appl.Math*, 16:468–487, 1968.
- [DidW17] Mareike Dressler, Sadik Ilman, and Timo de Wolff. A positivstellensatz for sums of nonnegative circuit polynomials. *SIAM Journal on Applied Algebra and Geometry*, 1(1):536–555, 2017.

- [DLP⁺07] Alexandre Deslauriers, Pierre L’Ecuyer, Jutta Pichitlamken, Armann Ingolfsson, and Athanassios N. Avramidis. Markov chain models of a telephone call center with call blending. *Computers & operations research*, 34(6):1616–1645, 2007.
- [DLRS10] Jesús A. De Loera, Jörg Rambau, and Francisco Santos. *Triangulations Structures for algorithms and applications*. Springer, 2010.
- [dQLMSG14] Marcus A. de Q.V. Lima, Paulo R. M. Maciel, Bruno Silva, and Almir P. Guimarães. Performability evaluation of emergency call center. *Performance Evaluation*, 80:27–42, 2014.
- [DY79] Eugene B. Dynkin and Aleksandr A. Yushkevich. *Controlled Markov processes*, volume 235. Springer, 1979.
- [EN94] Javier Esparza and Mogens Nielsen. Decidability issues for Petri nets. *BRICS Report Series*, 1(8), 1994.
- [Erl17] Agner K. Erlang. Solution of some problems in the theory of probabilities of significance in automatic telephone exchanges. *Post Office Electrical Engineer’s Journal*, 10:189–197, 1917.
- [Fei94] Eugene A. Feinberg. Constrained semi-Markov decision processes with average rewards. *Zeitschrift für Operations Research*, 39(3):257–288, 1994.
- [FG18] Shmuel Friedland and Stéphane Gaubert. Spectral inequalities for nonnegative tensors and their tropical analogues, 2018. arXiv:1804.00204.
- [FGQ11] Nadir Farhi, Maurice Goursat, and Jean-Pierre Quadrat. Piecewise linear concave dynamical systems appearing in the microscopic traffic modeling. *Linear Algebra and Appl.*, 2011.
- [FH15] Estibaliz Fraca and Serge Haddad. Complexity analysis of continuous Petri nets. *Fundamenta informaticae*, 137(1):1–28, 2015.
- [FN85] Gérard Florin and Stéphane Natkin. Les réseaux de Petri stochastiques. *Technique et Science Informatiques*, 4(1):143–160, 1985.
- [FP62] Miroslav Fiedler and Vlastimil Ptak. On matrices with non-positive off-diagonal elements and positive principal minors. *Czechoslovak Mathematical Journal*, 12(3):382–400, 1962.
- [GAA⁺20] Stéphane Gaubert, Marianne Akian, Xavier Allamigeon, Marin Boyet, Baptiste Colin, Théotime Grohens, Laurent Massoulié, David P. Parsons, Frédéric Adnet, Érick Chanzy, Laurent Goix, Frédéric Lapostolle, Éric Lecarpentier, Christophe Leroy, Thomas Loeb, Jean-Sébastien Marx, Caroline Télion, Laurent Tréluyer, and Pierre Carli. Understanding and monitoring the evolution of the Covid-19 epidemic from medical emergency calls: the example of the Paris area. *Comptes Rendus. Mathématique*, 358(7):843–875, 2020.
- [GG98] Stéphane Gaubert and Jeremy Gunawardena. The duality theorem for min-max functions. *Comptes Rendus de l’Académie des Sciences-Series I-Mathematics*, 326(1):43–48, 1998.

- [GG04a] Stéphane Gaubert and Jeremy Gunawardena. The Perron–Frobenius theorem for homogeneous, monotone functions. *Trans. of AMS*, 356(12):4931–4950, 2004.
- [GG04b] Bruno Gaujal and Alessandro Giua. Optimal stationary behavior for a class of timed continuous Petri nets. *Automatica*, 40(9):1505–1516, 2004.
- [Gil80] Richard D. Gill. Censoring and stochastic integrals. *Statistica Neerlandica*, 34(2):124–124, 1980.
- [GJ08] Bernd Gärtner and Martin Jaggi. Tropical support vector machines. Technical report, Citeseer, 2008.
- [GLS12] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2. Springer Science & Business Media, 2012.
- [GMR02] Ofer Garnett, Avishai Mandelbaum, and Martin Reiman. Designing a call center with impatient customers. *Manufacturing & Service Operations Management*, 4(3):208–227, 2002.
- [GP88] Jacob E. Goodman and Richard Pollack. Hadwiger’s transversal theorem in higher dimensions. *Journal of the American Mathematical Society*, pages 301–309, 1988.
- [Gro08] Donald Gross. *Fundamentals of queueing theory*. John Wiley & Sons, 2008.
- [GS18] Alessandro Giua and Manuel Silva. Petri nets and automatic control: A historical perspective. *Annual Reviews in Control*, 45:223–239, 2018.
- [Gus18] Klas Gustavsson. *Stochastic Modeling and Management of an Emergency Call Center: A Case Study at the Swedish Emergency CallCenter Provider, SOS Alarm Sverige AB*. PhD thesis, Mid Sweden University, 2018.
- [GY94] Paul Glasserman and David D. Yao. *Monotone structure in discrete-event systems*. John Wiley & Sons, Inc., 1994.
- [HA99] Mark Hartmann and Cristina Arguelles. Transience bounds for long walks. *Math. Oper. Res.*, 24(2):414–439, 1999.
- [Hac76a] Michel H. T. Hack. *Decidability questions for Petri Nets*. PhD thesis, Massachusetts Institute of Technology, 1976.
- [Hac76b] Michel H. T. Hack. *Petri net language*. Massachusetts Institute of Technology, 1976.
- [Hal74] Paul R. Halmos. *Measure theory*, volume 18. Springer, 1974.
- [HG11] Yonghui Huang and Xianping Guo. Finite horizon semi-Markov decision processes with application to maintenance systems. *European Journal of Operational Research*, 212(1):131–140, 2011.
- [HKNT98] Graham Horton, Vidyadhar G. Kulkarni, David M. Nicol, and Kishor S. Trivedi. Fluid stochastic Petri nets: Theory, applications, and solution techniques. *European Journal of Operational Research*, 105(1):184–201, 1998.

- [HMU01] John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. Introduction to automata theory, languages, and computation. *Acm Sigact News*, 32(1):60–65, 2001.
- [HOvdW05] Bernd Heidergott, Geert J. Olsder, and Jacob van der Woude. *Max-plus at Work*. Princeton Univ. Press, 2005.
- [How60] Ronald A. Howard. *Dynamic programming and Markov processes*. John Wiley, 1960.
- [HR04] Simon Hardy and Pierre N. Robillard. Modeling and simulation of molecular biology systems using Petri nets: modeling goals of various approaches. *Journal of bioinformatics and computational biology*, 2(04):619–637, 2004.
- [Hun07] John D. Hunter. Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.
- [ICWC10] Armann Ingolfsson, Fernanda Campello, Xudong Wu, and Edgar Cabral. Combining integer programming and the randomization method to schedule employees. *European Journal of Operational Research*, 202(1):153–163, 2010.
- [Jaś04] Anna Jaśkiewicz. On the equivalence of two expected average cost criteria for semi-Markov control processes. *Mathematics of Operations Research*, 29(2):326–338, 2004.
- [Jew63] William S. Jewell. Markov-renewal programming. I: Formulation, finite return models. *Operations Research*, 11(6):938–948, 1963.
- [JK01] Geurt Jongbloed and Ger Koole. Managing uncertainty in call centres using poisson mixtures. *Applied Stochastic Models in Business and Industry*, 17(4):307–318, 2001.
- [JPM20] François Javaudin, Yann Penverne, and Emmanuel Montassier. Organisation of prehospital care: the French experience. *European Journal of Emergency Medicine*, 27(6):404–405, 2020.
- [JSV17] Anders Jensen, Jeff Sommars, and Jan Verschelde. Computing tropical prevarieties in parallel. In *Proceedings of the International Workshop on Parallel Symbolic Computation*, pages 1–8, 2017.
- [JX04] Liu Jianyong and Zhao Xiaobo. On average reward semi-Markov decision processes with a general multichain structure. *Mathematics of Operations Research*, 29(2):339–352, 2004.
- [Kal97] Olav Kallenberg. *Foundations of modern probability*, volume 2. Springer, 1997.
- [KM27] William O. Kermack and Anderson G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115(772):700–721, 1927.
- [KM02] Ger Koole and Avishai Mandelbaum. Queueing models of call centers: An introduction. *Annals of Operations Research*, 113(1-4):41–59, 2002.

- [Koh80] Elon Kohlberg. Invariant half-lines of nonexpansive piecewise-linear transformations. *Math. Oper. Res.*, 5(3):366–372, 1980.
- [KP06] Ger Koole and Auke Pot. An overview of routing and staffing algorithms in multi-skill customer contact centers, 2006.
- [KV15] Panqanamala Ramana Kumar and Pravin Varaiya. *Stochastic systems: Estimation, identification, and adaptive control*. SIAM, 2015.
- [KW14] Song-Hee Kim and Ward Whitt. Are call center and hospital arrivals well modeled by nonhomogeneous Poisson processes? *Manufacturing & Service Operations Management*, 16(3):464–480, 2014.
- [LB10] Jean-Yves Le Boudec. *Performance evaluation of computer and communication systems*, volume 2. EPFL Press Lausanne, 2010.
- [LBT01] Jean-Yves Le Boudec and Patrick Thiran. *Network calculus: a theory of deterministic queuing systems for the Internet*. Springer, 2001.
- [LDF⁺20] Xavier Leroy, Damien Doligez, Alain Frisch, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. The OCaml system – release 4.11, 2020.
- [LGO18] Pierre L’Ecuyer, Klas Gustavsson, and Leif Olsson. Modeling bursts in the arrival process to an emergency call center. In M. Rabe, A. A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, editors, *Proceedings of the 2018 Winter Simulation Conference*, 2018.
- [LGV⁺20] Frédéric Lapostolle, Laurent Goix, Isabelle Vianu, Érick Chanzy, Carla De Stefano, Judith Gorlicki, Tomislav Petrovic, and Frédéric Adnet. Covid-19 epidemic in the Seine-Saint-Denis department of greater Paris: One month and three waves for a tsunami. *European Journal of Emergency Medicine*, 27(4):274–278, 2020.
- [LHSC07] Bruce G. Lewis, Ric D. Herbert, Peter. F Summons, and William J. Chivers. Agent-based simulation of a multi-queue emergency services call centre to evaluate resource allocation. In *MODSIM 2007, International Congress on Modelling and Simulation, Modelling and Simulation Society of Australia and New Zealand*, pages 11–17, 2007.
- [Lim05] Lek-Heng Lim. Singular values and eigenvalues of tensors: a variational approach. In *Proc. IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP ’05)*, volume 1, pages 129–132, 2005.
- [Lit61] John D. C. Little. A proof for the queuing formula: $L = \lambda W$. *Operations research*, 9(3):383–387, 1961.
- [Lit07] Grigori L. Litvinov. Maslov dequantization, idempotent and tropical mathematics: A brief introduction. *Journal of Mathematical Sciences*, 140(3):426–444, 2007.
- [LN12] Bas Lemmens and Roger Nussbaum. *Nonlinear Perron–Frobenius theory*, volume 189 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2012.

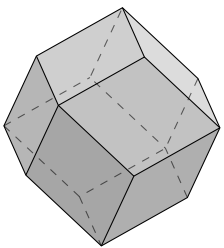
- [Lot02] Monsieur Lothaire. *Algebraic combinatorics on words*, volume 90. Cambridge university press, 2002.
- [LR96] Kurt Lautenbach and Hanno Ridder. The linear algebra of deadlock avoidance—a Petri net approach. *Univ. Koblenz, Inst. Comput. Sci., Germany, Tech. Rep*, 1996.
- [LS05] Bas Lemmens and Michael Scheutzow. On the dynamics of sup-norm non-expansive maps. *Ergodic Theory Dynam. Systems*, 25(3):861–871, 2005.
- [LS09] Jim Lawrence and Valeriu Soltan. The intersection of convex transversals is a convex polytope. *Contributions to Algebra and Geometry*, 50(1):283–294, 2009.
- [LvHK96] Ted Lewis, Balder von Hohenbalken, and Victor Klee. Common supports as fixed points. *Geometriae Dedicata*, 60(3):277–281, May 1996.
- [Mam86] John W. Mamer. Successive approximations for finite horizon, semi-Markov decision processes with application to asset liquidation. *Operations Research*, 34(4):638–644, 1986.
- [Mar89] Peter Martus. *Asymptotic properties of nonstationary operator sequences in the nonlinear case*. Phd thesis, Friedrich-Alexander Universität, Erlangen-Nürnberg, 1989.
- [Mey00] Carl D. Meyer. *Matrix analysis and applied linear algebra*, volume 71. Siam, 2000.
- [Mil68] Bruce L. Miller. Finite state continuous time Markov decision processes with a finite planning horizon. *SIAM Journal on Control*, 6(2):266–280, 1968.
- [MNS14] Glenn Merlet, Thomas Nowak, and Sergeï Sergeev. Weak CSR expansions and transience bounds in max-plus algebra. *Linear Algebra and its Applications*, 461:163–199, 2014.
- [MNS21] Glenn Merlet, Thomas Nowak, and Sergeï Sergeev. On the tightness of bounds for transients of weak CSR expansions and periodicity transients of critical rows and columns of tropical matrix powers. *Linear and Multilinear Algebra*, 0(0):1–31, 2021.
- [Mol82] Michael K. Molloy. Performance analysis using stochastic Petri nets. *IEEE Transactions on computers*, 31(09):913–917, 1982.
- [MP19] Romain Moutard and Yann Penverne. Mission de modernisation de l’accessibilité et de la réception des communications d’urgence pour la sécurité, la santé et les secours. Technical report, Ministère de l’Intérieur, Ministère de la Santé et des Solidarités, 2019.
- [MV98] Jean Mairesse and Laurent Vuillon. Optimal sequences in a heap model with two pieces. In *Paris*. Citeseer, 1998.
- [NCM⁺17] Nikolaos Nikolaou, Maaret Castrén, Koenraad G. Monsieurs, Diana Cimpoesu, Marios Georgiou, Violetta Raffay, Rudolph Koster, Silvija Hunyadi-Antičević, Anatolij Truhlář, Leo Bossaert, et al. Time delays to reach dispatch centres in different regions in Europe. Are we losing the window of opportunity?—The EUROCALL study. *Resuscitation*, 111:8–13, 2017.

- [Nev65] Jacques Neveu. *Mathematical foundations of the calculus of probability*. Holden-Day series in probability and statistics. Holden-day, 1965.
- [Nor98] James R. Norris. *Markov chains*, volume 2. Cambridge University Press, 1998.
- [NPJ98] Heikki E. Nikkanen, Claude Pougès, and Lenworth M. Jacobs. Emergency medicine in France. *Annals of emergency medicine*, 31(1):116–120, 1998.
- [NRA13] Kok M. Ng, Mamun B. I. Reaz, and Mohd A. M. Ali. A review on the applications of Petri nets in modeling, analysis, and control of urban traffic. *IEEE Transactions on Intelligent Transportation Systems*, 14(2):858–870, 2013.
- [Pal57] Rudolf C. A. Palm. *Research on telephone traffic carried by full availability groups*. Tele, 1957.
- [pDT20] The pandas Development Team. pandas-dev/pandas: Pandas, February 2020.
- [Per06] Benoît Perthame. *Transport equations in biology*. Springer Science & Business Media, 2006.
- [Pet62] Carl A. Petri. *Kommunikation mit automaten*. PhD thesis, Fachbereich Informatik, 1962.
- [Pet81] James L. Peterson. *Petri net theory and the modeling of systems*. Prentice Hall PTR, 1981.
- [Pet20] Eva Petitdemange. *SAMUFLUX: une démarche outillée de diagnostic et d’amélioration à base de doubles numériques: application aux centres d’appels d’urgence de trois SAMU*. PhD thesis, Ecole des Mines d’Albi-Carmaux, 2020.
- [PF91] Brigitte Plateau and Jean-Michel Fourneau. A methodology for solving Markov models of parallel systems. *Journal of parallel and distributed computing*, 12(4):370–387, 1991.
- [Put14] Martin L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [Ram73] Chander Ramchandani. *Analysis of asynchronous concurrent systems by timed Petri nets*. PhD thesis, Massachusetts Institute of Technology, 1973.
- [RB12] D. Roubos and S. Bhulai. Approximate dynamic programming techniques for skill-based routing in call centers. *Probability in the Engineering and Informational Sciences*, 26(4):581–591, 2012.
- [Rei13] Wolfgang Reisig. *Understanding Petri Nets: Modeling Techniques, Analysis Methods, Case Studies*. Springer Publishing Company, Incorporated, 2013.
- [Reu89] Christophe Reutenauer. *Aspects mathématiques des réseaux de Petri*. Elsevier Masson, 1989.
- [Rob13] Philippe Robert. *Stochastic networks and queues*, volume 52. Springer Science & Business Media, 2013.

- [Ros70] Sheldon M. Ross. Average cost semi-Markov decision processes. *Journal of Applied Probability*, 7(3):649–656, 1970.
- [RS01] Laura Recalde and Manuel Silva. Petri net fluidification revisited: Semantics and steady state. *European Journal of Automation, APII-JESA*, 35(4):435–449, 2001.
- [RTS99] Laura Recalde, Enrique Teruel, and Manuel Silva. Autonomous continuous P/T systems. In *International Conference on Application and Theory of Petri Nets*, pages 107–126. Springer, 1999.
- [Rud76] Walter Rudin. *Principles of mathematical analysis*, volume 3. McGraw-hill New York, 1976.
- [Sav69] Emanuel S. Savas. Simulation and cost-effectiveness analysis of New York’s emergency ambulance service. *Management science*, 15(12):B–608, 1969.
- [SC09] R. Malcolm Smith and Alasdair K.T. Conn. Prehospital care – scoop and run or stay and play? *Injury*, 40:S23–S26, 2009.
- [Sch80] Helmut Schellhaas. Markov renewal decision processes with finite horizon. *Operations-Research-Spektrum*, 2(1):33–40, 1980.
- [Sch92] Manfred Schäl. On the second optimality equation for semi-Markov decision models. *Mathematics of Operations Research*, 17(2):470–486, 1992.
- [Sen89] Linn I. Sennott. Average cost semi-Markov decision processes and the control of queueing systems. *Probability in the Engineering and Informational Sciences*, 3(2):247–272, 1989.
- [SF77] Paul J. Schweitzer and Awi Federgruen. The asymptotic behavior of undiscounted value iteration in Markov decision problems. *Mathematics of Operations Research*, 2(4):360–381, 1977.
- [SF78] Paul J. Schweitzer and Awi Federgruen. The functional equations of undiscounted Markov renewal programming. *Mathematics of Operations Research*, 3(4):308–321, 1978.
- [SF79] Paul J. Schweitzer and Awi Federgruen. Geometric convergence of value-iteration in multichain Markov decision problems. *Adv. Appl. Prob.*, 11:188–217, 1979.
- [Sif79] Joseph Sifakis. Use of Petri nets for performance evaluation. *Acta Cybernetica*, 4(2):185–202, 1979.
- [SS12] Sergeĭ Sergeev and Hans Schneider. CSR expansions of matrix powers in max algebra. *Transactions of AMS*, 364(11):5969–5994, May 2012.
- [Sti85] Shaler Stidham. Optimal control of admission to a queueing system. *IEEE Transactions on Automatic Control*, 30(8):705–713, 1985.
- [SW93] Winfried Stute and Jane-Ling Wang. The strong law under random censorship. *The Annals of statistics*, pages 1591–1607, 1993.

- [TB07] Gonca Tuncel and Gunhan M. Bayhan. Applications of Petri nets in production scheduling: A review. *The International Journal of Advanced Manufacturing Technology*, 34(7):762–773, 2007.
- [TCHR20] Johannes Trunk, Bertrand Cottenceau, Laurent Hardouin, and Joerg Raisch. Modelling and control of periodic time-variant event graphs in dioids. *Discrete Event Dynamic Systems*, 30(2):269–300, 2020.
- [TK93] Kishor S. Trivedi and Vidyadhar G. Kulkarni. FSPNs: fluid stochastic Petri nets. In *International Conference on Application and Theory of Petri Nets*, pages 24–31. Springer, 1993.
- [TLB16] Thuy Ta, Pierre L’Ecuyer, and Fabian Bastin. Staffing optimization with chance constraints for emergency call centers. In *MOSIM 2016-11th International Conference on Modeling, Optimization and Simulation*, 2016.
- [TMDC20] Caroline Telion, Jean-Sébastien Marx, Caroline Dautreppe, and Pierre Carli. Retour d’expérience sur la régulation au Samu de Paris pendant la crise de Covid-19. *Annales françaises de médecine d’urgence*, 10(4-5):202–211, 2020.
- [TVR99] John N. Tsitsiklis and Benjamin Van Roy. Optimal stopping of Markov processes: Hilbert space theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. *IEEE Transactions on Automatic Control*, 44(10):1840–1851, 1999.
- [vBKvdMB15] Martin van Buuren, Geert J. Kommer, Rob van der Mei, and Sandjai Bhulai. A simulation model for emergency medical services call centers. In *2015 winter simulation conference (WSC)*, pages 844–855. IEEE, 2015.
- [vBKvdMB17] Martin van Buuren, Geert J. Kommer, Rob van der Mei, and Sandjai Bhulai. EMS call center models with and without function differentiation: A comparison. *Operations Research for Health Care*, 12:16–28, 2017.
- [VdA98] Wil M. P. Van der Aalst. The application of Petri nets to workflow management. *Journal of circuits, systems, and computers*, 8(01):21–66, 1998.
- [Vir01] Oleg Viro. Dequantization of real algebraic geometry on logarithmic paper. In Carles Casacuberta, Rosa Maria Miró-Roig, Joan Verdera, and Sebastià Xambó-Descamps, editors, *European Congress of Mathematics*, pages 135–146, Basel, 2001. Birkhäuser Basel.
- [VMJS13] Carlos Vazquez, Cristian Mahulea, Jorge Júlvez, and Manuel Silva. Introduction to fluid Petri nets. *Lecture Notes in Control and Information Sciences*, 433:365–386, 01 2013.
- [VVR⁺17] Damien Viglino, Aurelien Vesin, Stephane Ruckly, Xavier Morelli, Rémi Slama, Guillaume Debaty, Vincent Danel, Maxime Maignan, and Jean-François Timsit. Daily volume of cases in emergency call centers: construction and validation of a predictive model. *Scandinavian journal of trauma, resuscitation and emergency medicine*, 25(1):1–7, 2017.

- [Wil98] Ruth J. Williams. Diffusion approximations for open multiclass queueing networks: sufficient conditions involving state space collapse. *Queueing systems*, 30(1):27–88, 1998.
- [WM10] Wes McKinney. Data Structures for Statistical Computing in Python. In Stéfan van der Walt and Jarrod Millman, editors, *Proceedings of the 9th Python in Science Conference*, pages 56 – 61, 2010.
- [WZ77] Richard Lee Wheeden and Antoni Zygmund. *Measure and integral*, volume 26. Dekker New York, 1977.
- [Ye07] Heng-Qing Ye. A paradox for admission control of multiclass queueing network with differentiated service. *Journal of applied probability*, 44(2):321–331, 2007.
- [Ye11] Yinyu Ye. The simplex and policy-iteration methods are strongly polynomial for the Markov decision problem with a fixed discount rate. *Mathematics of Operations Research*, 36(4):593–603, 2011.
- [YT91] Yuehwern Yih and Arne Thesen. Semi-Markov decision models for real-time scheduling. *The International Journal of Production Research*, 29(11):2331–2346, 1991.
- [Yus82] Aleksandr A. Yushkevich. On semi-Markov controlled models with an average reward criterion. *Theory of Probability & Its Applications*, 26(4):796–803, 1982.
- [Zuc78] Dror Zuckerman. Optimal stopping in a semi-Markov shock model. *Journal of Applied probability*, 15(3):629–634, 1978.





Titre : Systèmes dynamiques affines par morceaux appliqués à l'évaluation de performance de centres d'appels d'urgence

Mots clés : Évaluation de performance, systèmes dynamiques, géométrie tropicale, contrôle stochastique

Résumé : Nous développons des méthodes mathématiques pour l'évaluation de performance et le dimensionnement de centres d'appels d'urgence. Nous travaillons avec le formalisme des réseaux de Petri, et étudions en particulier leur dynamique décrite par des systèmes affines par morceaux.

Nous établissons une correspondance entre le comportement des réseaux de Petri monotones et l'étude de la fonction valeur des processus de décision semi-Markoviens. Cela nous permet de caractériser finement le débit des réseaux de Petri, et déduire des recommandations pratiques de dimensionnement pour les centres d'appels d'urgence. Au moyen d'une correspondance supplémentaire avec les problèmes de plus court chemin stochas-

tique, nous bornons aussi le temps d'absorption d'un pic d'appels. Nous étudions également des systèmes non-monotones, lorsque des règles de priorité sont utilisées, et en décrivons les phases de congestion par des complexes polyédraux.

Cette approche théorique est validée par des campagnes de simulation numérique se basant sur une étude approfondie d'appels d'urgence réels. Nous mettons à profit ces simulations pour comparer différentes architectures de centres d'appels et quantifier l'intérêt opérationnel de systèmes biniveaux par rapport à des systèmes mononiveaux. Nos travaux s'appuient sur des collaborations avec les centres d'appels d'urgence de la région parisienne : SAMU, pompiers et police.

Title: Piecewise affine dynamical systems applied to the performance evaluation of emergency call centers

Keywords: Performance evaluation, dynamical systems, tropical geometry, stochastic control

Abstract:

We develop mathematical methods for the performance analysis and dimensioning of emergency call centers. We work with the formalism of Petri nets and focus in particular on their dynamics expressed with piecewise affine systems.

We establish a formal correspondence between the behaviour of monotonic Petri nets and the study of the value function of semi-Markov decision processes. This enables us to precisely characterize the throughput of Petri nets, and deduce practical staffing recommendations for emergency call centers. Using an additional correspondence with stochastic shortest path problems, we bound the

time needed to absorb a peak of calls. We also study non-monotonic systems where priority rules come into play, and describe their congestion phases with polyedral complexes.

This theoretical approach is validated with numerical simulation campaigns based on an in-depth study of actual emergency calls. We build on these simulations to compare different layouts of call centers and quantify the interest of using bilevel (or two-tier) architectures compared to single-tier platforms. Our work has benefited from collaborations with emergency call centers of Paris area: SAMU (emergency health services), firefighters and police.

