



HAL
open science

Probabilistic model for distributed and concurrent systems. Limit theorems and application to statistical parametric estimation

Samy Abbes

► **To cite this version:**

Samy Abbes. Probabilistic model for distributed and concurrent systems. Limit theorems and application to statistical parametric estimation. Networking and Internet Architecture [cs.NI]. Université Rennes 1, 2004. English. NNT: . tel-00007150v2

HAL Id: tel-00007150

<https://theses.hal.science/tel-00007150v2>

Submitted on 19 Oct 2004

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.

Numéro d'ordre : 3026

THESE

présentée

devant l'Université de Rennes 1

pour obtenir le grade de

Docteur de l'Université de Rennes 1

Mention : Informatique

par

Samy ABBES

préparée à l'IRISA

Campus de Beaulieu
35 042 Rennes Cedex
<http://www.irisa.fr>

au sein de l'École Doctorale MATISSE

sous la direction de

Albert BENVENISTE
IRISA / INRIA

Modèle probabiliste de systèmes distribués et concurrents.
Théorèmes limite et application à l'estimation statistique de
paramètres.

rédigée en anglais

Soutenue le 14 octobre 2004, devant la Commission d'Examen :

<i>Président du Jury</i>	Jean-Pierre Banâtre	DR INRIA (IRISA)
<i>Rapporteur</i>	Glynn Winskel	Professeur (Univ. de Cambridge)
<i>Rapporteur</i>	Jean Mairesse	CR CNRS (LIAFA)
<i>Examineur</i>	Albert Benveniste	DR INRIA (IRISA)
<i>Examineur</i>	Joost-Pieter Katoen	Professeur (Univ. de Twente)
<i>Examineur</i>	Vadim Kaimanovitch	DR CNRS (IRMAR)
<i>Examineur</i>	Stefan Haar	CR INRIA (IRISA)

Thèse de Doctorat de l'Université de Rennes 1

Probabilistic model for distributed and concurrent systems. Limit theorems and application to statistical parametric estimation.

Modèle probabiliste de systèmes distribués et concurrents. Théorèmes limite et application à l'estimation statistique de paramètres.

Avec une Introduction en français

Samy ABBES

Direction des Recherches : **Albert BENVENISTE**

Septembre 2002–Juin 2004

La valeur de toute science a pour fondement des raisonnements de probabilités.


É. Borel, *Valeur pratique et philosophique des probabilités*, 1939 (Gauthier-Villars)

Remerciements

Je tiens à remercier chaleureusement M. Vincent Pantaloni, Professeur Agrégé, qui a attentivement relu ce document. *“Whosoever thinks a faultless piece to see, thinks what ne'er was, nor is, nor e'er shall be” (Pope).*

Note: The probabilistic background is treated in Chapter 2, Section 1, and in Chapter 4, Section 1.

Organisation of the Document



Contents

Organisation of the document	9
Introduction (<i>français</i>)	12
Introduction (<i>English</i>)	25
1 Preliminaries on Computational Models	37
2 Topological Event Structures	61
3 Geometry of Stopping Prefixes	83
4 Distributed Probabilities	121
5 Markov Nets	151
6 Recurrent Nets	185
7 Ergodicity and Limit theorems	203
8 Algorithms for Statistics	225
9 Conclusions and Perspectives	243
Bibliography	249
Index	254

Introduction

Ce document propose une analyse mathématique et algorithmique de certains modèles de concurrence¹ issus de l'Informatique, structures d'événements et réseaux de Petri. On décrit ci-dessous les contributions apportées, et le contexte où ils se placent.

Les résultats nouveaux sont indiqués par l'indentation à droite et à gauche du texte, comme pour ce paragraphe.

I—Motivation et contexte

L'étude des systèmes artificiels complexes a dégagé au cours des dernières décennies le concept de système distribué. Les approches et les modèles sont variés, cependant les conclusions et les conceptions convergent sur un certain nombre de points.

Une idée importante est que la complexité des systèmes conduit à renoncer à leur contrôle global. Tout d'abord l'accroissement démesuré du nombre d'états empêche l'analyse opérationnelle globale des systèmes. D'autre part la nature répartie des systèmes étudiés remet en cause la notion même d'état global de systèmes tels que les réseaux de télécommunication, et les réseaux de transmission d'information en général. Dans les grands réseaux de communication, les temps de synchronisation des composants rendent tout simplement inaccessible l'état global du système à un instant donné. La notion d'état global du système doit donc être remplacée par celle d'*état local*. Avec la notion d'état local, et en conséquence d'actions locales, apparaissent aussitôt des situations de *concurrence de processus*.

La théorie des systèmes s'est donc tournée vers les modèles de concurrence élaborés et étudiés depuis les années 1970, et dont le représentant emblématique est le modèle des réseaux de Petri (*Petri nets*). L'analyse des systèmes utilise

¹Dans ce document, "concurrence" est l'anglicisme passé dans le langage scientifique courant, correspondant au terme anglais "concurrency". *Concurrence* doit donc être compris au sens de simultanéité, et non pas au sens de compétition.

fréquemment des modèles probabilistes, permettant le traitement statistique des données. Des exemples classiques d'applications à la gestion et à la surveillance des systèmes sont l'identification de paramètres et le diagnostic de pannes basés sur des modèles de Markov cachés (HMM, *Hidden Markov Models*) avec les algorithmes EM et de Viterbi, mais aussi l'évaluation de performance des systèmes par exemple par des méthodes asymptotiques. Ainsi on a pu adapter la théorie des files d'attente à des systèmes connectés en réseau, en particulier la théorie des files d'attente markoviennes. La topologie des réseaux définit des opérateurs d'algèbre $(\max, +)$ conduisant à des équations d'évolution temporelles, d'où une analyse des performances du système. Le modèle qui se prête à cette approche est celui des réseaux de Petri à choix libres. La littérature est extrêmement fournie sur le sujet, voir entre autres [2, 4, 3, 30].

D'autre part, des modèles probabilistes ont été élaborés pour prendre en compte le développement des logiques temporelles et des langages de communication entre processus. Les algèbres de processus stochastiques par exemple étendent les algèbres de processus temporisés en donnant une valeur aléatoire au paramètre temporel [22, 25]. Les grammaires de processus concernent des processus de la forme (λ, p) , où p est un processus dans un langage de communication proche de CCS et λ est un paramètre probabiliste, par exemple le paramètre d'une loi exponentielle. Cette approche a aussi été appliquée aux *bundle event structures* [11], un modèle inspiré des structures d'événements de Winskel.

De nombreuses études de ces modèles sont motivées par des problèmes de sécurité de protocoles, exprimables comme problèmes de vérification. Larsen et Skou ont introduit des relations de simulation et de bisimulation, généralisant aux modèles probabilistes les relations d'équivalence observationnelle sur des processus étiquetés (*labelled transition systems*) [29]. Outre l'étude d'algorithmes probabilistes de tests, un thème d'étude est l'expressivité des modèles modulo simulation ou bisimulation, qu'on cherche à spécifier par une logique temporelle. Par exemple pour les systèmes de transition entièrement probabilistes (*fully probabilistic transition system*), une bisimulation faible sur les états du système coïncide avec l'équivalence provenant de $\text{PCTL} \setminus X$, fragment de PCTL (*Probabilistic Computation Tree Logic*) sans l'opérateur Next [5]. Il faut noter l'importance dans ces modèles de l'étiquetage (*labelling*) des événements par un alphabet d'actions. C'est par rapport à l'étiquetage qu'est définie la sémantique des processus probabilistes. La bisimulation isole des classes d'équivalence d'états telles que les transitions au sein d'une même classe du système se comportent comme des transitions internes, invisibles depuis les autres classes.

Ainsi, depuis les automates probabilistes de Rabin (1963), de nombreux modèles de calcul issus de l'Informatique ont été l'objet d'extensions probabilistes. Dans ces modèles l'état du système est représenté par un processus aléatoire $(X_t)_t$, où le paramètre t représente l'avancement du temps, suivant une ligne discrète ou continue. Il s'agit donc d'un temps *global à l'échelle du système*. En d'autres termes, et si le système modélise une structure en réseau, tous les composants du réseau sont synchronisés sur une même horloge globale. Or pour l'étude des systèmes distribués,

cette conception présente des inconvénients [8]. Car bien que tous les composants soient soumis au temps physique universel, il est raisonnable de penser que les synchronisations effectives n'interviennent que sur un ensemble discret d'événements, qui ne peut en aucune manière être fixé à l'avance. Entre les instants de synchronisation, l'asynchronie des composants se prête très mal à une représentation par un processus $(X_t)_t$ avec t avançant comme précédemment le long d'une ligne. La mauvaise adéquation entre le système et cette représentation est due à la nature *d'ordre total* de la ligne, qui ne correspond plus à l'ordonnement naturel des événements du système. Les événements asynchrones étant non reliés causalement, maintenir une relation de comparaison temporelle entre eux est superflu d'une part, et maladroit si cela entraîne une plus grande complexité de calculs.

Ainsi, en dehors de toute considération probabiliste, la représentation temporelle des systèmes répartis tire avantage d'un affaiblissement de la nature totalement ordonnée du temps, pour admettre la présence d'événements d'ordonnements *non comparables*. Cette représentation est précisément la représentation par *concurrence forte* (*true-concurrency*) des processus concurrents. Cette sémantique s'oppose à une sémantique d'entrelacement, qui distingue les entrelacements d'événements concurrents. L'économie de complexité provenant de la sémantique de concurrence forte a été mise à profit par exemple pour le *model-checking* de systèmes concurrents [32, 46, 19]. L'approche de concurrence forte a aussi été appliquée avec succès au diagnostic de pannes dans les réseaux de télécommunications, diagnostic centralisé et décentralisé [6, 8, 20].

Très récemment, on a commencé à s'intéresser à des modèles probabilistes pour la dynamique de processus fortement concurrents (*true-concurrent*) [45, 7, 43]. Le présent document propose une construction générale de processus probabilistes fortement concurrents ainsi que l'étude de certaines de leurs propriétés. On établit des propriétés asymptotiques de ces systèmes, notamment des propriétés de récurrence et la Loi forte des grands nombres, dans un contexte de concurrence forte. Les résultats sont appliqués à un problème statistique classique, l'estimation de paramètres, et on propose un algorithme d'estimation.

Les outils sont empruntés aux mathématiques et à l'informatique. Les outils analytiques sont ceux de la topologie et de la théorie de la mesure: espaces métriques, systèmes projectifs de probabilités. Les modèles et leur interprétation dynamique, ainsi que certains algorithmes, sont issus de l'informatique: réseaux de Petri, structures d'événements et réseaux d'occurrence. En dehors de contributions à l'analyse combinatoires de ces modèles, on montre en quelle mesure les outils analytiques mis en jeu, topologie et théorie de la mesure, apportent des éléments nouveaux pour l'appréhension des modèles.

II—Modèles acycliques : temps partiellement ordonné et probabilisation

Le modèle où nous nous plaçons est celui des dépliages de réseaux de Petri saufs (*safe Petri nets*). Étant donné un réseau de Petri sauf, l'idée est d'identifier les processus ab et ba si a et b sont des transitions *concurrentes*, c'est-à-dire n'ayant pas de ressources communes. C'est de cette manière qu'on *affaiblit* l'ordre chronologique total : les occurrences de a et b deviennent chronologiquement *non comparables*. La clôture transitive de cette relation est la relation d'*entrelacement* sur les suites de tirs dans le réseau. Par construction, la relation d'entrelacement identifie deux suites de tirs qui se déduisent l'une de l'autre par une suite d'échanges de transitions concurrentes. Les classes d'équivalences sont appelées processus *true-concurrent* du système, ou plus simplement *processus*. L'ensemble des processus constitue la dynamique du système, et chaque processus correspond à une exécution (partielle) du système. Ainsi une suite de transitions du réseau de Petri définit un processus, et différentes suites définissent le même processus si elles sont équivalentes modulo entrelacement. La concaténation de suites de transitions induit un ordre partiel sur les suites de transitions, qui induit à son tour un ordre partiel sur les processus. La relation $u \subseteq v$ pour deux processus correspond au fait que le processus u peut être continué jusqu'à atteindre v .

Le dépliage d'un réseau de Petri est un réseau acyclique qui recopie le réseau en différenciant les occurrences des transitions, ce qui revient à déplier indéfiniment les boucles du réseau d'origine. Mathématiquement, un dépliage est de manière équivalente une *structure d'événements* (*prime event structure*) ou un *réseau d'occurrences* (*occurrence net*). Ces deux modèles équivalents ont surtout la particularité d'être *acycliques*. Un processus du réseau de Petri d'origine est décrit par un unique ordre partiel d'événements du dépliage. On dit qu'on *relève* dans le dépliage le processus donné par une séquence de transitions du réseau. Comme il se doit, le relevé est indépendant de la séquence pourvu qu'elle reste dans la même classe d'équivalence modulo entrelacement. Les préfixes communs des différents relevés s'identifient dans le dépliage, d'où une structure d'ordre partiel sur les processus dans le dépliage. Cette relation est conjuguée à la relation d'inclusion naturelle sur les processus, ce qui montre que le dépliage capture très exactement la dynamique *true-concurrent* du réseau. Ces résultats sont ceux des travaux de Winskel sur les dépliages [47, 33, 50], et sont en bien des points comparables à certains résultats de la théorie du relèvement des chemins dans les variétés par exemple topologiques ou différentiables. De nombreuses extensions de ces travaux ont été effectuées, en particulier pour traiter le cas des réseaux non bornés, et pour relier cette théorie à celle des traces de Mazurkiewicz [23, 26].

Dans cette optique, que serait une probabilisation *true-concurrent* des réseaux de Petri saufs ? Peut-on définir une dynamique aléatoire non pas sur les séquences de transitions dans le réseau, mais sur les classes d'équivalence modulo entrelacement ? C'est ce qu'on appelle une probabilisation *true-concurrent*.

L'espace de probabilité adéquat est l'espace des processus maximaux, correspondant aux exécutions complètes du système. C'est bien cet espace qui est considéré lorsque le réseau de Petri est réduit à un système séquentiel comme une chaîne de Markov. L'espace de probabilité s'identifie avec le *bord à l'infini* du dépliage. Plus précisément, il s'agit du bord à l'infini des *processus du dépliage*, la distinction étant triviale pour les systèmes séquentiels. Par abus de langage, on parle de bord à l'infini du dépliage. La probabilité d'une séquence de transitions, *modulo entrelacement*, est alors la probabilité de l'*ombre* du relevé de la séquence dans le dépliage, ombre portée sur le bord du dépliage. L'ombre d'un processus dans le dépliage est illustrée par la Figure 1.




Figure 1: L'ombre d'un processus portée sur le bord définit la probabilité d'apparition du processus.

L'idée de probabiliser de cette manière la dynamique des réseaux de Petri est récente, due à Völzer [45]. Les auteurs qui ont mené des études sur le sujet se sont heurtés à des difficultés qui les ont amenés à réduire la classe de réseaux à probabiliser. Il s'agit des réseaux à choix libres étendus (*extended free choice*) pour Völzer [45], des réseaux sans confusion (*confusion-free*) pour Varracca [43] et pour Varracca-Winskel-Völzer [44] (la différence est mince entre choix libre étendu et sans confusion), et des réseaux dits "à choix compact" de Benveniste-Haar-Fabre [7].

Le but que se posent ces auteurs, au moins en un premier temps, est la construction d'une probabilité sur le bord par une technique d'extension. Indubitablement, il y a là un résultat analytique, dont la nature est indépendante de tout problème de concurrence. A. Benveniste a identifié avec raison ce résultat comme une limite projective de systèmes de probabilités, bien que les détails en soient négligés dans [7].

Dans le présent document, on montre qu'une clef pour l'étude analytique du bord à l'infini des modèles concurrents acycliques, c'est-à-dire l'ensemble des processus maximaux, consiste à plonger le bord à l'infini dans l'espace des processus (quelconques), ces deux espaces étant mu-

nis de topologies adéquates. On considère un système projectif, formé d'ensembles finis, pour chacun des deux espaces : espace des processus et espace des processus maximaux. Un premier résultat est l'adéquation entre l'espace des processus quelconques et la limite projective du système projectif associé. Bien que ce résultat soit assez élémentaire, il permet de définir de manière simple une topologie sur les processus, dérivée de la topologie projective. On retrouve cette topologie dans la littérature pour des modèles semblables: topologie de Lawson, topologie de traces infinies [28, 24]. On retrouve les résultats topologiques dus à la nature des espaces de processus, en particulier la compacité et la séparabilité, dans la structure de limite projective, sans introduire de métrique explicite.

Il n'en va pas exactement de même pour la topologie du bord à l'infini. J'ai été amené à réviser la notion de temps d'arrêt proposée par Benveniste *et al.* dans [7], pour étudier à la place les *préfixes intrinsèques* des structures d'événements et introduire la classe de structures d'événements *localement finies*, généralisant les arbres localement finis. Cette étude a grandement profité de la relation de *conflit dynamique* qui m'a été indiquée par D. Varacca et G. Winskel. Les préfixes intrinsèques finis d'une structure d'événements localement finie forment un semi-treillis supérieur dont la borne supérieure est la structure d'événements toute entière. Ils induisent une filtration partiellement ordonnée de sous-tribus du bord. L'étude des structures d'événements localement finies montre cette différence entre systèmes séquentiels et systèmes concurrents: le dépliage d'un système fini concurrent n'est pas nécessairement localement fini, alors qu'il en est toujours ainsi pour les systèmes séquentiels. Le bord à l'infini d'une structure d'événements localement finie s'identifie à une limite projective d'ensemble finis, d'où l'application du théorème d'extension de systèmes projectifs de probabilités (théorème de Prokhorov).

III—Probabilités et concurrence

Pour les structures d'événements localement finies, la construction d'une mesure de probabilité sur le bord se réduit à la construction d'un système projectif de probabilités sur des ensembles finis. L'extension de mesures pour les systèmes concurrents est donc semblable à l'extension de mesures pour les systèmes séquentiels.

Construire une probabilité sur le bord à l'infini d'un arbre, c'est-à-dire pour un système acyclique séquentiel, c'est construire une probabilité sur les chemins infinis de l'arbre. Cela revient à faire avancer un bateau sur la mer vers l'horizon par une

suite d'incrément, en prenant à chaque pas une décision aléatoire pour le faire dévier plutôt vers la droite ou plutôt vers la gauche, tout en sachant que quoi qu'il arrive, *le bateau va de l'avant*. C'est une traduction imagée du théorème de prolongement des systèmes projectifs de mesures de probabilité.

Maintenant pour un système concurrent, il s'agit de faire avancer plusieurs bateaux en même temps vers l'horizon. Les bateaux peuvent se rejoindre par synchronisation ou se séparer en plusieurs bateaux. Les bateaux ont *a priori* des paramètres propres, mais aussi des paramètres qui indiquent leurs influences respectives les uns sur les autres. Il est normal d'étudier si elles existent des probabilités qui annulent les influences mutuelles, rendant les bateaux indépendants en probabilité. C'est bien le sens de l'intuition de Benveniste *et al.* [7] (*l'indépendance conditionnelle des layers*), rejoignant la maxime de Winskel qui affirme que, au sens de la concurrence, "le parallélisme est une forme d'orthogonalité" [49].

On introduit certains objets géométriques liés aux structures d'événement permettant de définir les *probabilités distribuées*, qui possèdent les propriétés d'indépendance énoncées ci-dessus. On montre l'existence, et l'unicité sous certaines conditions naturelles, d'une probabilité distribuée construite à partir d'une famille de probabilités locales. L'opération qui associe une probabilité distribuée à une famille de probabilités locales est appelée *produit distribué*, et constitue une contribution principale de ce document. Le produit distribué généralise le produit de probabilités conditionnelles, par un traitement particulier de la concurrence. Naturellement, le caractère distribué des probabilités est trivial si la concurrence est absente. Autrement dit, toute probabilité sur un produit infini s'écrit comme un produit distribué.

Ce modèle probabiliste prend en compte la concurrence des systèmes de manière intrinsèque, sans recourir par exemple aux *schedulers* ou à un étiquetage supplémentaire, ou à toute autre variable non-déterministe extérieure comme on le fait par exemple pour les produits asynchrones d'automates probabilistes [40]. La construction du produit distribué est faite dans le cadre général des structures d'événements. À ma connaissance, cette construction est la méthode la plus générale pour construire une probabilité sur des processus *true-concurrent*, sans référence à un temps global réel.

La construction du produit distribué est fondé sur une décomposition des processus dans les structures d'événements. Les outils introduits avec cette décomposition forment ce que j'ai appelés des outils de *géométrie discrète* dans les structures d'événements. Pour les modèles acycliques étudiés dans ce document, structures d'événements et réseaux d'occurrences, j'appelle *géométrie* ce qui concerne les éléments qui composent le modèle, *i.e.* les événements et les conditions d'un réseau d'occurrences par exemples, par opposition aux deux autres objets associés au modèle : processus et processus maximaux, qui sont respectivement un ordre partiel topologique et un espace topologique, avec leurs tribus boréliennes respectives.

L'essentiel de l'étude consiste à relier les propriétés géométriques aux propriétés d'espaces mesurables et d'espaces de probabilités.

On décompose les processus maximaux par une construction géométrique itérative. Cette décomposition isole des processus locaux appelés *germes* du processus maximal qui, s'ils sont concurrents, sont indépendants les uns des autres. La décomposition n'est pas unique, due aux différents entrelacements possibles, mais la collection des germes en constitue un invariant, intrinsèque au processus maximal. Les germes sont les processus maximaux de sous-structures d'événements finies, appelées *cellules de branchement*. Les probabilités locales, appelées *probabilités de branchement*, probabilisent les germes des cellules de branchement. Dans la construction du produit distribué, on associe l'indépendance "horizontale" de germes concurrents à leur indépendance *en probabilité*.

Les cellules de branchement sont dynamiques: un même événement peut appartenir à différentes cellules de branchement, en fonction du contexte qui entoure l'événement. C'est donc une approche différente de celle des *clusters* (voir par exemple [19, 21]), qui sont statiquement définis dans un réseau de Petri.

IV—Systèmes aléatoires concurrents sans mémoire

En particulierisant la construction des probabilités distribuées des structures d'événements aux dépliages des réseaux de Petri saufs, on définit une dynamique aléatoire et *true-concurrent* des réseaux de Petri saufs.

De la théorie des systèmes dynamiques, on retient l'idée que les trajectoires infinies explorent les systèmes finis sans mémoire de manière régulière en moyenne. La formalisation de cette propriété est l'objet des théorèmes limites des probabilités, en particulier la Loi forte des grands nombres et le Théorème de la Limite Centrale. L'étude de théorèmes limites pour les systèmes concurrents suppose de définir les systèmes sans mémoires.

Pour cela j'ai introduit un outil d'une grande utilité dans ce travail, le cône du futur d'un processus dans les modèles acycliques. Par exemple dans un réseau d'occurrences, pour le cône du futur d'un processus v , on part de la simple observation que les processus contenant v peuvent être interprétés comme les processus d'un sous-réseau d'occurrences. C'est ce sous-réseau qui constitue le cône du futur de v , de sorte que v s'ajoute à tous les processus de son cône de futur. La propriété essentielle est que le bord à l'infini du cône du futur de v s'identifie à l'ombre de v . Si le bord du réseau d'occurrences est muni d'une probabilité \mathbb{P} , l'ombre de v , et donc le bord du cône du futur de v , est muni de la probabilité *conditionnelle* à v . Si

deux processus finis mènent au même marquage, leurs cones du futur sont isomorphes comme réseaux d'occurrences étiquetés, et donc leurs ombres sont isomorphes comme espaces topologiques et mesurables. Je dis qu'une probabilité \mathbb{P} définie sur le bord à l'infini du dépliage est *homogène*, ou *sans mémoire*, si l'isomorphisme d'espaces mesurables entre deux ombres isomorphes respecte les probabilités conditionnelles. Ainsi l'isomorphisme est un isomorphisme d'espaces de probabilités.

Je montre qu'une probabilité homogène et distribuée est déterminée par un nombre fini de paramètres, sous forme d'une famille finie de probabilités de branchement. Un réseau de Petri sauf et compact, *i.e.* dont le dépliage est localement fini, équipé d'une probabilité homogène et distribuée, est appelé *réseau de Markov distribué*. Les probabilités de branchements correspondent aux lignes de la matrice de transition d'une chaîne de Markov. La thèse de ce document est que **les réseaux de Markov distribués constituent une généralisation à des systèmes concurrents des chaînes de Markov finies**. À l'appui de cette thèse, je montre qu'un certain nombre de résultats bien connus pour les chaînes de Markov finies s'étendent au réseaux.

Une étude préliminaire consiste à reformuler et à établir la propriété de Markov forte dans un contexte *true-concurrent*. Le résultat implique la propriété de Markov habituelle pour les chaînes de Markov. Pour énoncer la propriété de Markov, j'ai été amené à proposer une généralisation des temps d'arrêts étudiés classiquement avec les processus stochastiques (chaînes de Markov, martingales, *etc*), sous la forme d'*opérateurs d'arrêt* (*stopping operator*). Contrairement aux temps d'arrêts pour réseaux d'occurrences de [7], les opérateurs d'arrêt admettent comme cas particuliers des opérateurs d'atteinte, correspondant aux temps d'atteinte (*hitting time*) étudiés pour les chaînes de Markov. Il s'ensuit une étude de la *réurrence des réseaux markoviens*, parallèle à l'étude de la récurrence des chaînes de Markov. On montre en particulier l'alternative 0-1 pour un marquage d'être récurrent ou transcient, et qu'un marquage atteignable depuis un marquage récurrent est lui-même récurrent. On distingue une récurrence globale et une récurrence locale, cette distinction étant triviale pour les chaînes de Markov.

La formulation de la propriété de Markov introduit la notion de *fonction homogène* pour palier cette difficulté intrinsèque aux systèmes récurrents, qu'une translation dans le temps ne transporte pas un espace vers un autre qui lui est isomorphe. À cause de la concurrence, l'espace des processus n'est pas homogène, d'où l'absence d'opérateur naturel de translation (*shift operator*) qui permettrait l'application directe de la théorie des systèmes dynamiques. L'étude de la propriété de Markov et de la récurrence est faite directement sur les réseaux, et les preuves généralisent les démonstrations classiques pour les résultats analogues sur les chaînes de Markov.

L'étude de la récurrence des réseaux est un premier exemple de caractéristique analytique des réseaux, qui profite de la souplesse du formalisme des probabilités. En effet, un réseau récurrent est, pour parler vite, un réseau qui revient infiniment

souvent dans son marquage d'origine, avec probabilité 1. La récurrence est un exemple typique de propriété vraie en probabilité, mais non totalement vraie. Il est très courant pour un réseau récurrent de contenir des exécutions qui ne reviennent pas au marquage d'origine : ces exécutions sont rares puisque toutes ensemble elles ont probabilité 0, mais elles existent. Ainsi une vérification complète échouerait à déterminer un tel réseau récurrent, il faut une formulation analytique des ensembles rares, telle que la fournit le formalisme de la théorie de la mesure, et en particulier les probabilités.

Certains résultats ergodiques, inspirés de résultats de la théorie des systèmes dynamiques, sont démontrés directement dans les réseaux. Pour en déduire la Loi forte des grands nombres, un ingrédient supplémentaire est apportée par l'étude d'une chaîne de Markov auxiliaire². En plus d'une condition de récurrence sur le réseau, on introduit une hypothèse pour contrôler l'amplitude de la concurrence au sein d'un réseau. L'amplitude est définie par une variable aléatoire appelée hauteur concurrente (*concurrent height*). De nouveau, l'apport de l'outil analytique est décisif pour formuler un contrôle sur la hauteur concurrente: autant il est beaucoup top contraignant de requérir que la hauteur concurrente soit bornée, autant il est naturel de demander qu'elle soit intégrable, *i.e.* d'espérance finie.

La Loi forte des grands nombres pose le problème de l'*unité de temps* pour les systèmes concurrents. Les outils géométriques décrits ci-dessus (§ III) montrent que si, pour les systèmes séquentiels et plus généralement pour les systèmes concurrents *sans confusion* [33], l'unité de temps est l'événement courant du système, il n'en est plus de même pour les systèmes concurrents plus généraux. Les *cellules de branchement* isolent des suites de décision qu'on est conduit à considérer comme atomiques, et qui sont comptabilisées comme telles pour une quantification du temps écoulé le long d'un processus : une unité de temps par cellule. La Loi forte des grands nombres donne la limite du rapport entre le nombre de réalisations d'une propriété locale au réseau, et le déroulement du temps comptabilisé de cette manière. Je mets en évidence la présence d'une densité asymptotique des cellules de branchement, qui correspond à la *mesure stationnaire* d'une chaîne de Markov ergodique.

Obtenir le Théorème de la Limite Centrale de manière entièrement intrinsèque aux réseaux peut sembler plus accessible que les théorèmes ergodiques. Je montre que les probabilités distribuées fournissent un cadre adapté pour l'application de la théorie des martingales. La propriété d'addition des espérances de carrés de

²Il faut se rappeler que la Loi forte des grands nombres pour les chaînes de Markov, et plus généralement pour les processus stationnaires, dérive en toute généralité du théorème ergodique de Birkhoff, dont la démonstration est fondée sur le lemme ergodique maximal. Il paraît peut-être trop ambitieux de vouloir démontrer un résultat analogue directement pour les systèmes concurrents.

martingales se retrouve couplée avec une addition “horizontale”, due aux processus concurrents. La composante horizontale est triviale (un singleton) pour les systèmes séquentiels. Cependant, le résultat final pour un Théorème de la Limite Centrale n’est pas donné dans ce document.

V—Méthodes algorithmiques

On aborde l’étude d’algorithmes utilisant les notions introduites plus haut.

On s’intéresse à la mise en œuvre opérationnelle d’une procédure d’estimation statistique des paramètres. Bien que les objets caractérisant les probabilités distribuées soient finis, leur calculabilité n’est pas évidente. Comme application, on propose une procédure pour l’estimation statistique de paramètres.

Pour la calculabilité des objets, les principales questions sont : peut-on décider si un réseau de Petri sauf est compact, c’est-à-dire si son dépliage est localement fini ? Si le réseau est compact, peut-on calculer les cellules de branchement initiales de son dépliage ? Je réduis partiellement ce problème à un problème d’atteignabilité, et je donne un algorithme qui finit toujours et décèle une sous-classe de réseaux non compacts. L’étude de la récurrence locale permet de dégager des suites de variables aléatoires indépendantes se prêtant donc parfaitement à l’estimation de paramètres, de la même façon qu’on étudie les estimateurs empiriques des chaînes de Markov finies. Cependant une différence due à la présence de la concurrence doit être prise en compte. Les données statistiques provenant du modèle sont par nature partiellement ordonnées, tandis que tout traitement opérationnel finit par atteindre une phase de traitement *séquentiel*. Pour modéliser ce décalage, je rajoute une contrainte d’observation au modèle sous la forme d’une variable non déterministe correspondant à la donnée d’une séquentialisation particulière d’un processus concurrent. En rapportant cette variables à d’autres variables, elles purement aléatoires, je montre que son effet est négligeable asymptotiquement, d’où une procédure opérationnelle d’estimation statistique.

L’étude de la composition de réseaux pour l’estimation distribuée est proposée pour des travaux futurs.

VI—Conclusion

Les réseaux de Markov distribués constituent une généralisation des chaînes de Markov à certains systèmes concurrents, les réseaux de Petri *compacts*. La dynamique de concurrence forte (*true-concurrent*) affaiblit la notion de temps global totalement ordonné. On introduit une construction par increments partiellement ordonnés pour une dynamique aléatoire de processus concurrents. Par comparaison avec l'espace des processus d'une chaîne de Markov, l'espace des processus concurrents d'un réseau n'est plus homogène par rapport aux translations dans le temps, ce qui oblige à reformuler un certain nombre de notions classiques pour l'étude des processus aléatoires, en particulier les temps d'arrêts, la propriété de Markov et la Loi forte des grands nombres. On montre la propriété de Markov concurrente, qui permet une étude de la récurrence des réseaux, et certains résultats ergodiques aboutissant à la Loi forte des grands nombres. On étudie la calculabilité des objets introduits, pour décrire une procédure opérationnelle d'estimation statistique.

Les méthodes introduites montrent que les outils analytiques permettent une approche originale et efficace des modèles de concurrence que sont les réseaux de Petri et les structures d'événements.

Introduction

The present introduction in English summarises the above Introduction in French.

This document proposes a mathematical and algorithmic analysis of concurrency models from Computer science: event structures and safe Petri nets. We describe above our contributions and their context. Centred parts of the text underline the description of new results.

I—Motivation and Context

Management of telecommunication networks is concerned by the development of tools for the analysis of distributed and asynchronous systems. In telecommunication networks the global state of a system has little meaning due to the asynchronism between the components of the system [6]. It is simply impossible to obtain a snapshot of the state of the system. Therefore the notion of global state of a system is not an appropriate notion for studying distributed systems, and is advantageously replaced by the notion of local state. Together with the notion of local state appears a notion of *concurrency* of processes. The interest of system management for probabilistic methods leads to study probabilistic extensions of models from Concurrency theory.

The development of temporal logics and the study of languages for communicating processes has encouraged the study of stochastic extensions of timed concurrency models like stochastic processes algebra [11, 40, 22]. Following the approach of Larsen and Skou [29], security issues are studied with probabilistic tests. Observational properties of models are studied through bisimulation equivalence, related to equivalence induced by probabilistic extensions of logics like PCTL [29, 5]. As an other probabilistic model for concurrent systems, stochastic Petri nets extend the queueing theory to networked systems, and are used, for instance, for performance evaluation of systems [4, 2, 3, 30].

These probabilistic models are based on the classical analysis of random processes with the form $(X_t)_t$, where t is a discrete or continuous time parameter. Typical

examples are given by Markovian processes such as Markov chains in discrete or continuous time. As a consequence, and for a networked system, this is like considering that all the components of the network synchronise on a *global clock*.

However an asynchronous execution of concurrent processes does not fit well this representation. If *local events* of concurrent processes are not causally related, keeping a chronological relation between them is unnecessary. Distinguishing distinct interleavings of concurrent events leads to useless computational complexity and should thus be avoided. On the contrary, identifying interleavings of concurrent events leads to the *true-concurrent* semantics for the dynamics of concurrency models [33], recently used for network management in [6, 8]. This new notion of process changes the representation of time, yielding a *partially ordered global* time. Hence local clocks are substituted to the global clock, and as a consequence the global time becomes partially ordered.

Very recently, probabilistic discrete-events models for the true concurrent dynamics of concurrency models has attracted interest [45, 7, 43]. This document proposes a general construction for true-concurrent probabilistic processes and a study of some of their properties. We establish asymptotic properties, in particular recurrence properties and the Strong law of large numbers, in the true-concurrent framework. We apply the results to the classical statistical problem of parametric estimation, and we propose related algorithmic procedures.

We use tools from Mathematics and from Computer science. Analytical tools come from Topology and Measure theory: metric spaces, projective systems of probabilities. The models and their semantics, together with some algorithms, come from Computer science: Petri nets, event structures and occurrence nets. This document contributes to the analysis of concurrency models in two ways: We propose new tools for a geometry on event structures, and we show that the analytical tools bring new elements for the study of concurrency.

II—Acyclic Models: Partially Ordered Time and Randomisation.

We study the model of *safe Petri nets* under their true-concurrent dynamics. Using the unfolding theory from Winskel [47, 33], the dynamics of a safe Petri net is captured by the dynamics of an acyclic net that represents the *set of phases* of the original Petri net. The acyclic net, called the *unfolding* of the Petri net, lies in the category of *occurrence nets*. A more abstract representation, still acyclic and equivalent from the dynamics point of view, is given by the *event structure* associated with the unfolding.

Acyclic models, and in particular the unfoldings of safe Petri nets, present a natural framework for randomisation of true-concurrent processes. The set of *maximal processes* is the adequate probability space to be considered. In this Introduction,

we call this space the *border at infinity* of the set of processes, or simply the border at infinity of the model. The probability for a finite process to occur is then given by the probability of the *shadow* of the finite process, as depicted in Figure 2.

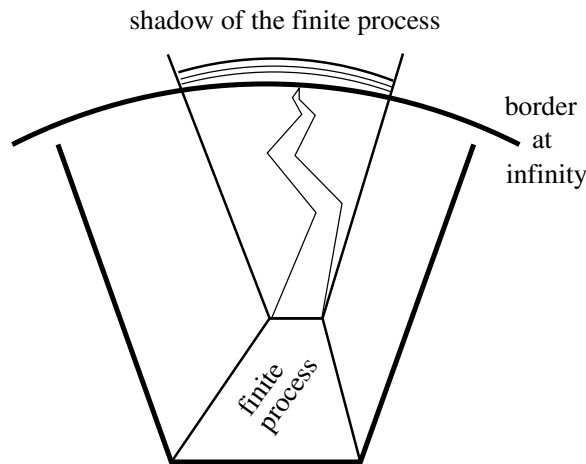


Figure 2: *The shadow of a finite process defines the probability for the process to occur.*

The introduction of this probability space is recent [45]. Authors have studied the way to extend probabilities defined on finite sets to a probability on the border at infinity (without this vocabulary) for some particular classes of safe Petri nets: extended free-choice Petri nets for Völzer [45], confusion-free nets for Varacca-Winkel-Völzer [44], and the so-called “choice-compact” nets of Benveniste-Haar-Fabre [7]. Benveniste *et al.* underline in [7] that the analytical basis for extension of probabilities in this framework is the extension of projective systems of probabilities, although the details are somewhat neglected.

In this document, I show that the partial order of processes of an acyclic model is given by a projective limit of finite sets, which induces a natural topology on the set of processes³. This is the key for the study of the border at infinity of the model. I show that the border at infinity can be realised as the limit of a sub-projective system if and only if the border at infinity is topologically closed in the set of processes, or equivalently if the border at infinity is compact. The extension of probabilities is then a consequence of Prokhorov’s extension Theorem (1930’s). The difference between sequential and concurrent systems is that the unfolding of a finite sequential system has always the compactness property, unlike finite concurrent systems.

³One finds this topology on closed models in the literature: Lawson topology, topology on traces [28, 27]. The definition as a projective topology has not been observed—although the metric constructions proposed are equivalent to the construction of the projective topology for metric spaces. Using the projective topology, compactness and separability of the space of processes follow from classical results.

I have replaced the notion of “stopping time for occurrence nets” of Benveniste *et al.* [7] with the notion of *intrinsic prefix*, and I have introduced the class of *locally finite* event structures, generalising the locally finite trees. This study has taken benefit from the notion of *dynamic conflict*, kindly indicated to me by D. Varacca and G. Winskel.

III—Probability and Concurrency

For locally finite event structures, the construction of a probability on the border at infinity reduces to the construction of a projective system of probabilities on finite sets. There is thus no fundamental difference in the extension process between sequential and concurrent systems.

For a *tree*, *i.e.* for a sequential and acyclic system, the border at infinity is the set of infinite paths of the tree. Constructing a probability on the border at infinity is like driving a boat on the sea, that goes *forward* towards the horizon. At each step, a random decision steers the boat to the left or to the right. In the concurrent framework, there are now several boats that go forward, and that can join by synchronisation or split to create new boats. Each boat has own probabilistic parameters, and other parameters reflect the mutual influences. Paying particular attention to probabilities that insures the *probabilistic independence* of concurrent boats is natural. This is the intuition of Benveniste *et al.* (conditional independence of layers), that agrees with Winskel’s maxim, according to which “parallelism is a form of orthogonality” [50].

I have introduced probabilities with the above independence properties under the name of *distributed probabilities*. Geometric objects related to event structures are introduced in order to properly define distributed probabilities. Local sub-event structures, called *branching cells*, are equipped with local probabilities, called *branching probabilities*. I show the existence and the uniqueness under natural conditions, of a distributed probability constructed from a family of branching probabilities. The construction of distributed probabilities describes an operation that associates a distributed probability with a family of branching probabilities. I call the operation the *distributed product*. The distributed product generalises the product of conditional probabilities, with a particular treatment for concurrency. When a system does not present concurrency, the distributed character of probabilities becomes trivial. Hence, every probability on an infinite product can be written as a distributed product.

This probabilistic model takes into account the concurrency of systems in an intrinsic manner, without using scheduler nor labelling as in asynchronous product of

probabilistic automata [40]. To the extend of my knowledge, the distributed product is the more general technique for randomising the dynamics of true-concurrent processes without reference to a global real time.

The construction of the distributed product in event structures is based on a decomposition of processes. The tools introduced with this decomposition constitute the basic tools for a *discrete geometry* on event structures. For acyclic models, I call *geometric* a property that concerns the elements of the model, by contrast with the two other objects associated with the model, the topological partial order of processes and the border at infinity. The general idea is to relate geometric properties of the model to topological and measurable properties for the spaces of processes and of maximal processes.

An iterative construction decomposes maximal processes. The decomposition isolates local processes, called *germs* of a maximal process. Concurrent germs of a same process have the independence property. Although the decomposition is not unique, due to the possible interleavings of germs, the collection of germs of a maximal process is uniquely defined. Germs are maximal processes of the branching cells, that are finite sub-event structures.

Branching cells are dynamic: for instance a same event can belong to different branching cells, according to the context in which the event is taken. It is thus different from *clusters* (see for instance [19, 21]). A branching probability randomises the germs of a branching cell. The distributed product associates the “horizontal” independence of concurrent germs with an independence in the probabilistic sense.

IV—Concurrent and Memory-less Random Systems

The construction of distributed probabilities applies in particular to define a true-concurrent randomisation for the dynamics of safe Petri nets. In order to obtain Limit theorems in this framework, we need to propose a notion of probabilistic memory-less dynamics for concurrent systems.

For this I have introduced a tool quite used in this work, the *cone of futures* of a process in acyclic models. The cone of future of a process v contains the possible events that can continue v . They form a sub-acyclic model, whose border at infinity identifies with the shadow of process v . If the border at infinity of the global acyclic model is equipped with a probability \mathbb{P} the border at infinity of the cone of v is then equipped with the *conditional probability* $\mathbb{P}(\cdot | v)$. If two processes v and v' in the unfolding of a safe Petri net lead to the same marking of the net, their cones are isomorphic as labelled occurrence nets, and this implies that the

shadows are isomorphic as measurable spaces. I say that a probability is *memory-less*, also called *homogeneous*, if this isomorphism of measurable spaces is actually an isomorphism of probability spaces, so that conditional probabilities are invariant. I show that a probability homogeneous and distributed is defined by a finite collection of parameters, each one ranging in a subset of a vector space of finite dimension.

I say that a safe Petri net is *compact* if its unfolding is locally finite. A compact net, whose unfolding is equipped with a probability homogeneous and distributed on the border at infinity is called a *distributed Markov net*. In this case the branching probabilities are in finite number and correspond to the rows of the transition matrix of a Markov chain. My aim is to show that **distributed Markov nets constitute a generalisation of finite Markov chains to concurrent systems**. For this, I show that several results well known for finite Markov chains extend to distributed Markov nets.

A preliminary study reformulates and establishes the Strong Markov property in the true-concurrent framework. The proof is self-contained, and the result implies the usual Strong Markov property for Markov chains. To state the Markov property, I propose a generalisation of stopping times classically studied with stochastic processes—also called optional times, introduced by Doob—under the form of *stopping operators*. Unlike stopping times of [7], stopping operators admit *hitting operators*, corresponding to hitting times of Markov chains. The new formulation of the Markov property is a work-around for the intrinsic difficulty with concurrent systems, that the space of processes is not invariant *w.r.t.* translations in the time space. The Markov property for concurrent systems allows a study of recurrence of nets, in the same way than the recurrence of finite Markov chains. I show the 0-1 alternative for a marking to be recurrent or transient, and that recurrence is a *conservative property*. Both results are inspired by Markov chains theory. I introduce an original *local recurrence* for nets. Local and global recurrences match for sequential systems, but not for concurrent systems.

Recurrence is a first example of an intrinsic property of nets expressed within the analytical formalism. Roughly speaking, a recurrent net comes back in its initial marking infinitely often with probability 1. Typically, a net may contain trajectories that do not return infinitely often to the initial marking, although they are very rare since all together they have probability zero. The analytical framework given by Measure theory, and in particular by Probability theory, is appropriate for dealing with negligible sets, whereas a complete verification would fail to detect recurrent nets. Recurrence also illustrates the following principle: properties concerning the states of a Markov chain are translated into a global property in memory-less nets. To obtain local results, we focus on distributed probabilities and take advantage from their local independence properties.

Some ergodic results, inspired by results from dynamical systems theory, are directly shown on the nets. With these results, I derive the Strong law

of large numbers for nets from the study of an auxiliary Markov chain. For the Strong law of large numbers to hold, and beside a recurrence condition imposed on nets, I introduce a condition to control the “concurrency range” of a net. A random variable, called *concurrent height*, modelises the concurrency range of the net. Again, the analytical formalism is crucial: requiring the variable to be bounded is too restrictive, whereas it is natural to impose that the concurrent height is integrable, *i.e.* has finite mean.

The Strong law of large numbers brings the problem of the *unit of time* for concurrent systems. The geometric tools introduced above (§ III), agree with the usual convention that the unit of time for a sequential system is the current event. The same holds for *confusion-free* concurrent systems [33], but not for more general concurrent systems. Branching cells isolate a partially ordered family of decisions that we consider as atomic. Hence we count branching cells to quantify the time elapsed along a process. The Strong law of large numbers gives the limit of the ratio between the number of occurrences of a local property in the net and the quantity of concurrent time elapsed. I show that distributed probabilities induce an asymptotic density of branching cells, that correspond to the stationary measure of an ergodic Markov chain.

Obtaining the Central Limit Theorem in a way fully intrinsic to nets seems to be possible using Martingale theory. Indications are given for the construction of partially ordered martingales for distributed probabilities. The final result for a Central Limit Theorem is not given in this document.

V—Algorithmic Methods

I study algorithms for computing some of the objects described above. Although the objects that characterise distributed probabilities are finite, their computability is not obvious. As an application, I propose a procedure for statistical parametric estimation.

The main computability questions are: can we decide if a safe net is compact, *i.e.* if its unfolding is locally finite? If the net is compact, can we compute the branching cells of the unfolding? I partially reduce the problem to a reachability problem. I give an algorithm that always ends and detects a class of non compact nets.

The local recurrence study shows that distributed probabilities define sequences of independent random variables from which the local prob-

abilistic parameters of the system can be retrieved. The statistical estimators associated correspond to the empirical estimator for Markov chains. Concurrency introduces however a specificity. Statistical data from the model are partially ordered, whereas I consider that any operational treatment reaches a phase of sequential computations (for instance, when data pass through channels and buffers). A non-deterministic (unknown but non random) sequentialisation of a concurrent process is thus introduced. I show that the effect of this variable is controlled by purely random variables asymptotically negligible, whence an operational statistical procedure. The problem of composition of nets for distributed estimation is proposed for a future work.

VI—Conclusion

Distributed Markov nets extend finite Markov chains to a class of concurrent systems, *compact* Petri nets. The true-concurrent dynamics weakens the notion of a global totally ordered time. As a consequence, the space of processes is not homogeneous *w.r.t.* translations in the time space. A random dynamics is constructed for partially ordered processes with partially ordered local increments. We define memory-less systems, and we show that they satisfy a Strong Markov property within a true-concurrent formulation. We apply the Markov property to the study of recurrent nets. Ergodic results yield the Strong law of large numbers, formulated with a concurrent unit of time. We study the computability of the objects introduced, and we describe a statistical estimation procedure.

The techniques introduced show that analytical tools provide an efficient and original approach for safe Petri nets and event structures.

Chapter 1

Preliminaries on Computational Models

In this first chapter, we present the computational models that we will study. Roughly speaking, there are two kinds of models: acyclic models, and finite machines.

A finite machine represents a “concrete device”, that we associate with an abstract acyclic model, representing the “set of phases” of the machine. An execution of the machine is then a trajectory in the acyclic model. For a sequential machine, like an automaton or a transition system, it is well known that the acyclic system is a tree, given by the covering of the graph of states and transitions of the system. Finite executions of the machine match finite paths in the tree starting from the root: we say that an execution is *lifted* into a path of the covering tree.

We study safe Petri nets as concurrency models. The concurrency of safe Petri nets expresses through the equivalence relation called *trace* equivalence or *interleaving* equivalence. The interleaving equivalence identifies two finite executions, seen as sequences of transitions, where two transitions occurring concurrently can be exchanged. In other words, we say that $ab = ba$ if transitions a and b are concurrent, and we close the relation on sequences of playing transitions by transitivity. What we call a *process*, or an *execution of the system*, or a *trace*, is an equivalence class of sequences of transitions, *modulo the trace equivalence*. This dynamics is also called the *true-concurrent* dynamics of the net, contrasting with the interleaving dynamics. For the reasons exposed in the Introduction, we are interested in the true-concurrent dynamics of nets.

It is remarkable that a representation of true-concurrent processes as “paths” of an acyclic system holds for safe Petri nets in a way formally equivalent to sequential systems. For a safe Petri net, the acyclic system analogous to the covering tree of a transition system is called the *unfolding* of the net. The unfolding of a safe net lies in the category of *occurrence nets*, and maps to the original net. True-concurrent processes in the net are *lifted* into true-concurrent processes of the unfolding. For this, a playing sequence of the net is first lifted into a playing sequence of the unfolding. Modulo trace equivalence in the unfolding, the lifted process is a congruence *w.r.t.* the trace equivalence in the net.

The concurrent models that we present are of three kinds, very closely related one with another. The safe and finite Petri nets constitute our finite concurrent machines. Their unfoldings lie in the category of labelled occurrence nets, or foldings. Event structures, and more precisely *prime* event structures, are an abstract model quasi equivalent to the model of occurrence nets.

Our aim is not to explicitly set up the elements from category theory that can be a basis for a presentation of concurrent models, although we insist on the relations between the different models. The goal of the chapter is rather to collect the different properties of models that will be used throughout the document, and to explain how we can go from a model to another.

Throughout the document, we will compare the results for concurrent systems with results for sequential systems: trees and transition systems. It is thus important to fix the notations and conventions that we will use for comparison. In particular, we will set the *sequential nets* that we use to simulate sequential systems. We also present the probabilistic framework associated with sequential systems, *i.e.* we present the basic framework of finite Markov chains.

The chapter begins in Section I, *Event structures*, with the model of prime event structures, which is the more abstract model, with the advantage of a simple formalism. Section II, *net models*, presents a quick overview of trace theory for safe nets. The unfolding theory is also shortly presented. The last Section, *Sequential systems*, fixes all our conventions concerning sequential systems. We also recall the basis of finite Markov chains theory, connected to simple probabilistic transition systems.

I—Event Structures

This Section recalls the basic definitions concerning event structures and their dynamics. We will only study *prime* event structures, thus we follow the classical presentation $(\mathcal{E}, \preceq, \#)$ ([33]). A presentation of general event structures is found in [48]. We will say *event structures* for short, *always* meaning prime event structures.

I-1 Partial Orders and Lattices.

We recall first basic well-known notions on partial orders and lattices. See for instance a presentation in [16].

I-1.1 Partial Orders. Let E be a set. A **relation** F on E is given by a subset of $E \times E$. We denote $x F y$ for $x, y \in E$ with $(x, y) \in F$. The inverse relation,

denoted by F^{-1} , is defined by:

$$x F^{-1} y \Leftrightarrow y F x .$$

The data (E, F) is said to be a **partial order** if F is reflexive transitive and satisfies:

$$\forall x, y \in E, \quad x F y \quad \& \quad y F x \Rightarrow x = y .$$

If (E, F) is a partial order, a sequence $(x_n)_n$ is said to be **non decreasing** if $x_n F x_m$ for all integers $n \leq m$. A non decreasing sequence is also called a **chain**. The sequence is said to be **increasing** if $n < m \Rightarrow x_n F x_m$ and $x_n \neq x_m$. The sequence is said to be **non increasing**, respectively **decreasing**, if F^{-1} is non decreasing, respectively increasing. A partial order (E, F) is said to be **well founded** if there is no infinite decreasing sequence.

A **morphism** of partial orders $f : (E, F) \rightarrow (E', F')$ is a mapping $f : E \rightarrow E'$ such that:

$$\forall x, y \in E, \quad x F y \Rightarrow f(x) F' f(y) .$$

Let A be a subset of a partial order (E, F) . The restriction $F|_A = F \cap (A \times A)$ defines the partial order $(A, F|_A)$.

I-1.2 Maximal Elements. Let (E, F) be a partial order. An element $x \in E$ is said to be **maximal** in E if:

$$\forall y \in E, \quad x F y \Rightarrow x = y .$$

x is said to be **minimal** if x is maximal in (E, F^{-1}) .

Let A be a subset of E . Let D^A be the subset of **upper bounds** of A defined by:

$$D^A = \{x \in E \mid \forall y \in A : y F x\} .$$

An element $a \in E$ is said to be a **least upper bound** of A if a is a minimal element of D^A , in which case a is unique. If such an element exists, in particular D^A is non empty.

We recall Zorn's Lemma, equivalent to the Axiom of choice, and that we will apply on twice in this document (I-2.12 below, and Ch. 3, III-1.2).

I-1.3 Theorem. (*Zorn's Lemma, [13]*) *Let (E, F) be a partial order. Assume that for every chain $A = \{x_i, i \in I\}$ of elements of E , with I any totally ordered set, admits an upper bound. Then E admits a maximal element.*

I-1.4 Lattices. A partial order (E, F) is said to be a **semi upper lattice** if every finite subset A of E admits a least upper bound in A . The semi lattice is said to be **complete** if every subset of E admits a least upper bound.

A partial order (E, F) is said to be a **lattice** if (E, F) and (E, F^{-1}) are two upper semi lattices.

A typical example of lattice is given by the powerset of a set O . If $E = \mathcal{P}(O)$ is the powerset of O , consisting of the subsets of O , (E, \subseteq) is a complete lattice. The prefixes of a configuration in an event structure are an other example, see I-2.8.

I-2 Event Structures.

Without other precision, a family of subsets of a same set is ordered by the inclusion \subseteq .

I-2.1 Definition. An **event structure** is a triple $(\mathcal{E}, \preceq, \#)$, where \mathcal{E} is a set at most countable, and $\preceq, \#$ are two binary relations satisfying the following axioms:

1. (\mathcal{E}, \preceq) is a partial order,
2. For every $e \in \mathcal{E}$, the set $\{x \in \mathcal{E} \mid x \preceq e\}$ is finite,
3. $\#$ is a symmetric and irreflexive relation on \mathcal{E} ,
4. (Inherited conflict) For all $x, y, e \in \mathcal{E}$, $e \# x$ and $x \preceq y \Rightarrow e \# y$.

The elements of \mathcal{E} are called the **events**, \preceq is the **causality relation**, $\#$ is the **conflict relation**.

The point 2. is fundamental from a computational point of view, since it means that an event in a computational execution has a finite number of ancestors. Remark that it implies that (\mathcal{E}, \preceq) is well founded.

I-2.2 The empty set with the empty relations is the empty event structure.

I-2.3 Graphical Representation of Event Structures. We represent event structures as in Figure 1.1. Events are drawn with a bullet, the causality relation is the transitive reflexive closure of the relation depicted by the oriented arcs. The conflict relation is generated by arcs drawn like: $\bullet \rightsquigarrow \bullet$, using the inherited conflict property. In Figure 1.1, the conflict relation is given by:

$$\# = \{(a, b), (a, c), (a, d)\}.$$

I-2.4 Relations. For $(\mathcal{E}, \preceq, \#)$ an event structure, we set the following relations, where Id denotes the identity relation $\{(x, x) \mid x \in \mathcal{E}\}$ on \mathcal{E} .

$$\prec = \preceq \setminus \text{Id}, \quad \succeq = (\preceq)^{-1}, \quad \succ = \succeq \setminus \text{Id}.$$

We define the **concurrency relation** \parallel by:

$$\forall x, y \in \mathcal{E} \quad x \parallel y \Leftrightarrow \neg(x \prec y) \text{ and } \neg(x \succ y) \text{ and } \neg(x \# y).$$




Figure 1.1: Representation of event structures. Generators of the causality and of the conflicts relations are drawn.

1-2.5 Extension of Relations to Sets. We extend the relations of conflict and of concurrency to subsets of \mathcal{E} as follows:

- Two subsets A, B are said to be concurrent, denoted by $A \parallel B$ if $a \parallel b$ for all $a \in A$ and for all $b \in B$.
- Two subsets A, B are said to be in conflict, denoted by $A \# B$ if there is a pair $(a, b) \in A \times B$ such that $a \# b$.

1-2.6 Sub-structures. Let A be a subset of an event structure \mathcal{E} . A is an event structure for the relations \preceq_A and $\#_A$, respectively defined as the restrictions to A of \preceq and $\#$. We will only consider the event structure $(A, \preceq_A, \#_A)$ for a subset $A \subseteq \mathcal{E}$.

1-2.7 Prefixes. A subset P of an event structure \mathcal{E} is said to be a **prefix** if P is \preceq -closed (downwards closed), *i.e.* if P satisfies:

$$\forall e \in P, \quad \forall x \in \mathcal{E}, \quad x \preceq e \Rightarrow x \in P.$$

The lattice of **finite prefixes** of \mathcal{E} is denoted by \mathcal{P}_0 . Using Axiom 2. in 1-2.1, every event e belongs to a finite prefix, the prefix $\{x \in \mathcal{E} \mid x \preceq e\}$.

1-2.8 Configurations. A prefix $v \subseteq \mathcal{E}$ is said to be a **configuration** of \mathcal{E} if no two nodes of v are in conflict. We denote by \mathcal{W} the set of configurations of \mathcal{E} , and by \mathcal{W}_0 the set of finite configurations. Both sets are ordered by inclusion, are stable under any intersection, and have \emptyset as unique minimal element.

The configurations included in a configuration w form a complete lattice, with union as least upper bound. They are called the **sub-configurations** of w .

Since the conflict relation is irreflexive and inherited (1-2.1), the finite prefix $\{x \in \mathcal{E} \mid x \preceq e\}$ is a configuration of \mathcal{E} for every $e \in \mathcal{E}$. We denote it by $[e]$. Obviously, $[e]$ is the smallest configuration that contains e . We denote by $[e[$ the configuration $[e[= [e] \setminus \{e\}$. If we have to specify the event structure \mathcal{E} , we use the notations $[e]^\mathcal{E}$ and $[e[^\mathcal{E}$.

1-2.9 Interpretation of Configurations. A configuration is interpreted as a *process* of a computational system. Events correspond to atomic actions of the system. A process is causally consistent: in a process, the causal predecessors of any event belong to the process, and the events of a same process do not exhibit any conflict.

The notion of compatibility of processes is of great importance.

1-2.10 Compatibility. Two configurations v, v' are said to be **compatible** if $v \cup v'$ is a configuration. Any family of pairwise compatible configurations is the family of sub-configurations of a configuration, the union of the family.

We say that an event e is **compatible** with a configuration v if $[e]$ and v are compatible in the above sense, *i.e.* if $[e] \cup v \in \mathcal{W}$.

1-2.11 Maximal Configurations. We denote by Ω the set of maximal elements of the partial order (\mathcal{W}, \subseteq) .

1-2.12 Lemma. *For each $v \in \mathcal{W}$, there is an element $\omega \in \Omega$ such that $\omega \supseteq v$.*

Proof— Let $\mathcal{W}(v)$ denote the set of configurations that contain v . Then any chain $(v_n)_{n \in I}$ of \mathcal{W}_v , where I is any totally ordered set, admits a least upper bound, that is $\bigcup_{n \in I} v_n$. By Zorn's Lemma (1-1.3), $\mathcal{W}(v)$ admits a maximal element, that is also maximal in \mathcal{W} . \square

1-2.13 Configurations in a Prefix. Let P be a prefix of an event structure \mathcal{E} . We denote by \mathcal{W}_P and by Ω_P respectively the set of configurations and the set of maximal configurations of P , as a sub-event structure (1-2.6). A conflict-free prefix of P is a conflict-free prefix of \mathcal{E} , so we have an injection $i : \mathcal{W}_P \hookrightarrow \mathcal{W}$. We will not mention the use of i , and state: a subset $v \subseteq P$ is a configuration of P if and only if v is a configuration of \mathcal{E} . In particular, for every $v \in \mathcal{W}$ and for every prefix P , $v \cap P \in \mathcal{W}_P$.

1-2.14 The Domain of Configurations. The structure of a partial order \mathcal{W} deriving from an event structure \mathcal{E} has been characterised in [47, 33]. We briefly recall the result.

Let (W, \sqsubseteq) be a partial order. An element $p \in W$ is said to be *complete prime* if for every subset $X \subseteq E$, if X admits a least upper bound a and if $p \sqsubseteq a$, then there is an element $x \in X$ such that $p \sqsubseteq x$. The partial order (W, \sqsubseteq) is said to be *prime algebraic* if for every element $v \in W$, if we set:

$$W_v = \{p \sqsubseteq v \mid p \text{ is complete prime}\},$$

then W_v admits v as least upper bound. A first result is the following:

Say that an event structure \mathcal{E} is elementary if the conflict relation is empty. Then the partial order of configurations (\mathcal{W}, \subseteq) of an elementary event structure is a prime algebraic complete lattice. Any prime algebraic complete lattice can be obtained this way.

\mathcal{E}, e	• an event structure, an event of \mathcal{E}
\mathcal{W}, w, v	• the partial order of configurations of \mathcal{E} , elements of \mathcal{W}
Ω, ω	• the set of maximal elements of \mathcal{W} , an element of Ω
$\mathcal{P}, \mathcal{P}_0$	• the sets of prefixes and of finite prefixes of \mathcal{E}
P	• a prefix or a finite prefix of \mathcal{E}
\mathcal{W}_P, Ω_P	• the set of configurations of prefix P , the set of maximal configurations of P .
$[e] = [e]^\mathcal{E},$ $[e] = [e]^\mathcal{E}$	• the smallest configuration containing event e , and $[e] \setminus e$

Table 1.1: *Basic Notations for Event Structures.*

The characterisation of partial orders of configurations \mathcal{W} for general event structures is treated in [33] using the condition of coherence from [31] as follows. Let (W, \sqsubseteq) be a partial order. A subset $A \subseteq W$ is *pairwise consistent* if any two of its elements have an upper bound in W . (W, \sqsubseteq) is said to be *coherent* if every pairwise consistent subset of W admits a least upper bound. Then we have:

Let $(\mathcal{E}, \preceq, \#)$ be an event structure. Then the partial order \mathcal{W} of configurations of \mathcal{E} is a prime algebraic coherent partial order. The complete primes of \mathcal{W} are the elements $[e]$, with e ranging over \mathcal{E} . Any prime algebraic coherent partial order can be obtained by this way.

II—Net Models

This Section describes the net models used throughout this document. Our basic objects are safe marked nets and their unfoldings. We present in II-1–II-2 the basis concerning safe nets and their dynamics from the trace theory point of view. We describe the model of occurrence nets in II-3. The unfolding theory, that relates safe nets to occurrence nets, is the topic of II-4.

II-1 Petri Nets and Safe Petri Nets.

Before we focus on safe Petri nets and on occurrence nets, it is useful to state the definition of nets in a slightly general way.

II-1.1 Definition. (*Nets*) Let P and T be two disjoint sets, both at most countable. Call P a set of **places**, and T a set of **transitions**. Let F be a relation on

$P \cup T$, called the **flow relation**. We assume that F connects places to transitions and transitions to places, hence F is given as a subset $F \subseteq (P \times T) \cup (T \times P)$. We assume further that no transition is isolated, *i.e.*:


$$\forall t \in T, \quad \exists b, b' \in P : b F t, t F b'.$$

We define a **net** as any triple $\mathcal{N} = (P, T, F)$, with P, T, F as above.

An element of the disjoint union $P \sqcup T$ is called a **node** of the net. For shortness we write $x \in \mathcal{N}$ to designate a node of \mathcal{N} .

Let e be a node of the net. We define the **preset** of e as the subset $\bullet e = \{x \in P \cup T \mid x F e\}$, and its **postset** by $e \bullet = \{x \in P \cup T \mid e F x\}$. The elements of the preset and of the postset of a node e are all of the same sort, places or transitions, which differs from the sort of node e .

Petri nets are graphically represented by circles for places and rectangles for transitions, and arcs for the flow relation, as follows:



We will now fix a net, and fill some places with tokens. The distribution of the tokens is thought of as the *state* of the system. A dynamics is introduced by a “rule of the game”, called the Petri net game, which allows to change the position of tokens. This rule introduces concurrency in the dynamics.

II-1.2 Multisets Notation. It is convenient to use a multiset notation. A multiset on set X is a mapping $A : X \rightarrow \mathbb{N}$. A subset A is represented by the multiset given by its characteristic function $\mathbf{1}_A : X \rightarrow \mathbb{N}$:

$$\mathbf{1}_A(x) = 1, \quad \text{if } x \in A, \quad \mathbf{1}_A(x) = 0, \quad \text{if } x \notin A.$$

The addition and subtraction are defined on multisets as functions $X \rightarrow \mathbb{N}$, as well as the order \leq .

II-1.3 Definition. (Markings, marked nets) Let M_0 be a multiset of places of a net \mathcal{N} , call M_0 a **marking**. We say that a place p is **isolated** *w.r.t.* the marking M_0 , if:

$$p \notin M_0, \quad \text{and} \quad \bullet p = \emptyset.$$

We define a **Petri net**, or a **marked net**, as any tuple $\mathcal{N} = (P, T, F, M_0)$, such that no place of \mathcal{N} is isolated *w.r.t.* M_0 . We call M_0 the **initial marking** of \mathcal{N} .

Very often, we omit the relation F , and we write $\mathcal{N} = (P, T, M_0)$ for a marked net.

II-1.4 Dynamics of Marked Nets (the Petri Net Game). Let $\mathcal{N} = (P, T, F)$ be a net, let M be a marking of \mathcal{N} , and let t be a transition of \mathcal{N} . We say that t is a **playing transition** of \mathcal{N} from marking M , or that \mathcal{N} can **play** transition t from M , if $\bullet t \subseteq M$. In this case, let M' be the marking given by:

$$M' = M - \bullet t + t \bullet .$$

We denote that \mathcal{N} can play t from M , together with the fact that $M' = M - \bullet t + t \bullet$, by:

$$\mathcal{N} : M \xrightarrow{t} M' .$$

If the net \mathcal{N} is clear from the context, we shortly write: $M \xrightarrow{t} M'$. The elements of $\bullet t$ are seen as *resources* that are consumed to *produce* $t \bullet$. The operation is illustrated in Figure 1.2. A marking is represented with tokens, the marking at right is the result of playing transition t from the marking at left.




Figure 1.2: *Playing transition t .*

We say that a sequence of transitions $r = (e_n)_{n \in I}$ of \mathcal{N} , indexed by a set with the form $I = \{1, \dots, N\}$ with $N \in \mathbb{N}$, or by $I = \mathbb{N}$, is a **playing sequence** of \mathcal{N} from M , if there is sequence of markings $(M_n)_{n \in I}$ such that we have, by setting $M_0 = M$:

$$\forall n \geq 1 \quad M_{n-1} \xrightarrow{e_n} M_n .$$

In this case the sequence $(M_n)_{n \in I}$ is unique. If the playing sequence r is finite indexed by $I = \{1, \dots, N\}$, we say that r **leads** from M to M_N . We denote it by $M \xrightarrow{r} M_N$. If \mathcal{N} is marked with initial marking M_0 , a marking M is said to be **reachable** if there is a playing sequence r such that $M_0 \xrightarrow{r} M$. A great simplification is brought as follows.

II-1.5 Definition. (*Safe Petri nets*) Let \mathcal{N} be a marked net. \mathcal{N} is said to be **safe**, or **contact-free**, or **1-bounded**, if for every reachable marking M of \mathcal{N} , $M \leq 1$.

Hence, for safe Petri nets, reachable markings identify with subsets of conditions, and in particular the initial marking.

II-2 True-Concurrent Dynamics of Safe Nets.

The trace dynamics on safe Petri nets is defined as follows [33, 39].

II-2.1 Definition. (*Interleaving relation, traces*) Let \mathcal{N} be a safe marked net, and let $\overline{\mathcal{R}}$ denote the set of playing sequences of \mathcal{N} . The **interleaving relation**, denoted by \sim , is the smallest equivalence relation on $\overline{\mathcal{R}}$ such that, for all transitions e, f :

$$\forall r = (\dots, e, f, \dots) \in \overline{\mathcal{R}}, \quad (\bullet e \cup e \bullet) \cap (\bullet f \cup f \bullet) = \emptyset \Rightarrow r \sim (\dots, f, e, \dots).$$

We denote by \mathcal{R} the quotient set $\mathcal{R} = \overline{\mathcal{R}}/\sim$. The equivalence class of a playing sequence is called its **trace modulo interleaving**. We denote by \mathcal{R}_0 the image in \mathcal{R} of finite playing sequences, and we call the elements of \mathcal{R}_0 the **finite traces** of \mathcal{N} .

The purpose of traces is to *not distinguish* between playing sequences, equivalent modulo interleaving of concurrent transitions.

II-2.2 True-Concurrent Dynamics and True-Concurrent Randomisation. Studying the true-concurrent dynamics of safe marked nets means studying its traces. Equivalently, a property on playing sequences is true-concurrent if it depends only on the traces of the playing sequences. We will explain later how this notion can be used for a true-concurrent randomisation.

II-2.3 Marking Associated with a Finite Trace. Let (\mathcal{N}, M_0) be a safe marked net. The mapping which associates with a finite playing sequence r the marking M such that $M_0 \xrightarrow{r} M$, is a \sim -congruence, and thus factorises through \mathcal{R}_0 . We say that a finite trace s **leads** to M if $M_0 \xrightarrow{r} M$ for any playing sequence r in the class s , and we denote it by: $M_0 \xrightarrow{s} M$.

II-2.4 Example. In the net of Figure 1.3, at left, the playing sequences ab and ba are equivalent since $\bullet a \cap \bullet b = \emptyset$. The resulting marking after playing their common trace is depicted at right.




Figure 1.3: *Transitions a and b are concurrent.*

II-2.5 Partial Order of Traces. For two finite traces s, s' , we define the relation $s \subseteq s'$ if and only if there are sequences r and r' belonging to the classes s and s' respectively, such that the sequence r is a prefix of r' . Then $(\mathcal{R}_0, \subseteq)$ is a partial order, that extends to \mathcal{R} as follows: for two traces $r, s \in \mathcal{R}$, $r \subseteq s$ if and only if there are non-decreasing sequences r_n and s_n of finite playing sequences, such that $\bigcup_n r_n$ and $\bigcup_n s_n$ belong to the classes r and s , and such that $r_n \subseteq s_n$ for all n .

The unfolding theory (II-4) realises \mathcal{R} as the partial order of configurations of an event structure. In this representation, the relation \subseteq just defined is conjugated to the inclusion relation on sets, and we will focus on this later model. The trace semantics for safe Petri nets give the “physical” meaning—since a finite Petri net is seen as a “physical” machine—of true-concurrent dynamics for more abstract objects such as event structures.

II-3 Occurrence Nets.

Occurrence nets are acyclic net models. We recall the basic material concerning occurrence nets: cuts, prefixes and configurations, and also the well-known relation between occurrence nets and event structures [47, 33].

II-3.1 Definition. (*Occurrence net*) Let $\mathcal{N} = (B, E, F)$ be a net, let $X = B \cup E$ denote the set of nodes of \mathcal{N} , and let \prec and \preceq denote respectively the transitive and the reflexive transitive closure of F on X . Define the **immediate conflict** relation $\#_1$ on transitions, and the **conflict** relation $\#$ on nodes by:

$$\begin{aligned} \forall e, e' \in E, \quad e \#_1 e' & \text{ iff } e \neq e', \bullet e \cap \bullet e' \neq \emptyset, \\ \forall x, x' \in X, \quad x \# x' & \text{ iff } \exists e \preceq x, \exists e' \preceq x' : e \#_1 e'. \end{aligned}$$

\mathcal{N} is an **occurrence net** if the following axioms are satisfied.

1. (X, \preceq) is partial order,
2. for every $x \in X$, $\{y \in X : y \preceq x\}$ is finite,
3. for every $b \in B$, $|\bullet b| \leq 1$,
4. $\#$ is irreflexive (no auto-conflict).

II-3.2 Terminology for Occurrence Nets. We adopt the usual convention according to which transitions $e \in E$ of an occurrence net are called **events**, and places $b \in B$ are called **conditions**. The relation \preceq is called the *causality* relation. We also set the following relations:

$$\prec = \preceq \setminus \text{Id}_X, \quad \succeq = (\preceq)^{-1}, \quad \succ = \succeq \setminus \text{Id}_X.$$

In Figure 1.4, examples of occurrence nets are depicted at right. Nets at left are not occurrence nets.

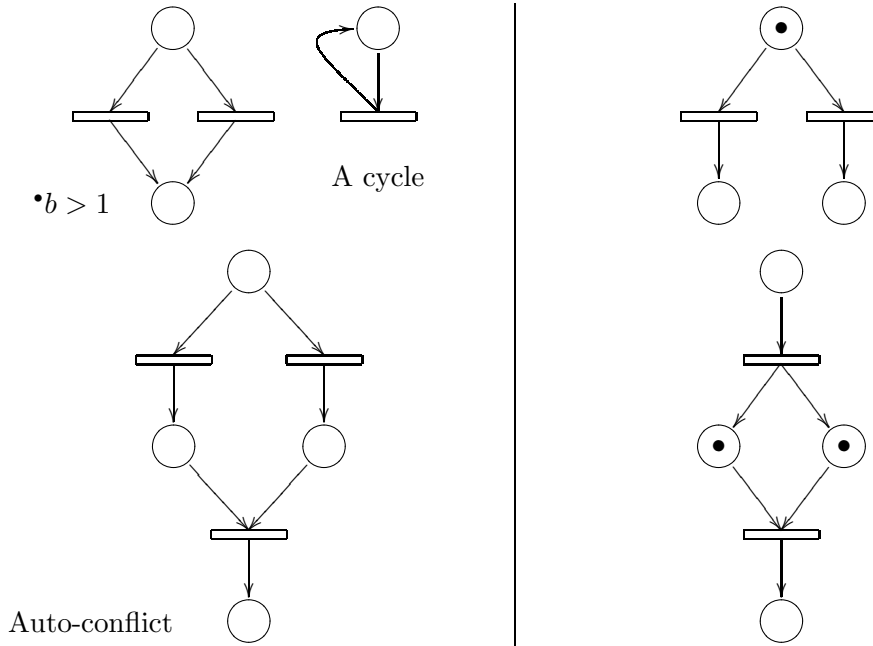


Figure 1.4: At left, non occurrence nets, at right, occurrence nets.

II-3.3 Concurrency and Concurrent Width. Define the *concurrency* relation \parallel on X by:

$$\parallel = X \times X \setminus (\prec \cup \succ \cup \#) .$$

We say that the occurrence net \mathcal{N} has *finite concurrent width* if every \parallel -clique of X is finite. This is equivalent to say that every \parallel -clique of conditions is finite, due to Axiom 3.

II-3.4 Cuts. Initial Cut. Let $\mathcal{N} = (B, E, F)$ be an occurrence net. A subset of conditions $A \subseteq B$ is said to be a *cut* if A is a \parallel -clique, maximal in B . In that case, A is also maximal in $X = E \cup B$.

The markings of Figure 1.4 at right indicate cuts of the occurrence nets.

We denote by $c_0 = \text{Min}_{\prec}(\mathcal{N})$ the set of minimal elements of \mathcal{N} , *w.r.t.* the causality relation. Since we assume that no event of the net is isolated (II-1.1), the elements of c_0 are the conditions b such that $\bullet b = \emptyset$. It follows from Axiom 2 in II-3.1 that the causality relation is well founded. Therefore c_0 is empty if and only if B is empty, *i.e.* if \mathcal{N} is the empty occurrence net. In all cases, c_0 is a cut of \mathcal{N} , which is called the *initial cut* of \mathcal{N} .

II-3.5 Open Subsets. Events of a Subset. Prefixes. Configurations. Let A be subset of X , where X denotes the set of nodes of an occurrence net \mathcal{N} .

We say that A is *open* if $\bullet e \subseteq A$ and $e^\circ \subseteq A$ for all events e of A .

We denote the set of events of A by $\overset{\circ}{A}$.

A is said to be a **prefix** of \mathcal{N} if:

- A is open,
- A contains the initial cut.
- A is **downward closed**, *i.e.*:

$$\forall x \in A, \quad \{y \in X \mid y \preceq x\} \subseteq A.$$

We say that a subset v of \mathcal{N} is a **configuration** of \mathcal{N} if v is a conflict-free prefix of \mathcal{N} , *i.e.* a prefix with no two nodes in conflict, or equivalently, with no two events in conflict. We insist that configurations contain the initial cut.

II-3.6 Relationship Occurrence Nets / Event structures. Let $\mathcal{N} = (B, E, F)$ be an occurrence net. Then $\mathcal{E} = (E, \preceq|_E, \#|_E)$ is an event structure, the canonic event structure of \mathcal{N} . The concurrency relation in \mathcal{E} coincides with the restriction to E of \parallel defined in \mathcal{N} .

There are natural isomorphisms of partial orders between the sets of prefixes, of finite prefixes, of configurations, of finite configurations, of \mathcal{N} on the one hand, and of \mathcal{E} on the other hand. The set of maximal configurations of \mathcal{N} and of \mathcal{E} are one-to-one, and all these mappings between subsets of \mathcal{N} and of \mathcal{E} are given by $A \mapsto \overset{\circ}{A} = A \cap E$. The inverse mappings are given by the operation that maps a subset $A \subseteq E$ to the smallest open subset of \mathcal{N} that contains A and the initial cut.

We mention that the converse operation, from event structures to occurrence nets, is also possible ([47, 33]). If \mathcal{E} is an event structure of finite concurrent width, the occurrence net \mathcal{N} that induces the event structure \mathcal{E} , resulting from the construction of [47, 33] is not of finite concurrent width in general since the initial cut obtained may be infinite, but a slight modification of the construction changes the occurrence net into an other one of finite concurrent width.

II-3.7 Notations. Due to the above isomorphisms, we keep the same notations that we have adopted for event structures (Table 1.1, page 43): \mathcal{W} and \mathcal{W}_P for the configurations and the configurations of a prefix P , *etc.*

A fundamental tool for studying configurations of occurrence nets is the following result, that relates cuts and configurations.

II-3.8 Proposition. *The mapping $\gamma : v \mapsto \gamma(v) = \text{Max}_{\preceq}(v)$ is one-to-one between the finite configurations of \mathcal{N} and the cuts of \mathcal{N} .*

II-3.9 Petri Net Game in Occurrence Nets. Let $\mathcal{U} = (E, B, F)$ be an occurrence net. There is a natural marked net associated to \mathcal{U} , that is $\mathcal{N} = (E, B, F, c_0)$, where c_0 denotes the initial cut of \mathcal{U} . If (e_1, \dots, e_n) is a finite playing sequence of \mathcal{U} , then:

$$c_0 + (e_1 + e_1^\bullet) + \dots + (e_n + e_n^\bullet),$$

is a configuration of \mathcal{U} . This association, from finite playing sequences to configurations, extends to infinite playing sequences, and is a trace congruence.

II-3.10 Proposition. *A marked occurrence net is safe. There is a natural isomorphism of partial orders Φ between finite traces of playing sequences, and finite configurations of the net. The marking associated with a trace s is given by $\gamma(\Phi(s))$.*

The isomorphism of partial orders naturally extends from finite traces to traces. The point is now that for any safe marked net \mathcal{N} , there is an occurrence net \mathcal{U} which realises the traces of \mathcal{N} as its own traces, *i.e.* as its own configurations. This is the topic of the unfolding theory.

II-4 Unfolding Theory.

As we have pointed out in the Introduction of this chapter, the unfolding theory is very closely related, at least formally, to the covering theory set up in various categories—topological, smooth or Riemannian manifolds, graphs. As in these classical frameworks, the unfolding of a safe net is characterised by a universal property. We thus need the notion of “covering”, *i.e.* a class of morphism. The appropriate definition of morphism for Petri nets is due to G. Winskel. Since we are only interested in unfoldings for the moment, we consider only the *foldings* ([49]) among the more general class of morphisms.

II-4.1 Definition. (*Labelled occurrence nets. Foldings.*) Let $\mathcal{N} = (P, T, F)$ be a safe net. A **labelled occurrence net** is a pair (\mathcal{U}, ρ) , where $\mathcal{U} = (B, E, G)$ is an occurrence net, and $\rho : \mathcal{U} \rightarrow \mathcal{N}$ is a mapping that respects the sort of the nodes. Hence we write $\rho = \beta \sqcup \eta$, with $\beta : B \rightarrow P$ and $\eta : E \rightarrow T$. We assume moreover that ρ satisfies the following property:

For each event $e \in E$, the restrictions $\beta|_{\bullet_e}$ and $\beta|_{e^\bullet}$ are two one-to-one mappings $\beta|_{\bullet_e} : \bullet_e \rightarrow \bullet(\eta(e))$ and $\beta|_{e^\bullet} : e^\bullet \rightarrow (\eta(e))^\bullet$.

Assume that M_0 is an initial marking of \mathcal{N} . Then (\mathcal{U}, ρ) is said to be a **folding** of (\mathcal{N}, M_0) if $\rho : \mathcal{U} \rightarrow \mathcal{N}$ is a labelled occurrence net, and if moreover:

The restriction $\beta|_{\text{Min}_{\preceq}(\mathcal{U})} : \text{Min}_{\preceq}(\mathcal{U}) \rightarrow M_0$ is a bijection between the initial markings of \mathcal{U} and of \mathcal{N} .

II-4.2 Definition. (*Morphisms of labelled occurrence nets and of foldings*) Let (\mathcal{U}, ρ) and (\mathcal{U}', ρ') be two occurrence nets, labelled by the same safe net \mathcal{N} . We say that a mapping $f : \mathcal{U} \rightarrow \mathcal{U}'$ is a **morphism of labelled occurrence nets** if (\mathcal{U}, f) is an occurrence net labelled by \mathcal{U}' , and satisfying $\rho' \circ f = \rho$.

Assume that (\mathcal{U}, ρ) and (\mathcal{U}', ρ') are foldings of (\mathcal{N}, M_0) . Then a mapping $f : \mathcal{U} \rightarrow \mathcal{U}'$ is said to be a **morphism of foldings** if f is just a morphism of labelled occurrence nets, and this implies that $f|_{\text{Min}_{\preceq}(\mathcal{U})} : \text{Min}_{\preceq}(\mathcal{U}) \rightarrow \text{Min}_{\preceq}(\mathcal{U}')$ is a bijection.

II-4.3 Theorem. (*Winskel*) Let \mathcal{N} be a safe marked net. There is a folding $(\mathcal{U}(\mathcal{N}), \rho)$ of \mathcal{N} , called **unfolding of \mathcal{N}** , satisfying the following universal property: For any folding (V, f) of \mathcal{N} , there is a unique morphism of foldings $g : V \rightarrow \mathcal{U}(\mathcal{N})$, such that the following diagram commutes:

$$\begin{array}{ccc}
 \mathcal{N} & \xleftarrow{\rho} & \mathcal{U}(\mathcal{N}) \\
 & \swarrow f & \uparrow \exists! g \\
 & & V
 \end{array}$$

The unfolding $\mathcal{U}(\mathcal{N})$ is unique, up to a unique isomorphism of foldings.

See [50, 18] for the construction of the unfolding.

II-4.4 Conservation of Trace Dynamics. The trace dynamics of the safe net is the same than the trace dynamics of its unfolding. Recall first that the initial markings M_0 and c_0 of the net and of its unfolding $\mathcal{U}(\mathcal{N})$ are one-to-one.

For $r = (e_i)_i$ a finite playing sequence of \mathcal{N} , there is a unique lifting of r in \mathcal{N} , that is a playing sequence $\bar{r} = (\bar{e}_i)_i$ of $\mathcal{U}(\mathcal{N})$, such that $\rho(\bar{e}_i) = e_i$ for all i : the lifted sequence projects into the original sequence. Define $v = \bar{\Psi}(r)$ as the configuration associated to \bar{r} as in II-3.9. Then $\bar{\Psi}$ is a congruence *w.r.t.* the interleaving.

The quotient mapping Ψ has the following property.

II-4.5 Proposition. ([33], Prop. 6) Let \mathcal{N} be a safe marked net, and let $(\mathcal{U}(\mathcal{N}), \rho)$ be the unfolding of \mathcal{N} . The mapping Ψ is an isomorphism of partial orders, which maps the finite traces of \mathcal{N} onto the finite configurations of $\mathcal{U}(\mathcal{N})$. For each finite trace r of \mathcal{N} , leading in \mathcal{N} to the marking M , if we set $v = \Psi(r)$, we have $M = \rho(\gamma(v))$, and the restriction $\rho|_{\gamma(v)} : \gamma(v) \rightarrow M$ is a bijection.

This result says that finite traces of playing sequences in \mathcal{N} are in bijection with the finite traces of playing sequences in $\mathcal{U}(\mathcal{N})$. And in an occurrence net, finite traces identify with finite configurations. Figure 1.5, left, depicts an example of safe marked net, with a prefix of the unfolding at right hand. In the representation of the unfolding, conditions and events are labelled with the names of corresponding places and transitions in the original net.




Figure 1.5: A safe marked net and a prefix of its unfolding

II-4.6 Concurrent Width of Unfoldings. Assume that \mathcal{N} is finite. It follows from II-3.8 and II-4.5 that the concurrent width of $\mathcal{U}(\mathcal{N})$ is finite. Indeed, any cut is one-to-one with a marking of the net \mathcal{N} , which is finite.

We will use later in this document the following Lemma.

II-4.7 Lemma. Let \mathcal{N} be a safe marked net. Let \mathcal{K} be the unfolding of \mathcal{N} , and let $m : \mathcal{O} \rightarrow \mathcal{K}$ be a morphism of labelled occurrence nets. Then m is injective.

Proof – It is enough to show that the restriction of m to conditions is injective. Since a conflict in \mathcal{O} induces a conflict in \mathcal{K} , it is enough to show that the restriction of m to configurations is injective. Let b, b' be conditions of a configuration $v \subseteq \mathcal{O}$ such that $m(b) = m(b')$. We may assume without generality that $b' \in \mathcal{K}^{[b]}$, since otherwise the converse would hold: $b' \in \mathcal{K}^{[b']}$. Hence there is a condition $p \in \gamma([b])$ such that $p \preceq b'$. It implies that $m(p) \preceq m(b') = m(b)$. On the other hand, since p and b belong to the same cut $\gamma(v)$ in \mathcal{O} , we have $p \parallel b$ and thus $m(p) \parallel m(b)$. It follows that $m(p) = m(b)$. But m is injective on cuts, hence $p = b$. Now we have $b \preceq b'$, so there is a chain of conditions from b to b' , that transports through m into a chain of conditions in \mathcal{K} , from $m(b)$ to $m(b') = m(b)$. Since \mathcal{K} is acyclic, the chain is trivial, which implies that $b = b'$. \square

III—Sequential Systems

Throughout this document, and in particular for probabilistic issues, we will try to generalise existing notions and results that are stated for sequential systems. We

will in particular refer to the theory of finite Markov chains.

Historically, Markov used first the probabilistic processes that he introduced at the beginning of the 20th Century to extend the range of application of the Weak law of large numbers to more general processes than independent sequences of real random variables. The study of *discrete* Markovian processes, and in particular of finite Markov chains, has also been an early topic of researches¹.

Markovian processes have become one of the most important probabilistic models, both in theoretical research and in practical applications of probabilities. In particular, the dynamics of sequential computational models is usually extended to a probabilistic dynamics based on finite Markov chains theory. Our goal is to provide an extension of discrete Markovian processes to concurrent systems.

III-1 Finite Markov Chains.

III-1.1 Canonical Markov Chain. Let E a finite set of states, and let μ be a probability measure on E , called **starting or initial probability measure** (we refer to Ch. 2, I-2 for background on probability). Let P be a transition matrix on E . By a transition matrix, we mean that P is a square matrix with elements $P_{x,y}$ indexed by $E \times E$, and such that:

$$\forall x, y \in E : P_{x,y} \in [0, 1], \quad \forall x \in E : \sum_{y \in E} P_{x,y} = 1.$$

The canonical Markov chain associated to the triple (E, P, μ) is defined as follows. Let \mathcal{A} denote the product space $\mathcal{A} = E^{\mathbb{N}}$, and let $X_n : \mathcal{A} \rightarrow E$ denote the canonical projections for $n \geq 0$. \mathcal{A} is equipped with the product σ -algebra \mathcal{F} , *w.r.t.* the discrete σ -algebra on E . Equivalently, we write $\mathcal{F} = \langle X_0, \dots \rangle$, meaning that \mathcal{F} is the smallest σ -algebra on \mathcal{A} that makes all the mappings X_n measurable. The Kolmogorov extension theorem implies that there is a unique probability measure \mathbb{P}_μ on \mathcal{A} , such that:

1. The law of X_0 in E is μ .
2. For all $x, y \in E$ and for all $n \geq 0$, for all arrays $(x_0, \dots, x_{n-1}) \in E^{n-1}$, we have:

$$\mathbb{P}_\mu(X_{n+1} = y \mid X_0 = x_0, \dots, X_{n-1} = x_{n-1}, X_n = x) = P_{x,y}.$$

This is also read as follows: The chain starts with initial distribution μ , and then recursively: given the present state x , the probability of jumping into state y does not depend on the previous history, but only on the present state x ; moreover this probability is given by the entry (x, y) in the transition matrix. We say that the

¹Markov and Kolmogorov have defined in the 1920's Markovian processes on words, introducing what is called today *speech recognition*.

process $(X_n)_n$ is memory less, since at each instant, its probability distribution in the future only depends on the present, and not on the past.

If the starting measure μ is a Dirac measure $\mu = \delta_{x_0}$, given by:

$$\delta_{x_0}(x) = 0 \quad \text{if } x \neq x_0, \quad \delta_{x_0}(x_0) = 1,$$

we say that the Markov chain starts with initial state x_0 . Indeed, we have $\mathbb{P}(X_0 = x_0) = 1$. We write $\mathbb{P}_{x_0} = \mathbb{P}_{\delta_{x_0}}$.

III-1.2 Dual Markov Chain. We introduce the notion of *dual Markov chain*, that will be of practical useful use for studying probabilistic transition systems.

Let $(X_n)_{n \geq 0}$ be a Markov chain on a finite set S , with transition matrix P , defined on the product space $\mathcal{A} = S^{\mathbb{N}}$. We define the **dual Markov chain** by:

$$\forall n \geq 1, \quad Y_n = (X_{n-1}, X_n).$$

The interpretation is that $(Y_n)_{n \geq 1}$ is the sequence of actions of the chain. With the point of view of a particle jumping from state to state, the chain X_n is the sequence of states, and the sequence Y_n is the set of jumps. Clearly, one defines the other. We formalise this below.

Let A denote the set $A = S \times S$, called set of **actions**. It is readily seen that $(Y_n)_{n \geq 1}$ is a Markov chain on A . If μ is the starting measure of $(X_n)_n$, the law of Y_1 in A is called the **dual starting measure**, and is given by:

$$\nu(x, y) = \mu(x)P_{x,y}.$$

The transition matrix Q of $(Y_n)_{n \geq 1}$ is called the **dual transition matrix**, and is given by:

$$Q_{(x,y),(x',y')} = \mathbf{1}_{\{y=x'\}}P_{x',y'}.$$

We set the following mappings, given by the projections on first and second components of actions $A = S \times S$:

$$\partial_-, \partial_+ : A \rightarrow S.$$

Let \mathcal{A}' be the product space $\mathcal{A}' = A^{\mathbb{N}}$. There is an injection $k : \mathcal{A}' \rightarrow \mathcal{A}$, $k(a_1, a_2, \dots) = (\partial_-(a_1), \partial_-(a_2), \dots)$, that satisfies the following property.

III-1.3 Proposition. *Let \mathbb{P}_μ be the probability measure on $\mathcal{A} = S^{\mathbb{N}}$ associated with the Markov chain $(X_n)_{n \geq 0}$ on S , with starting measure μ on S and transition matrix P . Let \mathbb{Q}_ν be the probability measure on $\mathcal{A}' = A^{\mathbb{N}}$, with $A = S \times S$, associated with the dual starting measure on A and the dual transition matrix on A . Then:*

$$(\mathcal{A}, \mathbb{P}_\mu) \rightarrow (\mathcal{A}', \mathbb{Q}_\nu) \quad (X_n)_{n \geq 0} \mapsto (Y_n)_{n \geq 1}, \quad Y_n = (X_{n-1}, X_n) \quad \forall n \geq 1,$$

is an isomorphism of probability spaces.

III-2 Probabilistic Transition Systems.

III-2.1 Basic Definitions. Let (E, A) be a finite oriented graph with E a set of nodes, and with set of arrows a family A of pairs $(x, y) \in E \times E$. We denote by $(\partial_-, \partial_+) : A \rightarrow E \times E$ the natural mappings that indicate the initial and final nodes of arrows. We say that E is the space of states, and that A is the space of actions or of transitions, and we consider an initial state s_0 . We define the triple (E, A, s_0) as a **transition system**.

For each state x , let Dx be the set of states y connected from x by an action $(x, y) \in A$, and let μ_x be a probability measure on Dx . Implicitly, we assume thus that Dx is non empty. We say that $(E, A, s_0, (\mu_x)_{x \in E})$ is a **probabilistic transition system**. Its dynamics is described by the trajectories of a particle starting from $x_0 = s_0$ at time $n = 0$. At instant $n \geq 0$, the particle jumps from its current state x_n to a state x_{n+1} according to the result of a coin, that has μ_x as probability distribution on Dx . Hence the particle follows the arrows in the graph, taking at each node a decision that is random, and independent of the previous decisions. Probabilistic transition systems are represented as in Figure 1.6 as a graph with weighted arcs.




Figure 1.6: A probabilistic transition system with three states. States are depicted by circles, and actions by arrows. An arrow is equivalent to the pair of its initial and final states. The initial state is filled with a token.

A special symbol Δ can be added to E , to carry the cases where no further decision can be taken, *i.e.* if $Dx = \emptyset$.

III-2.2 Markov Chain of a Probabilistic Transition System. Let $(E, A, x_0, (\mu_x)_{x \in S})$ be a probabilistic transition system. The successive states X_n of the particle form a Markov chain on E starting from x_0 . The transition matrix P on $E \times E$ is given by $P_{x,y} = \mu_x(y)$ if (x, y) is an arrow of the graph, $P_{x,y} = 0$ otherwise. It is like adding arrows, but with probabilistic weight zero. We call $(X_n)_{n \geq 0}$ the canonical Markov chain associated with the probabilistic transition system.

III-2.3 Dual Transition System. If $(X_n)_{n \geq 0}$ is the Markov chain that describes the states of a probabilistic transition system, the dual chain $(Y_n)_{n \geq 1}$ describes the transitions of the system. Figure 1.7 depicts the dual transition system of Figure 1.6.




Figure 1.7: Dual transition system of transition system depicted in Figure 1.6. Each state \bullet of the dual transition system is an arrow of the original system.

III-3 Sequential Net Associated with a Transition System.

We fix the rule that we use to associate a safe Petri net to a transition system. Let (E, A, x_0) be a transition system. We simply consider the net (E, A, F, x_0) , with same places than states in E , with set of transitions the set of actions A . The flow relation F is defined by (Cf. Figure 1.8):

$$\begin{aligned} \forall x \in E, \forall a \in A, \quad \partial_-(a) = x \Rightarrow xFa, \\ \forall x \in E, \forall a \in A, \quad \partial_+(a) = x \Rightarrow aFx. \end{aligned}$$

Hence, for each arrow $x \rightarrow y$ in the transition system, we put a transition t with $x \rightarrow t \rightarrow y$, and that satisfies: $\bullet t = x$ and $t \bullet = y$. Each reachable marking has a unique element (carries a unique token), and in particular the net is safe. A playing sequence $(r_n)_n$ in the net determines a unique sequence of transitions in the transition system, and we have observed that this determines a unique sequence of states $(X_n)_n$, with $X_0 = x_0$. The sequence of states corresponds to the successively reached markings in the net. This association is one-to-one.




Figure 1.8: From a transition system to the sequential net.

III-4 Trees and Probabilistic Trees.

Trees are the acyclic model associated with sequential systems. We define trees in two ways: as graphs, and as event structures ([48]). We only consider *oriented*

trees.

III-4.1 Trees as Graphs. Let T be a set of vertices, an *oriented arrow* is a pair $(x, y) \in T \times T$. An **oriented graph** is a pair (T, G) , where G is a set of oriented arrows on a set of vertices T . Let \leq denote the reflexive transitive closure of G . We say that (T, G) is a **tree** if the following conditions hold:

1. If T is non empty, \leq admits a unique minimal element, called the **root** of the tree.
2. (T, \leq) is a partial order.
3. $\forall x, y, z \in T, \quad xGz \ \& \ yGz \Rightarrow x = y$.

III-4.2 Trees of Events. We can also follow Winskel in [48] for the definition of trees, and define a **tree of events** as an event structure $(\mathcal{T}, \preceq, \#)$ satisfying this property: For all pairs (v, v') of configurations of \mathcal{T} , if v and v' are compatible then v and v' are comparable. That is:

$$\forall v, v' \in \mathcal{W}, \quad v \cup v' \in \mathcal{W} \Rightarrow v \subseteq v' \text{ or } v' \subseteq v, \quad (1.1)$$

where \mathcal{W} denotes the set of configurations of \mathcal{T} . It follows that every configuration of \mathcal{T} is a totally ordered subset of (\mathcal{T}, \preceq) . We describe below the relation between the two models of trees (see also Figure 1.9).

Let G be the binary relation on a tree of events \mathcal{T} , that connects any event to its immediate successors. Then G satisfies the points 2 and 3 of III-4.1, but does not satisfy point 1: \mathcal{T} has several minimal events in general. There is a slight difference: a tree of events \mathcal{T} is a disjoint union of trees in the above sense, each pair of distinct roots being in conflict. The roots of the different trees that compose \mathcal{T} are called the **roots** of \mathcal{T} . Every configuration of $(\mathcal{T}, \preceq, \#)$ is a path of (\mathcal{T}, G) , that can be empty, and that contains exactly one of the roots if non-empty.

Conversely, let (T, G) be a tree, we define a tree of events associated with (T, G) . We define the causality relation \preceq on T as the reflexive and transitive closure of G . The conflict relation $\#$ is defined on T as the smallest conflict relation on T that contains the pairs of distinct roots, and such that:

$$\forall x, y, z \in T, \quad xGy \ \& \ xGz \Longrightarrow y \# z \text{ or } y = z.$$

For simplicity, we assume that the tree is everywhere infinite: every $x \in T$ has a successor by G . Then the maximal configurations of $(T, \preceq, \#)$ as a tree of events match the infinite paths of (T, G) as a tree.

III-4.3 Covering of a Transition System. Let (S, A, x_0) be a transition system (we assume that all nodes have successors). We define a tree (T, G) by setting:

$$T = \{(x_0, \dots, x_n), n \geq 0 \mid (x_0, \dots, x_n)G(x_0, \dots, x_n, y) \Leftrightarrow (x_n, y) \in A\}.$$




Figure 1.9: From a tree as graph to a tree of events.

The covering of a transition system is illustrated by Figure 1.10. Let $p : T \rightarrow S$ be the mapping defined by $p : (x_0, \dots, x_n) \mapsto x_n$. Then the pair (T, p) is the universal covering of the graph (S, A) , starting from x_0 , which means the following:

For every finite sequence of states $(x_i)_i$ in the transition system, there is a unique path $(e_i)_i$ in the tree, starting from the root and with $p(e_i) = x_i$ for all i . The path $(e_i)_i$ is the *lifting* of $(x_i)_i$.




Figure 1.10: Covering of a transition system.

III-4.4 Probabilistic Covering. Assume moreover that $(\mu_x)_{x \in S}$ is a family of probabilities that makes $(S, A, x_0, (\mu_x)_x)$ a probabilistic transition system. Let \mathcal{K} denote the set of infinite paths in the tree (T, G) , and let \mathcal{A} denote the product space $\mathcal{A} = S^{\mathbb{N}}$. There is an injection $k : \mathcal{K} \hookrightarrow \mathcal{A}$. The space \mathcal{K} has \mathbb{P}_{x_0} probability 1 in \mathcal{A} . It implies that the injection k admits an inverse mapping \mathbb{P}_{x_0} -a.s defined, $k^{-1} : \mathcal{A} \rightarrow \mathcal{K}$. Equipped with the probability $k^{-1}\mathbb{P}_{x_0}$, \mathcal{K} is isomorphic as a probability space to $(\mathcal{A}, \mathcal{F}, \mathbb{P}_{x_0})$.

Probabilistic transition systems belong to the class of *sub-shifts of finite types*, a model studied in dynamical systems theory ([42]).

III-4.5 Unfolding of a Sequential Net. Let $\mathcal{N} = (S, A, F, x_0)$ be the safe marked net associated with the transition system (S, A, x_0) . Let (\mathcal{U}, ρ) be the unfolding of \mathcal{N} . Let H denote the set of conditions of \mathcal{U} . Then H , equipped with the graph structure induced by causality, and with the mapping $\rho|_H : H \rightarrow S$, is the

covering of the transition system S . In particular the initial cut of \mathcal{U} has a unique condition, labelled with the initial state x_0 . The cuts of the unfolding are given by any unique condition $b \in \mathcal{U}$.

The underlying event structure \mathcal{T} of \mathcal{U} is a tree of events (III-4.2). Let $a_i = \rho(r_i)$ be the label of the root r_i of T_i , with $\mathcal{T} = T_1 \sqcup \dots \sqcup T_n$ the decomposition of \mathcal{T} as a disjoint union of trees with a unique root. Then a_i ranges over the set $\partial_-^{-1}(x_0)$. Each T_i is the covering of the **dual transition system** of S , starting from different roots r_i . That is, the transition system with set of states A , and with arrows the pairs $(a, b) \in A \times A$ such that $\partial_+(a) = \partial_-(b)$.

III-4.6 Probabilistic Unfolding. The set Ω of maximal configurations of \mathcal{T} is one-to-one with the infinite paths in the covering H . Using III-4.4, Ω can thus be equipped with a probability that makes it isomorphic as a probability space to $(\mathcal{A}, \mathcal{F}, \mathbb{P}_{x_0})$, and we write $(\Omega, \mathcal{F}, \mathbb{P})$ to denote it.

This isomorphism, together with the isomorphism between a chain and its dual chain of III-1.3, allows to consider that the Markov chains $(X_n)_{n \geq 0}$ and $(Y_n)_{n \geq 1}$ are defined on the space $(\Omega, \mathcal{F}, \mathbb{P})$. They respectively give the successive markings and transitions of an execution of \mathcal{N} , and we have the equality of σ -algebras:

$$\forall n \geq 1, \quad \langle X_0, \dots, X_n \rangle = \langle Y_1, \dots, Y_n \rangle. \quad (1.2)$$

III-5 More General Transition Systems.

The model that we consider is simpler than probabilistic transition systems usually encountered in the literature. These models often consider that the set A of actions is not given as a set of arrows. Each action determines an arrow, but the mapping $(\partial_-, \partial_+) : A \rightarrow E \times E$ is not injective. Hence different actions can induce the same move of the particle in the transition system.

This 2-step semantics is intended for instance to set up synchronous products of transition systems, with application to the definition of grammars of processes for process algebra [22, 41]. If the choice of action a_n is random, depending only on the current state X_n , then the sequence $(X_n, a_n)_{n \geq 0}$ is a finite Markov chain, thus we are back to our simplified model. Asynchronous products of automata also consider non deterministic actions, in the sense: non determined, but non random. Nevertheless, the random part of the dynamics is always based on the model of finite Markov chains.

This justifies that we only use simple probabilistic transition systems as typical examples of sequential systems, since they can simulate any finite Markov chain. Adding labelling on nets for studying the probabilistic bisimulation of nets is an other step, not treated in this document.

Chapter 2

Topological Event Structures

Although we are interested in discrete models with discrete events, these models generate—in general—a non countable set of processes. Concepts and tools from mathematics are thus required to analyse the set of processes, in particular from topology and measure theory. We focus in this chapter on the topological properties of processes, measurability properties will follow.

Event structures are an acyclic model for concurrent systems, where processes are described by configurations. Among all configurations of an event structure, we pay particular attention to the set of *maximal configurations*. This space is the natural support for probability measures. Hence we are interested in the topological properties of the set of maximal processes, call it Ω , and not really of the set of all processes, let us call it \mathcal{W} . However, the topology on \mathcal{W} is the key for studying the topology on Ω . It appears that \mathcal{W} naturally identifies with the projective limit of a projective systems of finite sets. We equip thus \mathcal{W} with the topology inherited from the projective limit. This topology brings us into the familiar framework of compact metric spaces.

We have in mind the application to Ω of the Prokhorov extension theorem, that extends projective systems of probability measures. For this, we need to identify Ω with a projective limit. We define first from Ω a natural projective system of finite sets. Then we show that Ω identifies with the projective limit if and only if Ω is compact, *i.e.* closed in \mathcal{W} .

Can we see the compactness of Ω in the event structure? We give a geometric condition, *i.e.* a condition that concerns the event structure, insuring compactness: We introduce intrinsic prefixes, and the condition is that every event belongs to a finite intrinsic prefix. We also underline the computational interest of intrinsic prefixes: their own dynamics simulates the local dynamics of the event structure, which is not the case in general for arbitrary prefixes.

In particular, *locally finite event structures* are shown to satisfy this condition. I have studied the class of locally finite event structures to improve the study based on stopping times for occurrence nets of Benveniste *et al.* in [7]. The study has taken benefit from the definition of the *dynamic conflict relation*, kindly communicated to

me by D. Varacca and G. Winskel.

In Section I, *Background on Projective Systems*, we recall some notions from topology and measure theory, in order to present the theory of product spaces and projective limits. We also recall the vocabulary from probability theory: random variables, probability law, *etc.* We present the extension theorem that we will apply for projective systems of probability measures, a version of the Prokhorov theorem from Bourbaki. In Section II, *Defining topologies for event structures*, we successively define and study topologies on \mathcal{W} , the partial order of configurations, and on Ω , the set of maximal configurations—the boundary at infinity. We also introduce associated projective systems. The main result is the condition for Ω to be homeomorphic to a projective limit, the compactness of Ω . We apply this result with the Prokhorov extension theorem in Section III, *Extension of probability measures*. We introduce the intrinsic prefixes, and we show the compactness of Ω under the condition that finite intrinsic prefixes cover the whole event structure. We apply this result to stopping prefixes and to locally finite event structures, and we state the extension theorem for locally finite event structures.

Some remaining open questions are the topic of a discussion in IV.

I—Background on Projective Systems

We recall some notions from topology, in particular concerning products and projective limits. We introduce the basis of vocabulary from probability theory, and we give a version of the Prokhorov extension theorem from Bourbaki.

I-1 Basic Notations.

We denote by \mathbb{N} the set of integers, and we use $\overline{\mathbb{N}}$ to denote $\mathbb{N} \cup \{\infty\}$. We denote by \mathbb{R} the set of real numbers.

If X is a set, we identify an element $x \in X$ and the singleton $\{x\}$.

I-2 Topology and Probability.

I-2.1 Topological Spaces and Metric Spaces. A *topological space* is a pair (E, τ) , where E is a set and τ is a collection of subsets of E , called the collection of *open sets* of E , such that: τ contains E and the empty set, τ is stable under any union and under finite intersection. Any complement cU of U in E is by definition a *closed subset* of E . If τ' is an other topology on E such that $\tau' \subseteq \tau$, then τ' is said to be *weaker* than τ .

The space (E, τ) is said to be *Hausdorff* if for every disjoint elements x, y of E , there are disjoint open sets U, V such that $x \in U$ and $y \in V$. The *closure* \overline{F} of a subset $F \subseteq E$ is the smallest closed subset that contains F . F is *dense* in E if $\overline{F} = E$. E is said to be *separable* if there is a subset $F \subseteq E$ at most countable and dense in E .

For any subset $F \subseteq E$, we denote by $\tau|_F$ the *restriction* of τ to F , the topology on F given by the collection of sets $U \cap F$, where U ranges over τ .

A sequence $(x_n)_{n \geq 1}$ in a topological space (E, τ) is said to be *convergent* to an element $x \in E$ if:

$$\forall U \in \tau, x \in U, \quad \exists N \geq 1 : n \geq N \Rightarrow x_n \in U,$$

and the limit is unique if the space is Hausdorff.

If E is a set, a non-negative and symmetric function $d : E \times E \rightarrow \mathbb{R}$ is a *metric* if for every pair (x, y) of elements of E : $d(x, y) = 0 \Rightarrow x = y$, and if for every triple (x, y, z) , the triangular inequality is satisfied: $d(x, z) \leq d(x, y) + d(y, z)$. The pair (E, d) is said to be a *metric space*. We denote by $B(x, r)$ the open ball of centre x and of radius r . An open set of (E, d) is defined as any subset $U \subseteq E$ such that:

$$\forall x \in U, \quad \exists r > 0 : B(x, r) \subseteq U,$$

and then the collection of open sets defines a topology on E .

A subset $A \subseteq E$ of a metric space (E, d) is said to be *compact* if every sequence of A admits a subsequence that converges in A . Compactness is intrinsic: a subset $A \subseteq E$ is compact as a subset of E if and only if A is compact *w.r.t.* the induced topology on $(A, d|_A)$. A closed subset of a compact metric space is compact.

1-2.2 Measurable Spaces and Random Variables. A σ -algebra \mathcal{F} of a set E is a collection of subsets of E , containing \emptyset , and stable under: complement in E , countable union, and countable intersection. The pair (E, \mathcal{F}) is said to be a measurable space. The elements of \mathcal{F} are called the measurable subsets of (E, \mathcal{F}) , or of E for short. A mapping $X : E \rightarrow Y$ between two measurable sets (E, \mathcal{F}) and (Y, \mathcal{G}) is said to be measurable if for every measurable subset $A \subseteq Y$, $X^{-1}(A)$ is measurable in E . X is also called a random variable. We denote by $\langle X \rangle$ the σ -algebra generated by X , defined by:

$$\langle X \rangle = \{X^{-1}(A), A \in \mathcal{G}\}.$$

1-2.3 Borel σ -Algebra. If (X, τ) is a topological space, the Borel σ -algebra is defined as the smallest σ -algebra that contains the open subsets of X . The σ -algebras that we consider in this document are always given by Borel σ -algebras on topological space. For example, the set \mathbb{R} is equipped with its Borel σ -algebra.

1-2.4 Probability Space and Probability Law (Image Probability). Let (Ω, \mathcal{F}) be a measurable space. A probability measure, or a probability on (Ω, \mathcal{F}) is a set function $\mathbb{P} : \mathcal{F} \rightarrow [0, 1]$, satisfying $\mathbb{P}(\Omega) = 1$, and such that for every sequence $(B_n)_{n \geq 0}$ of disjoint measurable spaces:

$$\mathbb{P}\left(\bigcup_{n \geq 0} B_n\right) = \sum_{n=0}^{\infty} \mathbb{P}(B_n).$$

Let $f : \Omega \rightarrow E$ be a random variable, *i.e.* a measurable mapping with values in a measurable space (E, \mathcal{G}) . The **probability law** of f in E , or the law of f in E , or the **image** of \mathbb{P} under f , is the probability measure on F , denoted by $f\mathbb{P}$, and defined by:

$$\forall A \in \mathcal{G}, \quad f\mathbb{P}(A) = \mathbb{P}(f^{-1}(A)).$$

This action is indeed a left action, *i.e.* we have whenever the composition is well defined:

$$(f \circ g)\mathbb{P} = f(g\mathbb{P}).$$

Let $X : \Omega \rightarrow E$ be a random variable. Let \mathbb{P} be a probability on Ω , and let \mathbb{P}_X be the law of X in E (\mathbb{P}_X is a probability on E). Then we have the *transfer formula*, for every non-negative measurable function $h : E \rightarrow \mathbb{R}$:

$$\int_{\Omega} h(X) d\mathbb{P} = \int_E h(x) d\mathbb{P}_X(x).$$

I-2.5 Finite Probability Spaces. If Ω is finite, and unless otherwise specified, we consider the discrete σ -algebra \mathcal{F} given by the powerset of Ω . A probability \mathbb{P} is then an additive set function on the algebra of sets \mathcal{F} . Any function $f : \Omega \rightarrow \mathbb{R}$ such that:

$$\forall x \in \Omega, \quad f(x) \in [0, 1], \quad \sum_{x \in \Omega} f(x) = 1,$$

determines a unique probability such that $\mathbb{P}(\{x\}) = f(x)$ for all $x \in \Omega$. \mathbb{P} is given by: $\mathbb{P}(A) = \sum_{x \in A} f(x)$.

I-2.6 Notation of Subsets from Probability. We will not use the word event for measurable subsets of probability spaces, to avoid confusion with the events from event structures. However we will keep the usual notation that makes us write $\{\phi\}$ to denote the set of elements $\omega \in \Omega$ that satisfy the property $\phi(\omega)$. For instance, if $T : \Omega \rightarrow \mathbb{N}$ is an integer random variable, we write:

$$\begin{aligned} \{T = n\} &= \{\omega \in \Omega \mid T(\omega) = n\}, \\ \mathbb{P}(T = n) &= \mathbb{P}(\{\omega \in \Omega \mid T(\omega) = n\}). \end{aligned}$$

I-2.7 Properties Almost Sure (a.s). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. We say that a property $\Phi(\cdot)$ that depends on $\omega \in \Omega$, is true \mathbb{P} **almost surely**, abbreviated in \mathbb{P} -*a.s.*, if there is a measurable subset $A \subseteq \Omega$ with $\mathbb{P}(A) = 0$, such that $\Phi(\omega)$ is true for all $\omega \notin A$. In other words¹, Φ is true \mathbb{P} -*a.s.* if Φ holds with probability 1.

For example, we say: $X = Y$ \mathbb{P} -*a.s.*, for X and Y two random variables, if $\{X = Y\}$ has probability 1. We say that a sequence of random variables $(X_n)_{n \geq 1}$ converges \mathbb{P} -*a.s.* if $(X_n(\omega))_{n \geq 1}$ is \mathbb{P} -*a.s.* a convergent sequence.

Some more notions from Probability are introduced in Ch. 4, Section I.

I-3 Product Spaces.

I-3.1 Topology Generated by a Collection of Mappings. Let (E, τ) be a topological space. A basis of open sets of τ is a sub-collection $\tau' \subseteq \tau$, such that for every $U \in \tau$ there is a non empty $V \in \tau'$ with $V \subseteq U$.

Let E be a set, let I be an at most countable set of indices, and let $(E_i, \tau_i)_{i \in I}$ be a collection of topological spaces. For every $i \in I$, let $\pi_i : E \rightarrow E_i$ be a mapping. There is a weakest topology τ on E that makes all the mappings π_i continuous. A basis of τ is given by the collection of subsets with the form ([12]):

$$\bigcap_{j \in J} \pi_j^{-1}(A_j), \tag{2.1}$$

¹For simplicity, we will assume without more details that the σ -algebra \mathcal{F} is complete, *i.e.* contains all subsets contained in measurable sets of probability 0.

where J is a finite subset of I , and A_j is an open subset of E_j for every $j \in J$. We denote this topology by $\tau = \langle \pi_i, i \in I \rangle$. This construction applies to the product spaces as follows.

1-3.2 Topological Product Spaces. Let $(E_i, \tau_i)_{i \in I}$ be as above an at most countable collection of topological spaces. Let E be the product set of the family $(E_i, \tau_i)_{i \in I}$, and let $\pi_i : E \rightarrow E_i$ be the collection of projections. We define the product topology ([12],[38]) by $\tau = \langle \pi_i, i \in I \rangle$. It follows from (2.1) that a basis of open sets of (E, τ) is given by the collection of elementary cylinders:

$$U = \prod_{i \in I} A_i, \quad (2.2)$$

where A_i is an open subset of E_i for every $i \in I$, and $A_i = E_i$ for all but a finite number of i . If we have a basis of open sets for E_i , we can choose the sets $A_i \neq E_i$ in this basis.

Let \mathcal{F}_i be the Borel σ -algebra of τ_i for every $i \in I$, and let \mathcal{F} be the Borel σ -algebra associated to τ . Then \mathcal{F} is the smallest σ -algebra that contains $\langle \tau_i \rangle$ for every $i \in I$, and \mathcal{F} is generated by the subsets with the form (2.2), with A_i measurable for all $i \in I$, and $A_i = E_i$ for all but a finite number of i . Such subsets of E are called the elementary measurable cylinders of E .

The two following results are well known.

1-3.3 Proposition. *Assume that $(E_i)_{i \in I}$ is an at most countable collection of topological spaces, and let E denote the product space. If all E_i are Hausdorff, then E is Hausdorff. If all E_i are separable metric, then E is separable metric.*

We will only use the countable version of the Tychonoff theorem.

1-3.4 Theorem. ([38],[9]) *A countable product of compact spaces is compact in the product topology.*

1-4 Projective Systems.

1-4.1 Projective Systems. Let I be a partially ordered set of indices. We denote by Δ the set of pairs $(i, j) \in I \times I$ such that $i \preceq j$. Let $(E_i)_{i \in I}$ be a family of sets. Assume that, for every pair $(i, j) \in \Delta$, a mapping $\pi_{i,j} : E_j \rightarrow E_i$ is defined. We say that the family $(\pi_{i,j})_{(i,j) \in \Delta}$ is a filtration of $(E_i)_{i \in I}$, and we say that $(E_i)_{i \in I}$ is a projective system *w.r.t.* the filtration $(\pi_{i,j})_{(i,j) \in \Delta}$, if the two following conditions are satisfied ([9]):

1. $\forall i \in I: \pi_{i,i} = \text{Id}_{E_i},$
2. $\forall i, j, k \in I, \quad i \preceq j \preceq k \Rightarrow \pi_{i,k} = \pi_{i,j} \circ \pi_{j,k}.$

I-4.2 Projective Limit. Let E be the product of the family $(E_i)_{i \in I}$. The projective limit of the projective system $(E_i)_{i \in I}$, *w.r.t.* the filtration $(\pi_{i,j})_{(i,j) \in \Delta}$, is defined as the subset of E given by:

$$F = \{(x_i)_{i \in I} \in E \mid \forall i, j \in I, i \preceq j \Rightarrow x_i = \pi_{i,j}(x_j)\}. \quad (2.3)$$

The projective limit is denoted by $F = \varprojlim_{i \in I} E_i$. We still write π_i instead of $\pi_i|_F$ to denote the restrictions to F of the projections π_i .

I-4.3 Topological and Measurable Projective Systems. We say that the filtration is topological (respectively, measurable), and that the projective system is topological (respectively, measurable), if all E_i are equipped with a topology (respectively, a σ -algebra), such that all mappings $\pi_{i,j}$ are continuous (respectively, measurable) for all pairs (i, j) of Δ . For a topological projective system, we will then equip F with the restriction of the product topology. This topology is called the *projective topology*. The form (2.3) implies that $F = \varprojlim_{i \in I} E_i$ is a closed subset of the product, as an intersection of closed subsets. A basis of open sets of F is given by the collection of subsets U with the form:

$$U = F \cap \prod_{i \in I} A_i,$$

the trace in F of an elementary open cylinder.

Remark that, if I is equipped with the trivial partial order $\preceq = \text{Id}_{I \times I}$, then $\varprojlim_{i \in I} E_i$ coincides with the product space $\prod_{i \in I} E_i$.

I-4.4 Cofinal Sequences. We say that a sequence $(i_n)_{n \geq 0}$ is **cofinal** in the partial order (I, \preceq) if $(i_n)_{n \geq 0}$ is non-decreasing ($n \leq m \Rightarrow i_n \preceq i_m$), and if for every $i \in I$, there is an integer n such that $i \preceq i_n$.

I-4.5 Lemma. Let $(i_n)_{n \geq 1}$ be a sequence of indices cofinal in I , and let $(E_i, \pi_{i,j})$ be a projective system indexed by I . Then there is a homeomorphism:

$$\varprojlim_{i \in I} E_i \rightarrow \varprojlim_{n \geq 1} E_{i_n}, \quad (z_i)_{i \in I} \mapsto (z_{i_n})_{n \geq 1}.$$

I-5 Projective Systems of Probability Measures.

Let $(E_i, \mathcal{F}_i)_{i \in I}$ be a measurable projective system of measurable spaces, *w.r.t.* a filtration $(\pi_{i,j})_{i \preceq j}$. We assume that each σ -algebra \mathcal{F}_i is the Borel σ -algebra of a topology on E_i . The projective σ -algebra $\varprojlim_{i \in I} \mathcal{F}_i$ is the Borel σ -algebra of the projective topology.

1-5.1 Definition. (*Projective systems of probabilities*) Let $(E_i, \mathcal{F}_i)_{i \in I}$ be a measurable projective system of measurable spaces, *w.r.t.* a filtration $(\pi_{i,j})_{i \preceq j}$. We say that a family $(\mathbb{P}_i)_{i \in I}$, with \mathbb{P}_i a probability measure on (E_i, \mathcal{F}_i) , is a **projective system of probability measures**, if the following holds:

$$\forall i, j \in I, \quad i \preceq j \Rightarrow \mathbb{P}_i = \pi_{i,j} \mathbb{P}_j, \quad (2.4)$$

where $\pi_{i,j} \mathbb{P}_j$ denotes the image of \mathbb{P}_j in E_i under the measurable mapping $\pi_{i,j} : E_j \rightarrow E_i$ (1-2.4).

1-5.2 Projective System of Probabilities Coming from a Probability on the Projective Limit. With the previous notations, assume that \mathbb{P} is a probability measure on $\varprojlim_{i \in I} (E_i)$, equipped with the projective σ -algebra. Define $\mathbb{P}_i = \pi_i \mathbb{P}$ for all $i \in I$. Then $(\mathbb{P}_i)_{i \in I}$ is a projective system of probabilities on $(E_i, \pi_{i,j})$, naturally associated with \mathbb{P} . The converse operation, from projective systems of probabilities to probability on the projective limit, is the topic of an Extension theorem. Classical references describe the Prokhorov condition for the existence of projective limit of probability measures. In particular, the following statement holds ([9], Th. 2 p. 53).

1-5.3 Theorem. *Let $(E_i, \mathcal{F}_i, \pi_{i,j})$ be a measurable projective system indexed by a partial order I that admits a cofinal sequence. We assume that \mathcal{F}_i is the Borel σ -algebra of a Hausdorff topological space. Let $(\mathbb{P}_i)_{i \in I}$ be a projective system of probability measures. Then there is a unique probability \mathbb{P} on $\varprojlim_{i \in I} (E_i)$ such that $\mathbb{P}_i = \pi_i \mathbb{P}$ for all $i \in I$.*

II—Defining Topologies for Event Structures

We show that natural projective systems are associated with event structures. The projective formalism brings us to the definition of topologies on the set of configurations, and on the set of maximal configurations of an event structure. Although probabilistic constructions are concerned with the set of maximal configurations, the topology on the set of configurations is the key of our study. This way, we also re-obtain results concerning the relationship between the Scott and the Lawson topology on the domain of configurations.

The contribution of this section consists in the application of the projective formalism to the study of the compactness of the set of maximal configurations—that can be seen as the *border at infinity* of an acyclic concurrent model.

II-1 Definition of Mappings.

We recall that \mathcal{P} and \mathcal{P}_0 denote respectively the complete lattice of prefixes, and the lattice of finite prefixes of an event structure \mathcal{E} . \mathcal{W} denotes the set of configurations of \mathcal{E} , and for $P \in \mathcal{P}$, \mathcal{W}_P denotes the set of configurations of P , that coincides with the set of configurations of \mathcal{E} included in P .

II-1.1 Global Projections. We begin by defining mappings that will be used throughout the document. Let P be a prefix of \mathcal{E} . For each configuration v , $v \cap P$ is a conflict-free prefix of P , *i.e.* an element of \mathcal{W}_P . We define thus a surjective mapping π_P by setting:

$$\pi_P : \mathcal{W} \rightarrow \mathcal{W}_P, \quad w \mapsto w \cap P .$$

We recall that Ω denotes the set of maximal configurations of \mathcal{E} . We set $\Gamma_P = \pi_P(\Omega)$, and we still denote by π_P the restriction $\pi_P|_{\Omega}$:

$$\pi_P : \Omega \rightarrow \Gamma_P, \quad \omega \mapsto \omega \cap P .$$

We restrict our attention to the family $(\pi_P)_P$ where P ranges over \mathcal{P}_0 . The family $(\pi_P)_{P \in \mathcal{P}_0}$ separates \mathcal{W} , as stated by the following.

II-1.2 Lemma. *Let $w, w' \in \mathcal{W}$. Then $w = w'$ if and only if $\pi_P(w) = \pi_P(w')$ for all $P \in \mathcal{P}_0$.*

Proof – Since every event belongs to a finite prefix, every configuration w satisfies:

$$w = \bigcup_{P \in \mathcal{P}_0} \pi_P(w) ,$$

which implies the statement of the lemma. \square

II-1.3 Relative Projections. Let $P \subseteq P'$ be two prefixes of \mathcal{E} . We set as above the two following mappings:

$$\begin{aligned} \pi_{P,P'} : \mathcal{W}_{P'} &\rightarrow \mathcal{W}_P, & w &\mapsto w \cap P , \\ \pi_{P,P'} : \Gamma_{P'} &\rightarrow \Gamma_P, & w &\mapsto w \cap P . \end{aligned}$$

We have in particular $\pi_P = \pi_{P,\mathcal{E}}$ for every $P \in \mathcal{P}$. The family of mappings $\pi_{P,P'}$ satisfies obviously the two following properties:

- (a) $\forall P \in \mathcal{P}, \quad \pi_{P,P} = \text{Id}_{\mathcal{W}_P} ,$
- (b) $\forall P, P', P'' \in \mathcal{P}, \quad P \subseteq P' \subseteq P'' \Rightarrow \pi_{P,P''} = \pi_{P,P'} \circ \pi_{P',P''} .$

II-2 The Topological Space of Configurations.

As we are interested in controlling the processes restricted to finite prefixes, it is natural from a computational point of view to define the following topology on \mathcal{W} . The terminology of *projective topology* will be justified in the sequel. The definition has a sense according to I-3.1.

II-2.1 Definition. (*Projective topology*) For each $P \in \mathcal{P}_0$, we equip each \mathcal{W}_P with the discrete topology. We define the **projective topology** on \mathcal{W} as the topology generated by the collection of mappings $\pi_P : \mathcal{W} \rightarrow \mathcal{W}_P$, with P ranging over \mathcal{P}_0 . Equivalently, the projective topology is the weakest topology making all the mappings $\pi_P : \mathcal{W} \rightarrow \mathcal{W}_P$ continuous, with P ranging over \mathcal{P}_0 . In the sequel, \mathcal{W} is equipped with the projective topology.

II-2.2 Convergence in the Projective Topology. The convergence in the projective topology is addressed as follows. A sequence $(v_n)_{n \geq 0}$ of configurations converges to $w \in \mathcal{W}$ if and only if:

$$\forall P \in \mathcal{P}_0, \quad \exists N \geq 0 : n \geq N \Rightarrow v_n \cap P = w \cap P. \quad (2.5)$$

In particular, if $(v_n)_{n \geq 0}$ is non-decreasing, *i.e.* $i \leq j \Rightarrow v_i \subseteq v_j$ for all i, j , then $\lim_{n \rightarrow \infty} v_n = \bigcup_{n \geq 0} v_n$ always holds in the projective topology.

II-2.3 The Complete Projective System. We introduce a projective system whose limit realizes the topological space \mathcal{W} . According to properties (a) and (b) in II-1.3, the family of finite sets $(\mathcal{W}_P)_{P \in \mathcal{P}_0}$ forms a projective system *w.r.t.* to the family of mappings $\pi_{P,P'}$, trivially continuous for $P \subseteq P'$ ranging over \mathcal{P}_0 .

II-2.4 Remark. Since \mathcal{E} is at most countable, since every event e belongs to a finite prefix, and since \mathcal{P}_0 is a lattice, it follows that \mathcal{P}_0 admits a countable cofinal sequence (I-4.4). That is, there is a non decreasing sequence $(P_n)_{n \geq 1}$ of finite prefixes, such that for every finite prefix P , there is an integer n with $P_n \supseteq P$. Define for instance $\mathcal{E} = \{e_1, e_2, \dots\}$, choose $Q_i \in \mathcal{P}_0$ that contains e_i for each $i \geq 1$, and then consider the cofinal sequence $(P_n)_n$ defined by:

$$P_n = \bigcup_{i=1}^n Q_i.$$

II-2.5 Notation. Let Ξ denote the projective limit $\Xi = \varprojlim_{P \in \mathcal{P}_0} \mathcal{W}_P$. We recall that Ξ has the following expression as a subset of the product space $\prod_P \mathcal{W}_P$:

$$\Xi = \left\{ z = (z_P)_{P \in \mathcal{P}_0} \mid \begin{array}{l} \forall P \in \mathcal{P}_0 : z_P \in \mathcal{W}_P, \\ \forall P, P' \in \mathcal{P}_0 : \pi_{P,P'}(z_{P'}) = z_P \end{array} \right\}.$$

There is a natural mapping $\Phi : \mathcal{W} \rightarrow \Xi$, defined by $\Phi(v) = (\pi_P(v))_{P \in \mathcal{P}_0}$. Indeed, if $z = (\pi_P(v))_{P \in \mathcal{P}_0}$, we have for every $P \subseteq P'$:

$$\pi_{P,P'}(z_{P'}) = (v \cap P') \cap P = z_P ,$$

showing that z is an element of Ξ .

II-2.6 Proposition. *The mapping $\Phi : \mathcal{W} \rightarrow \Xi = \overleftarrow{(\mathcal{W}_P)_{P \in \mathcal{P}_0}}$ of II-2.5 is a homeomorphism. In particular, \mathcal{W} is compact, metric and separable.*

Proof – Injective mapping. Lemma II-1.2 implies that Φ is injective.

Continuous mapping. Let U be an elementary open set in Ξ . We denote by $\overline{\Xi}$ the product space $\prod_{P \in \mathcal{P}_0} \mathcal{W}_P$. U is given by $U = \overline{U} \cap \Xi$, with \overline{U} an elementary cylinder of $\overline{\Xi}$, given by:

$$\overline{U} = \prod_{P \in \mathcal{P}_0} A_P ,$$

with $A_P = \mathcal{W}_P$ for all but a finite number of P . \mathcal{P}_0 is a lattice, so there is a $P_0 \in \mathcal{P}_0$ such that:

$$\forall P \in \mathcal{P}_0, \quad A_P \neq \mathcal{W}_P \Rightarrow P \subseteq P_0 . \quad (2.6)$$

Equivalently $(z_P)_P \in U$ if and only if $z_P \in A_P$ for all $P \subseteq P_0$, which in turn is equivalent to $z_{P_0} \in A_{P_0}$. In particular, for all $v \in \mathcal{W}$, $\Phi(v) \in U$ if and only if $\pi_{P_0}(v) \in A_{P_0}$. Equivalently, $\Phi^{-1}(U) = \pi_{P_0}^{-1}(A_{P_0})$, which shows that $\Phi^{-1}(U)$ is open in the projective topology, and thus Φ is continuous.

Surjective mapping. Let $z = (z_P)_{P \in \mathcal{P}_0}$ be an element of Ξ . We set the following prefix of \mathcal{E} :

$$v = \bigcup_{P \in \mathcal{P}_0} z_P .$$

Assume that v contains a conflict, *i.e.* two events e, e' with $e \# e'$. There are finite prefixes P_0, P_1 such that $e \in z_{P_0}$ and $e' \in z_{P_1}$. Let $P_2 \in \mathcal{P}_0$ containing P_0 and P_1 . As z is an element of Ξ , z satisfies:

$$\pi_{P_0, P_2}(z_{P_2}) = z_{P_0} , \quad \pi_{P_0, P_1}(z_{P_2}) = z_{P_1} ,$$

equivalently:

$$z_{P_2} \cap P_0 = z_{P_0} , \quad z_{P_2} \cap P_1 = z_{P_1} .$$

It implies that z_{P_2} contains the conflict $e \# e'$, a contradiction. Thus v is a conflict-free prefix of \mathcal{E} , *i.e.* a configuration. We show that $\Phi(v) = z$, *i.e.* that $v \cap P = z_P$ for all $P \in \mathcal{P}_0$. For any $Q \in \mathcal{P}_0$, let $Q' \in \mathcal{P}_0$ that contains both Q and P . We have $z_Q = \pi_{Q, Q'}(z_{Q'})$, and therefore $z_Q \cap P = z_{Q'} \cap Q \cap P$, whence we get:

$$z_Q \cap P = z_P \cap Q . \quad (2.7)$$

Taking union *w.r.t.* Q in (2.7), we get: $v \cap P = z_P$, as was to be shown. Therefore Φ is bijective, and we denote by $\Phi^{-1} : \Xi \rightarrow \mathcal{W}$ the inverse mapping.

Open mapping (Φ^{-1} is continuous). By Definition II-2.1, a basis of open sets for the projective topology is given by the family $\pi_P^{-1}(v)$, with (P, v) such that P ranges over \mathcal{P}_0 and v ranges over \mathcal{W}_P . Let $V = \pi_P^{-1}(v)$ be such an elementary open set in \mathcal{W} , we show that $\Phi(V)$ is open. Let X be the subset of Ξ given by:

$$X = \Xi \cap \prod_{P \in \mathcal{P}_0} A_P ,$$

with $A_{P_0} = \pi_{P_0}(v)$ and $A_P = \mathcal{W}_P$ for every $P \neq P_0$. Then X is open and we obviously have that $\Phi(V) \subseteq X$. We show the converse inclusion. Let $z \in X$, and let $w = \Phi^{-1}(z)$. Then $w \cap P_0 = z_{P_0} \in A_{P_0}$, from which follows that $w \in \pi_{P_0}^{-1}(v) = V$. We have shown that $\Phi(V) = X$ is open, thus Φ is an open mapping.

With the four previous points, we have shown that Φ is a homeomorphism $\mathcal{W} \rightarrow \Xi$. By the classical results, as the \mathcal{W}_P are finite and thus compact metric, Ξ is compact metric whenever \mathcal{P}_0 contains a cofinal sequence. We have seen in II-2.4 that such a sequence exists. Therefore \mathcal{W} is a compact and separable metric space. \square

II-3 The Space of Maximal Configurations.

We now examine the set Ω of maximal configurations of \mathcal{E} . As a subset of \mathcal{W} , Ω is equipped with the restriction of the projective topology.

II-3.1 Definition. (*Operational topology, finite shadows*) We define the **operational topology** on Ω as the restriction to Ω of the projective topology.

We will always use the following basis of open sets for the topology on Ω .

II-3.2 Lemma. *A countable base of open sets of the operational topology is given by the family of finite shadows of Ω , where a **finite shadow** is any subset of Ω with the form:*

$$\Omega(v) = \{\omega \in \Omega \mid \omega \supseteq v\} , \quad v \in \mathcal{W}_0 .$$

Proof— Let τ denote the restriction to Ω of the projective topology in \mathcal{W} , and let τ' denote the topology generated by the collection of finite shadows. Any $\Omega(v) \in \tau'$ is written as $\Omega(v) = \pi_P^{-1}(v) \cap \Omega$, with $P = v$ and $\pi_P : \mathcal{W} \rightarrow \mathcal{W}_P$. Hence $\Omega(v)$ is open in the projective topology, and thus $\tau' \subseteq \tau$.

Conversely, let $P \in \mathcal{P}_0$, $v \in \mathcal{W}_P$, and set $A = \pi_P^{-1}(v) \cap \Omega$, with $\pi_P : \mathcal{W} \rightarrow \mathcal{W}_P$. The collection of such subsets A generates τ . To show that $A \in \tau'$, let $\omega \in A$. For each $e \in P \setminus v$, e is in conflict with an event $y \in \omega$, otherwise ω would not be maximal. Let $y(e)$ be such an event. Since P is finite, there is a finite prefix Q that contains v and all the events $y(e)$ for $e \in P \setminus v$. We set $v' = \omega \cap Q$. Then the finite shadow $\Omega(v')$ satisfies: $\omega \in \Omega(v') \subseteq A$. This shows that A is open in τ , and thus $\tau = \tau'$. \square

II-3.3 The Operational Projective System. Until a certain point, we can make the same construction as we did for the space of configurations. Recall that Γ_P denotes, for P a finite prefix, the set: $\Gamma_P = \{\omega \cap P, \omega \in \Omega\}$. Then we have, for the family of mappings $\pi_{P,P'} : \Gamma_{P'} \rightarrow \Gamma_P$ with $P \in \mathcal{P}_0$:

- (a) $\forall P \in \mathcal{P}_0: \pi_{P,P} = \text{Id}_{\Gamma_P}$,
 (b) $\forall P, P', P'' \in \mathcal{P}_0, P \subseteq P' \subseteq P'' \Rightarrow \pi_{P,P''} = \pi_{P,P'} \circ \pi_{P',P''}$.

Hence the family $(\Gamma_P)_{P \in \mathcal{P}_0}$ is a projective system *w.r.t.* the filtration $\pi_{P,P'}$ with $P, P' \in \mathcal{P}_0$ and $P \subseteq P'$. Let Γ denote the projective limit $\Gamma = \varprojlim_{P \in \mathcal{P}_0} \Gamma_P$. For every $\omega \in \Omega$, the family $z = (z_P)_{P \in \mathcal{P}_0}$ with $z_P = \pi_P(\omega)$ satisfies:

$$\forall P, P' \in \mathcal{P}_0 \quad P \subseteq P' \Rightarrow z_P = \pi_{P,P'}(z_{P'}) ,$$

hence we define a mapping $\Psi : \Omega \rightarrow \Gamma$ by setting $\Psi(\omega) = z$. We have the following commutative diagram, with natural continuous injections on the sides:

$$\begin{array}{ccc} \mathcal{W} & \xrightarrow{\Phi} & \Xi \\ \uparrow i & & \uparrow j \\ \Omega & \xrightarrow{\Psi} & \Gamma \end{array} \quad (2.8)$$

i is continuous since the topology on Ω is the restriction of the topology on \mathcal{W} . It is readily checked by hand that $j : \Gamma \rightarrow \Xi$ is continuous.

From the diagram follows that Ψ is injective. We also have that Ψ is continuous: Let V be an open set in Γ . Then there is an open set U in Ξ such that $j(V) = j(\Gamma) \cap U$, and we have:

$$\Psi^{-1}(V) = (\Phi \circ i)^{-1}(j(\Gamma) \cap U) = (\Phi \circ i)^{-1}(U) ,$$

showing that Ψ is continuous.

II-3.4 Lemma. We have $\Phi^{-1}(\Gamma) = \overline{\Omega}$, where $\overline{\Omega}$ denotes the closure of Ω in \mathcal{W} , and where we identify Γ and its image $j(\Gamma)$ in Ξ .

Proof— We have $\Phi(\Omega) \subseteq \Gamma$. Γ is a closed subset of Ξ , as a countable intersection of closed subsets of Ξ :

$$\Gamma = \bigcap_{P \in \mathcal{P}_0} \pi_P^{-1}(\Gamma_P) .$$

Since Φ is continuous, it follows that $\Phi(\overline{\Omega}) \subseteq \Gamma$, and thus $\overline{\Omega} \subseteq \Phi^{-1}(\Gamma)$.

To show the converse inclusion, let $v \in \Phi^{-1}(\Gamma)$. There is an element $(z_P)_{P \in \mathcal{P}_0} \in \Gamma$ such that $v \cap P = z_P$ for all $P \in \mathcal{P}_0$. For each $P \in \mathcal{P}_0$, let $\omega_P \in \Omega$ be an element such that $z_P = \omega_P \cap P$. Choose $(P_n)_{n \geq 0}$ a cofinal sequence of \mathcal{P}_0 (I-4.4), *i.e.* a non decreasing sequence of finite prefixes, that contain every finite prefix. We set $\omega^n = \omega_{P_n}$, and we show that $\lim_{n \rightarrow \infty} \omega^n = v$ holds in \mathcal{W} (Cf. Remark II-2.2 on the convergence in Ξ).

We have $v \cap P = z_P$ for all $P \in \mathcal{P}_0$, and in particular:

$$\forall n \geq 0, \quad v \cap P_n = z_{P_n} = \omega^n \cap P_n. \quad (2.9)$$

Fix a prefix $P \in \mathcal{P}_0$, and choose N an integer such that $n \geq N \Rightarrow P \subseteq P_n$. Then (2.9) implies: $n \geq N \Rightarrow v \cap P = \omega^n \cap P$, which is the convergence $\lim_{n \rightarrow \infty} \omega^n = v$. This shows that $\Phi^{-1}(\Gamma) \subseteq \overline{\Omega}$. \square

II-3.5 Theorem. *The mapping $\Phi|_{\overline{\Omega}} : \overline{\Omega} \rightarrow \Gamma$ is a homeomorphism.*

Proof – According to II-3.4 and II-2.6, $\Phi|_{\overline{\Omega}}$ is a one-to-one continuous mapping $\Phi|_{\overline{\Omega}} : \overline{\Omega} \rightarrow \Gamma$. $\overline{\Omega} \subseteq \mathcal{W}$ is compact since \mathcal{W} is compact. By a classical result, it implies that $\Phi|_{\overline{\Omega}}$ is a homeomorphism. \square

II-3.6 Corollary. *Let \mathcal{E} be an event structure. With the above notations, we have equivalence between the following propositions:*

1. $\Psi : \Omega \rightarrow \Gamma$ is a homeomorphism,
2. Ω is compact.

Proof – $1 \Rightarrow 2$. We have seen above in the proof of II-3.4, that Γ is closed in \mathcal{W} . It implies that Γ is compact since \mathcal{W} is compact (II-2.6). Hence Ω is compact.

$2 \Rightarrow 1$. We have that $\Phi|_{\overline{\Omega}} : \overline{\Omega} \rightarrow \Gamma$ is a homeomorphism (II-3.5). $\Omega = \overline{\Omega}$ since Ω is compact, hence $\Psi = \Phi|_{\overline{\Omega}}$ is a homeomorphism $\Omega \rightarrow \Gamma$. \square

II-3.7 Example. It can actually hold that $\Psi : \Omega \rightarrow \Gamma$ is not a homeomorphism. It is readily checked that Ψ is always open (maps an open set onto an open set), and therefore the only possibility for Ψ not to be a homeomorphism is to be not surjective.

Consider the event structure of Figure 2.1. Let ω_∞ and ω_n denote the maximal configurations given by:

$$\omega_\infty = g \oplus e_1 \oplus e_2 \oplus \cdots, \quad \forall n \geq 1, \quad \omega_n = e_1 \oplus \cdots \oplus e_n \oplus f_n.$$

Consider the following cofinal sequence of finite prefixes:

$$\forall n \geq 1, \quad P_n = \{g, e_1, \dots, e_n, f_1, \dots, f_n\}.$$

Since $(P_n)_n$ is cofinal, the projective limit Γ is given by $\Gamma = \varprojlim_{n \geq 1} \Gamma_n$, with $\Gamma_n = \Gamma_{P_n}$ (Lemma I-4.5). For each n , the configuration $v_n = e_1 \oplus \cdots \oplus e_n$ belongs to Γ_n since $v_n = \omega_{n+1} \cap P_n$. Moreover the sequence $(v_n)_{n \geq 1}$ is coherent, *i.e.* satisfies:

$$m \geq n \Rightarrow v_m \cap P_n = v_n.$$

Now we observe that the element $z = (v_n)_{n \geq 1}$, which is thus an element of Γ , cannot be written as $z = \Phi(\omega)$, with $\omega \in \Omega$. Indeed, assume that there is such an element ω . Then $\omega \cap P_1 = v_1$ implies that $g \notin \omega$, and $\omega \cap P_n \supseteq v_n$ for all n implies that $\omega = \omega_\infty$, a contradiction. This shows that Φ is not surjective.

Remark that the operational topology makes Ω homeomorphic to $\bar{\mathbb{N}} = \mathbb{N} \cup \{\infty\}$, with the *discrete* topology. In other words, ω_∞ is not the Alexandrov point at infinity of $\{\omega_n, n \geq 1\}$. Indeed the finite shadow $\Omega(g)$ only contains ω_∞ , so ω_∞ is isolated. Observe that it is always true that a finite and maximal configuration is isolated, and in particular here the $\omega_n, n \geq 1$.




Figure 2.1: $\Psi : \Omega \rightarrow \Gamma$ is not surjective.

II-3.8 Remark. This example can be realised as the unfolding of a safe Petri net. Indeed, we recognise the unfolding of the net depicted in Figure 1.5. The construction of probability measures for this example, that does not fit the extension technique presented below, is discussed in Chapter 9.

III—Extension of Probability Measures

This last section is devoted to the application of the above results to the extension of probability measures. Using the form of a projective limit for the space Ω associated with an event structure \mathcal{E} , we apply the Prokhorov extension theorem to extend projective systems of probabilities, defined on *finite* spaces. For this we introduce

the class of intrinsic prefixes, and we discuss the computational interest of this class of prefixes.

We apply this result to locally finite event structures. This class of event structures will be studied in more details in Chapter 3. We only use here a basic property of locally finite event structures.

III-1 Probabilistic Event Structures and Intrinsic Prefixes.

III-1.1 Definition. (*Operational σ -algebra, probabilistic event structure*) Let \mathcal{E} be an event structure, with Ω the topological space of maximal configurations of \mathcal{E} . We define the **operational σ -algebra** on Ω as the Borel σ -algebra \mathcal{F} of Ω (I-2.3). Hence \mathcal{E} defines the measurable space (Ω, \mathcal{F}) .

A **probabilistic event structure** is a pair $(\mathcal{E}, \mathbb{P})$, where \mathcal{E} is an event structure and \mathbb{P} is a probability measure on (Ω, \mathcal{F}) .

Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure. For every finite prefix P , the mapping $\pi_P : \Omega \rightarrow \Gamma_P$ induces a probability $\mathbb{P}_P = \pi_P \mathbb{P}$ on Γ_P . Remark that \mathbb{P}_P is not defined on Ω_P in general, and hence does not define a probabilistic event structure on P . However $\Omega_P \subseteq \Gamma_P$ always holds. Indeed, let $v \in \Omega_P$, and let $\omega \in \Omega$ such that $\omega \supseteq v$, then $\pi_P(\omega)$ is a configuration of P that contains v . Since v is maximal in P , $v = \pi_P(\omega)$, and thus $v \in \Gamma_P$.

In order to induce probabilistic event structures on prefixes, we are thus led to introduce the following definition.

III-1.2 Definition. (*Intrinsic prefixes*) Let P be a prefix of \mathcal{E} . We denote by Ω_P the set of maximal configurations of the sub-event structure P . We say that P is **intrinsic to \mathcal{E}** , or **intrinsic**, if $\Omega_P = \Gamma_P$. If $(\mathcal{E}, \mathbb{P})$ is a probabilistic event structure, $(P, \pi_P \mathbb{P})$ defines an **induced** probabilistic event structure for every intrinsic prefix P .

III-1.3 Example. In the event structure of Figure 2.2, left, prefix $P = \{a\}$ is not intrinsic in the event structure. Indeed, with $\omega = \{b\}$ we have $P \cap \omega = \emptyset$, that is not a maximal configuration of P . Hence the dynamics of the event structure at right is not faithful to the restricted dynamics of the event structure at left.




Figure 2.2: Non intrinsic prefix at left hand, intrinsic at right hand.

III-1.4 Remark. Clearly, the union of two intrinsic prefixes is an intrinsic prefix. Intrinsic prefixes are not stable in general under intersection, as shown by the

event structure of Figure 2.3. P_1 and P_2 are intrinsic with $P_1 = \{a, b\}$ and $P_2 = \{b, c\}$, but $P_1 \cap P_2 = \{b\}$ is not: take $\omega = \{a, c\}$, then $\omega \cap \{b\} = \emptyset$, not maximal in $\{b\}$.




Figure 2.3: Non intrinsic intersection of intrinsic prefixes.

III-1.5 Example. In the event structure \mathcal{E} of Figure 2.1, the only intrinsic prefix that contains the event g is \mathcal{E} itself.

Proof – Let P be an intrinsic prefix that contains g . Let $N = \max\{i \geq 1 \mid e_i \in P\}$. Assume that $N < \infty$. Then $\omega = e_1 \oplus \dots \oplus e_N \oplus f_{N+1}$ is maximal in \mathcal{E} and satisfies $\omega \cap P = e_1 \oplus \dots \oplus e_N$, a configuration of P that can be completed in P by $g \oplus e_1 \oplus \dots \oplus e_N$. This contradicts that P is intrinsic. We have shown that P contains all e_i , and we show in a similar way that P contains all f_i . Hence $P = \mathcal{E}$. \square

III-2 The Extension Theorem.

Since we have given a condition for Ω to be homeomorphic to a projective limit, we can apply under this condition the Extension theorem of Prokhorov to extend projective systems of probability measures (I-5). From a computational point of view, the advantage is that the projective system consists of finite spaces. However this is not enough, in particular for the two following reasons.

1. The condition to apply the Extension theorem is abstractly given on Ω , this information is not easily computable. A first step towards a better computability will be to give a condition on \mathcal{E} that implies the compactness of Ω .
2. If P is a finite prefix, the information contained in (P, Γ_P) is strictly larger than P in general. In other words, given a finite prefix P , we do not know how Ω intersects P through the configurations $\omega \cap P$. The use of intrinsic prefixes is thus fully justified: in this case $(P, \Gamma_P) = (P, \Omega_P)$ is determined by the only data P .

We will show how the use of intrinsic prefixes brings a solution to both problems. In particular, the class of locally finite event structures provides a convenient framework *w.r.t.* both restrictions.

The first result that we get from our above study is the ability to apply the Prokhorov extension theorem. Recall the projective system of finite spaces

$(\Gamma_P)_{P \in \mathcal{P}_0}$. From equation (2.4), 1-5, a family $(\mathbb{P}_P)_{P \in \mathcal{P}_0}$ with \mathbb{P}_P a probability on Γ_P is a projective system of probability measures if the following condition holds:

$$\forall P, P' \in \mathcal{P}_0, \quad P \subseteq P' \Rightarrow \mathbb{P}_P = \pi_{P, P'} \mathbb{P}_{P'} .$$

III-2.1 Theorem. *Let \mathcal{E} be an event structure, and let $(\mathbb{P}_P)_{P \in \mathcal{P}_0}$ be a projective system of probability measures w.r.t. the projective system $(\Gamma_P)_{P \in \mathcal{P}_0}$. Assume that Ω is compact w.r.t. the operational topology. Then there is a unique probability measure \mathbb{P} on Ω such that:*

$$\forall P \in \mathcal{P}_0, \quad \pi_P \mathbb{P} = \mathbb{P}_P .$$

Proof – Let $\Gamma = \overleftarrow{(\Gamma_P)_{P \in \mathcal{P}_0}}$, and let $\rho_P : \Gamma \rightarrow \Gamma_P$ denote the natural projections for $P \in \mathcal{P}_0$. Since \mathcal{P}_0 admits a countable cofinal sequence, the Prokhorov extension theorem applies (1-5.3). There is a unique probability \mathbb{Q} on Γ such that:

$$\forall P \in \mathcal{P}_0, \quad \rho_P \mathbb{Q} = \mathbb{P}_P .$$

Let $\Psi : \Omega \rightarrow \Gamma$ be the homeomorphism given by 11-3.6, since we assume that Ω is compact, and let $\mathbb{P} = \Phi^{-1} \mathbb{Q}$. Since π_P and ρ_P are conjugated by $\pi_P = \rho_P \circ \Phi$ for all $P \in \mathcal{P}_0$, we obtain by the chain rule for image probabilities (1-2.4):

$$\pi_P \mathbb{P} = \pi_P(\Phi^{-1} \mathbb{Q}) = (\pi_P \circ \Phi^{-1}) \mathbb{Q} = \rho_P \mathbb{Q} = \mathbb{P}_P .$$

This shows the existence of \mathbb{P} , and by similar arguments, the uniqueness follows from the uniqueness in the Prokhorov theorem. \square

As we have noted above, we need a more computable condition than the compactness of Ω to apply the extension theorem. We give a geometric condition on \mathcal{E} that insures the compactness of Ω .

III-2.2 Lemma. *Let \mathcal{E} be an event structure. Assume that for each event $e \in \mathcal{E}$ there is a finite intrinsic prefix P such that $e \in P$. Then Ω is compact.*

Proof – We show that Ω is closed in \mathcal{W} . Let $(\omega^n)_{n \geq 1}$ be a sequence of Ω , convergent to v in \mathcal{W} . Assume that $v \notin \Omega$. Then there is an event $e \notin v$, and such that $v \cup e \in \mathcal{W}$. Let P be a finite intrinsic prefix that contains e , such a prefix exists by hypothesis. Then e is compatible with $v \cap P$. Since $\lim_{n \rightarrow \infty} \omega^n = v$, there is an integer $n \geq 1$ such that $\omega^n \cap P = v \cap P$. Since P is intrinsic, $\omega^n \cap P$ is maximal in P , and compatible with e . Therefore e belongs to $\omega^n \cap P$, and thus $e \in v$, a contradiction. This shows that Ω is closed in \mathcal{W} , and thus compact w.r.t. the operational topology. \square

Remark that, since finite unions of intrinsic prefixes are intrinsic, the condition of III-2.2 is equivalent to requiring the existence of a cofinal sequence (1-4.4) in \mathcal{P}_0 of finite intrinsic prefixes.

III-2.3 Example. Let $\mathcal{E} = \mathcal{T}$ is a tree of events with a unique root, and let G be the graph relation on \mathcal{T} (Cf. Ch. 1, III-4.2). Then it is well known that Ω is compact if and only if the sets $D(x) = \{y \in \mathcal{T} \mid xGy\}$ are all finite, for x ranging over \mathcal{T} . Such trees are called *locally finite* in the literature.

To see that, for trees, the assumption of Lemma III-2.2 is fulfilled if and only if Ω is compact, remark that intrinsic prefixes are exactly the prefixes P that satisfy:

$$\forall x \in \mathcal{T}, \quad \begin{cases} D(x) \cap P = \emptyset, \\ \text{or: } D(x) \cap P = D(x). \end{cases}$$

Hence finite prefixes are unions of sets $D(x)$, and every event belongs to such a finite prefix if and only if all $D(x)$ are finite: the statement of lemma III-2.2 holds with an “if and only if” for trees. In these intrinsic prefixes, one recognises the “stopping times for occurrence nets” of Benveniste *et al.* ([7]). In particular, if \mathcal{T} is the covering of a finite transition system, then Ω is always compact.

III-2.4 Example. Consider the event structure depicted in Figure 2.1, p. 75, with $\Psi : \Omega \rightarrow \left(\Gamma_P\right)_{P \in \mathcal{P}_0}$ non surjective. We have shown in III-1.5 that every intrinsic prefix containing the event g is \mathcal{E} itself. The condition of Lemma III-2.2 is not fulfilled.

In this example, we check that the result of Lemma III-2.2 does not occur: Ω is not compact. For this observe that all elements $\omega \in \Omega$ are isolated in Ω . That is, for every $\omega \in \Omega$, there is a finite shadow $\Omega(v)$ such that $\Omega(v) = \{\omega\}$. This is clearly the case for elements $\omega_n = e_1 \oplus \dots \oplus e_n \oplus f_n$ —every finite maximal configuration is always isolated. For ω_∞ , take $v = \{g\}$.

I expect that the Lemma III-2.2 can be improved to an “if and only if” statement for all event structures. As we see, it works with this example.

III-3 Locally Finite Event Structures.

We will introduce and study in details in Chapter 3 the class of locally finite event structures. For the time being, the reader is asked to accept that one can define a class of event structures called *locally finite event structures*, with the following property:

A locally finite event structure comes with a lattice \mathcal{S}_0 of finite intrinsic prefixes, called finite stopping prefixes. Every event $e \in \mathcal{E}$ belongs to a finite stopping prefix.

In particular, by Lemma III-2.2, the space Ω is compact and the extension theorem III-2.1 applies to locally finite event structures. The lattice \mathcal{S}_0 contains a sequence cofinal in \mathcal{P}_0 , therefore the mapping:

$$\Omega \rightarrow \left(\Omega_B\right)_{B \in \mathcal{S}_0}, \quad \omega \rightarrow (\omega \cap B)_{B \in \mathcal{S}_0};$$

is a homeomorphism. Whence the following result.

III-3.1 Theorem. *Let \mathcal{E} be a locally finite event structure, with \mathcal{S}_0 the lattice of stopping prefixes of \mathcal{E} . Then $(\Omega_B)_{B \in \mathcal{S}_0}$ is a projective system w.r.t. the filtration $\pi_{B, B'}$. Assume that $(\mathbb{P}_B)_{B \in \mathcal{S}_0}$ is a projective system of probability measures associated with the projective system $(\Omega_B)_{B \in \mathcal{S}_0}$, i.e.:*

$$\forall B, B' \in \mathcal{S}_0, \quad B \subseteq B' \Rightarrow \mathbb{P}_B = \pi_{B, B'} \mathbb{P}_{B'}. \quad (2.10)$$

Then there is a unique probability \mathbb{P} on Ω such that $\mathbb{P}_B = \pi_B \mathbb{P}$ for all $B \in \mathcal{S}_0$.

IV—Conclusion.

Summary. In this chapter, we have defined topological and probabilistic event structures. We have shown how the projective formalism applies to describe the dynamics of event structures.

We have given a condition on event structure that insures the compactness of Ω . This condition applies to locally finite event structure, the main class of event structures studied in the rest of this document. In particular, for locally finite event structures, the construction of a probability is equivalent to the construction of a projective system of probabilities on *finite sets*. We have shown that the local finiteness, due to the intrinsic character of stopping prefixes, is a reasonable assumption from the computational point of view.

Extensions. We have seen that, for the extension theorem of probability measures to hold, it is enough that every event belongs to a finite intrinsic prefix. Locally finite event structures satisfy this property, but the converse result has not been established. Nevertheless, it will not be needed. I also expect that the class of locally finite event structures coincides with the class of event structures with Ω compact, so that we should have the equivalences:

$$\begin{array}{ccccc} \Omega \text{ is} & \iff & \text{every event of } \mathcal{E} \text{ belongs} & \iff & \text{every event of } \mathcal{E} \text{ belongs} \\ \text{compact} & & \text{to a finite intrinsic prefix} & & \text{to a finite stopping prefix} \end{array}$$

The question of computability of local finiteness itself will be relevant when dealing with unfoldings of finite nets instead of abstract event structures, to have a finite input. Some results on this topic are given in Chapter 8.

Chapter 4 is devoted to the construction of projective systems of probabilities. First, Chapter 3 introduces some needed material concerning event structures.

Geometry of Stopping Prefixes

In the previous chapter, we have shown that, for locally finite event structures, the construction of a probability on the set Ω of maximal configurations reduces to the construction of finite probabilistic event structures (B, \mathbb{P}_B) , where B ranges over the lattice of finite stopping prefixes. There is a coherence condition on the family $(\mathbb{P}_B)_B$: it must be a projective system of probabilities.

Consider the case of a locally finite tree, *i.e.* a tree with finite branching. The boundary at infinity is reached by the canonical sequence of prefixes B_n , that contain events connected to the root with less than n events. For event structures, these prefixes have little meaning from the dynamics point of view, since they are not intrinsic in general. We use instead the partially ordered filtration given by the finite stopping prefixes, as depicted in Figure 3.1.



Figure 3.1: *Filtrations of sequential and concurrent acyclic systems.*

Continuing the analogy with sequential systems, one may wish to inductively construct projective systems of probability measures by conditional increments:

$$\mathbb{P}(t_1, \dots, t_{n+1}) = \underbrace{\mathbb{P}(t_1, \dots, t_n)}_{\text{already computed}} \underbrace{\mathbb{P}(t_1, \dots, t_{n+1} \mid t_1, \dots, t_n)}_{\text{to be defined}} .$$

Although increments between stopping prefixes of concurrent systems can be defined, a direct probabilistic transcription fails. For concurrent systems, the class of

configurations reached by the stopping prefixes lacks a compositionality property: the class is not closed under concatenation. This is our main motivation for introducing a new class of configurations.

With probabilistic applications in mind, and in particular as we want a notion of Markovian process, the closure of processes under concatenation is expected. Indeed, we want to have invariance properties *w.r.t.* the future of finite processes. Geometrically, *i.e.* in the event structure, the future of processes is obtained by the concatenation of configurations. We need thus to manipulate a class of configurations closed under concatenation.

We propose therefore to consider a class that have these two properties: stability under concatenation and containing the stopped configurations, *i.e.* the configurations with the form $\omega \cap B$ with $\omega \in \Omega$ and B a finite stopping prefix. The configurations obtained by recursively concatenating stopped configurations are called well-stopped, and they form the minimal class with the required properties.

Each maximal configuration admits a partially ordered decomposition through well-stopped configurations. Although the decomposition is not unique, the *minimal increments* of well-stopped configurations that compose a maximal configuration ω are uniquely defined. They form what we call the *germs* of ω . The germs are configurations *locally maximal*: they are indeed maximal configurations of a finite sub-event structure, called a *branching cell*.

The germs of an element ω are partially ordered. Some are causally related and some are concurrent. The parallelism of germs is stronger than the simple concurrency of events. The parallelism of germs implies their *independence* from a set point of view: all combinations of parallel germs are allowed, hence the state space is a free product of local state spaces. This property has a natural probabilistic counterpart as a probabilistic independence, but this is the topic of next chapter. We underline the three following properties of the decomposition through germs:

1. *Universal w.r.t. the past.* The decomposition through germs fits the global puzzle drawn by the stopping prefixes. The restriction of a decomposition to a stopping prefix coincides with the decomposition of the restriction.
2. *Universal w.r.t. the future.* The decomposition of the tails of a maximal configuration coincides with the restriction of the maximal configuration to a *cone of future*—to be defined later.
3. *Parallelism and independence.* If v is a configuration formed by a finite number of germs, the germs that continue v are both *concurrent* and *independent*.

Although this construction has been achieved with probabilistic applications in mind, I hope that it presents some interest for itself. It is exposed together with a very useful tool, the cone of future of a configuration. The cone of future allows a proper definition of concatenation and of subtraction of configurations. It will be of constant use throughout the document.

Concerning event structures as a computational model, the contribution of this document consists in the introduction of these two notions: the well-stopped con-

figurations (germs and branching cells, *etc*), and the cone of future. How the cone of future can be inserted in the categorical framework for Petri nets introduced by Winskel will be detailed in Chapter 5. The notion of stopping prefix comes back to “stopping times for occurrence nets”, introduced by A. Benveniste *et al.* in [7], and which have been improved by D. Varacca and G. Winskel by the use of the *dynamic conflict*. However, the whole study presented here is original.

As we need some technical results, the chapter is rather long. Section I, *Dynamic conflict and stopping prefixes*, gives the basic definitions concerning stopping prefixes and stopped configurations. Section II, *Cone of future*, studies this object that, as we said, will be used throughout the document. Its study consists mainly in examining its relationships with other objects: prefixes, stopping prefixes, induced topologies. As a first introduction to well stopped decompositions and well-stopped configurations, Section III presents the *Normal decomposition* of maximal configurations—the whole study of well-stopped configurations consists in variants on this theme. Section III also introduces the important *initial stopping prefix*, and the *initial branching cells*. In Section IV, *Well-stopped configurations*, we introduce the well-stopped configurations and their basic properties.

The compositional properties of well-stopped configurations are treated in Sections V and VI. Section V, *Branching cells*, introduces the definition of branching cells, generalising the initial branching cells. It collects then a certain number of technical results, leading to a lemma that we call an *exchange lemma*. Section VI, *The dynamic puzzle*, states the most useful results about well-stopped configurations, that will be used in the rest of the document. Finally, Section VII precises what become the notions introduced when the event structure is a tree of events.

Most of the “back-end” results on well-stopped configurations used later in the document are contained in Section VI. The notations are collected in Table 3.1, page 120.

I—Dynamic Conflict and Stopping Prefixes

Stopping prefixes and stopped configurations are basic tools for the analysis of the dynamics of event structures. They rely on the *dynamic conflict* relation.

I-1 Dynamic Conflict.

I-1.1 Definition. (*Dynamic conflict*) We define the **dynamic conflict** relation $\#_d^\mathcal{E}$ on \mathcal{E} by:

$$\forall x, y \in \mathcal{E} \quad x \#_d^\mathcal{E} y \Leftrightarrow (x \# y) \text{ and } ([x] \cup [y] \in \mathcal{W}) \text{ and } ([x] \cup [y] \in \mathcal{W}).$$

We denote the dynamic conflict relation by $\#_d$ for short.

Events are in dynamic conflict if they are “simultaneously” enabled but in conflict.

I-1.2 Lemma. *Let v, v' be two incompatible configurations. Then there is a pair of events $(e, e') \in v \times v'$ such that $e \#_d e'$.*

Proof— For every configuration w and every event x , we set:

$$C_x(w) = \{y \in w \mid x \# y\}.$$

For $v, v' \in \mathcal{W}$, we set $C_v(v') = \{e \in v \mid C_e(v') \neq \emptyset\}$. Now for the proof, let $v, v' \in \mathcal{W}$ be incompatible. Then $C_v(v')$ is non empty, and thus admits at least a minimal element $e \in v$. We have $C_e(v') \neq \emptyset$, thus $C_e(v')$ admits a minimal element $e' \in v'$. As e is minimal in $C_v(v')$ we have:

$$\forall x \in \mathcal{E}, \quad x \prec e \Rightarrow x \text{ compatible with } v',$$

therefore $[e[$ is compatible with v' . In particular, as $e' \in v'$, $[e[$ and $[e'$ are compatible. Now as e' is minimal in $C_e(v')$, we have that $[e]$ and $[e'[$ are compatible. As $e \# e'$, this shows that $e \#_d e'$, and we have as requested: $e \in v$ and $e' \in v'$. \square

I-2 Stopping Prefixes.

I-2.1 Definition. (*Stopping prefix*) We say that a prefix B of \mathcal{E} is a **stopping prefix** if B is $\#_d$ -closed, *i.e.* if B satisfies:

$$\forall e \in B, \quad \forall x \in \mathcal{E}, \quad e \#_d x \Rightarrow x \in B.$$

We denote by \mathcal{S} the complete lattice of stopping prefixes of \mathcal{E} , and by \mathcal{S}_0 the lattice of finite stopping prefixes.

I-2.2 Smallest Stopping Prefix Containing an Event or a Set of Events.

Since \mathcal{S} is a complete lattice, and as \mathcal{E} itself is a stopping prefix, there is for every event $e \in \mathcal{E}$ a smallest stopping prefix that contains e . We denote it by $B(e)$. Note that $B(e)$ is not necessarily finite. For the same reasons, for every subset $A \subseteq \mathcal{E}$, there is a smallest stopping prefix $B(A)$ containing A , and we have $B(A) = \bigcup_{e \in A} B(e)$.

I-2.3 Definition. (*Locally finite event structure*) We say that the event structure \mathcal{E} is **locally finite** if for every event $e \in \mathcal{E}$, there is a finite stopping prefix B such that $e \in B$.

Equivalently, \mathcal{E} is locally finite if and only if $B(e)$ is finite for every $e \in \mathcal{E}$, if and only if $B(A)$ is finite for every finite subset $A \subseteq \mathcal{E}$.

I-2.4 Example. (*Locally finite trees*) If \mathcal{T} is a tree of events, the intrinsic prefixes described in Chapter 2, III-2.3, are finite stopping prefixes. Hence a tree of event is locally finite as an event structure if and only if it is locally finite as a graph, in the usual sense of finite branching. In particular the covering of a finite transition systems is locally finite.

I-2.5 Example. Figure 3.2 shows examples of stopping prefixes. The stopping prefix $B(e)$ is framed with dotted lines, with e the indicated event. Figure 3.3 depicts partially ordered stopping prefixes.




Figure 3.2: Stopping prefix $B(e)$ with $e = \boxed{\bullet}$.




Figure 3.3: Partially ordered stopping prefixes.

Stopping prefixes have the following property, that will be fundamental in the sequel. The result of the lemma is clear on Figure 3.3.

l-2.6 Lemma. *Let B be a stopping prefix of \mathcal{E} , and let D be a prefix of \mathcal{E} . If $B \cap D = \emptyset$ then $B \parallel D$.*

Proof— We have to show that $x \parallel y$ for every pair $(x, y) \in B \times D$. Let $(x, y) \in B \times D$. As B and D are two prefixes with $B \cap D = \emptyset$, we cannot have $x \preceq y$, nor $x \succeq y$. Assume that $x \# y$ holds. Then by Lemma l-1.2, there exists $x' \in [x]$ and $y' \in [y]$ such that $x' \#_d y'$. But then $x' \in B$ and $y' \in D$, and B is $\#_d$ -closed, thus $y' \in B \cap D$, a contradiction with $B \cap D = \emptyset$. This shows that $x \parallel y$. \square

l-2.7 Remark. The result of Lemma l-2.6 fails if B and D are two prefixes in general. Take for instance the prefixes defined by the two events of the event structure: $\bullet \rightsquigarrow \bullet$.

l-3 Stopped Configurations.

The following lemma motivates the definition l-3.2 of *stopped configurations*. Recall that a prefix P of an event structure \mathcal{E} is said to be *intrinsic to \mathcal{E}* if every maximal configuration ω intersects P through a maximal configuration of P , which defines the mapping $\pi_P : \Omega \rightarrow \Omega_P$, $\pi_P(\omega) = \omega \cap P$.

l-3.1 Lemma.

1. Any stopping prefix of \mathcal{E} is intrinsic to \mathcal{E} .
2. If $B \subseteq B'$ are two stopping prefixes, then B is a stopping prefix of B' . In particular for $B \subseteq B'$, B is intrinsic to B' , so that we have the following commutative diagram:

$$\begin{array}{ccc} \Omega & \xrightarrow{\pi_{B'}} & \Omega_{B'} \\ & \searrow \pi_B & \downarrow \pi_{B, B'} \\ & & \Omega_B \end{array}$$

Proof— 1. Let B be a stopping prefix of \mathcal{E} , and let $\omega \in \Omega$ be a maximal configuration of \mathcal{E} , we have to show that $\omega \cap B$ is a maximal configuration of B . Assume that $\omega_B = \omega \cap B$ is not maximal in B . Then there is an event $e \in B$ such that $e \notin \omega_B$, and such that $\omega_B \cup \{e\}$ is a configuration of B .

e belongs to B and not to $\omega \cap B$, thus e does not belong to ω . We claim that $e \# \omega$, i.e. there is an event $x \in \omega$ with $x \# e$. Otherwise $[e] \cup \omega$ is a conflict-free prefix that strictly contains ω , contradicting that ω is maximal. Hence $e \# \omega$, and it follows from Lemma l-1.2 that there is a pair (f, y) of events such that:

$$a) f \preceq e, \quad b) y \in \omega, \quad c) f \#_d y.$$

Since $e \in B$, and since B is a prefix, a) implies that $f \in B$. As B is $\#_d$ -closed, c) implies that $y \in B$, and with b), we get: $y \in \omega_B$. But then the configuration $\omega_B \cup \{e\}$ contains the conflict $f \# y$, a contradiction. We have shown that ω_B is maximal in B .

2. Obvious. \square

1-3.2 Definition. (*Stopped configurations*) We say that a configuration w is **stopped** in \mathcal{E} , or that w is a **stopped configuration** of \mathcal{E} , if there is a stopping prefix B such that $w \in \Omega_B$. We say that w is **finitely stopped** if there is a *finite* stopping prefix B such that $w \in \Omega_B$. The elements of Ω_B , with $B \in \mathcal{S}$, are called the ***B*-stopped** configurations of \mathcal{E} .

In particular, since \mathcal{E} is a stopping prefix, every maximal maximal configuration is stopped. From Lemma 1-3.1, point 2, the stopped configurations of a stopping prefix B are the stopped configurations of \mathcal{E} included in B .

1-3.3 If v is stopped in \mathcal{E} , then v is $B(v)$ -stopped. Indeed let $B \in \mathcal{S}$ such that $v \in \Omega_B$. Then $B(v) \subseteq B$, and thus $v \in \Omega_{B(v)}$ by Lemma 1-3.1, point 2.

As a consequence, in a locally finite event structure, every finite stopped configuration is finitely stopped: if v is B stopped and if v is finite, then v is maximal in $B(v)$, which is finite since the event structure is locally finite. The non locally finite event structure of Ch. 2, Figure 2.1, page 75, admits finite stopped configurations that are not finitely stopped.

1-3.4 Example. In a tree, every configuration is stopped (Cf. VII). The same holds for confusion-free event structures (Cf. Ch. 5). In general event structures, this is not the case as shown by the example below on the left hand. In this finite event structure \mathcal{E} , the configuration v consisting of the framed event is not stopped. Otherwise, v would be maximal in $B(v) = \mathcal{E}$, but this is not the case. Compare with the example below on the right hand: by adding a conflict the event structure becomes a tree and the configuration becomes stopped. There is no more concurrency.



II—Cone of Future

We introduce the cone of future of a configuration, a notion that we will use throughout the document. We present the definition in II-1. We can then define in II-2 a concatenation and a cancellation of configurations with this formalism. Then we

analyse the relations between the cone of future and other objects that we have defined for event structures: the induced topology is examined in II-3. The dynamic conflict and the stopping prefixes *in a cone of future* are studied in II-4, whereas the cone of future *in a prefix* is the topic of II-5.

II-1 The Fundamental Isomorphism of Partial Orders.

Let v be a configuration of \mathcal{E} . We denote by $\mathcal{W}(v)$ and $\Omega(v)$ the following subsets of \mathcal{W} and Ω :

$$\mathcal{W}(v) = \{w \in \mathcal{W} \mid w \supseteq v\}, \quad \Omega(v) = \{\omega \in \Omega \mid \omega \supseteq v\}. \quad (3.1)$$

We have already encountered $\Omega(v)$ in Ch. 2, II-3.1, we call it the **shadow** of v . We say that the shadow is *finite* if v is finite. These two subsets can be realised as, respectively, the set of configurations and of maximal configurations of a sub-event structure of \mathcal{E} .

II-1.1 Definition. (*Cone of future of a configuration*) Let v be a configuration of \mathcal{E} . We define the **cone of future** of v as the sub-event structure \mathcal{E}^v given by:

$$\begin{aligned} \mathcal{E}^v &= \{e \in \mathcal{E} \mid \exists w \in \mathcal{W} : w \supseteq \{e\} \cup v\} \setminus v \\ &= \{e \in \mathcal{E} \setminus v \mid e \text{ is compatible with } v\}. \end{aligned}$$

We denote by \mathcal{W}^v and Ω^v respectively the set of configurations and of maximal configurations of \mathcal{E}^v .

Let w be a configuration that contains v . Then $w \cap \mathcal{E}^v$ is conflict-free in \mathcal{E}^v since it is conflict-free in \mathcal{E} (Cf. Ch. 1, I-2.6, for sub-event structures). Clearly, $w \cap \mathcal{E}^v$ is a prefix of \mathcal{E}^v , thus we have $w \cap \mathcal{E}^v \in \mathcal{W}^v$. So we define a mapping by setting:

$$\bar{\Phi}_v : \mathcal{W}(v) \rightarrow \mathcal{W}^v, \quad w \mapsto w \cap \mathcal{E}^v.$$

II-1.2 Lemma. For every event $e \in \mathcal{E}^v$, we have $[e]^{\mathcal{E}^v} = [e]^{\mathcal{E}} \setminus v = [e]^{\mathcal{E}} \cap \mathcal{E}^v$.

Proof – Let $y(e) = [e]^{\mathcal{E}} \setminus v$. Using that $[e]^{\mathcal{E}^v} = \{x \in \mathcal{E}^v \mid x \preceq e\}$, we see that $[e]^{\mathcal{E}^v} \subseteq y(e)$. As e is compatible in \mathcal{E} with v , every $f \in [e]^{\mathcal{E}}$ is compatible in \mathcal{E} with v . This implies that $y(v) \subseteq [e]^{\mathcal{E}^v}$, whence the equality. The part $[e]^{\mathcal{E}} \setminus v = [e]^{\mathcal{E}} \cap \mathcal{E}^v$ is obvious. \square

II-1.3 Proposition. The mapping $\bar{\Phi}_v : \mathcal{W}(v) \rightarrow \mathcal{W}^v$ is an isomorphism of partial orders, that induces by restriction a one-to-one mapping:

$$\Phi_v : \Omega(v) \rightarrow \Omega^v.$$

Proof— The mapping $\bar{\Phi}_v : w \mapsto w \cap \mathcal{E}^v$ is increasing. To show that $\bar{\Phi}_v$ is an isomorphism, we exhibit its inverse mapping. We set $\bar{\Psi}_v(z) = v \cup z$ for every configuration z of \mathcal{E}^v . We have for every $z \in \mathcal{W}^v$:

$$z = \bigcup_{e \in z} [e]^{\mathcal{E}^v},$$

thus $\bar{\Psi}_v(z) = \bigcup_{e \in z} (v \cup [e]^{\mathcal{E}^v})$. By Lemma I-3.1, we have $v \cup [e]^{\mathcal{E}^v} = v \cup [e]^{\mathcal{E}}$ for every $e \in z$. Thus $\bar{\Psi}_v(z)$ is a configuration of \mathcal{E} , as a union of pairwise compatible configurations. Clearly, we have $\bar{\Psi}_v \circ \bar{\Phi}_v = \text{Id}_{\mathcal{W}(v)}$ and $\bar{\Phi}_v \circ \bar{\Psi}_v = \text{Id}_{\mathcal{W}^v}$, so $\bar{\Phi}_v$ is invertible and $(\bar{\Phi}_v)^{-1} = \bar{\Psi}_v$. In particular, $\bar{\Phi}_v$ maps the maximal elements of $\mathcal{W}(v)$ onto the maximal elements of \mathcal{W}^v , *i.e.* onto Ω^v . We conclude by observing that the maximal elements of $\mathcal{W}(v)$ are the elements of $\Omega(v)$. \square

Figure 3.4 depicts a geometric representation for the cone of future of a configuration.




Figure 3.4: *Cone of future of a configuration. A configuration v cuts off three regions: events in v , events in the cone of future \mathcal{E}^v , and events in conflict with v .*

II-1.4 Notation. For every configuration v , $\bar{\Phi}_v$ and $\bar{\Psi}_v$ are denoted by the arrows $\mathcal{W}(v) \rightarrow \mathcal{W}^v$ and $\Omega(v) \rightarrow \Omega^v$.

II-1.5 Remark. (*Trivial cases*) Observe the following equalities:

$$\mathcal{E}^\emptyset = \mathcal{E}, \quad \forall \omega \in \Omega, \quad \mathcal{E}^\omega = \emptyset.$$

For every configuration v , $\mathcal{E}^v = \emptyset$ implies that v is maximal.

II-2 Concatenation and Cancellation.

The above isomorphism allows to properly define the concatenation, and also the cancellation, a sort of subtraction of configurations.

II-2.1 Definition. Let v be a configuration, and let $y \in \mathcal{W}^v$. We define the **concatenation** of v and y , in this order, by:

$$\oplus : \{(v, y) \mid v \in \mathcal{W}, y \in \mathcal{W}^v\} \rightarrow \mathcal{W}, \quad v \oplus y = \overline{\Phi}_v^{-1}(y) = v \cup y,$$

and we always have $v \oplus y \in \mathcal{W}$.

We define the **left-cancellation**, or shortly the **cancellation**, as follows. For $v \subseteq v'$ two configurations, $v' \ominus v$ is the unique configuration y of \mathcal{E}^v such that $v' = v \oplus y$, and y is given by $y = v' \cap \mathcal{E}_v = v' \setminus v$. We have thus:

$$\ominus : \{(v', v) \in \mathcal{W} \times \mathcal{W} \mid v \subseteq v'\} \rightarrow \mathcal{W}, \quad v' \ominus v = \overline{\Phi}_v(v') = v' \setminus v,$$

and we always have that $v' \ominus v \in \mathcal{W}^v$.

If $e \in \mathcal{E}^v$ satisfies $\{e\} = [e]^{\mathcal{E}^v}$, i.e. if e is minimal in \mathcal{E}^v , we shortly write $v \oplus e$ to denote $v \oplus [e]^{\mathcal{E}^v} = v \cup \{e\}$. Figure 3.5 illustrates the concatenation and the cancellation.




Figure 3.5: Concatenation and cancellation.

II-2.2 Composition. If $v \in \mathcal{W}$ and $y \in \mathcal{W}^v$, we have $v \oplus \Omega^v(y) = \Omega(v \oplus y)$ and:

$$\overline{\Phi}_{v \oplus y} = \overline{\Phi}_y^v \circ \overline{\Phi}_v, \quad (3.2)$$

where $\overline{\Phi}_y^v$ denotes the isomorphism $\mathcal{W}^v(y) \rightarrow (\mathcal{W}^v)^y$ constructed *w.r.t.* the event structure \mathcal{E}^v .

II-2.3 Lemma. (*Concatenation of cones*) For any $v \in \mathcal{W}$ and $y \in \mathcal{W}^v$, we have:

$$(\mathcal{E}^v)^y = \mathcal{E}^{v \oplus y}, \quad (3.3)$$

and for any configurations $v \subseteq v'$ of \mathcal{E} :

$$(\mathcal{E}^v)^{v' \ominus v} = \mathcal{E}^{v'}. \quad (3.4)$$

We have for any pair v, v' of configurations: $v \subseteq v' \Rightarrow \mathcal{E}^{v'} \subseteq \mathcal{E}^v$. This inclusion induces through its action on sets the canonical injection: $\Omega^{v'} \rightarrow \Omega^v$, conjugated to the inclusion mapping:

$$\Omega(v') \rightarrow \Omega(v) .$$

Proof – It all follows from (3.2). \square

II-3 Compositionality of Topologies.

Let v be a finite configuration of \mathcal{E} . As a subset of \mathcal{W} , $\mathcal{W}(v)$ is equipped with the restriction of the projective topology on \mathcal{W} . On the other hand, \mathcal{W}^v is equipped with the projective topology inherited from the event structure \mathcal{E}^v . In the same manner, $\Omega(v)$ is equipped with the restriction to $\Omega(v)$ of the operational topology on Ω , whereas Ω^v is equipped with the operational topology inherited from the event structure \mathcal{E}^v .

II-3.1 Notation. We denote by \mathcal{P}_0^v the set of finite prefixes of \mathcal{E}^v . Then we have the following elementary result.

II-3.2 Lemma. *Let v be a finite configuration of \mathcal{E} . For every $P \in \mathcal{P}_0$, $P \cap \mathcal{E}^v \in \mathcal{P}_0^v$. For every $Q \in \mathcal{P}_0^v$, $v \cup Q \in \mathcal{P}_0$.*

Proof – The first part is trivial. For the second part, let $Q \in \mathcal{P}_0^v$. We write $v \cup Q$ as a union of finite prefixes:

$$v \cup Q = \bigcup_{e \in Q} (v \cup [e]^{\mathcal{E}^v}) = \bigcup_{e \in Q} (v \cup [e]^{\mathcal{E}}),$$

the later by lemma II-1.2. So $v \cup Q \in \mathcal{P}$. \square

As a consequence, we have:

II-3.3 Proposition. *Let v be a finite configuration of \mathcal{E} . The mappings $\mathcal{W}(v) \rightarrow \mathcal{W}^v$ and $\Omega(v) \rightarrow \Omega^v$ are two homeomorphisms.*

Proof – We show for instance that $\mathcal{W}(v) \rightarrow \mathcal{W}^v$ is continuous. An elementary open subset of \mathcal{W}^v has the form:

$$U = \{x \in \mathcal{W}^v \mid Q \cap x = z\} ,$$

with Q a finite prefix of \mathcal{E}^v and z a configuration of Q (from Ch. 2, II-2.1). According to lemma II-3.2, $P = v \cup Q$ is a finite prefix of \mathcal{E} . Then we have:

$$\begin{aligned} \overline{\Phi}_v^{-1}(U) &= \{v \oplus x, x \in U\} \\ &= \{y \in \mathcal{W} \mid y \cap P = v \oplus z\} , \end{aligned}$$

an open subset of $\mathcal{W}(v)$. We show in a similar manner that $\mathcal{W}^v \rightarrow \mathcal{W}(v)$ is continuous. Hence $\mathcal{W}(v) \rightarrow \mathcal{W}^v$ is an isomorphism of partial orders and a homeomorphism. It implies that the restriction of the mapping to the maximal elements is a homeomorphism. Since the maximal elements of $\mathcal{W}(v)$ are indeed the elements of $\Omega(v)$, we obtain that $\Omega(v) \rightarrow \Omega^v$ is a homeomorphism. \square

II-4 Dynamic Conflict in a Cone of Future.

For a configuration v , we denote by \mathcal{S}^v the lattice of stopping prefixes of \mathcal{E}^v , and by \mathcal{S}_0^v the lattice of finite stopping prefixes of \mathcal{E}^v . We denote by $\#_d^v$ the dynamic conflict relation defined in \mathcal{E}^v . For e an event of \mathcal{E}^v , we denote by $[e]^v$ the smallest configuration of \mathcal{E}^v that contains e , and similarly for $[e]^v$. As the conflict in \mathcal{E}^v is by definition the restriction of $\#$ to \mathcal{E}^v (sub-structure, Ch. 1, I-2.6), the relation $\#_d^v$ on \mathcal{E}^v is given by:

$$e \#_d^v e' \Leftrightarrow (e \# e') \text{ and } ([e]^v \cup [e']^v \in \mathcal{W}^v) \text{ and } ([e]^v \cup [e']^v \in \mathcal{W}^v). \quad (3.5)$$

II-4.1 Lemma. *The dynamic conflict relations $\#_d$ and $\#_d^v$ satisfy:*

$$\#_d^v = \#_d \cap (\mathcal{E}^v \times \mathcal{E}^v). \quad (3.6)$$

Proof— Applying Lemma II-1.2, we have for every event $e \in \mathcal{E}^v$:

$$[e]^v = [e] \cap \mathcal{E}^v, \quad [e']^v = [e'] \cap \mathcal{E}^v. \quad (3.7)$$

$v \oplus [e]^v$ is a configuration of \mathcal{E} that contains e , thus $v \oplus [e]^v \supseteq v \cup [e]$. From (3.7) we have $v \oplus [e]^v \subseteq v \cup [e]$, and thus $v \oplus [e]^v = v \cup [e]$. It implies: $v \cup [e]^v = v \cup [e]$.

Now to show (3.6), we begin with the inclusion \subseteq . Let $e, e' \in \mathcal{E}^v$ such that $e \#_d^v e'$. Then $[e]^v \cup [e']^v \in \mathcal{W}^v$, and thus:

$$\begin{aligned} v \oplus ([e]^v \cup [e']^v) &\in \mathcal{W} \\ &= (v \oplus [e]^v) \cup (v \oplus [e']^v) \\ &= (v \cup [e]) \cup (v \cup [e']) \\ &\supseteq [e] \cup [e']. \end{aligned}$$

So $[e]$ and $[e']$ are compatible. We show in the same way that $[e]$ and $[e']$ are compatible, and as $e \#_d^v e' \Rightarrow e \# e'$, we have $e \#_d e'$. This shows \subseteq in (3.6). For the \supseteq part, take $e, e' \in \mathcal{E}^v$ with $e \#_d e'$. Then (3.7) shows that $e \#_d^v e'$. \square

II-4.2 Lemma. *For every configuration v , the association $B \mapsto B \cap \mathcal{E}^v$ defines mappings from \mathcal{S} to \mathcal{S}^v , and from \mathcal{S}_0 to \mathcal{S}_0^v .*

Proof— Follows immediately from II-4.1 and II-3.2. \square




Figure 3.6: Stopping prefix B splits into two stopping prefixes in the cone of future \mathcal{E}^{i_1} .

II-4.3 Remark. (Stopped configurations are not closed under concatenation)
 Lemma II-4.2 shows that $B \cap \mathcal{E}^v$ is a stopping prefix of \mathcal{E}^v . However not every stopping prefix of \mathcal{E}^v is obtained this way in general. Consider for instance the event structure \mathcal{E} depicted in Figure 3.6, with stopping prefix $B = \mathcal{E}$. Then B splits into two stopping prefixes in the cone of future \mathcal{E}^{i_1} .

In particular, observe that $v = i_1 \oplus i_3$ is not a stopped configuration (otherwise v would be maximal in $B(v) = \mathcal{E}$), although v is the concatenation of two stopped configurations. This example shows that *the class of stopped configurations is not closed under concatenation*.

II-5 Cone of Future in a Prefix.

If P is a prefix of \mathcal{E} , we recall that we do not distinguish between configurations of \mathcal{E} included in P , and configurations of P . For v a configuration of P , P^v denotes the cone of future of v in P .

II-5.1 Lemma. *Let v be configuration included in a prefix P . Then we have: $P^v = \mathcal{E}^v \cap P$.*


Proof – Obvious. \square

II-5.2 Lemma. *Let B be a stopping prefix of \mathcal{E} . Let w be a configuration of \mathcal{E} , and let $v = w \cap B$. Then we have: $B^v = \mathcal{E}^w \cap B$.*

Proof – We write $w = v \oplus z$, with z a configuration of \mathcal{E}^v , satisfying $z \cap B = \emptyset$. Let $e \in B^v$, and assume that $e \notin \mathcal{E}^w$. Since $e \notin w$, it implies that $e \# w$, and thus $e \#^{\mathcal{E}^v} z$, since $w = v \oplus z$. Applying Lemma I-1.2 in the event structure \mathcal{E}^v , we find $f \in \mathcal{E}^v$ with $f \preceq e$ and $g \in z$ such that $f \#_d^v g$, and thus $f \#_d g$. As B is $\#_d$ -closed, and since $f \in B$, it implies that $g \in B$, which contradicts that $z \cap B = \emptyset$. This shows that $e \in \mathcal{E}^w$, and thus $B^v \subseteq \mathcal{E}^w \cap B$. The converse inclusion is obvious. \square

II-5.3 Remark. The condition “ B is a stopping prefix” is necessary in Lemma II-5.2. Indeed, consider the event structure \mathcal{E} below, with $w = \{b\}$ and

with $B = \{a, c\}$, which is not a stopping prefix. Then we have $v = \emptyset$, and $B^v \neq \emptyset$, whereas $\mathcal{E}^w \cap B = \emptyset$.



III—Normal Decomposition

We introduce now the *normal decomposition* of maximal and of stopped configurations. A more extensive study of decompositions similar to the normal decomposition is the topic of paragraph IV. We also introduce the class of event structures of *finite concurrent width*¹.

III-1 Finite Concurrent Width.

We introduce the following class of event structures. They model systems where finitely many events can concurrently appear, a reasonable computational assumption.

III-1.1 Definition. (*Finite concurrent width*) We say that an event structure \mathcal{E} is of **finite concurrent width** if every \parallel -clique of \mathcal{E} is finite.

Obviously, if \mathcal{E} has finite concurrent width, the same holds for every sub-event structure of \mathcal{E} , and in particular for \mathcal{E}^v for every $v \in \mathcal{W}$.

Let \mathcal{S} denote the lattice of stopping prefixes of an event structure \mathcal{E} . Let \mathcal{S}^* denote the set of non empty stopping prefixes of \mathcal{E} . Then $(\mathcal{S}^*, \subseteq)$ is a partial order. In the sequel we can avoid the use of the following lemma, it is however instructive (notice the use of Zorn's Lemma).

III-1.2 Lemma. *Assume that \mathcal{E} has finite concurrent width. Then every $B \in \mathcal{S}^*$ is supset of a minimal element of \mathcal{S}^* .*

Proof— We will use the following elementary remark: since \mathcal{E} is well founded, every prefix $P \subseteq \mathcal{E}$ is well founded, thus P admits minimal elements $e \in P$, which are all minimal in \mathcal{E} .

¹In Ch. 7, we will introduce a concurrent height.

Let $B \in \mathcal{S}$. We show that every chain $(B_n)_{n \in I}$ admits a lower bound in \mathcal{S}^* , where I is a totally ordered set and B_n , for $n \in I$, are non empty stopping prefixes of B . Since \mathcal{S} is a lattice, the bound exists in \mathcal{S} , given by $C = \bigcap_{n \in I} B_n$, and we show that C is non empty. Assume that $C = \emptyset$, and consider the sequence of events constructed as follows. Fix ω a maximal configuration of \mathcal{E} . Then for all $n \in I$, $\omega \cap B_n$ is maximal in B_n , and since we assume $B_n \neq \emptyset$, $\omega \cap B_n$ is non empty.

Choose e_0 a minimal event of ω , and thus of \mathcal{E} . Assume that e_0, \dots, e_i have been constructed until rank $k \geq 0$, with e_i a minimal event of ω for each $i \geq 0$, and $e_i \neq e_j$ for all $i \neq j$. Since $\bigcap_{n \in I} \downarrow B_n = \emptyset$, there is an index $n \in I$ such that $B_n \cap \{e_0, \dots, e_k\} = \emptyset$. As $\omega \cap B_n$ is non empty, we choose e_{k+1} minimal in $\omega \cap B_n$, and the induction is complete. Then all e_i are distinct, and for all $i, j \geq 0$, $\neg(e_i \preceq e_j)$ holds since e_j is minimal in \mathcal{E} , and similarly $\neg(e_i \succeq e_j)$. Since e_i and e_j belong to the same configuration ω , $\neg(e_i \# e_j)$ also holds, and thus $e_i \parallel e_j$. This contradicts that \mathcal{E} has finite concurrent width. Hence $C \neq \emptyset$.

Zorn's lemma implies that the set of non empty stopping prefixes included in B admits a minimal element. \square

III-1.3 Example. We show in Figure 3.7 an example of an event structure \mathcal{E} of infinite concurrent width, and such that \mathcal{S}^* admits no minimal element. Indeed every $B_k = \{e_k, e_{k+1}, \dots\}$ is a stopping prefix, and every non empty stopping prefix B that contains e_k strictly contains B_{k+1} . Remark the infinite \parallel -clique $\{e_{2k+1}, k \geq 0\}$.




Figure 3.7: \mathcal{S}^* admits no minimal element.

III-1.4 Remark. Let \mathcal{E} be an event structure locally finite and of finite concurrent width. A prefix P does not necessarily even intersect a minimal element of \mathcal{S}^* (Figure 3.8, $P = \{e\}$). If e is minimal in \mathcal{E} , the stopping prefix $B(e)$, minimal to contain e , is not necessarily minimal in \mathcal{S}^* (Figure 3.8).

Local finiteness does not imply finite concurrent width: consider for instance the infinite event structure \mathcal{E} without conflict: $\bullet \bullet \bullet \dots$. \mathcal{E} is locally finite and has infinite concurrent width. And finally, finite concurrent width does not imply local finiteness, as show by the example of non locally finite event structure given in Chapter 2, Figure 2.1, p. 75.




Figure 3.8: Although e is minimal in \mathcal{E} , $B(e)$ (external frame) is not minimal in \mathcal{S}^* since $B(e)$ contains a smaller stopping prefix (internal frame).

III-2 Initial Stopping Prefix and Initial Branching Cells.

We keep the notation \mathcal{S}^* to denote the non void stopping prefixes. We define the *initial branching cells* as the minimal elements of \mathcal{S}^* , if they exist.

III-2.1 Definition. (*Initial branching cells, initial stopping prefix, full-initial configurations.*) We set $\Delta^\perp(\mathcal{E}) = \min(\mathcal{S}^*)$, and the elements of $\Delta^\perp(\mathcal{E})$ are called the **initial branching cells** of \mathcal{E} . We define the **initial stopping prefix** of \mathcal{E} by:

$$B^\perp(\mathcal{E}) = \bigcup_{\lambda \in \Delta^\perp(\mathcal{E})} \lambda. \tag{3.8}$$

We say that a configuration w is **full-initial** in \mathcal{E} if w is a maximal configuration of the initial stopping prefix $B^\perp(\mathcal{E})$.

Figure 3.9 depicts the initial branching cells of an event structure. Figure 3.10 depicts a more abstract representation.




Figure 3.9: Initial branching cells.




Figure 3.10: Initial branching cells and initial stopping prefix.

III-2.2 Lemma. Let \mathcal{E} be an event structure of finite concurrent width. Then the initial stopping prefix $B^\perp(\mathcal{E})$ is empty if and only if \mathcal{E} is empty. For every $B \in \mathcal{S}$ we have: $B^\perp(B) = B^\perp(\mathcal{E}) \cap B$.

Proof– The first part follows from III-1.2, the second part is obvious. \square

III-2.3 Proposition. *Let \mathcal{E} be an event structure of finite concurrent width. Then for every $v \in \mathcal{W}_0$, $\Delta^\perp(\mathcal{E}^v)$ is a finite collection of sets. If moreover \mathcal{E} is locally finite, then all $\lambda \in \Delta^\perp(\mathcal{E}^v)$ are finite. In particular the initial stopping prefix $B^\perp(\mathcal{E})$ is finite.*

Proof – Since \mathcal{E}^v has finite concurrent width for every $v \in \mathcal{W}_0$, we can assume without loss of generality that $v = \emptyset$, and thus $\mathcal{E}^v = \mathcal{E}$. We use an independent result, stated below in IV-3.1: all $\lambda \in \Delta^\perp(\mathcal{E})$ are concurrent. Since every $\lambda \in \Delta^\perp(\mathcal{E})$ is non empty by definition, we choose a collection $(e_\lambda)_{\lambda \in \Delta^\perp(\mathcal{E})}$ with $e_\lambda \in \lambda$ for all $\lambda \in \Delta^\perp(\mathcal{E})$, which is thus a \parallel -clique one-to-one with $\Delta^\perp(\mathcal{E})$ (the construction applies even if \mathcal{E} is the empty event structure). Since \mathcal{E} has finite concurrent width, it implies that $\Delta^\perp(\mathcal{E})$ is finite.

Assume that \mathcal{E} is locally finite. For every $\lambda \in \min(\mathcal{S}^*)$, λ contains an event e , and thus $\lambda \supseteq B(e)$ (I-2.2), which implies $\lambda = B(e)$. $B(e)$ is finite since \mathcal{E} is locally finite. In particular, $B^\perp(\mathcal{E})$ is a finite collection of finite sets, thus $B^\perp(\mathcal{E})$ is finite. \square

III-3 Normal decomposition.

We show that every maximal configuration is the concatenation of recursive full-initial configurations, and this decomposition is unique, hence the terminology *normal decomposition*.

From now on, in this chapter, we assume that \mathcal{E} is an event structure **locally finite, and of finite concurrent width**.

III-3.1 Proposition-Definition. (*Normal decomposition*) For $\omega \in \Omega$ fixed, there is a unique infinite sequence $(v_j, z_j)_{j \geq 1}$ such that, with $v_0 = \emptyset$ we have for all $j \geq 1$:

$$v_j = v_{j-1} \oplus z_j, \quad \text{with } v_j \subseteq \omega \text{ and } z_j \in \Omega_{B^\perp(\mathcal{E}^{v_{j-1}})}, \quad (3.9)$$

i.e. z_j is full-initial in $\mathcal{E}^{v_{j-1}}$. We define the sequence $(v_j, z_j)_{j \geq 1}$, completed with $v_0 = \emptyset$, as the **normal decomposition** of ω .

Proof – *Existence.* By induction, assume that (v_j, z_j) has been constructed for $1 \leq j \leq n$. We have $\omega \supseteq v_j$, so that $\omega \in \Omega(v_j)$. Thus $\xi_j = \omega \ominus v_j$ is maximal in \mathcal{E}^{v_j} by Proposition II-1.3. We set $z_{j+1} = \xi_j \cap B^\perp(\mathcal{E}^{v_j})$, and z_{j+1} is maximal in the stopping prefix $B^\perp(\mathcal{E}^{v_j})$ as $B^\perp(\mathcal{E}^{v_j})$ is intrinsic to \mathcal{E}^{v_j} . Then by setting $v_{j+1} = v_j \oplus z_{j+1}$, we have $v_{j+1} \subseteq \omega$. This completes the induction.

Uniqueness. Let $(v_j, z_j)_{j \geq 1}$ and $(v'_j, z'_j)_{j \geq 1}$ be two sequences satisfying (3.9). By definition $v_0 = v'_0 = \emptyset$. Assume that $v_n = v'_n$, with $n \geq 0$, then z_{n+1} and z'_{n+1} are two configurations compatible and maximal in $B^\perp(\mathcal{E}^{v_n})$, so they are equal, and thus $v_{n+1} = v'_{n+1}$. This shows the uniqueness. \square

III-3.2 Other Formulation. For every $\omega \in \Omega$, the normal decomposition of ω is given by:

$$v_0 = \emptyset, \quad z_{n+1} = \omega \cap B^\perp(\mathcal{E}^{v_n}), \quad v_{n+1} = v_n \oplus z_{n+1}.$$

If we set $\xi_n = \omega \ominus v_n$ for all $n \geq 0$, then we have for all $n \geq 0$:

$$\xi_n \in \Omega^{v_n}, \quad z_{n+1} = \xi_n \cap B^\perp(\mathcal{E}^{v_n}).$$

Figure 3.11 illustrates a normal decomposition.




Figure 3.11: Three first steps of a normal decomposition.

Now we show that the normal decomposition of ω entirely characterises ω .

III-3.3 Proposition. Let $(v_n, z_n)_{n \geq 1}$ be the normal decomposition of an element $\omega \in \Omega$. Then $\bigcup_{n \geq 1} v_n = \omega$. Equivalently, $\lim_{n \rightarrow \infty} v_n = \omega$ in the projective topology.

Proof – Let $v = \bigcup_{n \geq 1} v_n$. We show that v is maximal. Assume that v is not maximal, then there is an event $e \in \mathcal{E}$ such that $v \oplus e \in \mathcal{W}$. It implies that $[e] \subseteq v$, so there is an integer n such that $e \in v_n$. Then e is a minimal event of \mathcal{E}^{v_n} , and it implies that there is a $\lambda \in \Delta^\perp(\mathcal{E}^{v_n})$ such that $e \in \lambda$. Let $z_{n+1} = v_{n+1} \ominus v_n$. e is compatible with v , so e is compatible with z_{n+1} , which is maximal in λ . It implies that $e \in z_{n+1}$, and contradicts that $e \ni v$. Thus v is maximal, and as ω is maximal too, $v = \omega$. \square

Remark that if ω is finite, $v_n = \omega$ for n large enough, and then $\mathcal{E}^{v_n} = \emptyset$, $B^\perp(\mathcal{E}^{v_n}) = \emptyset$ and $z_n = \emptyset$.

III-3.4 Proposition. (Induced normal decomposition) Let B be a stopping prefix of \mathcal{E} , and let $\omega_B = \omega \cap B$, with ω an element of Ω . Let $(v_n, z_n)_{n \geq 1}$ and $(u_n, r_n)_{n \geq 1}$ be the normal decompositions of ω and of ω_B . Then we have for all $n \geq 1$:

$$u_n = v_n \cap B, \quad r_n = z_n \cap B.$$

Proof— We show the following by induction on $n \geq 1$:

$$H_n : \begin{cases} r_n = z_n \cap B, & u_n = v_n \cap B, \\ B^{u_n} = \mathcal{E}^{v_n} \cap B. \end{cases}$$

We have $r_1 = \omega_B \cap B^\perp(B)$ and $B^\perp(B) = B^\perp(\mathcal{E}) \cap B$ by Lemma III-2.2, and thus $r_1 = z_1 \cap B$. It follows that $u_1 = v_1 \cap B$. We apply Lemma II-5.2 to obtain that $B^{u_1} = \mathcal{E}^{v_1} \cap B$. This shows H_1 (observe that H_0 is not defined).

Assume that H_n holds. Then we have:

$$r_{n+1} = \omega_B \cap B^\perp(B^{u_n}) = \omega \cap B^\perp(B^{u_n}). \quad (3.10)$$

We have $B^{u_n} = \mathcal{E}^{v_n} \cap B$ by induction, thus B^{u_n} is a stopping prefix of \mathcal{E}^{v_n} by Lemma II-4.2. Then we apply Lemma III-2.2 in the event structure \mathcal{E}^{v_n} , and we get:

$$\begin{aligned} B^\perp(B^{u_n}) &= B^\perp(\mathcal{E}^{v_n}) \cap B^{u_n} \\ &= B^\perp(\mathcal{E}^{v_n}) \cap B. \end{aligned}$$

With (3.10) it implies:

$$r_{n+1} = \omega \cap B^\perp(\mathcal{E}^{v_n}) \cap B = z_{n+1} \cap B.$$

It follows that $u_{n+1} = v_{n+1} \cap B$. We apply Lemma II-5.2 in the event structure \mathcal{E}^{v_n} to get:

$$\begin{aligned} (B^{u_n})^{r_n} &= (\mathcal{E}^{v_n})^{z_n} \cap B^{u_n} \\ B^{u_{n+1}} &= \mathcal{E}^{v_{n+1}} \cap B^{u_n} = \mathcal{E}^{v_{n+1}} \cap B. \end{aligned}$$

This shows H_{n+1} and completes the induction, which implies the statement of the lemma. \square

IV—Well-Stopped Configurations

We still consider locally finite event structures, of finite concurrent width. The class of stopped configurations is not rich enough. In particular, we have seen in II-4.3 that the concatenation of two stopped configurations is not stopped. From a dynamics point of view, concatenation however is crucial. Thus we close the class of finitely stopped configurations under concatenation, and the resulting class of configurations has good compositional properties. We call them *well-stopped* configurations.

We collect in this Section the basic properties of well-stopped configurations. We present the definition and the immediate properties of well-stopped configurations in IV-1. We show in IV-2 that they have good properties *w.r.t.* the restriction to stopping prefixes. In IV-3 and IV-4, we analyse the properties of the initial stopping prefix. We describe in particular its well-stopped configurations, in terms of *germs*. Finally, IV-5 gives a useful characterisation of well-stopped configurations.

IV-1 Well-Stopped Sequences and Well-Stopped Configurations.

To close the class of finitely stopped configurations under concatenation, we proceed as follows.

IV-1.1 Definition. (*Well-stopped sequences, well-stopped configurations*) Let I be a finite or infinite interval of \mathbb{N} containing 0, finite or infinite. Let $(v_n)_{n \in I}$ be a non decreasing sequence of configurations of \mathcal{E} . We set $\mathcal{E}_0 = \mathcal{E}$, and for every non zero $n \in I$, we set $\mathcal{E}_n = \mathcal{E}^{v_n}$ and $z_n = v_n \ominus v_{n-1}$. Hence z_n is the unique configuration of \mathcal{E}_{n-1} such that:

$$v_n = v_{n-1} \oplus z_n .$$

We say that $(v_n)_{n \in I}$ is a **well-stopped sequence** if z_n is a finite stopped configuration of \mathcal{E}_{n-1} for every non zero integer $n \in I$. We say that a configuration $v \in \mathcal{W}$ is **well-stopped** if v is limit of a well-stopped sequence of configurations, *i.e.* if:

$$v = \bigcup_{n \in I} v_n ,$$

with $(v_n)_{n \in I}$ a well-stopped sequence. We say that $(v_n)_{n \in I}$ **leads** to v .

Figure 3.12 illustrates the construction of well-stopped configurations.




Figure 3.12: *Construction of well-stopped configurations. Successive cones of futures are drawn and successive stopping prefixes are filled in with gray.*

IV-1.2 Length and Finite Well-Stopped Configurations. A well-stopped sequence $(v_i)_{i \in I}$ is said to be of **finite length** if there is an integer $N \in I$ such that $v_n = v_N$ for all $n \in I$ with $n \geq N$. The **length** of the sequence is then the smallest such integer N . A configuration v is said to be **finite well-stopped** if $v = \bigcup_{n \in I} v_n$, with $(v_n)_{n \in I}$ a well-stopped sequence of finite length.

IV-1.3 Notations. We denote by \mathcal{X} the set of well-stopped configurations, and we denote by \mathcal{X}_0 the set of finite well-stopped configurations. For $v \in \mathcal{W}$, we denote by \mathcal{X}^v the set of well-stopped configurations of \mathcal{E}^v , and by \mathcal{X}_0^v the set of finite well-stopped configurations of \mathcal{E}^v .

IV-1.4 We can always assume that a well-stopped sequence is infinite, by completing the sequence by $z_n = \emptyset$ for n large if it is of finite length—otherwise nothing has to be done—, which still gives a well-stopped sequence and the same union.

IV-1.5 Lemma. *Assume that \mathcal{E} is a locally finite event structure with finite concurrent width. Then every maximal configuration is a well-stopped configuration, i.e. $\Omega \subseteq \mathcal{X}$.*

Proof— Let $(v_n, z_n)_{n \geq 1}$ be the normal decomposition of an element $\omega \in \Omega$ (III-3.1). For every integer $n \geq 1$, z_n is maximal in $B^\perp(\mathcal{E}^{v_{n-1}})$, and $B^\perp(\mathcal{E}^{v_{n-1}})$ is finite by Proposition III-2.3, as \mathcal{E} has finite concurrent width. Thus z_n is finitely stopped in $\mathcal{E}^{v_{n-1}}$. \square

IV-1.6 Lemma. (Concatenation) *If $v \in \mathcal{X}_0$ and if $y \in \mathcal{X}_0^v$, then $v \oplus y \in \mathcal{X}_0$.*

Proof— Let $(v_n^1, z_n^1)_{n \geq 1}$ be a well-stopped sequence of \mathcal{E} , of finite length N , and leading to $v = v_N$. Let $(v_n^2, z_n^2)_{n \geq 1}$ be a well-stopped sequence of \mathcal{E}^v . Then the concatenation of $(v_n^1, z_n^1)_{1 \leq n \leq N}$ and of $(v_{n-N}^2, z_{n-N}^2)_{n \geq N+1}$ is a well-stopped sequence of \mathcal{E} , by Definition IV-1.1. \square

IV-1.7 Lemma. (Extraction of well-stopped sequences). *Let $(v_k, z_k)_{k \geq 1}$ be a well-stopped sequence of \mathcal{E} . Let n be an integer, and for each $k \geq 1$ let:*

$$v'_k = v_{n+k} \ominus v_n, \quad z'_k = z_{n+k}.$$

Then $(v'_k, z'_k)_{k \geq 1}$ is a well-stopped sequence of \mathcal{E}^{v_n} .

Proof— Let $\mathcal{T} = \mathcal{E}^{v_n}$. We have:

$$v'_{k+1} \ominus v'_k = (v_{n+k+1} \setminus v_{n+k}) \setminus v_n = z_{n+k+1} \setminus v_n.$$

We have $z_{n+k+1} \setminus v_n = z_{n+k+1}$ since $z_{n+k+1} \subseteq \mathcal{E}^{v_n}$, and thus $v'_{k+1} \ominus v'_k = z'_k$. Moreover we have for each $k \geq 1$, using Lemma II-2.3:

$$\mathcal{E}^{v_{n+k}} = \mathcal{E}^{v_n \oplus v'_k} = \mathcal{T}^{v'_k}.$$

Thus $z'_k = z_{n+k}$ is finitely stopped in $\mathcal{T}^{v'_{k-1}}$. \square

IV-2 Well-Stopped Configurations in a Stopping Prefix.

As stopped configurations, well-stopped ones are intrinsic *w.r.t.* stopping prefixes, in the following sense. The lemma is illustrated by Figure 3.13.




Figure 3.13: Lemma IV-2.1

IV-2.1 Lemma. Let B be a stopping prefix of \mathcal{E} . A configuration v of B is well-stopped in B if and only if v is well-stopped in \mathcal{E} . The decompositions are obtained one from the other by:

$$(v_n, z_n)_n \rightarrow (v_n \cap B, z_n \cap B)_n, \quad \text{and} \quad (v_n, z_n)_n \rightarrow (v_n, z_n)_n .$$

Proof— This is a consequence of the following lemma. \square

IV-2.2 Lemma. Let B be a stopping prefix of \mathcal{E} and let v be a configuration of B . Then we have:

- (a) D stopping prefix of $B^v \Rightarrow D$ stopping prefix of \mathcal{E}^v .
- (b) D stopping prefix of $\mathcal{E}^v \Rightarrow D \cap B$ stopping prefix of B^v .

Proof— According to Lemma II-5.2 we have: $B^v = \mathcal{E}^v \cap B$, and thus B^v is a stopping prefix of \mathcal{E}^v since B is a stopping prefix of \mathcal{E} (II-4.2). Now for (a), if D is a stopping prefix of B^v , then D is a stopping prefix of \mathcal{E}^v by Lemma I-3.1 applied in the event structure \mathcal{E}^v . For (b), if D is a stopping prefix of \mathcal{E}^v , we also have by Lemma I-3.1 applied in the event structure \mathcal{E}^v , that $D \cap B$ is a stopping prefix of B^v . \square

IV-3 Concurrent Decomposition of the Initial Stopping Prefix.

The following simple lemma is a key result for our study.

IV-3.1 Lemma. Two distinct initial branching cells are disjoint and concurrent.

Proof— Since branching cells are by definition minimal non void stopping prefixes, and since stopping prefixes are stable under intersection, distinct branching cells are

disjoint. That distinct branching cells are concurrent follows then from Lemma I-2.6. \square

The following result establishes the relation between initial branching cells, *i.e.* elements of $\Delta^\perp(\mathcal{E})$, and stopping prefixes of the initial stopping prefix $B^\perp(\mathcal{E})$. The proof is straightforward.

IV-3.2 Lemma. *A subset B of $B^\perp(\mathcal{E})$ is a stopping prefix of \mathcal{E} if and only if B is a union $B = \bigcup_{\lambda \in I} \lambda$, where I is a subset of $\Delta^\perp(\mathcal{E})$.*

IV-3.3 Definition. (*Initial λ -Germs, stopped-initial configurations*) Each initial branching cell λ is a sub-event structure of \mathcal{E} , and $\mathcal{W}_\lambda, \Omega_\lambda$ denote respectively the set of configurations and the set of maximal configurations of λ . Any element $z_\lambda \in \Omega_\lambda$ is called an **initial germ** of \mathcal{E} , or a **λ -germ**. We say that a configuration v is **stopped-initial** if v is stopped in the initial prefix $B^\perp(\mathcal{E})$.

IV-3.4 Product Space Associated with the Initial Decomposition of \mathcal{E} . With Lemma IV-3.1, the geometric decomposition $B^\perp(\mathcal{E}) = \bigcup_{\lambda \in \Delta^\perp(\mathcal{E})} \lambda$ implies a decomposition of $\Omega_{B^\perp(\mathcal{E})}$ into independent components associated with the branching cells. We set the following product spaces:

$$K^\perp(\mathcal{E}) = \prod_{\lambda \in \Delta^\perp(\mathcal{E})} \mathcal{W}_\lambda, \quad \Pi^\perp(\mathcal{E}) = \prod_{\lambda \in \Delta^\perp(\mathcal{E})} \Omega_\lambda. \quad (3.11)$$

For every $\lambda \in \Delta^\perp(\mathcal{E})$, we have the mapping $\pi_\lambda : \mathcal{W} \rightarrow \mathcal{W}_\lambda, w \mapsto w \cap \lambda$. Let $\mathcal{W} \rightarrow \Pi^\perp(\mathcal{E})$ be the direct product of the family $(\pi_\lambda)_{\lambda \in \Delta^\perp(\mathcal{E})}$. We thus have the following commutative diagrams, with a functorial injection from nodes at right hand to nodes at left hand.

$$\begin{array}{ccc} \mathcal{W} & & \Omega \\ \downarrow & \searrow & \downarrow \\ K^\perp(\mathcal{E}) & \longleftarrow & \mathcal{W}_{B^\perp(\mathcal{E})} \end{array} \quad \begin{array}{ccc} \Omega & & \Omega_{B^\perp(\mathcal{E})} \\ \downarrow & \searrow & \downarrow \\ \Pi^\perp(\mathcal{E}) & \longleftarrow & \Omega_{B^\perp(\mathcal{E})} \end{array} \quad (3.12)$$

We have seen that, since we assume that \mathcal{E} has finite concurrent width, $\Delta^\perp(\mathcal{E})$ is a finite collection of sets, and these sets are finite since \mathcal{E} is locally finite (Prop. III-2.3). Hence $K^\perp(\mathcal{E})$ and $\Pi^\perp(\mathcal{E})$ are finite. The following result is a direct consequence of Lemma IV-3.2.

IV-3.5 Lemma. *In Diagram (3.12), the mappings $\mathcal{W}_{B^\perp(\mathcal{E})} \rightarrow K^\perp(\mathcal{E})$ and $\Omega_{B^\perp(\mathcal{E})} \rightarrow \Pi^\perp(\mathcal{E})$ are two bijections.*

In particular, a configuration v is stopped-initial if and only if v can be written as a union:

$$v = \bigcup_{\lambda \in I} v_\lambda, \quad v_\lambda \in \Omega_\lambda,$$

where I is a subset of $\Delta^\perp(\mathcal{E})$. The decomposition is unique. v is full-initial if and only if $I = \Delta^\perp(\mathcal{E})$.

IV-4 Well-Stopped Configurations of the Initial Prefix.

We can now precisely describe the well-stopped configurations of the initial stopping prefix.

IV-4.1 Lemma. *If v is a stopped configuration of \mathcal{E} , then $\mathcal{E}^v \cap B(v) = \emptyset$.*

Proof— Assume that $B(v) \cap \mathcal{E}^v$ contains an event x . Then x is compatible with v . We have seen in I-3.3 that v is maximal in $B(v)$. It implies that $x \in v$, and contradicts that $v \cap \mathcal{E}^v = \emptyset$. \square

IV-4.2 Lemma. *Assume that \mathcal{E} satisfies $B^\perp(\mathcal{E}) = \mathcal{E}$, and let $(\emptyset, v_1, \dots, v_n)$ be a well-stopped sequence of \mathcal{E} . Then $(\emptyset, v_2 \ominus v_1, \dots, v_n \ominus v_1)$ is also a well-stopped sequence of \mathcal{E} .*

Proof— Let $B_1 = \mathcal{E} \setminus B(v_1)$, where $B(v_1)$ denotes the smallest stopping prefix that contains v_1 . We show that $B_1 = \mathcal{E}^{v_1}$. As v_1 is stopped, we have $\mathcal{E}^{v_1} \cap B(v_1) = \emptyset$ by Lemma IV-4.1, hence $\mathcal{E}^{v_1} \subseteq B_1$. By Lemma IV-3.2, and since $B^\perp(\mathcal{E}) = \mathcal{E}$, B_1 itself is a stopping prefix of \mathcal{E} , union of initial branching cells $\lambda \in \Delta^\perp(\mathcal{E})$ such that $\lambda \cap B(v_1) = \emptyset$. By Lemma IV-3.1 every event $e \in B_1$ is compatible with v_1 , and $e \notin v_1$, thus $e \in \mathcal{E}^{v_1}$. Finally $\mathcal{E}^{v_1} = B_1$. For each $n \geq 1$ let $z_n = v_n \ominus v_{n-1}$.

As $(\emptyset, v_1, \dots, v_n)$ is well-stopped in \mathcal{E} , it follows from Lemma IV-1.7 that $(\emptyset, v_2 \ominus v_1, \dots, v_n \ominus v_1)$ is well-stopped in $\mathcal{E}^{v_1} = B_1$. Since B_1 is a stopping prefix of \mathcal{E} , $(\emptyset, v_2 \ominus v_1, \dots, v_n \ominus v_1)$ is well-stopped in \mathcal{E} (Lemma IV-2.1). \square

IV-4.3 Proposition. *Assume that \mathcal{E} is locally finite and of finite concurrent width, and let v be a configuration of $B^\perp(\mathcal{E})$. Then v is well-stopped if and only if v is stopped.*

Proof— Let $B = B^\perp(\mathcal{E})$. B is finite according to Proposition III-2.3. Therefore $\Omega_B \subseteq \mathcal{X}_0$.

Conversely, by Lemma IV-2.1, we can assume without loss of generality that $\mathcal{E} = B$. Let $(v_n, z_n)_n$ be a well-stopped sequence of finite length of \mathcal{E} , leading to $v = v_N$. Using Lemma IV-4.2, we show by induction that z_k is stopped in \mathcal{E} for all k . It is easy to check that a union of compatible finitely stopped configurations, is a finitely stopped configuration. Thus $v = \bigcup_k z_k$ is finitely stopped in \mathcal{E} . \square

IV-5 Step-by-Step Decomposition.

We introduce a characterisation of well-stopped configurations, that can be taken as an equivalent definition.

IV-5.1 Definition. (*Step-by-step decomposition*) Let $(v_n, z_n)_{n \in I}$ be a well-stopped sequence of \mathcal{E} (IV-1.1). We say that the sequence is **step-by-step** if for every $n \geq 1$, z_n is stopped-initial (IV-3.3) in $\mathcal{E}^{v_{n-1}}$. We say that an element $v \in \mathcal{X}$ admits a step-by-step decomposition if there is a step-by-step sequence $(v_n, z_n)_{n \in I}$ such that $v = \bigcup_{n \in I} v_n$.

IV-5.2 Proposition. *In a locally finite event structure of finite concurrent width, every well-stopped configuration admits a step-by-step decomposition.*

Proof – We observe first that every stopped configuration admits a finite step-by-step decomposition. Indeed, if $v \in \Omega_B$ with $B \in \mathcal{S}_0$, v admits by Lemma IV-1.5 a normal form in B , which is step-by-step in \mathcal{E} by Lemma IV-2.1. The decomposition has finite length since B is finite and thus v is finite.

Now let $(v_n, z_n)_n$ be a well-stopped sequence decomposing v . Consider a finite step-by-step decomposition of each z_k in $\mathcal{E}^{v_{k-1}}$. Then the concatenation of these decomposition gives a step-by-step decomposition of v . \square

V—Branching Cells

Branching cells generalise the initial branching cells introduced in III-2.1. The idea is the following: what happens *just after* a finite configuration v can be seen in the initial stopping prefix of the cone of future \mathcal{E}^v . Hence the study of the initial stopping prefix applies in general. In particular, *branching cells* generalise the initial branching cells.

In V-1, we present the definition and the composition properties of branching cells. We expose combinatorial lemmas in V-2. The goal is to show Proposition V-2.4. Some other lemmas are summed up in V-3 into a lemma that appears as an *exchange lemma*.

V-1 Branching Cells.

We generalise the notion of initial branching cell introduced in Definition III-2.1.

V-1.1 Notation. For v a configuration of \mathcal{E} , $\Delta_{\mathcal{E}}^+(v)$, or $\Delta^+(v)$ for short, denotes the set of initial branching cells of the cone of future \mathcal{E}^v :

$$\Delta_{\mathcal{E}}^+(v) = \Delta^+(\mathcal{E}^v) .$$

We have in particular: $\Delta^+(\mathcal{E}) = \Delta_{\mathcal{E}}^+(\emptyset)$.

V-1.2 Definition. (Branching cells) Any element $\lambda \in \Delta^+(v)$, with $v \in \mathcal{X}_0$, is called a **branching cell** of \mathcal{E} . The elements $\lambda \in \Delta^+(v)$ are called branching cells at v . We denote by $\Lambda_{\mathcal{E}}$ the collection of branching cells of \mathcal{E} .

Branching cells are illustrated by Figure 3.14.




Figure 3.14: *Branching cells in $\Delta^+(v)$.*

V-1.3 Lemma. (Concatenation and restriction to a stopping prefix) For every $B \in \mathcal{S}$ we have:

$$\Lambda_B = \{\lambda \in \Lambda_{\mathcal{E}} \mid \lambda \cap B \neq \emptyset\} = \{\lambda \in \Lambda_{\mathcal{E}} \mid \lambda \subseteq B\} .$$

For every $v \in \mathcal{X}_0$ and $y \in \mathcal{X}_0^v$, we have:

$$\Delta_{\mathcal{E}^v}^+(y) = \Delta_{\mathcal{E}}^+(v \oplus y) .$$

In particular, $\Lambda_{\mathcal{E}^v} \subseteq \Lambda_{\mathcal{E}}$ if $v \in \mathcal{X}_0$.

Proof – The second part comes from II-2.3: $(\mathcal{E}^v)^y = \mathcal{E}^{v \oplus y}$. For the first part, let $B \in \mathcal{S}$, and $\lambda \in \Lambda_B$. Then there is a configuration v , well-stopped in B , such that $\lambda \in \Delta^+(v)$. By IV-2.1, v is well-stopped in \mathcal{E} , and λ is a stopping prefix of \mathcal{E}_v by IV-2.2 (a). It is a minimal non empty stopping prefix of \mathcal{E}^v by IV-2.2 (b), so we have $\lambda \in \Delta^+(v)$ and thus $\lambda \in \Lambda_{\mathcal{E}}$ and $\lambda \subseteq B$. With the same argument we show that any $\lambda \in \Lambda_{\mathcal{E}}$ with $\lambda \subseteq B$ is a branching cell of B , so we get:

$$\Lambda_B = \{\lambda \in \Lambda_{\mathcal{E}} \mid \lambda \subseteq B\} .$$

It follows from II-4.2 that, if $\lambda \in \Lambda_{\mathcal{E}}$ satisfies $\lambda \cap B \neq \emptyset$, then $\lambda \subseteq B$. This completes the proof. \square

V-2 Auxiliary Results.

Recall that stopped-initial configurations are stopped in the initial stopping prefix (IV-3.3).

V-2.1 Lemma. *Let $v \in \mathcal{W}$, and let A be a subset of \mathcal{E}^v . Assume that A is a prefix of \mathcal{E} . Then $A \parallel v$.*

Proof— We have to show that $x \parallel y$ for every pair $(x, y) \in A \times v$. Let $(x, y) \in A \times v$. As A is a prefix of \mathcal{E} , we have $x \succeq y \Rightarrow y \in A$, and $y \in A$ contradicts $v \cap \mathcal{E}^v = \emptyset$, thus we have $\neg(x \succeq y)$. As v is a prefix of \mathcal{E} , we have $x \preceq y \Rightarrow x \in v$, and $x \in v$ also contradicts $v \cap \mathcal{E}^v = \emptyset$, and thus $\neg(x \preceq y)$ holds. As $x \in \mathcal{E}^v$, x is compatible with v , thus $\neg(x \# y)$ holds, and finally $x \parallel y$. \square

V-2.2 Lemma. *Let v be a configuration of \mathcal{E} , and let $B \in \mathcal{S}$. Then we have: $B \cap v = \emptyset \Rightarrow B \subseteq \mathcal{E}^v$.*

Remark that the result does not hold in general if B is not a stopping prefix. Take for instance two minimal events in conflict $e_1 \# e_2$. Then $e_1 \cap e_2 = \emptyset$ but $e_1 \notin \mathcal{E}^{e_2}$.

Proof— As B and v are two prefixes of \mathcal{E} that do not intersect, no pair of events $(e, e') \in B \times v$ are causally related. If $e \# e'$, then by Lemma I-1.2 there are events x, y with $x \preceq e$ and $y \in v$ with $x \#_d y$. Then $x \in B$ as B is a prefix, and thus $y \in B$ as B is $\#_d$ -closed, contradicting $v \cap B = \emptyset$. Therefore we have $B \parallel v$ and $B \cap v = \emptyset$, which implies $B \subseteq \mathcal{E}^v$. \square

V-2.3 Lemma. *Let v be a stopped-initial configuration of \mathcal{E} , and let λ be an initial branching cell of \mathcal{E} . Then $v \cap \lambda = \emptyset \Rightarrow \lambda \in \Delta^+(v)$.*

Proof— According to Lemma V-2.2, $\lambda \subseteq \mathcal{E}^v$. Hence, since λ is a stopping prefix of \mathcal{E} , λ is a non empty stopping prefix of \mathcal{E}^v . It remains thus to show that λ is minimal in \mathcal{E}^v .

For this let γ be a stopping prefix of \mathcal{E}^v that contains an event in $\gamma \cap \lambda$. We set the following subset of \mathcal{E} :

$$B = \gamma \cap \lambda = (v \cup \gamma) \cap \lambda,$$

and it is clear on the latter that B is a prefix of \mathcal{E} . We show that B is $\#_d$ -closed in \mathcal{E} . Let $e \in B$, and let $f \in \mathcal{E}$ such that $e \#_d f$. Then $f \in \lambda$ since λ is $\#_d$ -closed in \mathcal{E} . Since $\lambda \subseteq \mathcal{E}^v$, it follows that $f \in \mathcal{E}^v$. γ is $\#_d^v$ -closed, which implies that $f \in \gamma$, and finally $f \in \lambda \cap \gamma$. Hence B is a stopping prefix of \mathcal{E} that intersects λ , thus $B \supseteq \lambda$. It follows that $\gamma \supseteq \lambda$. This shows that λ is an initial branching cell of \mathcal{E}^v . \square

V-2.4 Proposition. Let $\lambda \in \Delta^\perp(\mathcal{E})$, and let v be a well-stopped configuration of \mathcal{E} . Then $v \cap \lambda = \emptyset$ or $v \cap \lambda \in \Omega_\lambda$.

Proof— $v \cap \lambda$ is a well-stopped configuration of λ . Since λ is minimal non void, it is then obvious that $v \cap \lambda = \emptyset$ or $v \cap \lambda$ is maximal in λ . \square

V-3 An Exchange Lemma.

We now show a result that will be useful to conclude the study of well-stopped configurations. We use the notion of initial germ (IV-3.3).

V-3.1 Lemma. Let $v_0 \in \mathcal{X}_0$, and let x be an initial germ of \mathcal{E} compatible with v_0 . Then $v_0 \cap x \neq \emptyset \Rightarrow x \subseteq v_0$.

Proof— Let $\lambda \in \Delta^\perp(\mathcal{E})$ such that $x \in \Omega_\lambda$, and assume that $v_0 \cap x \neq \emptyset$. By V-2.4, $v_0 \cap \lambda \in \Omega_\lambda$. x and $v_0 \cap \lambda$ are two maximal and compatible configurations of λ , so they are equal, and thus $x \subseteq v_0$. \square

V-3.2 Lemma. Let x be an initial germ of \mathcal{E} , let z be stopped-initial in \mathcal{E} , and assume that $x \subseteq z$. Then $z \ominus x$ is stopped-initial in \mathcal{E}^x .

Proof— By Lemma IV-3.5, z decomposes itself as the following disjoint union:

$$z = x \sqcup z_{\lambda_1} \sqcup \dots \sqcup z_{\lambda_p}, \quad \lambda_i \in \Delta^\perp(\mathcal{E}),$$

with z_{λ_i} a λ_i -germ for all i . Thus $z \ominus x = z_{\lambda_1} \sqcup \dots \sqcup z_{\lambda_p}$. For each i , $x \cap \lambda_i = \emptyset$, and so $\lambda_i \in \Delta^+(x)$ by V-2.3. Thus, each z_{λ_i} is an initial germ of \mathcal{E}^x , and $z \ominus x$ is a union of compatible initial germs of \mathcal{E}^x . Applying IV-3.5 in the event structure \mathcal{E}^x , $z \ominus x$ is stopped-initial in \mathcal{E}^x . \square

V-3.3 Lemma. (First Exchange Lemma) Let x be an initial germ of \mathcal{E} , let $v_0 \in \mathcal{X}_0$, and let $z \in \mathcal{W}^{v_0}$. Assume that $v_0 \oplus z$ is compatible with x . We set the configuration $v = v_0 \cup x$, and we set $v' \in \mathcal{W}$ and $z' \in \mathcal{W}^v$ such that:


$$v' = (v_0 \oplus z) \cup x = (v_0 \cup x) \oplus z'. \quad (3.13)$$

Then z is stopped-initial in $\mathcal{E}^{v_0} \Rightarrow z'$ is stopped-initial in \mathcal{E}^v .

Figure 3.15 illustrates Lemma V-3.3.

Proof— It is enough to show the implication: z is an initial germ of $\mathcal{E}^{v_0} \Rightarrow z'$ is stopped-initial in \mathcal{E}^v , since v does not depend on z , and a union of compatible stopped-initial configurations is stopped-initial (IV-3.5). So we assume that z is an initial germ of \mathcal{E}^{v_0} .

First case: $v_0 \cap x \neq \emptyset$. Then, by V-3.1, $x \subseteq v_0$. Thus $v = v_0$ and $z' = z$. The result is trivial.

Figure 3.15: *The exchange lemma V-3.3.*

Second case: $v_0 \cap x = \emptyset$. Then it is a consequence of V-2.3 that x is an initial germ of \mathcal{E}^{v_0} .

- a) $z \cap x \neq \emptyset$. Then we apply V-3.1 in the event structure \mathcal{E}^{v_0} , to get that $x \subseteq z$. Since x and z_0 are two initial germs of \mathcal{E}^{v_0} , it implies $z = x$ and $z' = \emptyset$.
- b) $z \cap x = \emptyset$. Let $\lambda \in \Delta^\perp(\mathcal{E})$ such that $x \in \Omega_\lambda$. By V-2.4, if $(v_0 \oplus z) \cap \lambda \neq \emptyset$, then $(v_0 \oplus z) \cap \lambda$ is maximal in λ ; since it is compatible with x , it coincides with x , contradicting $(v_0 \oplus z) \cap x = \emptyset$. Thus we have $(v_0 \oplus z) \cap \lambda = \emptyset$, and by V-2.2 applied to λ , it implies $\lambda \subseteq \mathcal{E}^{v_0}$. In particular $x \subseteq \mathcal{E}^{v_0}$. Then (3.13) becomes:

$$(v_0 \oplus z) \oplus x = (v_0 \oplus x) \oplus z' ,$$

and thus $z' = z$ since the \oplus are disjoint unions. Let $\gamma \in \Delta^+(v_0)$ such that $z \in \Omega_\gamma$. Since we have seen that x is an initial germ of \mathcal{E}^{v_0} , $x \cap \gamma = \emptyset$. Applying V-2.3 in \mathcal{E}^{v_0} , we get:

$$\gamma \in \Delta_{\mathcal{E}^{v_0}}^+(x) = \Delta^\perp(\mathcal{E}^v) ,$$

the later by V-1.3. Since $z' \in \Omega_\gamma$, we have the requested result.

□

VI—The Dynamic Puzzle.

This section is intended to take advantage of the previous technical results, to state more general and practical statements. We introduce the partition of a well-stopped configuration through its germs. The associated branching cells form the *dynamic puzzle* around the configuration.

In VI-1, we introduce another characterisation of well-stopped configurations, through the so-called *germ-decomposition*. Then VI-2 is devoted to the result that compatible well-stopped configurations form a lattice. The dynamic puzzle associated to well-stopped configurations is the topic of VI-3.

VI-1 Germ-Decompositions.

Recall that an initial germ is maximal in an initial branching cell. We extend the definition as follows.

VI-1.1 Definition. (*Germ decomposition, germs*) A well-stopped sequence $(v_n, z_n)_{n \geq 1}$ is said to be a **germ-decomposition** if z_n is an initial germ of $\mathcal{E}^{v_{n-1}}$ for all integer $n \geq 1$. If the germ decomposition has finite length N , it is said to lead to v_N . We say that a well-stopped configuration v admits a germ-decomposition if there is a germ-decomposition leading to v . Any element z_n is called a **germ** of v .

Clearly, the concatenation (IV-1.6) of finite germ-decompositions is a germ-decomposition.

VI-1.2 Lemma. *Assume that \mathcal{E} is locally finite and has finite concurrent width. Then every finite well-stopped configuration of \mathcal{E} admits a germ-decomposition.*

Proof— Let $v \in \mathcal{X}_0$, and let $(v_n, z_n)_n$ be a step-by-step decomposition of v (IV-5.2). If each z_n admits a finite germ-decomposition in $\mathcal{E}^{v_{n-1}}$, then the concatenation of their germ-decomposition is a germ-decomposition of v . As $\mathcal{E}^{v_{n-1}}$ is locally finite and has finite concurrent width, it is enough to show:

Every v stopped-initial in \mathcal{E} admits a finite germ-decomposition.

For this, $\Delta^\perp(\mathcal{E})$ is finite as \mathcal{E} as finite concurrent width (III-2.3). As v is stopped-initial, v has the following decomposition into disjoint unions:

$$v = z_{\lambda_1} \sqcup \dots \sqcup z_{\lambda_p}, \quad z_{\lambda_i} \in \Omega_{\lambda_i}.$$

Let $v_0 = \emptyset$ and $v_k = v_{k-1} \cup z_k$ for $1 \leq k \leq p$. Clearly, all v_k are stopped-initial in \mathcal{E} , as unions of compatible germs (IV-3.5). For every k , $v_{k-1} \cap \lambda_k = \emptyset$ since v_{k-1} is stopped and z_k is maximal in λ_k . Applying V-2.3, we get that $\lambda_k \in \Delta^+(v_{k-1})$. It implies that $(v_k, z_{\lambda_k})_{1 \leq k \leq p}$ is a germ-decomposition, leading to $v_p = v$. \square

VI-2 The Lattice of Well-Stopped Sub-Configurations.

We begin by stating a result improving our first exchange lemma V-3.3.

VI-2.1 Lemma. (*Exchange Lemma*) *Let v, v' be two well-stopped configurations of \mathcal{E} , and assume that v and v' are compatible. Then $(v \cup v') \ominus v'$ is well-stopped in $\mathcal{E}^{v'}$.*

Proof— We assume without loss of generality that v and v' have decompositions of finite length. We assume first that $v' = x$ is an initial germ of \mathcal{E} . Let $(v_n, z_n)_n$ be a step-by-step decomposition of v . We set:

$$w_k = v_k \cup x, \quad z'_k = w_k \ominus w_{k-1}.$$

Then we have:

$$w_k = (v_{k-1} \oplus z_k) \cup x = (v_{k-1} \cup x) \oplus z'_k .$$

By the exchange lemma V-3.3, z'_k is stopped-initial in $\mathcal{E}^{w_{k-1}}$ since z_k is stopped-initial in $\mathcal{E}^{v_{k-1}}$. We have shown that $(w_k, z'_k)_k$ is a well-stopped sequence of \mathcal{E} , leading to $w_n = v \cup x$. Now let $y_k = w_k \ominus x$. Then we have $y_k = y_{k-1} \oplus z'_k$, and z'_k is stopped-initial in $\mathcal{E}^{v_{k-1}} = (\mathcal{E}^x)^{y_{k-1}}$. As $(y_k, z'_k)_k$ leads to $(v \cup x) \ominus x$, we have shown the result if $v' = x$ is an initial germ of \mathcal{E} .

For the general case, let $(v'_n)_{n \geq 1}$ be a germ-decomposition of v' . Such a decomposition is given by Lemma VI-1.2. By induction on n , it follows from the previous case that $(v \cup v'_n) \ominus v'_n$ is well-stopped in $\mathcal{E}^{v'_n}$. The result of the lemma follows, by considering an integer n large enough. \square

VI-2.2 Corollary. *Let $v \in \mathcal{X}_0$ and $v' \in \mathcal{X}$. If $v \subseteq v'$ then $v' \ominus v$ is well-stopped in \mathcal{E}^v .*

Proof— We can assume without generality that $v' \in \mathcal{X}_0$. Then we apply VI-2.1 with $v' \cup v = v'$. \square

VI-2.3 Theorem. *Let $w \in \mathcal{X}$. The set of well-stopped sub-configurations of w forms a lattice.*

Proof— It is enough to show that the following subset is a lattice:

$$H = \{v \in \mathcal{X}_0 \mid v \subseteq w\} .$$

Let $v, v' \in H$, we show that $v \cup v' \in H$. By VI-2.1, $v \cup v'$ is finitely well-stopped in v' . By concatenation (IV-1.6), $v \cup v'$ is well-stopped in \mathcal{E} .

Now we show that $y = v \cap v' \in H$. We set:

$$y_0 = \emptyset, \quad z_1 = y \cap B^\perp(\mathcal{E}) .$$

Observe that $z_1 = \emptyset$ if and only if $y = \emptyset = y_0$, since each prefix of \mathcal{E} encounters an initial branching cell (III-1.2). For any initial branching cell λ , $v \cap \lambda$ and $v' \cap \lambda$ are two compatible configurations, and each one is maximal in λ if non empty (V-2.4). Thus z_1 has the property:

$$\forall \lambda \in \Delta^\perp(\mathcal{E}), \quad z_1 \cap \lambda \neq \emptyset \Rightarrow z_1 \cap \lambda \in \Omega_\lambda .$$

It follows that z_1 is stopped-initial in \mathcal{E} . According to VI-2.1, it implies that $v \ominus z_1$ and $v' \ominus z_1$ are well-stopped in \mathcal{E}^{z_1} . We repeat inductively this construction and get a sequence $(y_n, z_n)_{n \geq 1}$ such that:

- i. $(y_n, z_n)_{n \geq 1}$ is a step-by-step well-stopped sequence of \mathcal{E} ,
- ii. For all $n \geq 1$, $y_n \subseteq y$, and $v \ominus y_n, v' \ominus y_n$ are well-stopped in \mathcal{E}^{y_n} ,
- iii. For all $n \geq 1$, $z_n = \emptyset$ if and only if $y = y_{n-1}$.

As y is finite, there is an integer n such that $z_n = \emptyset$, and then $y = y_{n-1}$, which shows that y is well-stopped, *i.e.* $y \in H$. \square

VI-3 Dynamic Puzzle.

Consider a lattice of compatible well-stopped configurations. Then consider the elementary increments between the well-stopped configurations: they form a collection of germs. These germs are maximal configurations of a collection of branching cells. We show that these branching cells do not intersect, hence they form what we call a puzzle. However, and as a typical effect of concurrency, the whole collection of branching cells defined by an event structure may intersect. Therefore we say that the puzzles drawn by well-stopped configurations are dynamic.

VI-3.1 Dynamic puzzle. For v a well-stopped configuration of \mathcal{E} , we define the following collection of branching cells of \mathcal{E} :

$$\bar{\Lambda}(v) = \{\lambda \in \Delta^+(y) \mid y \subseteq v, y \in \mathcal{X}_0\},$$

If v is finite and well-stopped, we define the **dynamic puzzle around** v by the collection of branching cells:

$$\Lambda(v) = \bar{\Lambda}(v) \setminus \Delta^+(v). \tag{3.14}$$

We note $\Lambda_{\mathcal{E}}(v)$ and $\bar{\Lambda}_{\mathcal{E}}(v)$ to precise the event structure \mathcal{E} .

VI-3.2 Remark. If v is finite and maximal, then $\mathcal{E}^v = \emptyset$ thus $\Delta^+(\mathcal{E}^v) = \emptyset$, and $\Lambda(v) = \bar{\Lambda}(v)$. The “boundary effect” is cancelled.

Figure 3.16 illustrates the dynamic puzzle.




Figure 3.16: *The dynamic puzzle.*

VI-3.3 Lemma. For any $v \in \mathcal{X}_0$, we have:

$$\Lambda(v) = \{\lambda \in \bar{\Lambda}(v) : v \cap \lambda \neq \emptyset\} = \{\lambda \in \bar{\Lambda}(v) : v \cap \lambda \in \Omega_\lambda\}. \tag{3.15}$$

Proof – Let $\lambda \in \overline{\Lambda}(v)$: there is a $y \in \mathcal{X}_0$ with $y \subseteq v$ and $\lambda \in \Delta^+(y)$. We set $s = v \ominus y$, and we have $s \in \mathcal{X}_0^y$ by VI-2.2. Since $\lambda \in \Delta^+(\mathcal{E}^y)$, we have by V-2.4:

$$s \cap \lambda \neq \emptyset \Leftrightarrow s \cap \lambda \in \Omega_\lambda .$$

Since $v \cap \lambda = s \cap \lambda$, the equality between the two right sets in (3.15) follows. For every $\lambda \in \overline{\Lambda}(v)$, we have $\lambda \in \Delta^+(v) \Rightarrow \lambda \cap v = \emptyset$, which shows the following inclusion:

$$\{\lambda \in \overline{\Lambda}(v) : \lambda \cap v \neq \emptyset\} \subseteq \Lambda(v) .$$

It remains to show the converse inclusion. Let $y \in \mathcal{X}_0$ with $y \subseteq v$, and let $\lambda \in \Delta^+(y)$. It is enough to show:

$$\lambda \cap v = \emptyset \Rightarrow \lambda \in \Delta^+(v) .$$

For this, let $s = v \ominus y \in \mathcal{X}_0^y$. We have $\lambda \in \Delta^+(\mathcal{E}^y)$, hence by applying V-2.3:

$$\lambda \cap v = \emptyset \Rightarrow \lambda \cap s = \emptyset \Rightarrow \lambda \in \Delta_{\mathcal{E}^y}^+(s) = \Delta_{\mathcal{E}}^+(y \oplus s) = \Delta^+(v) .$$

□

VI-3.4 Lemma. *Let $v, v' \in \mathcal{X}_0$, and assume that v, v' are compatible. Let $\lambda \in \Delta^+(v)$ and $\lambda' \in \Delta^+(v')$. Then $\lambda \cap \lambda' \neq \emptyset \Rightarrow \lambda = \lambda'$.*

Proof – We assume that $\lambda \cap \lambda' \neq \emptyset$, and we show that $\lambda = \lambda'$. Assume first that $v \cap v' = \emptyset$. Then we have:

$$v' = (v \cup v') \ominus v . \quad (3.16)$$

According to VI-2.1, it implies that v' is well-stopped in \mathcal{E}^v . (3.16) also implies that $\lambda' \in \Delta_{\mathcal{E}}^+(v') = \Delta_{\mathcal{E}^v}^+(v')$. Then $\lambda \cap \mathcal{E}^{v'}$ is a stopping prefix of $(\mathcal{E}^v)^{v'}$, that intercepts λ' by hypothesis. Since λ' is minimal, it implies that $\lambda' \subseteq \lambda$. Exchanging the role of v and v' leads to the converse inclusion, hence $\lambda = \lambda'$.

For the general case, we set:

$$v_1 = v \ominus (v \cap v'), \quad v_2 = v' \ominus (v \cap v'), \quad u = v \cap v' .$$

Then, by VI-2.3, u is well-stopped, and thus v_1 and v_2 are well-stopped in \mathcal{E}^u by Lemma VI-2.1. We also have: $\lambda \in \Delta_{\mathcal{E}}^+(v) = \Delta_{\mathcal{E}^u}^+(v_1)$, and in the same way $\lambda' \in \Delta_{\mathcal{E}^u}^+(v_2)$. Since $v_1 \cap v_2 = \emptyset$, the problem reduces to the previous case, and we conclude that $\lambda \cap \lambda' \neq \emptyset \Rightarrow \lambda = \lambda'$. □

The following result shows that the branching cells of $\overline{\Lambda}(v)$ are disjoint, hence in particular the branching cells in the dynamic puzzle are disjoint. Moreover branching cells of $\overline{\Lambda}(v)$ can be obtained from any step-by-step decomposition of v . It is thus an *invariant* of these decompositions.

VI-3.5 Theorem. *Let v be a well-stopped configuration of \mathcal{E} . All the branching cells of $\overline{\Lambda}(v)$ are disjoint.*

If $(v_n, z_n)_{n \geq 1}$ is any step-by-step decomposition of v , we have:

$$\overline{\Lambda}(v) = \bigcup_{n \geq 0} \Delta^+(v_n) . \tag{3.17}$$

Proof— The first part of the theorem follows from VI-3.4. To show (3.17), we can assume without loss of generality that $v \in \mathcal{X}_0$. Let $\Gamma = \bigcup_{n \geq 0} \Delta^+(v_n)$, we have by definition that $\Gamma \subseteq \overline{\Lambda}(v)$. To show the converse inclusion, let $\lambda \in \overline{\Lambda}(v)$, and let $v' \subseteq v$ such that $\lambda \in \Delta^+(v')$. We distinguish two cases.

- *First case.* Assume that there is an integer $p \geq 1$ such that $v_p \cap \lambda \neq \emptyset$. Let n be the smallest of such integers. Then λ contains an event e in $B^\perp(\mathcal{E}^{v_{n-1}})$, that belongs to a branching cell $\lambda' \in \Delta^+(v_{n-1})$. Since v' and v are compatible, it follows from VI-3.4 that $\lambda = \lambda'$, and thus $\lambda \in \Gamma$.
- *Second case.* Otherwise $v \cap \lambda = \emptyset$. Let $y = v \ominus v'$. Then y is well-stopped in $\mathcal{E}^{v'}$ (VI-2.2). We have $\lambda \in \Delta^\perp(\mathcal{E}^{v'})$, and $y \cap \lambda = \emptyset$ since $v \cap \lambda = \emptyset$. We apply V-2.3 to get that $\lambda \in \Delta_{\mathcal{E}^{v'}}^+(y) = \Delta^+(v)$. Since v is finite, there is an integer q such that $v = v_q$, and thus $\lambda \in \Gamma$.

□

VI-3.6 Remark. Branching cells of a same configuration do not intersect. However branching cells intersect in general. Figure 3.17 shows the branching cells that decompose the two maximal configurations of a same event structure. The unique branching cell of $\Delta^+(a)$ is framed on the left hand, and the unique branching cell of $\Delta^+(b)$ is framed on the right hand. These distinct branching cells intersect. However branching cells are disjoint for confusion-free event structures, and *a fortiori* for trees (Cf. VII for trees, and Ch. 5 for confusion-free event structures).




Figure 3.17: *Intersecting branching cells. $\Delta^+(a)$ at left, $\Delta^+(b)$ at right.*

The compositionality of dynamic puzzles is established by the following result.

VI-3.7 Proposition. (*Concatenation of dynamic puzzles*) Let $v \in \mathcal{X}_0$, and let $v' \in \mathcal{X}_0^v$. Then we have the following disjoint union:

$$\Lambda_{\mathcal{E}}(v \oplus v') = \Lambda_{\mathcal{E}}(v) \sqcup \Lambda_{\mathcal{E}^v}(v'). \quad (3.18)$$

We have $v \cap \lambda = (v \oplus v') \cap \lambda$ if $\lambda \in \Lambda_{\mathcal{E}}(v)$ and $v' \cap \lambda = (v \oplus v') \cap \lambda$ if $\lambda \in \Lambda_{\mathcal{E}^v}(v')$.

Proof— Assume that the following property is satisfied:

◇ For every $v \in \mathcal{X}_0$ and z a λ_0 -germ of \mathcal{E}^v , $\Lambda_{\mathcal{E}}(v \oplus z) = \Lambda_{\mathcal{E}}(v) \sqcup \{\lambda_0\}$.

Then (3.18) follows by considering a germ decomposition of v' in \mathcal{E}^v (VI-1.1). Now we show ◇. The inclusion \supseteq follows immediately from VI-3.3. For the converse inclusion, let $\lambda \in \Lambda_{\mathcal{E}}(v \oplus z)$. There is a $y \in \mathcal{X}_0$, with: $y \subseteq v \oplus z$ and $\lambda \in \Delta^+(y)$, and we have, according to VI-3.4: $\lambda \cap (v \oplus z) \neq \emptyset$. We distinguish the following cases:

1. $y \cap z \neq \emptyset$. Then $y \subseteq v$. We either have one of the followings:
 - a. $\lambda \cap z \neq \emptyset \Rightarrow \lambda = \lambda_0$, using VI-3.4.
 - b. $\lambda \cap z = \emptyset \Rightarrow \lambda \cap v \neq \emptyset$, and thus $\lambda \in \Lambda_{\mathcal{E}}(v)$, using the characterisation of VI-3.3.
2. $y \cap z = \emptyset$. Then $z \subseteq y$, and we either have one of the followings:
 - a. $\lambda \cap v \neq \emptyset \Rightarrow \lambda \in \Lambda_{\mathcal{E}}(v)$,
 - b. $\lambda \cap v = \emptyset \Rightarrow \lambda \cap z \neq \emptyset \Rightarrow \lambda = \lambda_0$, using VI-3.4.

This shows the \subseteq inclusion, and thus the equality in ◇.

The second assertion of VI-3.7 follows from VI-3.4. □

VII—The Case of Trees of Events

In this section, we examine the different notions introduced (stopping prefixes, well-stopped configurations, germs and branching cells) for trees, *i.e.* for event structures without concurrency. In Chapter 5, a similar study with similar results is done for *confusion-free* event structures.

VII-1 Stopping Prefixes and Cone of Future.

Assume that the event structure $\mathcal{E} = \mathcal{T}$ is a non empty tree of events (Ch. 1, III-4.2). Denote by G the graph relation on \mathcal{T} that connects an event to its successors.

Configurations of \mathcal{T} coincide with paths in (\mathcal{T}, G) . Hence every non void finite configuration v has the form:

$$v = e_1 \oplus \cdots \oplus e_n = [e_n],$$

with e_1 a root of \mathcal{T} and $e_i G e_{i+1}$ for all i .

VII-1.1 Stopping Prefixes. For each event $x \in \mathcal{T}$ we denote by $D(x)$ the immediate successors of x :

$$D(x) = \{y \in \mathcal{T} \mid x G y\}.$$

We also define:

$$F(e) = \{y \in \mathcal{T} \mid x \#_d y\}.$$

Let e be an event of \mathcal{T} , and let $v = [e]$. Then we have: $F(e) = \text{Min}_{\leq}(\mathcal{T}^v)$, the minimal events of \mathcal{T}^v . If $v \neq \emptyset$, v has the form $v = [f]$, and then:

$$F(e) = D(f).$$

Clearly, the smallest stopping prefix $B(e)$ that contains an event e with $[e] = e_1 \oplus \cdots \oplus e_n$, is given by:

$$B(e) = \bigcup_{i=1}^n F(e_i).$$

It follows that:

1. A tree of events \mathcal{T} is locally finite if and only if \mathcal{T} has finitely many roots, and if all of the trees as graphs that constitute the tree of events are locally finite in the usual sense (finite branching).
2. Every configuration of a tree is a stopped configuration. Indeed, we have $[e] = \omega \cap B(e)$ for every $\omega \in \Omega$ that contains e . This implies in particular that $\mathcal{X}_0 = \mathcal{W}_0$.

VII-1.2 Cone of Future. For every finite configuration $v = [e]$, the cone of future of v coincides with the set of successors of e : $\mathcal{E}^{[e]} = \{x \in \mathcal{E} \mid e \prec x\}$. Each cone of future \mathcal{E}^v is a tree of events.

VII-1.3 Branching Cells. A branching cell of \mathcal{T} has the form $\lambda = F(e)$, with e any event in λ . If $\Lambda_{\mathcal{T}}$ denotes the set of branching cells of \mathcal{T} , we have thus:

$$\forall \lambda, \lambda' \in \Lambda_{\mathcal{T}}, \quad \lambda \cap \lambda' \neq \emptyset \Rightarrow \lambda = \lambda'.$$

The “dynamic” puzzle is not dynamic.

VII-1.4 Lack of Concurrency. There is a unique minimal non void stopping prefix, consisting of the roots of \mathcal{T} . The lack of concurrency implies that $\Delta^\perp(\mathcal{E})$ is a singleton. The later condition is not sufficient however since there can be concurrent events inside a branching cell.

The initial stopping prefix coincides thus with this unique initial branching cell.

Since the cones of future are trees of events, $\Delta^+(v)$ consists for every configuration v of at most one branching cell. It follows that $([e_1], \dots, [e_n])$ is a germ-decomposition of $e_1 \oplus \dots \oplus e_n$.

VII-2 Unfolding of a Sequential Net.

Assume that \mathcal{T} is the tree of events given by the unfolding of a sequential net \mathcal{N} , with \mathcal{N} associated with a transition system (S, A, x_0) as in Ch. 1, III-3. The finite branching condition of VII-1.1 is satisfied. Hence unfoldings of sequential nets are locally finite.

Every branching cell λ projects one-to-one into the set of arrows a that all have the same origin x , for a certain state x of the system. It follows that the projections of branching cells in the transition system are disjoint.

VIII—Conclusion

In this chapter, we have introduced a class of configurations for locally finite event structures with finite concurrent width, the well-stopped configurations. Their study is motivated by their compositional properties. We have collected various results on well-stopped configurations that we use in particular in the next chapter to define projective systems of probabilities.

We have shown that well-stopped configurations have various recursive decompositions. We have defined the germs of a well-stopped configuration, that are its minimal increments. Germs are maximal configurations of finite sub-event structures. These finite event structures constitute the branching cells. Branching cells associated with a well-stopped configuration v are uniquely defined. Their collection is partially ordered, and they cut off without intersecting each other a neighbourhood around v . This collection is a factor of concurrency of the configuration, intrinsic to the configuration. In general, there are still concurrent elements *inside* branching cells.

To define and manipulate well-stopped configurations, we have also introduced the cone of future of a configuration. The cone of future of v is simply the sub-event structure that realises the shadow $\Omega(v)$ as its own set of maximal configurations.

$\#_d, \mathcal{S}$	• dynamic conflict (1-1.1), the lattice of stopping prefixes (1-2.1)
$\mathcal{S}_0, \mathcal{S}^*$	• the lattice of finite stopping prefix, the poset of non empty stopping prefixes
\mathcal{E}^v	• cone of future of v , a configuration of \mathcal{E}
$\#^v = \#\mathcal{E}^v, \#_d^v = \#_d^{\mathcal{E}^v}$	• conflict and dynamic conflict relations in \mathcal{E}^v
\mathcal{W}^v, Ω^v	• configurations and maximal configurations of \mathcal{E}^v
$\mathcal{W}(v), \Omega(v)$	• configurations and maximal configurations of \mathcal{E} containing v
$B^\perp(\mathcal{E}), \Delta^\perp(\mathcal{E})$	• initial stopping time of \mathcal{E} , collection of initial branching cells of \mathcal{E}
$\Delta_{\mathcal{E}}^+(v) (= \Delta^+(v))$	• collection of initial branching cells of \mathcal{E}^v
$\mathcal{X}, \mathcal{X}_0$	• set of well-stopped and finite well-stopped configurations of \mathcal{E}
$(v_n, z_n)_{n \geq 1}$	• a well-stopped sequence, <i>i.e.</i> satisfying $v_0 = \emptyset$, $z_n = v_{n+1} \ominus v_n$, and z_n finitely stopped in \mathcal{E}^v . Other restricting conditions on z_n make the sequence a step-by-step or a germ decomposition.
$\Lambda(v), \bar{\Lambda}(v)$	• the dynamic neighbourhood of $v \in \mathcal{X}$, $\bar{\Lambda}(v) = \Lambda(v) \sqcup \Delta^+(v)$ if $v \in \mathcal{X}_0$ (a disjoint union)
$\Lambda_{\mathcal{E}}$	• the collection of branching cells of \mathcal{E}

Table 3.1: Notations for Well-Stopped Configurations

Chapter 4

Distributed Probabilities

Chapter 2 has presented an analytical tool for the construction of a probability measure on the set Ω of maximal configurations of an event structure. Modulo a geometric condition—the local finiteness of the event structure—the probability is the limit of a projective system of probability measures on finite sets. In this chapter, we use the decompositions presented in Chapter 3 to explicitly construct projective systems of probabilities.

The projective systems that we construct define what we call *distributed probabilities*. The main feature is to join in some sense a probabilistic independence to the concurrency of local processes. The intuition according to which “parallelism is a form of orthogonality” [50] is confirmed. This study is a new interpretation of the idea of Benveniste *et al.* [7], that concurrency should fit probabilistic independence.

Chapter 3 has introduced the family of branching cells of an event structure. Each maximal configuration ω of a locally finite event structure is composed of *germs*, seen as local processes. A probability \mathbb{P} on Ω *induces* a collection of probabilities on germs, called *branching probabilities*. As the germs are local processes, branching probabilities are local probabilities, *i.e. local probabilistic parameters*. The main problem is then the reverse operation: given a collection of branching probabilities, can we construct a probability that induces this collection as induced branching probabilities? Is such a probability unique? The answer depends on the concurrency properties of the system.

Consider an event structure given by a *tree*. Randomising Ω consists in driving a boat that goes forward in the tree. Each choice is the result of the toss of a coin, whose probability law may vary in any way through the different choices. There is no concurrency, only one boat is in course. Any probability on Ω , *w.r.t.* the natural σ -algebra, is obtained this way—although it is an unusual way to formulate the extension theorem for projective limits.

For the concurrent picture, you now drive several boats, and they still go forward. The number of boats vary: they can join—synchronisation—or split—adding concurrency. There are now many joysticks: one for each boat that makes local choices, and also *additional parameters* related to the influence of each boat on the others. In a sequential system, the additional parameter is necessarily trivial. In concurrent systems, it is an additional assumption to require that the mutual influence is trivial.

This is what we do with distributed probabilities.

From a practical point of view, distributed probabilities are of high interest for computational reasons. They can be coded through a countable family of local probabilities on finite sets. If the event structure is given by the unfolding of a safe finite net, we can reduce the countable collection to a finite collection. They are intended to modelise distributed systems, where interactions are “minimal”: only local parameters must be specified.

Distributed probabilities also bring a new object: the distributed product of branching probabilities. The distributed product is the operation that associates a probability with a family of branching probabilities. As we have already noticed, its range of application is quite large since every probability on an infinite product $\Omega = E^{\mathbb{N}}$ can be obtained this way. Following the discussion engaged in the Introduction of Chapter 3, the distributed product, denoted with the symbol $\mathbb{P} = \bigotimes_{\mathcal{E}}^d \mathbb{P}^\lambda$ (distributed product on \mathcal{E} of the collection $(\mathbb{P}^\lambda)_\lambda$), satisfies the following properties:

1. *Universal w.r.t. the past.* For every finite stopping prefix B , the image in Ω_B of a d -product $\mathbb{P} = \bigotimes_{\mathcal{E}}^d \mathbb{P}^\lambda$ satisfies:

$$\pi_B \mathbb{P} = \bigotimes_B^d \mathbb{P}^\lambda .$$

2. *Universal w.r.t. the future.* For every finite well-stopped configuration, the probabilistic future $\mathbb{P}^v = \mathbb{P}(\cdot \mid \Omega(v))$ of a d -product is a d -product w.r.t. the cone of future \mathcal{E}^v :

$$\mathbb{P}^v = \bigotimes_{\mathcal{E}^v}^d \mathbb{P}^\lambda . \quad (4.1)$$

3. *Parallelism and probabilistic independence.* We have seen in Chapter 3 the association between the concurrency of branching cells and the independence of the associated germs. Distributed probabilities bring as a counterpart the probabilistic independence of the random variables defined by the concurrent germs.

Point 2. will be fundamental for the construction of *memory-less probabilities* in Chapter 5, connecting our constructions with the familiar framework of finite Markov chains.

The present chapter has 5 sections. The background from probability theory related to independence is presented in Section I, *Independence and conditional probabilities*. On the way, we collect the notions concerning conditional expectation and conditional probability used in next chapters. Section II, *General probabilistic framework* quickly recalls the vocabulary and notations concerning probabilistic event structures. In Section III, *Distributed probabilities*, we define the distributed probabilities and related objects: branching probabilities, random germs. The existence of distributed probabilities is the topic of Section IV, *The distributed product*.

We establish the theorem that states the existence and the compositional properties of the distributed product. Section V studies *Two examples of distributed product*. In the first example, a purely concurrent case leads to a direct product of probabilities. The second example is the example of trees, a purely sequential case; we show that every probability on a countable product (the boundary of a tree) is a distributed product.

I—Background: Independent Random Variables and Expectation

This section states the background material concerning independence of random variables needed in this chapter (I-1). We also recall some results on expectation and on conditional expectation that will be used later in the document (I-2–I-3), and we give an example of conditional expectation in the context of probabilistic event structures. Finally we recall two classical results from probability theory, the Borel-Cantelli lemma and the Strong law of large numbers for independent random variables.

I-1 Conditional Probability and Independence of Random Variables.

I-1.1 Conditional Probability. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. If A is a subset of Ω with $\mathbb{P}(A) > 0$, we denote by $\mathbb{P}(\cdot | A)$ the probability on Ω defined by:

$$\forall B \in \mathcal{F}, \quad \mathbb{P}(B | A) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(A)}.$$

We shall also denote by $\mathbb{P}(\cdot | A)$ the probability on A defined by:

$$\forall B \in \mathcal{F}, B \subseteq A, \quad \mathbb{P}(B | A) = \frac{\mathbb{P}(B)}{\mathbb{P}(A)}.$$

$\mathbb{P}(\cdot | A)$ is the probability *conditional on* A .

I-1.2 Independent Subsets. Two measurable subsets $A, B \subseteq \Omega$ are said to be independent (*w.r.t.* probability \mathbb{P}), if $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B)$.

With the vocabulary of conditional probabilities, and if $\mathbb{P}(A) > 0$, A and B are independent if and only if $\mathbb{P}(B | A) = \mathbb{P}(B)$.

I-1.3 Independence of σ -Algebras. Two σ -algebras $\mathcal{G}, \mathcal{G}' \subseteq \mathcal{F}$ are said to be independent if A and B are independent subsets, for all pairs $(A, B) \in \mathcal{G} \times \mathcal{G}'$. More generally, if I is a finite set and \mathcal{F}_i are sub- σ -algebras of \mathcal{F} , $(\mathcal{F}_i)_{i \in I}$ is said to be a family of independent σ -algebras if we have:

$$\mathbb{P}\left(\bigcap_{i \in I} A_i\right) = \prod_{i \in I} \mathbb{P}(A_i),$$

for every tuple of subsets $(A_i)_i \in \prod_i \mathcal{F}_i$.

I-1.4 Independence of Random Variables. Let $(\Omega, \mathcal{F}, \mathcal{P})$ be a probability space, and let $(E_i, \mathcal{F}_i)_{i \in I}$ be a finite collection of measurable spaces. Let $(X_i)_{i \in I}$

be a family of random variables $X_i : \Omega \rightarrow E_i$. The family $(X_i)_{i \in I}$ is said to be independent if the family of σ -algebras $\langle X_i \rangle_{i \in I}$ is independent. Equivalently:

$$\forall (A_i)_{i \in I} \in \prod_i \mathcal{F}_i, \quad \mathbb{P}\left(\bigcap_i (X_i \in A_i)\right) = \prod_i \mathbb{P}(X_i \in A_i).$$

For instance, assume that $X_i : \Omega \rightarrow \mathbb{N}$ are integer random variables, with $I = \{1, \dots, k\}$. Then $(X_i)_i$ is independent if and only if for every k -tuple $(n_i)_i$:

$$\mathbb{P}(X_1 = n_1, \dots, X_k = n_k) = \mathbb{P}(X_1 = n_1) \cdots \mathbb{P}(X_k = n_k).$$

More generally, the following result is of constant use for studying independent random variables. Recall that, if $(E_i, \mathcal{F}_i, \mu_i)_i$ is a finite family of probability spaces, the product spaces $\prod_i E_i$ is equipped with the product σ -algebra $\mathcal{F} = \bigotimes_i \mathcal{F}_i$, generated by the squares:

$$\prod_i A_i,$$

with $A_i \in \mathcal{F}_i$. The **product** probability $\mu = \bigotimes_i \mu_i$ is the unique probability on $(\prod_i E_i, \bigotimes_i \mathcal{F}_i)$ satisfying:

$$\forall (A_i)_{i \in I} \in \prod_i \mathcal{F}_i, \quad \mu\left(\prod_i A_i\right) = \prod_i \mu_i(A_i).$$

I-1.5 Theorem. *Let $(X_i)_{i \in I}$ be a family of random variables $X_i : \Omega \rightarrow E_i$. Let E denote the product space $E = \prod_i E_i$. Then the family $(X_i)_i$ is independent under \mathbb{P} if and only if the probability law \mathbb{P}^X of $X = (X_i)_i \in E$ is the product:*

$$\mathbb{P}^X = \bigotimes_{i \in I} \mathbb{P}^i,$$

with \mathbb{P}^i the probability law of X_i in E_i .

I-2 Expectation and Conditional Expectation.

We denote by L^1 the set of integrable real random variables $X : \Omega \rightarrow \mathbb{R}$. We denote by L^∞ the set of real random variables, bounded on a set of probability 1. For $X \in L^1$, the expectation of X is defined as its integral, and we write:

$$\mathbb{E}(X) = \int_{\Omega} X(\omega) \, d\mathbb{P}(\omega).$$

If A is measurable subset of Ω with $\mathbb{P}(A) > 0$, the expectation of X conditional on A is the integral of X w.r.t. the probability measure $\mathbb{P}(\cdot | A)$, i.e.:

$$\mathbb{E}(X | A) = \frac{1}{\mathbb{P}(A)} \int_A X \, d\mathbb{P}.$$

Theorem I-1.5 has the following consequence:

I-2.1 Theorem. Let X, Y be two independent random variables. For every pair (g, h) of integrable real random variables, with f X -measurable and g Y -measurable, we have:

$$\mathbb{E}(fg) = \mathbb{E}(f)\mathbb{E}(g) .$$

I-3 Conditional Expectation w.r.t. a σ -Algebra.

For \mathcal{G} a sub- σ -algebra of \mathcal{F} , we write $g \in \mathcal{G}$ if g is a \mathcal{G} -measurable random variable. \mathcal{G} -measurable functions are seen as test functions.

Let \mathcal{G} be a sub- σ -algebra of \mathcal{F} , and let $X \in L^1$. There is a random variable $Z \in L^1$, \mathcal{G} -measurable, and such that:

$$\forall g \in L^\infty, \quad g \in \mathcal{G} \Rightarrow \mathbb{E}(Xg) = \mathbb{E}(Zg) . \quad (4.2)$$

If Z' is an other \mathcal{G} -measurable random variable satisfying (4.2), then $Z = Z'$ \mathbb{P} -a.s. Z is called the **conditional expectation** of g w.r.t. \mathcal{G} , denoted by:

$$Z = \mathbb{E}(X | \mathcal{G}) .$$

I-3.1 Example. (*Conditional expectation w.r.t. a discrete random variable*) We write $\mathbb{E}(X | Y) = \mathbb{E}(X | \langle Y \rangle)$, where $\langle Y \rangle$ is the σ -algebra generated by a random variable Y . Assume that Y takes its values in a discrete set of values $y_1, \dots \in \mathbb{R}$. Let $Z = \mathbb{E}(X | Y)$. Z is $\langle Y \rangle$ -measurable, hence there is a mapping $z : \{y_1, \dots\} \rightarrow \mathbb{R}$ such that $Z = z(Y)$, and z is given by:

$$z(y) = \mathbb{E}(X | Y = y) , \quad (4.3)$$

if $\mathbb{P}(Y = y) > 0$, and any real number otherwise. In (4.3), the expectation is taken conditionally on the event $\{Y = y\} = \{\omega \in \Omega | Y(\omega) = y\}$, that is w.r.t. the conditional probability $\mathbb{P}(\cdot | Y = y)$.

I-3.2 Example. Assume that X and Y are independent real random variables. Then conditional expectation and expectation coincide:

$$\mathbb{E}(X | Y) = \mathbb{E}(X) .$$

Indeed for every random variable g Y -measurable, we have by Th. I-2.1:

$$\mathbb{E}(gX) = \mathbb{E}(g)\mathbb{E}(X) = \mathbb{E}(g\mathbb{E}(X)) .$$

The constant $\mathbb{E}(X)$ is Y -measurable and satisfies the characterisation of conditional expectation.

I-3.3 Example. (*Conditional expectation w.r.t. a stopping prefix*) This example illustrates the notion of probabilistic future, defined below in II-3.

Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure, and let B be a finite stopping prefix of \mathcal{E} . We denote as usual $\pi_B : \Omega \rightarrow \Omega_B$, and $\omega_B = \pi_B(\omega) = \omega \cap B$. Let \mathcal{F}_B denote the σ -algebra $\mathcal{F}_B = \langle \pi_B \rangle$. For any integrable function $h : \Omega \rightarrow \mathbb{R}$, we have:

$$\mathbb{E}(h | \mathcal{F}_B) = \int_{\Omega^{\omega_B}} h(\omega_B \oplus \xi) d\mathbb{P}^{\omega_B}(\xi).$$

$\mathbb{E}(h | \mathcal{F}_B)$ represents the mean value of h , given that $\omega \supseteq \omega_B$. We can define it as zero or as any function on $\mathbb{P}_B(\omega_B) = 0$.

Proof— Let ϕ be a \mathcal{F}_B -measurable function. There is a function $\dot{\phi} : \Omega_B \rightarrow \mathbb{R}$ with: $\phi(\omega) = \dot{\phi}(\omega_B)$. Since B is finite, Ω_B is finite, hence:

$$\begin{aligned} \mathbb{E}(h\phi) &= \sum_{w \in \Omega_B} \int_{\Omega(w)} h\phi d\mathbb{P} \\ &= \sum_{w \in \Omega_B} \dot{\phi}(w) \int_{\Omega(w)} h(\omega) d\mathbb{P}(\omega) \\ &= \sum_{w \in \Omega_B} \mathbb{P}(\Omega(w)) \dot{\phi}(w) \int_{\Omega^w} h(w \oplus \xi) d\mathbb{P}^w(\xi), \end{aligned}$$

where $\mathbb{P}^w = \frac{1}{\mathbb{P}(\Omega(w))} \mathbb{P}(\cdot \cap \Omega(w))$ denotes the probabilistic future of \mathcal{U}^w (see below, II-3). In the later sum, there is no need to define \mathbb{P}^w if $\mathbb{P}(\Omega(w)) = 0$. The sum can be written as:

$$\mathbb{E}(h\phi) = \mathbb{E}_B(\dot{Y}\dot{\phi}),$$

with $\dot{Y} : \Omega \rightarrow \mathbb{R}$ defined by:

$$\dot{Y} : \Omega_B \rightarrow \mathbb{R}, \quad \dot{Y}(w) = \int_{\Omega^w} h(w \oplus \xi) d\mathbb{P}^w(\xi) \quad \text{if } \mathbb{P}(\Omega(w)) > 0,$$

$\dot{Y} = 0$ otherwise. We set $Y : \Omega \rightarrow \mathbb{R}$, $Y(\omega) = \dot{Y}(\omega_B)$ to get:

$$\mathbb{E}(h\phi) = \mathbb{E}(Y\phi).$$

Y is \mathcal{F}_B -measurable and has the characteristic property of conditional expectation, hence $\mathbb{E}(h | \mathcal{F}_B) = Y$. \square

I-3.4 Properties of Conditional Expectation. The conditional expectation satisfies the following properties, for $\mathcal{G} \subseteq \mathcal{F}$ two σ -algebras:

$$\forall h \in \mathcal{F}, \quad \mathbb{E}(\mathbb{E}(h | \mathcal{G})) = \mathbb{E}(h), \tag{4.4}$$

$$\forall h \in \mathcal{F}, \forall g \in \mathcal{G}, \quad \mathbb{E}(gh | \mathcal{G}) = g\mathbb{E}(h | \mathcal{G}), \quad g \text{ is "constant" w.r.t. } \mathcal{G}. \tag{4.5}$$

I-4 Independent and Identically Distributed (i.i.d.) Random Variables.

We recall the statement of the Strong law of large numbers for *i.i.d.* sequences [10].

I-4.1 *i.i.d.* Sequences. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let (E, \mathcal{G}) be a measurable space. A sequence $(X_n)_{n \geq 1}$ of random variables $X_n : \Omega \rightarrow E$ is said to be *identically distributed* if the law \mathbb{P}^{X_n} of X_n in E is independent of n . If X_i are real random variable identically distributed, it implies that the X_n are integrable all together, and if they are integrable $\mathbb{E}(X_i) = \mathbb{E}(X_j)$ for all i, j .

The sequence $(X_n)_{n \geq 1}$ is said to be *independent* if for every n , the finite family (X_1, \dots, X_n) is independent (I-1.4). “*i.i.d.*” is an abbreviation for independent and identically distributed.

I-4.2 Remark. If $(X_n)_{n \geq 1}$ is *i.i.d.* and if X_n takes a finite number of values, then $(X_n)_{n \geq 1}$ is a finite Markov chain. The transition matrix has identical rows. Any row is given by the finite probability vector that gives the probability law of X_1 .

I-4.3 Theorem. Let $(X_n)_{n \geq 1}$ be an *i.i.d.* sequence of real random variables $X_n : \Omega \rightarrow \mathbb{R}$, that we assume integrable. Then for every non negative function $f : \mathbb{R} \rightarrow \mathbb{R}$, we have the convergence:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n f(X_k) = \mathbb{E}(f(X_1)), \quad \mathbb{P}\text{-a.s.}$$

If μ denotes the common probability law \mathbb{P}^{X_i} of X_i in \mathbb{R} we have:

$$\mathbb{E}(f(X_1)) = \int_{\mathbb{R}} f(x) d\mu(x).$$

I-5 Limit Sup of Subsets and the Borel-Cantelli Lemma.

I-5.1 Limit Sup of Subsets. Let (Ω, \mathcal{F}) be a measurable space, and let $(A_n)_{n \geq 1}$ be a sequence of measurable subsets of Ω . The *limit sup* of the sequence $(A_n)_n$ is the measurable subset of Ω defined by:

$$\limsup_{n \rightarrow \infty} A_n = \bigcap_{N \geq 1} \bigcup_{n \geq N} A_n.$$

The elements $\omega \in \limsup A_n$ are those elements $\omega \in \Omega$ that belong to infinitely many A_n . We also say that A_n *holds infinitely often*, abbreviated by *i.o.* Thus we write:

$$\limsup_{n \rightarrow \infty} A_n = \{A_n \text{ i.o.}\}.$$

We recall the following elementary result.

I-5.2 Theorem. (Borel-Cantelli Lemma) Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, and let $(A_n)_{n \geq 1}$ be a sequence of measurable subsets of Ω .

1. If $\sum_{n \geq 1} \mathbb{P}(A_n) < \infty$, then: $\mathbb{P}(A_n \text{ i.o.}) = 0$.
2. If the sequence $(A_n)_n$ is an independent sequence of subsets, and if $\sum_{n \geq 1} \mathbb{P}(A_n) = \infty$, then: $\mathbb{P}(A_n \text{ i.o.}) = 1$.

II—General Probabilistic Framework

We consider an event structure \mathcal{E} , that we assume *locally finite* and of *finite concurrent width* (Chapter 3, I-2.3, III-1.1).

II-1 Probability Space and Random Variables for Event Structures.

We denote by Ω the set of maximal configurations of \mathcal{E} . We have shown in Chapter 2 that the topology generated by the finite shadows:

$$\Omega(v) = \{\omega \in \Omega \mid \omega \supseteq v\}, \quad v \in \mathcal{W}_0,$$

makes Ω a separable metric space. We have defined a probabilistic event structure as a pair $(\mathcal{E}, \mathbb{P})$, with \mathbb{P} a probability measure on Ω , equipped with the σ -algebra \mathcal{F} generated by the collection of finite shadows.

We recall that for each finite prefix P of \mathcal{E} , the *finite* set $\Gamma_P = \{\omega \cap P \mid \omega \in \Omega\}$ is equipped with the discrete σ -algebra, that contains all its subsets. We denote by π_P the mapping:

$$\pi_P : \Omega \rightarrow \Gamma_P, \quad \omega \mapsto \pi_P(\omega) = \omega \cap P,$$

and we denote by $\mathcal{F}_P = \langle \pi_P \rangle$ the sub- σ -algebra of \mathcal{F} generated by π_P . \mathcal{F}_P is the collection of π_P -saturated subsets $A \subseteq \Omega$, *i.e.* such that:

$$\forall \omega \in A, \quad \forall \xi \in \Omega, \quad \pi_P(\omega) = \pi_P(\xi) \Rightarrow \xi \in A.$$

For any measurable space E , a random variable $Y : \Omega \rightarrow E$ is $\langle \pi_P \rangle$ -measurable if and only if there is a random variable $f : \Gamma_P \rightarrow E$ such that $Y = f \circ \pi_P$. Ω_P denotes the set of maximal configurations of P . Every stopping prefix $B \in \mathcal{S}$ is intrinsic to \mathcal{E} , *i.e.* satisfies $\Omega_B = \Gamma_B$.

Since \mathcal{E} is locally finite, Ω is isomorphic as a topological space, and hence as a measurable space, to the projective limit:

$$(\Omega, \mathcal{F}) \simeq \varprojlim_{B \in \mathcal{S}_0} (\Omega_B), \quad (4.6)$$

with Ω_B equipped with the discrete σ -algebra (Ch.2, III-3).

In particular, \mathcal{F} coincides with the σ -algebra:

$$\mathcal{F} = \langle \mathcal{F}_B, B \in \mathcal{S}_0 \rangle .$$

As a consequence, dealing with finite stopping prefixes instead of finite prefixes is allowed without loss of generality, from a measurable (and topological) point of view. Moreover, Theorem III-3.1 of Ch. 2 states that every projective system $(\mathbb{P}_B)_{B \in \mathcal{S}_0}$ of probability measures on $(\Omega_B)_{B \in \mathcal{S}_0}$, *w.r.t.* the natural filtration $\pi_{B,B'}$, can be extended to a unique probability measure \mathbb{P} on Ω such that $\mathbb{P}_B = \pi_B \mathbb{P}$ for all $B \in \mathcal{S}_0$. The diagram of Ch. 3, I-3.1, becomes the following commutative diagram of probability spaces, for $B, B' \in \mathcal{S}_0$ with $B \subseteq B'$:

$$\begin{array}{ccc} (\Omega, \mathcal{F}, \mathbb{P}) & \xrightarrow{\pi_{B'}} & (\Omega_{B'}, \mathcal{F}_{B'}, \mathbb{P}_{B'}) \\ & \searrow \pi_B & \downarrow \pi_{B,B'} \\ & & (\Omega_B, \mathcal{F}_B, \mathbb{P}_B) \end{array}$$

II-2 Likelihood of Configurations.

We recall that every shadow $\Omega(w)$, with $w \in \mathcal{W}$ a configuration, is measurable. Indeed, $\Omega(w)$ is the countable intersection of measurable subsets:

$$\Omega(w) = \bigcap_{\substack{v \subseteq w \\ v \in \mathcal{W}_0}} \Omega(v) .$$

We define the **likelihood** of \mathbb{P} as the function $p_{\mathcal{E}}$ with values in the real interval $[0, 1]$, given by:

$$p_{\mathcal{E}} : \mathcal{W} \rightarrow [0, 1], \quad w \mapsto p_{\mathcal{E}}(w) = \mathbb{P}(\Omega(w)) .$$

Let $B \in \mathcal{S}$. The probabilistic event structure (B, \mathbb{P}_B) defines a likelihood $p_B : \mathcal{W}_B \rightarrow [0, 1]$. For every configuration $v \subseteq B$, we have:

$$\pi_B^{-1}(\Omega_B(v)) = \Omega(v) ,$$

and therefore:

$$p_B(v) = p_{\mathcal{E}}(v) \tag{4.7}$$

is independent of $B \in \mathcal{S}$ such that $v \subseteq B$. Therefore, we simply note p for the likelihood.

II-3 Probabilistic Future of a Finite Configuration.

Let v be a finite configuration of \mathcal{E} , and let \mathcal{E}^v be the cone of future of v in \mathcal{E} (Ch. 3, II-1.1). The mapping $\Omega(v) \rightarrow \Omega^v$ is a homeomorphism (Ch. 3, II-3.3), and thus an isomorphism of measurable spaces, where:

- $\Omega(v)$ is equipped with the restriction of \mathcal{F} to $\Omega(v)$,
- Ω^v is equipped with the operational σ -algebra *w.r.t.* the event structure \mathcal{E}^v .

We denote with the same symbol \mathcal{F}^v , both σ -algebras on $\Omega(v)$ and on Ω^v . If v satisfies $p(v) > 0$, then we equip $(\Omega(v), \mathcal{F}^v)$ with the conditional probability \mathbb{P}^v defined by:

$$\mathbb{P}^v(\cdot) = \mathbb{P}(\cdot | \Omega(v)), \quad \mathbb{P}^v(A) = \frac{\mathbb{P}(A)}{p(v)} \quad \forall A \subseteq \Omega(v), A \in \mathcal{F}^v.$$

We denote with the same symbol the probability on Ω^v , image of \mathbb{P}^v under the isomorphism $\Omega(v) \rightarrow \Omega^v$, and given by:

$$\forall A \subseteq \Omega^v, A \in \mathcal{F}^v, \quad \mathbb{P}^v(A) = \frac{1}{p(v)} \mathbb{P}(v \oplus A). \quad (4.8)$$

II-3.1 Definition. (*Probabilistic future*) Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure, and let $v \in \mathcal{W}_0$ such that $p(v) > 0$. We define the **probabilistic cone of future** of v , or shortly the **probabilistic future** of v , as the probabilistic event structure $(\mathcal{E}^v, \mathbb{P}^v)$.

II-3.2 Likelihood in the Future. If $v \in \mathcal{W}_0$ satisfies $p(v) > 0$, we denote by p^v the likelihood of $(\mathcal{E}^v, \mathbb{P}^v)$, and we have from (4.8):

$$\forall y \in \mathcal{W}^v, \quad p^v(y) = \frac{p(v \oplus y)}{p(v)}. \quad (4.9)$$

II-4 \star -Regular Probabilities.

We introduce the following class of probabilities, that is convenient for our purpose. It is analogous to a statistical model *dominated along a filtration* [14].

II-4.1 Definition. (*\star -regular probability*) Let (X, d) be a separable metric space, let \mathcal{F} be the Borel σ -algebra generated by d , and let \mathbb{P} be a probability on (X, \mathcal{F}) . We say that \mathbb{P} is **\star -regular** if $\mathbb{P}(U) > 0$ for every non empty open set U .

II-4.2 Remark. Two \star -regular probabilities might not to be regular¹ one *w.r.t.* the other. I give below two counter-examples.

- Let $X = [0, 1]$ and let p, q be two different numbers of $]0, 1[$. Consider \mathbb{P} and \mathbb{Q} the probabilities on $X \cong \{0, 1\}^{\mathbb{N}}$ given by the coin game on $\{0, 1\}$, respectively with probability p and q respectively on 0. Then by the Strong Law of Large Numbers applied to the natural process $(X_n)_{n \geq 0}$ of *i.i.d* random variables, \mathbb{P} is concentrated on the set of sequences $(X_n)_{n \geq 0}$ such that:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \text{Card}\{k \leq n \mid X_k = 0\} = p$$

and similarly with q for \mathbb{Q} . Therefore \mathbb{P} and \mathbb{Q} are concentrated on disjoint sets, and thus none is regular *w.r.t.* the other, although both are \star -regular *w.r.t.* the Euclidean topology of X .

- The following example has been kindly communicated to me by Tanguy Briançon. Let $X = [0, 1]$ and let $(x_n)_{n \geq 1}$ be a sequence dense in X . Let $k > 0$ be a constant such that $\sum_{n \geq 1}^{\infty} \frac{2k}{n^2} < 1$. Let U be the open set of X given by:

$$U = \bigcup_{n \geq 1} B\left(x_n, \frac{k}{n^2}\right),$$

where $B(x, r)$ denotes the open interval of centre x and radius r . Then $\overline{U} = X$ since $(x_n)_n$ is dense. Let m denote the Lebesgue measure on X , and let $H = X \setminus U$. We have $m(H) \geq 1 - \sum_{n \geq 1}^{\infty} \frac{2k}{n^2} > 0$. Let m_U denote the restriction of m to U , given by $m_U(\cdot) = m(\cdot \cap U)$. Then m and m_U are both \star -regular, but $m(H) > 0$ and $m_U(H) = 0$. Therefore m is not regular *w.r.t.* m_U .

The class of \star -regular probabilities is stable under product and projective limit, which has the following expression in our context.

II-4.3 Lemma. *Let $(\mathcal{E}, \mathbb{P})$ be a locally finite probabilistic event structure, with likelihood p . The following propositions are equivalent:*

1. \mathbb{P} is \star -regular,
2. \mathbb{P}_B is \star -regular for every $B \in \mathcal{S}_0$,
3. $p(v) > 0$ for every $v \in \mathcal{W}_0$.

In this case, \mathbb{P}^v is defined for every $v \in \mathcal{W}_0$, and \mathbb{P}^v is \star -regular.

Proof— 1 \Leftrightarrow 3. Since the finite shadows form a basis of open sets of Ω , \mathbb{P} is \star -regular if and only if $\mathbb{P}(\Omega(v)) = p(v) > 0$ for every $v \in \mathcal{W}_0$.

2 \Leftrightarrow 3. As \mathcal{E} is locally finite, every finite configuration is subset of a $B \in \mathcal{S}_0$, and (4.7) gives then the result.

It follows from (4.9) that \mathbb{P}^v is then \star -regular. \square

¹A probability \mathbb{P} is said to be regular *w.r.t.* a probability \mathbb{Q} , which is denoted by $\mathbb{P} \ll \mathbb{Q}$, if $\mathbb{Q}(A) = 0 \Rightarrow \mathbb{P}(A) = 0$.

III—Distributed Probabilities

In this Section we introduce the definition of distributed probabilities, together with other related notions: branching probabilities, induced branching probabilities and random germs. We begin the study *a priori* of distributed probabilities, their effective construction is the topic of Section IV.

We begin in III-1 with the definition of *branching probability* and of *induced branching probability*. This comes together with the notion of *absolute random germ* associated with a branching cell. The notion of *conditional random germ* is presented in III-2. The definition of distributed probabilities, presented in III-3, establishes a relation between conditional and absolute random germs. We conclude the section by studying in III-4 the image of distributed probabilities in stopping prefixes.

III-1 Branching Probabilities.

We use the notions of branching cells and of well-stopped configurations introduced in Chapter 3. We recall that Table 3.1, page 120, presents a summary of the notations that we use concerning well-stopped configurations and branching cells. Since we assume that \mathcal{E} is locally finite, every branching cell is finite (Ch. 3, III-2.3).

III-1.1 Definition. (*Branching probability*) Let λ be a branching cell of \mathcal{E} . A **branching probability on λ** is a probability μ on the finite set Ω_λ . The branching probability μ is said to be **positive** if $\mu(z) > 0$ for every $z \in \Omega_\lambda$.

III-1.2 Induced Branching Probability. We show that a probability on Ω induces a collection of branching probabilities. For each $\lambda \in \Lambda_\mathcal{E}$, we set the following subset of Ω :


$$\mathcal{H}_\mathcal{E}^\lambda = \{\omega \in \Omega \mid \lambda \in \bar{\Lambda}(\omega)\},$$

where $\bar{\Lambda}(\omega)$, defined in Ch. 3, VI-3.1, denotes the collection of branching cells of \mathcal{E} that form the neighbourhood of ω . Thus $\mathcal{H}_\mathcal{E}^\lambda$ is the set of maximal configurations of \mathcal{E} that contain λ as a branching cell in their decomposition. $\mathcal{H}_\mathcal{E}^\lambda$ is a “thick” shadow, illustrated in Figure 4.1. We denote $\mathcal{H}_\mathcal{E}^\lambda$ by \mathcal{H}^λ if there is no ambiguity on the event structure \mathcal{E} .

III-1.3 Lemma. For every $\lambda \in \Lambda_\mathcal{E}$, \mathcal{H}^λ is a non empty open subset of Ω .

Proof – By definition of $\bar{\Lambda}(\omega)$, we have:

$$\mathcal{H}^\lambda = \bigcup_{v \in \mathcal{X}_0 : \lambda \in \Delta^+(v)} \Omega(v).$$

Figure 4.1: Thick shadow \mathcal{H}^λ .

Therefore \mathcal{H}_λ is open as a union of open sets. It is not empty since by definition of branching cells, there is a $v \in \mathcal{X}_0$ with $\lambda \in \Delta^+(v)$. \square

In particular \mathcal{H}^λ is measurable, and $\mathbb{P}(\mathcal{H}^\lambda) > 0$ for all $\lambda \in \Lambda_{\mathcal{E}}$ whenever \mathbb{P} is \star -regular. In order to deal with non \star -regular probabilities, we introduce the following definition. For simplicity, the reader may assume that all probabilities are \star -regular.

III-1.4 Definition. (*Positive trace of a probability, positive branching cell*) Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure. We define the positive trace of \mathbb{P} as the set of finite well-stopped configurations v such that $p(v) > 0$:

$$\mathcal{R}(\mathbb{P}) = \{v \in \mathcal{X}_0 \mid p(v) > 0\}.$$

We say that a branching cell λ of \mathcal{E} is \mathbb{P} -positive, or **positive** for short, if $\mathbb{P}(\mathcal{H}^\lambda) > 0$.

Remark that it is totally independent for a branching cell λ to be positive, and for a branching probability μ on λ to be positive.

Assume that λ is a positive branching cell. Then we equip \mathcal{H}^λ with the conditional probability measure $\mathbb{P}(\cdot \mid \mathcal{H}^\lambda)$, given by:

$$\forall A \subseteq \mathcal{H}^\lambda, A \in \mathcal{F}, \quad \mathbb{P}(A \mid \mathcal{H}^\lambda) = \frac{1}{\mathbb{P}(\mathcal{H}^\lambda)} \mathbb{P}(A). \quad (4.10)$$

III-1.5 Definition. (*Absolute random germs*) Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure. For each positive branching cell λ , we define the following random variable as the **absolute random germ** of λ :

$$X^\lambda : \mathcal{H}^\lambda \rightarrow \Omega_\lambda, \quad \omega \mapsto \omega \cap \lambda,$$

where \mathcal{H}^λ is equipped with the conditional probability (4.10).

Proof— We have to show that $\omega \cap \lambda \in \Omega_\lambda$ for every $\lambda \in \Lambda_\mathcal{E}$ and all $\omega \in \mathcal{H}^\lambda$. Fix $\lambda \in \Lambda_\mathcal{E}$, $\omega \in \mathcal{H}^\lambda$, and let $v \in \mathcal{X}_0$ such that $v \subseteq \omega$ and $\lambda \in \Delta^+(v)$. Then $\omega \ominus v$ is maximal in \mathcal{E}^v (Ch. 3, VI-2.2) and λ is a stopping prefix of \mathcal{E}^v , thus $(\omega \ominus v) \cap \lambda \in \Omega_\lambda$. As $\lambda \in \Delta^+(v)$, we have $v \cap \lambda = \emptyset$, and thus $(\omega \ominus v) \cap \lambda = \omega \cap \lambda$, from which follows that $\omega \cap \lambda \in \Omega_\lambda$. \square

III-1.6 Definition. (*Induced branching probability*) Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure, and let λ be a positive branching cell of \mathcal{E} . We define the branching probability on λ **induced by** \mathbb{P} , as the probability law of the random variable $X^\lambda : \mathcal{H}^\lambda \rightarrow \Omega_\lambda$. We denote this probability on Ω_λ by \mathbb{P}^λ .

For every $x \in \Omega_\lambda$, $\mathbb{P}^\lambda(x)$ is given by:

$$\mathbb{P}^\lambda(x) = \mathbb{P}(X^\lambda = x | \mathcal{H}^\lambda) = \frac{\mathbb{P}(\omega \in \mathcal{H}^\lambda \text{ and } \omega \cap \lambda = x)}{\mathbb{P}(\mathcal{H}^\lambda)}. \quad (4.11)$$

III-1.7 Lemma. Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure, and let B be a stopping prefix. Let $\lambda \in \Lambda_B$, i.e. a branching cell of \mathcal{E} subset of B . Then λ is positive for $(\mathcal{E}, \mathbb{P})$ if and only if λ is positive for (B, \mathbb{P}_B) , and in this case both induce the same branching probability in λ .

Proof— We have:

$$\mathcal{H}_B^\lambda = \{\omega_B \in \Omega_B \mid \lambda \in \overline{\Lambda}(\omega_B)\},$$

from which follows that $\pi_B^{-1}(\mathcal{H}_B^\lambda) = \mathcal{H}_\mathcal{E}^\lambda$. Therefore: λ is \mathbb{P} -positive \Leftrightarrow λ is \mathbb{P}_B -positive. Assume that λ is a positive branching cell. For every $\omega \in \mathcal{H}^\lambda$, we have $\omega \cap \lambda = \pi_B(\omega) \cap \lambda$ since $\lambda \subseteq B$, from which follows:

$$\begin{aligned} \forall z \in \Omega_\lambda, \quad \mathbb{P}_B^\lambda(z) &= \frac{\mathbb{P}_B(\omega_B \in \mathcal{H}_B^\lambda \text{ and } \omega_B \cap \lambda = z)}{\mathbb{P}_B(\mathcal{H}_B^\lambda)} \\ &= \frac{\mathbb{P}(\omega \in \mathcal{H}^\lambda \text{ and } \omega \cap \lambda = z)}{\mathbb{P}(\mathcal{H}^\lambda)} = \mathbb{P}^\lambda(z). \end{aligned}$$

\square

III-2 Concurrent Conditional Random Germs.

The absolute random germ X^λ depends only on λ , and thus is intrinsic to λ , as a branching cell of \mathcal{E} . We now define an other random germ with value in Ω_λ , that is not intrinsic to λ , but with the advantage that concurrent random germs can be considered on the same probability space.

We fix a finite well-stopped configuration $v \in \mathcal{X}_0$. For each $\lambda \in \Delta^+(v)$, we define the random variable Z_v^λ by:

$$Z_v^\lambda : \Omega(v) \rightarrow \Omega_\lambda, \quad \omega \mapsto \omega \cap \lambda .$$

It is immediate that $Z_v^\lambda \in \Omega_\lambda$. Indeed for $\omega \in \Omega(v)$, $\omega \ominus v$ is maximal in \mathcal{E}^v , and λ is a stopping prefix of \mathcal{E}^v , thus $\omega \cap \lambda = (\omega \ominus v) \cap \lambda \in \Omega_\lambda$. We recall that we denote by $\Pi(v)$ the product:

$$\Pi(v) = \prod_{\lambda \in \Delta^+(v)} \Omega_\lambda .$$

III-2.1 Definition. (*Conditional random germs*) Let v be a finite well-stopped configuration of a probabilistic event structure $(\mathcal{E}, \mathbb{P})$, and assume that $p(v) > 0$. We define the **conditional random germs** of v as the random variables Z_v^λ for $\lambda \in \Delta^+(v)$, defined on the same probability space $\Omega(v)$. We set the following product random variable:

$$Z_v : \Omega(v) \rightarrow \Pi(v), \quad Z_v = (Z_v^\lambda)_{\lambda \in \Delta^+(v)} .$$

As it was for absolute random germs, the restriction of conditional germs to stopping prefixes keeps the laws invariant, as stated by the following result. Recall the notation $\mathcal{R}(\mathbb{P})$ for the positive trace of \mathbb{P} , III-1.4.

III-2.2 Lemma. *Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure. Let $v \in \mathcal{R}(\mathbb{P})$, and let $B \in \mathcal{S}$ such that $\Lambda(v) \subseteq \Lambda_B$. Let Y_v^λ denote the conditional λ -germ induced by (B, \mathbb{P}_B) , and let $Y_v = (Y_v^\lambda)_{\lambda \in \Delta_B^+(v)}$. Then Z_v and Y_v have the same law in $\Pi(v)$.*

Proof – Remark first that Y_v and Z_v take their value in the same space $\Pi(v)$ since $\Lambda(v) \subseteq B$, and thus $\Delta_B^+(v) = \Delta_{\mathcal{E}}^+(v) \subseteq B$. Let $y = (y_\lambda)_{\lambda \in \Delta^+(v)}$. The law of Y_v evaluated at y is given by:

$$\mathbb{P}_B(Y_v = y \mid \Omega_B(v)) = \frac{1}{p_B(v)} \mathbb{P}_B(\omega_B \supseteq v, Y_v(\omega_B) = y),$$

where p_B denotes the likelihood in (B, \mathbb{P}_B) . For every $\omega \in \Omega(v)$, $Y_v(\omega \cap B) = Z_v(\omega)$, and since $p_B(v) = p(v)$ as seen in (4.9), we get:

$$\begin{aligned} \mathbb{P}_B(Y_v = y \mid \Omega_B(v)) &= \frac{1}{p(v)} \mathbb{P}(\omega \supseteq v, Z_v(\omega) = y) \\ &= \mathbb{P}(Z_v = y \mid \Omega(v)) . \end{aligned}$$

□

For every $v \in \mathcal{R}(\mathbb{P})$ and $\lambda \in \Delta^+(v)$, λ is a positive branching cell. In this case, the two random variables X^λ and Z_v^λ , with values in Ω_λ , define probability laws on Ω_λ . In general, these laws need not be equal (we derive a counter-example from a law on $\{0, 1\} \times \{0, 1\}$ that is not the product of its marginal laws). This observation leads to the definition of distributed probabilities.

III-3 Distributed Probabilities.

We define a particular class of probabilities, that we call distributed probabilities². We have seen that the germs that compose well-stopped configurations can be concurrent. Moreover, from a set point of view, concurrent germs are independent. Distributed probabilities translate this property into an independence in the probabilistic sense.

We give the definition, and then we show that the likelihood of distributed probabilities has a simple form (Cf. I-1 for background on independent random variables).

III-3.1 Definition. (*Distributed probability*) Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure. We say that the probability \mathbb{P} is ***distributed*** *w.r.t.* \mathcal{E} , or shortly that \mathbb{P} is distributed, if for every $v \in \mathcal{R}(\mathbb{P})$, we have:

1. the random variables $(Z_v^\lambda)_{\lambda \in \Delta^+(v)}$ form an independent family,
2. for every $\lambda \in \Delta^+(v)$, the law of Z_v^λ in Ω_λ is the law of X^λ .

Equivalently, \mathbb{P} is distributed if and only if for every $v \in \mathcal{X}_0$ with $p(v) > 0$, the law of Z_v in $\Pi(v) = \prod_{\lambda \in \Delta^+(v)} \Omega_\lambda$ is given by the direct product of probabilities:

$$\bigotimes_{\lambda \in \Delta^+(v)} \mathbb{P}^\lambda, \quad (4.12)$$

where \mathbb{P}^λ is the branching probability induced by \mathbb{P} on λ (III-1.6). Remark that this product is finite since we assume that \mathcal{E} has finite concurrent width. The product probability is the unique probability that gives probability:

$$\prod_{\lambda \in \Delta^+(v)} \mathbb{P}^\lambda(z^\lambda),$$

to a tuple $(z^\lambda)_\lambda \in \Pi(v)$.

III-3.2 Concurrency and Probabilistic Independence. Point 1 in Definition III-3.1 of distributed probabilities is the more intuitive part: the conditional random germs are independent in the probabilistic sense. This is also the sense of (4.12).

The next proposition shows that the likelihood of distributed probabilities has a simple form on \mathcal{X}_0 .

²The term “distributed” in the expression “distributed probability” refers to the distributed systems themes, and not to the theory of distributions nor to the distribution of the probability, *i.e.* to its law.

III-3.3 Counting Branching Cells. We introduce a tool for inductions on finite well-stopped configurations. We define the following function, that counts the number of branching cells in the decomposition of an element $v \in \mathcal{X}_0$:

$$N : \mathcal{X}_0 \rightarrow \mathbb{N}, \quad v \in \mathcal{X}_0 \mapsto \langle N, v \rangle = \text{Card}(\Lambda(v)). \quad (4.13)$$

The notation $\langle N, v \rangle$ will be justified in Chapter 6.

III-3.4 Proposition. *Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure, with \mathbb{P} a distributed probability. Then the likelihood p is given on \mathcal{X}_0 by:*

$$\forall v \in \mathcal{R}(\mathbb{P}), \quad p(v) = \prod_{\lambda \in \Lambda(v)} \mathbb{P}^\lambda(v \cap \lambda). \quad (4.14)$$

Proof— We show the result by induction on $\langle N, v \rangle = \text{Card} \Lambda_{\mathcal{E}}(v)$ (III-3.3). If empty, the product equals 1 by convention, so $p(\emptyset) = 1$ holds. Let $v' = v \oplus z$, with $z \in \Omega_{\lambda_0}$ and $\lambda_0 \in \Delta^+(v)$, and assume that $v \oplus z \in \mathcal{R}(\mathbb{P})$. Then $\Lambda(v') = \Lambda(v) \sqcup \{\lambda\}$, a disjoint union (Ch. 3, VI-3.7), and thus $\langle N, v' \rangle = \langle N, v \rangle + 1$. Obviously, $p(v) > 0$, thus we apply the chain rule to get:

$$p(v') = p(v) p^v(z) = p(v) \mathbb{P}^v(Z_v^{\lambda_0} = z). \quad (4.15)$$

The subset $\{Z_v^{\lambda_0} = z\}$ of $\Omega(v)$ projects to $\Pi(v)$ onto the product:

$$\{z\} \times \prod_{\substack{\lambda \in \Delta^+(v) \\ \lambda \neq \lambda_0}} \Omega_\lambda.$$

Using that \mathbb{P} is distributed, the law of Z^v has the form (4.12), so we get $p^v(z) = \mathbb{P}^{\lambda_0}(z) \times 1 = \mathbb{P}^{\lambda_0}(v' \cap \lambda_0)$. Using the induction hypothesis, we get from (4.15):

$$p(v') = \mathbb{P}^{\lambda_0}(v' \cap \lambda_0) \prod_{\lambda \in \Lambda(v)} \mathbb{P}^\lambda(v \cap \lambda).$$

It follows from Ch. 3, VI-3.3, that $v' \cap \lambda = v \cap \lambda$ for every $\lambda \in \Lambda(v)$, and this completes the induction. \square

III-3.5 Corollary. *Let \mathbb{P} and \mathbb{Q} be two probabilities distributed w.r.t. an event structure \mathcal{E} (locally finite and of finite concurrent width). Assume that the following holds:*

$$\forall \lambda \in \Lambda_{\mathcal{E}}, \quad \text{if } \lambda \text{ is } \mathbb{P}\text{-positive and } \mathbb{Q}\text{-positive, then } \mathbb{P}^\lambda = \mathbb{Q}^\lambda.$$

Then $\mathbb{P} = \mathbb{Q}$.

Proof— *First step.* We show that $\mathcal{R}(\mathbb{P}) = \mathcal{R}(\mathbb{Q})$. Let $v \in \mathcal{R}(\mathbb{P})$. We set:

$$v' = \sup\{y \in \mathcal{X}_0, y \subseteq v \mid y \in \mathcal{R}(\mathbb{Q})\}.$$

Assume that $v' \subsetneq v$. Since the well-stopped sub-configurations of v form a lattice (Ch. 3, VI-2.3), v' is well-stopped. Hence there is a $\lambda_0 \in \Delta^+(v')$ such that $\lambda \in \Lambda(v)$, i.e., $\lambda \notin \Delta^+(v)$. Let p and q denote the likelihoods of \mathbb{P} and \mathbb{Q} . We have $q(v') > 0$ by construction, so λ_0 is \mathbb{Q} -positive. Let z be the λ_0 germ of v . Then $v' \oplus z \subseteq v$. Since $p(v) > 0$, we have that λ_0 is \mathbb{P} -positive, and that $\mathbb{P}^{\lambda_0}(z) > 0$. Thus λ_0 is \mathbb{P} -positive and \mathbb{Q} -positive, and by the assumption it implies that $\mathbb{Q}^{\lambda_0}(z) = \mathbb{P}^{\lambda_0}(z)$. But then $q(v' \oplus z) > 0$, which contradicts that v' is maximal among configurations $y \in \mathcal{X}_0$ satisfying $y \in \mathcal{R}(\mathbb{Q})$ and $y \subseteq v$. It follows that $\mathcal{R}(\mathbb{P}) \subseteq \mathcal{R}(\mathbb{Q})$, and by symmetry: $\mathcal{R}(\mathbb{P}) = \mathcal{R}(\mathbb{Q})$.

Second step. Let $B \in \mathcal{S}_0$, and let $v \in \Omega_B$. If $p(v) = 0$, it follows from the above point that $q(v) = 0$, and thus $p(v) = q(v)$. If $p(v) > 0$, then $q(v) > 0$. It follows that all the branching cells of $\Lambda(v)$ are \mathbb{P} -positive and \mathbb{Q} -positive, and thus satisfy $\mathbb{P}^\lambda = \mathbb{Q}^\lambda$. Since \mathbb{P} and \mathbb{Q} are distributed, it follows from III-3.4 that $p(v) = q(v)$. We have shown that $\mathbb{P}_B = \mathbb{Q}_B$ for all $B \in \mathcal{S}_0$. By the uniqueness in the extension theorem (Ch. 2, III-3.1), it implies that $\mathbb{P} = \mathbb{Q}$. \square

III-4 Restriction of Distributed Probabilities to Stopping Prefixes.

Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure, and assume that \mathbb{P} is a distributed probability. We first look for conservation properties *w.r.t.* restriction to stopping prefixes. The following result shows that the distributed probabilities are stable under projective limit. The case of the probabilistic future is more delicate and is examined in the next section.

III-4.1 Proposition. *Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure. Then $(\mathcal{E}, \mathbb{P})$ is distributed if and only if (B, \mathbb{P}_B) is distributed for every $B \in \mathcal{S}_0$. In this case, for every $B \in \mathcal{S}_0$, (B, \mathbb{P}_B) and $(\mathcal{E}, \mathbb{P})$ induce the same collection of branching probabilities on Λ_B .*

Proof – The first part follows from III-1.7 and III-2.2.

Conversely, assume that (B, \mathbb{P}_B) is distributed for every $B \in \mathcal{S}_0$, we show that \mathbb{P} is distributed *w.r.t.* \mathcal{E} . Let $v \in \mathcal{R}(\mathbb{P})$. Since \mathcal{E} is locally finite, $\Delta^+(v)$ is a finite collection of finite sets (Ch. 3, III-2.3). There is thus a $B \in \mathcal{S}_0$ such that $\lambda \subseteq B$ for all $\lambda \in \Delta^+(v)$. Since we have (Ch. 3, V-1.3) $\Lambda_B = \{\lambda \in \Lambda_{\mathcal{E}} \mid \lambda \subseteq B\}$, it follows that $\Delta_{\mathcal{E}}^+(v) = \Delta_B^+(v)$. Then applying III-1.7 and III-2.2, the law of Z_v in $\Pi(v)$ is the product:

$$\bigotimes_{\lambda \in \Delta_B^+(v)} \mathbb{P}_B^\lambda = \bigotimes_{\lambda \in \Delta_{\mathcal{E}}^+(v)} \mathbb{P}^\lambda.$$

This shows that \mathbb{P} is distributed. \square

IV—Distributed Product

This Section states the existence of distributed probabilities. For this we propose a construction that reverses the operation $\mathbb{P} \rightarrow (\mathbb{P}^\lambda)_\lambda$. The goal is to establish theorem IV-2.2, that states both the existence of the distributed product and its compositional properties *w.r.t.* the past and *w.r.t.* the future. The proof is based on an induction. We use the analytical result of extension of measures shown in Chapter 2.

IV-1 Preliminary.

We have seen in III-3.4 that distributed probabilities, if they exist, must have their likelihood given by (4.14) on finite configurations. We will now check that this expression actually leads to a projective system.

Before that, we introduce the following random variable, defined for $\lambda \in \Lambda_\mathcal{E}$:

$$V_\mathcal{E}^\lambda : \mathcal{H}^\lambda \rightarrow \mathcal{W}, \quad \omega \mapsto V_\mathcal{E}^\lambda(\omega) = \min\{v \in \mathcal{X}_0 \mid v \subseteq \omega, \lambda \in \Delta^+(v)\}.$$

The set at right hand is non empty because of $\omega \in \mathcal{H}^\lambda$, and is a lattice by Ch. 3, VI-2.3, thus $V_\mathcal{E}^\lambda$ is well defined.

IV-1.1 Lemma. *The random variable $V_\mathcal{E}^\lambda$ satisfies the following property:*

$$\forall \omega, \omega' \in \Omega, \quad \omega \supseteq V_\mathcal{E}^\lambda(\omega') \Rightarrow V_\mathcal{E}^\lambda(\omega) = V_\mathcal{E}^\lambda(\omega'). \quad (4.16)$$

For all $B \in \mathcal{S}$ such that $\lambda \in \Lambda_B$, i.e. such that $\lambda \subseteq B$, $V_\mathcal{E}^\lambda$ is \mathcal{F}_B -measurable, given by $V_\mathcal{E}^\lambda = V_B^\lambda \circ \pi_B$.

Proof— We denote shortly $V = V_\mathcal{E}^\lambda$. Let $\omega, \omega' \in \Omega$. Let $y = V(\omega')$, and assume that $\omega \supseteq y$. Then y is finite well-stopped, and satisfies $\lambda \in \Delta_\mathcal{E}^+(y)$, thus $V(\omega) \subseteq y$, i.e.: $V(\omega) \subseteq V(\omega')$.

It implies, since $y \subseteq \omega'$, that $V(\omega) \subseteq \omega'$. Exchanging the role of ω, ω' and applying the same reasoning, we get that $V(\omega') \supseteq V(\omega)$, and hence $V(\omega) = V(\omega')$. The identity $V_\mathcal{E}^\lambda = V_B^\lambda \circ \pi_B$ for a $B \in \mathcal{S}_0$ that contains λ follows from $\Lambda_B = \Lambda_\mathcal{E} \cap B$, in the sense as given by Lemma V-1.3, Ch. 3. It implies that $V_\mathcal{E}^\lambda$ is \mathcal{F}_B -measurable. \square

IV-1.2 Lemma. *Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure. Assume that there is a collection of branching probabilities $(\mu_\lambda)_{\lambda \in \Lambda_\mathcal{E}}$, such that for every $v \in \mathcal{R}(\mathbb{P})$, the law of Z_v in $\Pi(v)$ is given by the product law $\bigotimes_{\lambda \in \Delta^+(v)} \mu_\lambda$. Then we have, for all $\lambda \in \Lambda_\mathcal{E}$:*

$$\mathbb{P}(\mathcal{H}^\lambda) > 0 \Rightarrow \mathbb{P}^\lambda = \mu_\lambda.$$

Proof— Let $\lambda \in \Lambda_{\mathcal{E}}$, assume that $\mathbb{P}(\mathcal{H}^\lambda) > 0$, and let $z \in \Omega_\lambda$. We denote by V the random variable $V = V_{\mathcal{E}}^\lambda$, and by T the set of values of V . Since \mathcal{E} is locally finite, there is a $B \in \mathcal{S}_0$ such that $\lambda \subseteq B$, and then V is \mathcal{F}_B -measurable (IV-1.1). It implies that T is a finite set, and in particular:

$$\mathbb{P}(\mathcal{H}^\lambda, X^\lambda = z) = \sum_{u \in T} \mathbb{P}(V = u, X^\lambda = z).$$

Equation (4.16) implies that $\{V = u\} = \Omega(u)$ for any $u \in T$. Therefore:

$$\mathbb{P}(\mathcal{H}^\lambda, X^\lambda = z) = \sum_{u \in T} \mathbb{P}(\Omega(u)) \mathbb{P}(X^\lambda = z | \Omega(u)) = \sum_{u \in T} p(u) \mathbb{P}^u(\xi \cap \lambda = z),$$

where ξ denotes the variable in Ω^u . Since any $u \in T$ is well-stopped, it follows from the hypothesis that:

$$\forall u \in T, \quad p(u) > 0 \Rightarrow \mathbb{P}^u(\xi \cap \lambda = z) = \mu_\lambda(z),$$

and thus: $\forall u \in T \quad p(u) \mathbb{P}^u(\xi \cap \lambda = z) = p(u) \mu_\lambda(z)$. Therefore:

$$\mathbb{P}(\mathcal{H}^\lambda, X^\lambda = z) = \left(\sum_{u \in T} p(u) \right) \mu_\lambda(z) = \mathbb{P}(\mathcal{H}^\lambda) \mu_\lambda(z).$$

We have thus: $\mathbb{P}(\mathcal{H}^\lambda) > 0 \Rightarrow \mathbb{P}^\lambda(z) = \mu_\lambda(z)$, and this holds for all $z \in \Omega_\lambda$. \square

IV-2 Construction of the Distributed Product.

IV-2.1 Definition. (*Probability consistent with a family of branching probabilities*) Let \mathcal{E} be a locally finite event structure of finite concurrent width, and let $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$ be a family of branching probabilities. We say that a probability \mathbb{P} on Ω is consistent with $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$, if we have:

$$\forall \lambda \in \Lambda_{\mathcal{E}}, \quad \mathbb{P}(\mathcal{H}^\lambda) > 0 \Rightarrow \mathbb{P}^\lambda = \mu_\lambda.$$

IV-2.2 Theorem. *Let \mathcal{E} be a locally finite event structure, of finite concurrent width. Let $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$ be a family of branching probabilities. There is a unique distributed probability measure \mathbb{P} consistent with $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$. We call this probability the distributed product of the family $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$, and we denote it by:*

$$\mathbb{P} = \bigotimes_{\lambda \in \Lambda_{\mathcal{E}}}^d \mu_\lambda.$$

\mathbb{P} is \star -regular if and only if every branching probability μ_λ is positive (III-1.1), and there is a one-to-one mapping between \star -regular distributed probabilities and families of positive branching probabilities $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$.

The distributed product satisfies the following composition formula:

$$\forall B \in \mathcal{S}, \quad \mathbb{P}_B = \bigotimes_{\lambda \subseteq B}^d \mathbb{P}^\lambda, \quad (4.17)$$

$$\forall v \in \mathcal{X}_0 \text{ with } p(v) > 0, \quad \mathbb{P}^v = \bigotimes_{\lambda \subseteq \mathcal{E}^v}^d \mathbb{P}^\lambda. \quad (4.18)$$

Proof – Existence. Fix the event structure \mathcal{E} . For every finite event structure \mathcal{K} such that $\Lambda_{\mathcal{K}} \subseteq \Lambda_{\mathcal{E}}$, let $m_{\mathcal{K}} : \Omega_{\mathcal{K}} \rightarrow [0, 1]$ be the real-valued function defined by:

$$\forall w \in \Omega_{\mathcal{K}}, \quad m_{\mathcal{K}}(w) = \prod_{\lambda \in \Lambda_{\mathcal{K}}(w)} \mu_{\lambda}(w \cap \lambda). \quad (4.19)$$

For such an event structure \mathcal{K} , we define the height of \mathcal{K} as the integer:

$$N(\mathcal{K}) = \max\{\langle N, w \rangle, w \in \Omega_{\mathcal{K}}\},$$

where $\langle N, w \rangle$ counts the branching cells of w (III-3.3). We show by induction on $N(\mathcal{K})$ that $m_{\mathcal{K}}$ defines a probability on $\Omega_{\mathcal{K}}$, *i.e.* it sums to 1 over $\Omega_{\mathcal{K}}$.

This is trivial for $\mathcal{K} = \emptyset$, since $\Omega_{\emptyset} = \{\emptyset\}$. Let \mathcal{K} be finite and non empty with $\Lambda_{\mathcal{K}} \subseteq \mathcal{E}$, and assume that the property holds for every event structure of height $< N(\mathcal{K})$. Denote by $B^0 = B^\perp(\mathcal{K})$ the initial stopping prefix of \mathcal{K} . For every $w \in \Omega_{\mathcal{K}}$, denote by v the full-initial germ of w , given by $v = w \cap \Omega_{B^0}$, and denote by \mathcal{K}^v the cone of future of v . We apply Ch. 3, VI-3.7, to the concatenation $w = v \oplus (w \ominus v)$ to get the disjoint union:

$$\Lambda(w) = \Delta^\perp(\mathcal{K}) \sqcup \Lambda_{\mathcal{K}^v}(w \ominus v). \quad (4.20)$$

In particular, if $v \neq \emptyset$, we have:

$$N(\mathcal{K}^v) \leq N(\mathcal{K}) - 1.$$

Since $\mathcal{K} \neq \emptyset$, and since $\emptyset \in \Omega_{B^0} \Rightarrow \mathcal{K} = \emptyset$, we actually have that $v \neq \emptyset$. Moreover, as v is stopped in \mathcal{K} , we have according to Ch. 3, V-1.3: $\Lambda_{\mathcal{K}^v} \subseteq \Lambda_{\mathcal{K}} \subseteq \Lambda_{\mathcal{E}}$, thus the induction hypothesis applies to \mathcal{K}^v .

It follows from (4.20) that the function $m_{\mathcal{K}}$ has the following expression:

$$\begin{aligned} m_{\mathcal{K}}(w) &= \left(\prod_{\lambda \in \Delta^\perp(\mathcal{K})} \mu_{\lambda}(v \cap \lambda) \right) \left(\prod_{\lambda \in \Lambda_{\mathcal{K}^v}(w \ominus v)} \mu_{\lambda}((w \ominus v) \cap \lambda) \right) \\ &= m_{B^0}(v) m_{\mathcal{K}^v}(w \ominus v). \end{aligned}$$

Since $w \in \Omega_{\mathcal{K}} \mapsto w \cap B^0$ is onto Ω_{B^0} , and since $w \in \Omega(v) \mapsto w \ominus v \in \Omega^v$ is a bijection (Ch. 3, II-1.3), $w \in \Omega_{\mathcal{K}} \mapsto (v, w \ominus v)$ is a bijection onto the following disjoint union:

$$\Omega_{\mathcal{K}} \mapsto \bigsqcup_{v \in \Omega_{B^0}} \{v\} \times \Omega^v.$$

We sum the above expression of $m_{\mathcal{K}}(w)$ to obtain:

$$\sum_{w \in \Omega_{\mathcal{K}}} m_{\mathcal{K}}(w) = \sum_{v \in \Omega_{B^0}} m_{B^0}(v) \left(\sum_{s \in \Omega_{\mathcal{K}^v}} m_{\mathcal{K}^v}(s) \right) = \sum_{v \in \Omega_{B^0}} m_{B^0}(v), \quad (4.21)$$

the later by applying the induction hypothesis to \mathcal{K}^v . Given the identification $\Omega_{B^0} \rightarrow \Pi^{\perp}(\mathcal{K})$ (Ch. 3, IV-3.5), it is clear that m_{B^0} identifies with the product probability $\bigotimes_{\lambda \in \Delta^{\perp}(\mathcal{K})} \mu_{\lambda}$, and hence sums to 1. Then it follows from (4.21) that $m_{\mathcal{K}}$ sums to 1, and this completes the induction.

In particular for B a finite stopping prefix of \mathcal{E} , we have $\Lambda_B \subseteq \Lambda_{\mathcal{E}}$ (Ch. 3, V-1.3), and thus m_B defines a probability on Ω_B . We show that $(m_B)_{B \in \mathcal{S}_0}$ is a projective system. Let $B, B' \in \mathcal{S}_0$ such that $B \subseteq B'$. Since Ω_B and $\Omega_{B'}$ are finite, we have to check that:

$$\forall v \in \Omega_B, \quad \sum_{w \in \Omega_{B'}(v)} m_{B'}(w) = m_B(v). \quad (4.22)$$

Let $v \in \Omega_B$. For every $w \in \Omega_{B'}(v)$, we have by Ch. 3, VI-3.7:

$$\Lambda_{B'}(w) = \Lambda_B(v) \sqcup \Lambda_{B^v}(w \ominus v).$$

Hence we get:

$$\begin{aligned} \sum_{w \in \Omega_{B'}(v)} m_{B'}(w) &= \left(\prod_{\lambda \in \Lambda_B(v)} \mu_{\lambda}(v \cap \lambda) \right) \sum_{s \in \Omega_{B'}^v} \prod_{\lambda \in \Lambda_{B^v}(s)} \mu_{\lambda}(s \cap \lambda) \\ &= m_B(v) \sum_{s \in \Omega_{B'}^v} m_{B^v}(s). \end{aligned}$$

Since B^v is a finite sub-event structure of \mathcal{E} with branching cells a sub-collection of $\Lambda_{\mathcal{E}}$, we have shown that the last sum equals 1, and (4.22) follows.

As $(m_B)_{B \in \mathcal{S}_0}$ is a projective system of probability measures, and since \mathcal{E} is locally finite, there is a unique probability \mathbb{P} on Ω such that $m_B = \pi_B \mathbb{P}$ for every $B \in \mathcal{S}_0$ (Ch. 2, III-3.1). The likelihood p of \mathbb{P} is given, on finitely B -stopped configurations, by (4.19) with $\mathcal{K} = B$. Now we show that p is given on \mathcal{X}_0 by:

$$\forall v \in \mathcal{X}_0, \quad p(v) = \prod_{\lambda \in \Lambda(v)} \mu_{\lambda}(v \cap \lambda). \quad (4.23)$$

For this, let $v \in \mathcal{X}_0$, and let $B \in \mathcal{S}_0$ that contains v . Then we have $p(v) = m_B(\Omega_B(v))$. Every $w \in \Omega_B(v)$ admits the decomposition:

$$w = v \oplus s, \quad s \in \Omega_B^v.$$

Since v is well-stopped in \mathcal{E} , v is well-stopped in B (Ch. 3, IV-2.1). As w is maximal and thus well-stopped in B , $s = w \ominus v$ is well-stopped in B^v (Ch. 3, VI-2.2). Applying

Ch. 3, VI-3.7, we have: $\Lambda_B(w) = \Lambda_B(v) \sqcup \Lambda_{B^v}(s)$, from which follows:

$$\begin{aligned} p(v) &= \sum_{s \in \Omega_B^v} p(v \oplus s) \\ &= \left(\prod_{\lambda \in \Lambda_B(v)} \mu_\lambda(v \cap \lambda) \right) \left(\sum_{s \in \Omega_B^v} m_{B^v}(s) \right). \end{aligned}$$

We have shown that the sum at right equals 1, which implies (4.23).

Now we can compute the law of Z_v in $\Pi(v)$, for $v \in \mathcal{X}_0$ with $p(v) > 0$. Let $u = (u_\lambda)_{\lambda \in \Delta^+(v)}$ be an element of $\Pi(v)$, we have, using (4.23):

$$\begin{aligned} \mathbb{P}^v(Z_v = u) &= \frac{1}{p(v)} p\left(v \oplus \bigcup_{\lambda \in \Delta^+(v)} u_\lambda\right) \\ &= \prod_{\lambda \in \Delta^+(v)} \mu_\lambda(u_\lambda) = \left(\bigotimes_{\lambda \in \Delta^+(v)} \mu_\lambda \right)(u). \end{aligned}$$

Hence the law of Z_v in $\Pi(v)$ is the product law $\bigotimes_{\lambda \in \Delta^+(v)} \mu_\lambda$, and this holds for every $v \in \mathcal{R}(\mathbb{P})$. According to IV-1.2, it implies:

$$\forall \lambda \in \Lambda_{\mathcal{E}}, \quad \mathbb{P}(\mathcal{H}^\lambda) > 0 \Rightarrow \mathbb{P}^\lambda = \mu_\lambda.$$

Therefore $(\mathcal{E}, \mathbb{P})$ is distributed and consistent with $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$. This completes the proof of existence of a distributed probability consistent with the family $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$. Actually, we have a bit more, since we have explicitly constructed a distributed product of $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$.

Uniqueness. If \mathbb{P} and \mathbb{Q} are consistent with $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$, then they satisfy the condition of III-3.5. Hence $\mathbb{P} = \mathbb{Q}$.

★-Regularity. The fact that \mathbb{P} is ★-regular if and only if all μ_λ are positive comes from (4.23). Hence for each family of positive branching probabilities $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$ we construct a unique consistent distributed and ★-regular probability. Conversely, let \mathbb{P} be distributed and ★-regular. \mathbb{P} admits the unique following form as a distributed product:

$$\mathbb{P} = \bigotimes_{\lambda \in \Lambda_{\mathcal{E}}}^d \mathbb{P}^\lambda,$$

and each \mathbb{P}^λ is a positive branching probability.

Composition formulae. By construction, the distributed product on any $B \in \mathcal{S}_0$ is given by $m_B = \pi_B \mathbb{P}$. Now let $B \in \mathcal{S}$ not necessarily finite. Let $(B_n)_{n \geq 0}$ be a sequence cofinal in \mathcal{S}_0 (Ch. 2, I-4.4). Then $B \cap B_n$ is cofinal in B . Making $n \rightarrow \infty$ in: $m_{B_n \cap B} = \pi_{B_n \cap B} \mathbb{P}$, we get:

$$\bigotimes_{\lambda \subseteq B}^d \mu_\lambda = \pi_B \mathbb{P}.$$

This shows (4.17).

Now we show (4.18). We write $B_n \rightarrow \mathcal{E}$ to denote that $(B_n)_n$ is a sequence of \mathcal{S}_0 , cofinal in \mathcal{S}_0 . Let $v \in \mathcal{X}_0$, and let $B \in \mathcal{S}_0$ such that $v \subseteq B$. For any $w \in \Omega_B(v)$, we have seen that, using the decomposition $w = v \oplus (w \ominus v)$, we get:

$$p(w) = p(v)m_{\mathcal{E}^v}(w \ominus v),$$

and thus $\mathbb{P}_B^v = m_{B^v}$, modulo the identification $\Omega_B(v) \rightarrow \Omega_{B^v}^v$. Consider a sequence $B_n \rightarrow \mathcal{E}$. Then $B_n \cap \mathcal{E}^v \rightarrow \mathcal{E}^v$, and we recall that we have $B_n \cap \mathcal{E}^v = B_n^v$ (Ch. 3, II-5.1), hence $B_n^v \rightarrow \mathcal{E}^v$. Making $n \rightarrow \infty$ in $\mathbb{P}_{B_n}^v = m_{B_n^v}$, we obtain:

$$\mathbb{P}^v = \bigotimes_{\lambda \subseteq \mathcal{E}^v}^d \mu_\lambda.$$

□

IV-2.3 Corollary. *Let $(\mathcal{E}, \mathbb{P})$ be a distributed probabilistic event structure. Let $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$ be any family of branching probabilities such that $\mu_\lambda = \mathbb{P}^\lambda$ for all \mathbb{P} -positive branching cells λ . Then we have:*

$$\mathbb{P} = \bigotimes_{\lambda \subseteq \mathcal{E}}^d \mu_\lambda.$$

Proof – Let $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$ be such a family, and let $\mathbb{Q} = \bigotimes_{\lambda \subseteq \mathcal{E}}^d \mu_\lambda$. Then \mathbb{P} and \mathbb{Q} are distributed and consistent with $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{E}}}$. Hence they are equal. □

V—Two Examples of Distributed Product.

We analyse two “extremal” examples of distributed products: one with full concurrency, and one without concurrency. With full concurrency, we find the direct product of measures. Without concurrency, *i.e.* for trees of events, we find without surprise that all probabilities are distributed. The distributed character of probabilities is trivial for trees.

V-1 Concurrent Product.

Assume that \mathcal{E} coincides with its initial stopping prefix, *i.e.*: $\mathcal{E} = B^\perp(\mathcal{E})$. Then $\Omega = \Omega_{B^\perp(\mathcal{E})}$, and any $\omega \in \Omega$ decomposes itself as a unique disjoint union:

$$\omega = \bigsqcup_{\lambda \in \Delta^\perp(\mathcal{E})} z^\lambda, \quad z^\lambda = \omega \cap \lambda.$$

The study of Ch. 2, IV-4, shows that the collection of branching cells of \mathcal{E} is given by $\Lambda_{\mathcal{E}} = \Delta^{\perp}(\mathcal{E})$. Hence, if $(\mu_{\lambda})_{\lambda \in \Delta^{\perp}(\mathcal{E})}$ is a collection of branching probabilities, the distributed product $\mathbb{P} = \bigotimes_{\lambda \in \Delta^{\perp}(\mathcal{E})}^{\text{d}} \mu_{\lambda}$ is just the direct product of measures: $\mathbb{P} = \bigotimes_{\lambda \in \Delta^{\perp}(\mathcal{E})} \mu_{\lambda}$.

V-2 Distributed Probabilities on Trees.

The construction of distributed probabilities is based on an analysis of concurrency properties. Since the concurrency relation is trivial in trees of events, we expect that the distributed character of probabilities are trivial. Indeed, every probability on Ω is distributed, and can be written as a distributed product. We begin by a quick review of the different probabilistic notions introduced, for trees.

V-2.1 Absolute and Conditional Germs. Let \mathcal{T} be a tree of events (Ch. 1, III-4.2). Let G denote the graph relation on \mathcal{T} :

$$\forall x, y \in \mathcal{T}, \quad xGy \Leftrightarrow y \in \text{Min}_{\preceq}(z \in \mathcal{T} \mid z \succ x),$$

such that the reflexive and transitive closure of G coincides with \preceq . We have seen in Ch. 3, VII, that every finite configuration is finitely stopped. We recall that the branching cells of \mathcal{T} are disjoint, and given by the collection:

$$\lambda_x = \{y \in \mathcal{T} \mid xGy\}, \quad x \in \mathcal{T},$$

excepted for the unique initial branching cell, formed by the finite set of roots of the tree of events. We also have for all $x \in \mathcal{T}$: $\Delta^+([x]) = \{\lambda_x\}$. There is thus a unique conditional germ $Z_{[x]}^{\lambda_x}$ (III-2.1). We write:

$$Z_x = Z_{[x]}^{\lambda_x}, \tag{4.24}$$

and Z_{\emptyset} for the unique initial germ.

Let x be an event of \mathcal{T} . We write $\Omega(x)$ to denote the finite shadow $\Omega([x])$. We trivially have the inclusion (III-1.2): $\Omega(x) \subseteq \mathcal{H}^{\lambda_x}$. Conversely, let $\omega \in \mathcal{H}^{\lambda_x}$. That is, there is finite configuration $v = [y]$ such that $y \subseteq \omega$ and $\lambda_x \in \Delta^+(v)$. Then we have $\lambda_x = \lambda_y$ and thus $x = y$, so we have: $\mathcal{H}^{\lambda_x} \subseteq \Omega(x)$. Hence we obtain:

$$\forall x \in \mathcal{T}, \quad \mathcal{H}^{\lambda_x} = \Omega(x). \tag{4.25}$$

It follows that the absolute random germ X^{λ} (III-1.5) coincides with the unique conditional random germ Z_x defined by (4.24).

V-2.2 Branching Probabilities. Let \mathbb{P} be any probability on (Ω, \mathcal{F}) . Let $\lambda = \lambda_x$ be a branching cell of \mathcal{T} . Assume that x has positive likelihood, *i.e.* $\mathbb{P}(\Omega(x)) > 0$. The branching probability induced by \mathbb{P} in λ is given by:

$$\forall z \in \Omega_{\lambda}, \quad \mathbb{P}^{\lambda}(z) = \mathbb{P}(X^{\lambda} = z \mid \mathcal{H}^{\lambda}) = \mathbb{P}(Z_x = z \mid \Omega(x)). \tag{4.26}$$

V-2.3 Distributed Probabilities. Let \mathbb{P} be a probability on Ω . Let $v = [x]$ be a finite (trivially well-stopped) configuration. The family of conditional germs at v reduces to the singleton (Z_x) , and is thus trivially an independent family. Moreover for λ the branching cell associated with v , the law of Z_x in Ω_λ coincides with the law of X^λ by (4.26). It follows that \mathbb{P} is a distributed probability (Definition III-3.1). Using the form of distributed product of IV-2.3, we have the following result.

V-2.4 Theorem. *For every probability \mathbb{P} on (Ω, \mathcal{F}) , $(\mathcal{T}, \mathbb{P})$ is distributed. There is a family of branching probabilities $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{T}}}$ such that $\mathbb{P} = \bigotimes_{\lambda \subseteq \mathcal{T}}^d \mu_\lambda$.*

VI—Conclusion

In this chapter we have constructed a product of branching probabilities, called distributed product. The distributed product associates a probabilistic event structure with the data consisting of:

- a locally finite event structure with finite concurrent width,
- a collection of branching probabilities, defined on the branching cells of the event structure.

The class of probabilities reached by this product considers jointly probability and concurrency: *distributed* probabilities associate a probabilistic independence with the concurrency of local processes. Local processes are defined by the germs of maximal configurations.

The distributed product has remarkable composition properties. Restriction to the past, and conditioning *w.r.t.* the future are two operations on probabilities that combine with operations on the event structure and on the processes. Remark the triple level of operations: event structures, processes and maximal processes, and probabilities. We obtain the following table:

Past	Stopping prefix B	$\begin{cases} \Omega \rightarrow \Omega_B \\ \omega \mapsto \omega \cap B \end{cases}$	$\begin{aligned} \mathbb{P}_B &= \pi_B \mathbb{P} \\ &= \bigotimes_{\lambda \subseteq B}^d \mathbb{P}^\lambda \end{aligned}$
Future	Cone of future \mathcal{E}^v	$\begin{cases} \Omega(v) \rightarrow \Omega^v \\ \omega \mapsto \omega \cap \mathcal{E}^v \end{cases}$	$\begin{aligned} \mathbb{P}^v &= \mathbb{P}(\cdot \mid \Omega(v)) \\ &= \bigotimes_{\lambda \subseteq \mathcal{E}^v}^d \mathbb{P}^\lambda \end{aligned}$

$\mathbb{P}, p_{\mathcal{E}} = p$	<ul style="list-style-type: none"> • a probability \mathbb{P} on Ω, the associated likelihood $p_{\mathcal{E}} : \mathcal{W} \rightarrow \mathbb{R}$, defined by $p_{\mathcal{E}}(\cdot) = \mathbb{P}(\Omega(\cdot))$ (II-2).
$\mathcal{R}(\mathbb{P})$	<ul style="list-style-type: none"> • the trace of \mathbb{P} in \mathcal{W}_0, the set of positive finite configurations (III-1.4)
$(\mathcal{E}^v, \mathbb{P}^v)$	<ul style="list-style-type: none"> • the probabilistic future of configuration v, where \mathcal{E}^v is the cone of future of v, and $\mathbb{P}^v = \mathbb{P}(\cdot \Omega(v))$ (II-3).
p^v	<ul style="list-style-type: none"> • the likelihood associated with \mathbb{P}^v (II-3.2)
\star -regular probability $\mathcal{H}_{\mathcal{E}}^{\lambda} = \mathcal{H}^{\lambda}$	<ul style="list-style-type: none"> • II-4 • the open subset of Ω associated with a branching cell λ (III-1)
X^{λ}	<ul style="list-style-type: none"> • the absolute random germ associated with branching cell λ, a random variable $X^{\lambda} : H^{\lambda} \rightarrow \Omega_{\lambda}$ (III-1.5)
\mathbb{P}^{λ}	<ul style="list-style-type: none"> • the branching probability in the positive branching cell λ induced by \mathbb{P}. \mathbb{P}^{λ} is the law of X^{λ} in Ω_{λ}
Z_v^{λ}	<ul style="list-style-type: none"> • the conditional λ-germ associated with finite configuration v and branching cell $\lambda \in \Delta^+(v)$. $Z_v^{\lambda} : \Omega^v \rightarrow \Omega_{\lambda}$ (III-2)
Z_v	<ul style="list-style-type: none"> • the collection $Z_v = (Z_v^{\lambda})_{\lambda \in \Delta^+(v)}$ (III-2.1)
$\langle N, v \rangle$	<ul style="list-style-type: none"> • the number of branching cells in the decomposition of v: $\langle N, v \rangle = \text{Card}(\Lambda(v))$ (III-3.3)

Table 4.1: Summary of notations for distributed probabilities, Chapter 4.

Chapter 5

Markov Nets and Distributed Markov Nets

So far we have studied the abstract model of probabilistic event structures. We now apply our results to particular event structures, those coming from the unfolding of *safe* and finite Petri nets.

The trace theory identifies the dynamics of a safe net with the dynamics of its unfolding, an occurrence net labelled by the safe net. By the well-known equivalence between the model of occurrence nets and the model of event structures, the randomisation of trace dynamics reduces to the construction of a probabilistic event structure. This idea is due to H. Völzer, and this is also the point of view of [7].

From a theoretical point of view as well as for practical applications, the memory-less—or Markovian—randomisation of systems is of major interest in the study of random systems. Informally, for memory-less systems, processes starting from the same state have the same futures in probability. We say that processes forget their past. The most popular model of such a system in discrete time is the model of finite (and homogeneous) Markov chains. It constitutes for instance the basis of probabilistic automata [34, 41].

Our first challenge is to bring a definition for memory-less and true-concurrent random systems, in the model of safe Petri nets. Given a safe Petri net, we constat the isomorphism between the cones of futures of two processes leading to the same marking. This formalises that the dynamics of Petri nets is intrinsically memory-less. This geometric mapping induces an isomorphism of measurable spaces on the shadows of processes. We require for a memory-less, or *homogeneous* probability, that this isomorphism respects the conditional probabilities defined on the shadows. From different states—markings—the probabilistic futures have the same probability law. This defines an equivalent of *homogeneous Markov chains*, whence the terminology of *homogeneous probabilistic nets*—or *homogeneous probability*—that we adopt. Homogeneous probabilities are characterised by an invariance property: the probabilistic future only depends on the current marking.

In the analysis of Markovian processes, one of the major tool is the so-called *Strong Markov property*. It is based on the notion of stopping time, that are particular *random* times. The important point is that stopping times, for sequential

systems, generalise the *constant* times given by the integers $0, 1, \dots$. The Markov property is also related to the notion of *shift operator*. Roughly speaking, a stopping time cuts up two pieces from each maximal process: a beginning—until the random time—and the tail. The shift operator associated with the stopping time is the punctual transformation on the canonical space Ω that forgets the beginning of the process, and keeps the tail.

Stating the Markov property for concurrent and memory-less systems is our second challenge. For concurrent systems, *constant times* have little meaning. Until a certain point, constant times can be replaced by stopping prefixes, and we introduce stopping operators to replace stopping times. Many properties of stopping times hold for stopping operators—but not all¹. We also define *shift operators* associated with stopping operators, with the same idea of keeping the tail of processes. Then we show the Strong Markov property for homogeneous probabilistic nets. We underline the formal analogy between the two statements, in the sequential and in the concurrent frameworks. The proof of the concurrent Markov property is self-contained, and implies the usual Markov property for Markov chains with usual stopping times.

The difficulty for concurrent systems is that the space of processes is not homogeneous, *w.r.t.* time translations. As a consequence, for concurrent systems, there is no natural shift operator $\Omega \rightarrow \Omega$ as for Markov chains. Hence even the only formulation of the Markov property for concurrent systems is not clear. We propose a formulation based on *homogeneous functions*, that extend the class of test functions $\Omega \rightarrow \mathbb{R}$ used with the Markov property for sequential systems.

For this study to be useful, we need to show the existence of homogeneous nets—of homogeneous probabilities as we say. The Markov chain theory gives a first example: the sequential nets. However, the more general construction that we propose does not bring homogeneous probabilities in general. Our construction is based on the distributed product of Chapter 4, which is in turn based on the decomposition of processes through *well-stopped* configurations. Due to the compositionality formula *w.r.t.* the future of the distributed product, the invariance property for homogeneity only holds, in general, for well-stopped configurations, not for all configurations. This distinction is trivial for sequential systems, but not for concurrent systems.

Modulo this restriction, all the notions introduced—homogeneous probabilities, homogeneous functions, stopping operators, shift operators—admit weaker versions, concerning well-stopped processes. This technical distinction has thus no consequence on the global way to proceed. We state the adapted version of the Markov property, that we call the *well-stopped Markov property*. A first application is found in this chapter through the study of the *embedded Markov chain* of a net. We will have several occasions in Chapters 6 and 7 to apply the Markov property for concurrent systems, following some classical applications of the Markov property for sequential systems.

The chapter contains 5 sections. Section 1, *True-concurrent randomisation of*

¹An important property of stopping times for sequential systems, the stability under $\min(\cdot, \cdot)$, does not hold for stopping operators, but holds for stopping prefixes. The translation is thus a bit ambiguous, and we shall say very carefully that stopping operators replace stopping times.

nets, formalises the randomisation of traces of safe Petri nets through probabilistic event structures. We also give the version of the cone of future adapted to occurrence nets—it was defined for event structures in Chapter 3. For simplicity, the notions related to concurrent Markovian processes are presented in Section II, *Homogeneous nets*, without reference to well-stopped configurations. Section III, *d-Homogeneous nets*, presents the weaker versions, where the invariance property of probabilities only concerns well-stopped configurations. No fundamental difference appears. Section IV, *Distributed Markov nets*, presents a construction of *d*-homogenous nets, using the distributed product introduced in Chapter 4. Finally, Section V, *Examples of distributed Markov nets*, presents explicit examples of Markov nets: we give examples of parametric families, and we analyse the case of *confusion-free* event structures.

I—True-Concurrent Randomisation of Nets

This short section states the definition and the general framework for probabilistic Petri nets. The true-concurrent randomisation of traces is defined in I-1. We have defined in Ch. 3 the cone of future of event structures. Here, in I-2, we define it for occurrence nets.

I-1 True-Concurrent Randomisation.

We consider a safe marked net \mathcal{N} , with initial marking M_0 . (\mathcal{U}, ρ) denotes the unfolding of \mathcal{N} . $(\mathcal{E}, \preceq, \#)$ denotes the event structure of finite concurrent width, canonically associated to \mathcal{U} .

We recall that the partial order of traces of executions of the net \mathcal{N} is isomorphic to the partial order of configurations of \mathcal{E} . Hence, a true-concurrent randomisation of the dynamics of \mathcal{N} is modelised by a probability measure \mathbb{P} on Ω , the set of maximal configurations of \mathcal{E} . For r a playing sequence of the net, its “probability” is given by what we called the *likelihood* of v , where v is the configuration associated to the sequence r . This probability is independent of r modulo interleaving, and defines thus what can be called a true-concurrent randomisation of the dynamics of \mathcal{N} . The probability of r to occur is given by the probability of the shadow:

$$p(v) = \mathbb{P}(\Omega(v)).$$

This definition, adapted from sequential systems to concurrent systems, is due to H. Völzer ([45]).

Since we have studied in Chapters 3 and 4 the class of locally finite event structures, we set accordingly the following definition for safe nets.

I-1.1 Definition. (*Compact nets*) A safe marked net \mathcal{N} is said to be **compact** if the unfolding \mathcal{U} of \mathcal{N} is locally finite.

(*Probabilistic net, positive net*) A **probabilistic net** is a pair $(\mathcal{N}, \mathbb{P})$, where \mathbb{P} is a probability measure on the set Ω of maximal configurations of \mathcal{U} . \mathcal{N} , or $(\mathcal{N}, \mathbb{P})$, is said to be **positive** if \mathbb{P} is \star -regular, *i.e.* if every finite trace has positive likelihood.

I-1.2 Example. (*Sequential nets*) Let \mathcal{N} be the sequential net associated to a probabilistic transition system $(S, A, x_0, (\mu_s)_{s \in S})$ (Ch. 1, III-3). Let \mathcal{U} be the unfolding of \mathcal{N} , and let \mathcal{T} be the event structure associated to \mathcal{U} , which is a tree of events. We have seen in Ch.1, III-4.6, that Ω is equipped with a probability \mathbb{P} such that the canonical Markov chain $(X_n)_{n \geq 0}$ and its dual $(Y_n)_{n \geq 1}$ are defined on Ω , and given as follows: For every $\omega \in \Omega$, if (x_0, x_1, \dots) and (t_1, t_2, \dots) are respectively the sequences of markings (states) and of transitions (actions) of the execution ω , then:

$$\forall n \geq 0, X_n(\omega) = x_n, \quad \forall n \geq 1, Y_n(\omega) = t_n.$$

We define $(\mathcal{N}, \mathbb{P})$ as the canonical probabilistic net associated with the probabilistic transition system S .

If all μ_s satisfy $\mu_s(x) > 0$ for all x connected to s in the transition system, then the likelihood p satisfies $p(e_1, \dots, e_n) > 0$ for every finite configuration (e_1, \dots, e_n) in \mathcal{T} . Hence $(\mathcal{N}, \mathbb{P})$ is a positive probabilistic net.

I-2 Cone of Future for Occurrence Nets.

I-2.1 Notations: Cuts and Markings. Recall that we denote by $\gamma(v)$ the cut $\gamma(v) = \max(v)$ associated with a finite configuration v of the unfolding (\mathcal{U}, ρ) . We denote by $m(v)$ the marking of the net \mathcal{N} associated with v , defined by $m(v) = \rho \circ \gamma(v)$. The configuration v leads to the marking $m(v)$.

I-2.2 Cone of Future. All the affirmations stated below follow easily from our study of the cone of future in event structures. We recall that $\gamma(v)$ denotes the cut associated with a finite configuration v .

Let \mathcal{U} be an occurrence net, and let \mathcal{E} be the event structure associated. For any cut c of \mathcal{U} , let \mathcal{U}^c denote the following open subset of \mathcal{U} :

$$\mathcal{U}^c = \left\{ x \in \mathcal{U} \mid \begin{array}{l} \forall b \in c, \quad \neg(x \# b), \\ \& \exists b \in c : \quad b \preceq x \end{array} \right\}. \quad (5.1)$$

For any finite configuration v of \mathcal{U} , we write \mathcal{U}^v to denote \mathcal{U}^c , where $c = \gamma(v)$ is the cut given by the maximal elements of v . We call \mathcal{U}^v or \mathcal{U}^c the **cone of future** of v or, equivalently, of $c = \gamma(v)$. The cone \mathcal{U}^v is equipped with the restricted flow relation $F|_{\mathcal{U}^v} = F \cap (\mathcal{U}^v \times \mathcal{U}^v)$, and (\mathcal{U}^v, F) is an occurrence net satisfying:

$$v \cup \mathcal{U}^v = \{x \in \mathcal{U} \mid \neg(x \# v)\}, \quad (5.2)$$

$$v \cap \mathcal{U}^v = \gamma(v) = \text{Min}_{\preceq}(\mathcal{U}^v), \quad (5.3)$$

$$\overset{\circ}{\mathcal{U}}^v = \mathcal{E}^v, \quad (5.4)$$

where \mathcal{E}^v denotes the cone of future for event structures, as defined in Ch. 3, II-1.1.

I-2.3 Remark. We insist that the cone \mathcal{U}^v contains all conditions of $\gamma(v)$, which is not implied by $\overset{\circ}{\mathcal{U}}^v = \mathcal{E}^v$.

I-2.4 Lemma. *Let (\mathcal{N}, M_0) be a safe marked net. For every reachable marking M , let $(\mathcal{U}M, \rho^M)$ denote the unfolding of (\mathcal{N}, M) . Let c be a cut of $\mathcal{U}M_0$, and let $M_1 = \gamma(c)$ be the marking of \mathcal{N} associated to c . Then there is a unique isomorphism of labelled occurrence nets:*

$$(\mathcal{U}M_0^c, \rho^{M_0}|_{\mathcal{U}M_0^c}) \rightarrow (\mathcal{U}M_1, \rho^{M_1}),$$

where $\mathcal{U}M_0^c$ denotes the cone of future of $\mathcal{U}M_0$ associated with the cut c .

Proof – We set $\mathcal{U} = \mathcal{U}M_0$ and $(\mathcal{V}, \theta) = (\mathcal{U}M_0^c, \rho^{M_0}|_{\mathcal{U}M_0^c})$. We denote by v the unique finite configuration of \mathcal{U} such that $\gamma(v) = c$. \mathcal{V} is an occurrence net, satisfying, using (5.3): $c = \text{Min}_{\preceq}(\mathcal{U}^v)$, and $\theta|_c$ is one-to-one onto M_1 . It follows that (\mathcal{V}, θ) is a folding of (\mathcal{N}, M_1) .

Conversely, let (V, r) be a folding of (\mathcal{N}, M_1) , we show that (V, r) maps to (\mathcal{V}, θ) . We have that $\text{Min}_{\preceq}(V)$ is one-to-one with M_1 . We also have that $\gamma(v) = \text{Max}_{\preceq}(v)$ in \mathcal{U} is one-to-one with M_1 . We can thus consider the concatenation of $(v, \rho^{M_0}|_v)$ and of (V, r) , obtained by taking first their disjoint union, and then identifying $\text{Max}_{\preceq}(v)$ and $\text{Min}_{\preceq}(V)$. Hence, we glue v and V along the cut c , and we denote the result by $v \sqcup_c V$.

It is straightforward to check that $v \sqcup_c V$ is a folding of (\mathcal{N}, M_0) . There is thus a unique morphism of foldings $f : v \sqcup_c V \rightarrow \mathcal{U}$ (Ch. 1, Th. II-4.3). Every node $f(x)$ is compatible with v , hence $f(x) \in v \cup \mathcal{U}^v$ by (5.2). We have $f(v) = v$ and thus $f(\text{Min}_{\preceq}(V)) = c$. According to Lemma II-4.7 of Ch. 1, the morphism f is injective. It implies that $f(V) \subseteq \mathcal{U}^v$, since $v \cap \mathcal{U}^v = \emptyset$ by (5.3). This shows that we get by restriction a folding: $f|_V : V \rightarrow \mathcal{U}^v = \mathcal{V}$.

We have shown that (\mathcal{V}, θ) satisfies the universal property of $\mathcal{U}M_1$, and by the uniqueness of the unfolding (Ch.1, II-4.3), it implies the statement of the lemma. \square

II—Homogeneous Nets.

In this section, we give a formalisation for memory-less true-concurrent processes, and we show the Markov property in this framework. Following as much as we can the theory of Markov chains, we define stopping operators for concurrent systems, a notion closely related to the classical notion (due to Doob) of stopping times for stochastic processes, also called optional times.

In II-1, we define an equivalence relation on cuts of labelled occurrence nets that will be useful throughout the whole study of Markov nets. We define the memory-less, or homogeneous probabilities in II-2. An other notion for stating the Markov property is the notion of stopping operator, introduced in II-3. It is then time in II-4 to compare our definitions with their analogous for Markov chains. After recalling the statement of the Markov property for Markov chains, we establish the analogous for concurrent systems in II-5.

II-1 Equivalence of \parallel -Cliques in Unfoldings.

II-1.1 Definition. (*Equivalence of \parallel -cliques and of cuts.*) We assume that (\mathcal{U}, ρ) is the unfolding of a safe finite marked net \mathcal{N} . Let c, c' be two \parallel -cliques of conditions of \mathcal{U} . We say that c and c' are **equivalent**, and we denote it by:

$$c \cong c',$$

if $\rho(c) = \rho(c')$.

The restriction $\rho|_d$ of ρ to a cut d is one-to-one onto $\rho(d)$: this is a consequence of Ch 1, II-4.5 and Ch. 1, II-3.8. Hence the same holds for \parallel -cliques. It follows that, if $c \cong c'$, there is an isomorphism of labelled occurrence nets:

$$(c, \rho|_c) \rightarrow (c', \rho|_{c'}). \quad (5.5)$$

If the clique is actually a cut, *i.e.* if c is a maximal \parallel -clique, the isomorphism can be extended to the cones of future as follows.

II-1.2 Lemma. *Let (\mathcal{U}, ρ) be the unfolding of a safe marked net (\mathcal{N}, M_0) . Let c and c' be two equivalent cuts of \mathcal{U} . Then there is a unique isomorphism of labelled occurrence nets:*

$$(\mathcal{U}^c, \rho|_{\mathcal{U}^c}) \rightarrow (\mathcal{U}^{c'}, \rho|_{\mathcal{U}^{c'}}),$$

that extends (5.5).

Proof— Let $M = \rho(c) = \rho(c')$ be the marking of \mathcal{N} associated with c and c' . Let \mathcal{V} denote the unfolding of (\mathcal{N}, M) . According to I-2.4, there are unique isomorphisms of foldings $f_1 : \mathcal{U}^c \rightarrow \mathcal{V}$ and $f_2 : \mathcal{U}^{c'} \rightarrow \mathcal{V}$, thus a unique isomorphism $g : \mathcal{U}^c \rightarrow \mathcal{U}^{c'}$. \square

Intuitively, the lemma is clear: if a same marking is reached by two processes, the possible continuations of the processes are the same. A continuation of one process fits the other process. This formalises the absence of memory of the dynamics of Petri nets. We will now associate this fact with invariance properties of probability laws.

II-1.3 Convention. Let (\mathcal{U}, ρ) and (\mathcal{U}', ρ') be two isomorphic labelled occurrence nets. For any configuration v of \mathcal{U} with image v' in \mathcal{U}' , it follows directly from (5.2) and (5.3) that \mathcal{U}^v and $\mathcal{U}'^{v'}$ are isomorphic, and I-2.4 implies that the isomorphism is unique. Hence taking the cone of future is a congruence *w.r.t.* isomorphism of labelled occurrence nets.

Therefore we will not distinguish between isomorphic labelled occurrence nets, and the operation of future is defined for the class of isomorphic labelled occurrence nets. With this convention, II-1.2 becomes: If v, v' are two finite configurations of the unfolding \mathcal{U} of a safe net, we have:

$$\gamma(v) \cong \gamma(v') \Rightarrow \mathcal{U}^v = \mathcal{U}^{v'}.$$

This implies in particular, for v, v' finite in \mathcal{U} with $\gamma(v) \cong \gamma(v')$:

$$\mathcal{W}^v = \mathcal{W}^{v'}, \quad \Omega^v = \Omega^{v'}, \quad \text{etc.} \quad (5.6)$$

II-2 Homogeneous Probabilities.

Let $(\mathcal{N}, M_0, \mathbb{P})$ be a marked probabilistic net. \mathcal{U} and \mathcal{E} denote the unfolding and its event structure. Let p denote the likelihood of \mathbb{P} , and let v be a finite configuration with $p(v) > 0$. The probabilistic future $(\mathcal{E}^v, \mathbb{P}^v)$ is well-defined (Ch. 4, II-3.1). Hence, if M denotes the marking reached by v in \mathcal{N} from M_0 , $(\mathcal{N}, M, \mathbb{P}^v)$ is a probabilistic net. The probability \mathbb{P}^v is defined on Ω^v , whose elements are the maximal configurations of \mathcal{U}^v . Recall that \mathcal{U}^v is defined modulo isomorphism of labelled occurrence nets. Let v' be an other finite configuration with $p(v') > 0$. As stated by (5.6), if $\gamma(v) \cong \gamma(v')$, then \mathbb{P}^v and $\mathbb{P}^{v'}$ are defined on the same probability space Ω^v . In general \mathbb{P}^v and $\mathbb{P}^{v'}$ are different.

II-2.1 Definition. (*Homogeneous probability*) Let $(\mathcal{N}, \mathbb{P})$ be a probabilistic net. Let \mathcal{U} denote the unfolding of \mathcal{N} , and p the likelihood of \mathbb{P} . We say that $(\mathcal{N}, \mathbb{P})$ is **homogeneous**, or **strongly homogeneous**, or shortly that \mathbb{P} is homogeneous, if \mathbb{P} satisfies the following property:

$$\forall v, v' \in \mathcal{W}_0, \quad p(v), p(v') > 0, \quad \gamma(v) \cong \gamma(v') \Rightarrow \mathbb{P}^v = \mathbb{P}^{v'}. \quad (5.7)$$

Equivalently, without the identification modulo isomorphism, homogeneity is formulated by: \mathbb{P} is homogeneous if and only if for all finite configurations v, v' such that $\gamma(v) \cong \gamma(v')$, the two random variables:

$$\left\{ \begin{array}{l} \Omega(v) \rightarrow \Omega^v \\ \omega \mapsto \omega \ominus v \end{array} \right\}, \quad \left\{ \begin{array}{l} \Omega(v') \rightarrow \Omega^{v'} \\ \omega \mapsto \omega \ominus v' \end{array} \right\},$$

induce probability laws in Ω^v and $\Omega^{v'}$, that are images one from the other under the natural isomorphism $\Omega^v \rightarrow \Omega^{v'}$.

This definition formalises the notion of memory-less randomisation of traces of safe nets. The construction of homogeneous probabilities is the topic of next sections in this chapter. For the moment, we study their properties *a priori*.

II-2.2 Probability from a Reachable Marking. If a probability \mathbb{P} on Ω is homogeneous, the probability space $(\Omega^v, \mathcal{F}^v, \mathbb{P}^v)$ only depends on the class of $\gamma(v)$ modulo the relation \cong , *i.e.* on the marking $m(v) = \rho \circ \gamma(v)$.

It implies that, if M is a marking satisfying $M = m(v)$, with v a finite configuration with positive likelihood, the probability space $(\Omega^v, \mathcal{F}^v, \mathbb{P}^v)$ does not depend on v . We denote it by $(\Omega^M, \mathcal{F}^M, \mathbb{P}^M)$. We have thus:

$$\forall v \in \mathcal{W}_0, \quad (\Omega^v, \mathcal{F}^v, \mathbb{P}^v) = (\Omega^{\gamma(v)}, \mathcal{F}^{\gamma(v)}, \mathbb{P}^{\gamma(v)}) = (\Omega^{m(v)}, \mathcal{F}^{m(v)}, \mathbb{P}^{m(v)}).$$

II-3 Stopping Operators.

We now define a notion closely related to the notion of stopping times, classical in the study of stochastic processes.

II-3.1 Definition. (Stopping Operator) We say that a mapping $V : \Omega \rightarrow \mathcal{W}$ is a **stopping operator** if V is measurable, and if V satisfies the two following properties:

1. $\forall \omega \in \Omega, \quad V(\omega) \subseteq \omega,$
2. $\forall \omega, \omega' \in \Omega, \quad \omega \supseteq V(\omega') \Rightarrow V(\omega) = V(\omega').$

II-3.2 Example. Let P be an intrinsic prefix of \mathcal{E} (Ch. 2, III-1.2). Then $\pi_P : \Omega \rightarrow \mathcal{W}, \omega \mapsto \omega \cap P$ is a stopping operator. In particular, π_B is a stopping operator for every stopping prefix B .

II-3.3 Example. Not every stopping operator has the form $V(\omega) = P \cap \omega$. The event structure \mathcal{E} of Figure 5.1 has two maximal configurations, $\omega_1 = a \oplus c$ and $\omega_2 = a \oplus b$. Let $V(\omega_1) = \omega_1$ and $V(\omega_2) = b$. If V is defined through a prefix P , then $P = V(\omega_1) \cup V(\omega_2) = \mathcal{E}$. It implies that $V = \text{Id}_\Omega$, which is not.




Figure 5.1: Stopping operator not defined through a prefix.

II-3.4 Example. (Local hitting operator) Let λ be a branching cell of \mathcal{E} . We have defined in Ch. 4, IV-1.1, the random variable $V^\lambda : \mathcal{H}^\lambda \rightarrow \mathcal{X}_0$ by:

$$\forall \omega \in \mathcal{H}^\lambda, \quad V^\lambda(\omega) = \min\{y \in \mathcal{X}_0, y \subseteq \omega \mid \lambda \in \Delta^+(y)\}.$$

We extend V^λ to Ω by setting $V^\lambda(\omega) = \omega$ if $\omega \notin \mathcal{H}^\lambda$, i.e. if $\lambda \notin \Lambda(\omega)$. We have seen that the restriction $V^\lambda|_{\mathcal{H}^\lambda}$ is measurable, thus $V^\lambda : \Omega \rightarrow \mathcal{W}$ is measurable. We have $V^\lambda(\omega) \subseteq \omega$ by construction. Point 2 of II-3.1 has been shown in Ch. 4, IV-1.1, on \mathcal{H}^λ , and extends to Ω . We call V^λ the hitting operator associated with λ .

II-4 The Strong Markov Property for Homogeneous Markov Chains.

We quickly recall the definition and properties of stopping times (also called *optional times*) for sequential systems. In particular, we recall the Strong Markov property for finite Markov chains, following [37]. We show that, for sequential systems, stopping operators are one-to-one with stopping times.

II-4.1 Stopping Times. Let P be a stochastic matrix on a finite set S , and let $(X_n)_{n \geq 0}$ be the canonical homogeneous Markov chain defined on the probability space $\Omega = \prod_{n \geq 0} S$ (Ch. 1, III-1). We denote by \mathcal{F}_n the σ -algebra $\mathcal{F}_n = \langle X_0, \dots, X_n \rangle$. We denote by \mathbb{P}_ν the probability on Ω associated with an initial probability ν on S . We write \mathbb{P}_x for \mathbb{P}_{δ_x} , where δ_x is the Dirac measure of an element $x \in S$. We write $\mathbb{E}_\nu(\cdot)$ and $\mathbb{E}_x(\cdot)$ to denote the expectations *w.r.t.* probabilities \mathbb{P}_ν and \mathbb{P}_x .

We denote by $\overline{\mathbb{N}}$ the set $\mathbb{N} \cup \{\infty\}$. A measurable random variable $T : \Omega \rightarrow \overline{\mathbb{N}}$ is said to be a **stopping time** *w.r.t.* the filtration $(\mathcal{F}_n)_{n \geq 0}$, if the following holds:

$$\forall n \geq 0, \quad \{T = n\} \in \mathcal{F}_n.$$

We recall that $\{T = n\}$ stands for the subset of Ω : $\{\omega \in \Omega \mid T(\omega) = n\}$. We denote by \mathcal{F}_T the sub- σ -algebra of subsets $A \in \mathcal{F}$ such that:

$$\forall n \geq 0, \quad A \cap \{T = n\} \in \mathcal{F}_n. \quad (5.8)$$

For $T : \Omega \rightarrow \overline{\mathbb{N}}$ a stopping time, we denote by $X_T : \Omega \rightarrow S$ the random variable defined $X_T(\omega) = X_{T(\omega)}(\omega)$. X_T gives the value of the process X “at time T ”.

II-4.2 Shift Operator. We define the shift operator as the pointwise transformation $\theta : \Omega \rightarrow \Omega$ by $\theta(s_0, s_1, \dots) = (s_1, s_2, \dots)$. We denote by θ_n the n^{th} power of θ : $\theta_0 = \text{Id}_\Omega$, $\theta_n = \theta \circ \theta_{n-1}$ for $n \geq 1$.

We denote by θ_T the transformation $\Omega \rightarrow \Omega$ given by $\theta_T(\omega) = \theta_{T(\omega)}(\omega)$.

II-4.3 The Strong Markov Property. ([37], Th. 3.5, p. 23) For every positive random variable h on (Ω, \mathcal{F}) , initial measure ν and stopping time T :

$$\mathbb{E}_\nu(h \circ \theta_T \mid \mathcal{F}_T) = \mathbb{E}_{X_T}(h), \quad \mathbb{P}_\nu\text{-a.s.}$$

By convention the two members vanish on $\{T = \infty\}$. The right member is the composite mapping of $\omega \mapsto X_T(\omega)$ and $x \rightarrow \mathbb{E}_x(h)$. Both members are \mathcal{F}_T -measurable mappings.

II-4.4 Commentary on the Markov Property. The meaning of the Markov property is more clear if $T : \Omega \rightarrow \overline{\mathbb{N}}$ is constant, say equal to n . It becomes: $\mathbb{E}(h \circ \theta_n \mid \mathcal{F}_n) = \mathbb{E}_{X_n}(h)$. Recall that $\mathcal{F}_n = \langle X_1, \dots, X_n \rangle$ contains the information

until time n , with final state X_n . The left hand member of the Markov equality represents the mean value of the function h on the truncated variable (X_n, X_{n+1}, \dots) , conditional on (X_1, \dots, X_n) , that is:

$$\mathbb{E}(h(X_n, X_{n+1}, \dots) \mid X_1, \dots, X_n) .$$

The right hand member of the Markov equality is the mean value of $h(Z_0, Z_1, \dots)$, where $(Z_p)_{p \geq 0}$ is a Markov chain on the same space and with same transition matrix than $(X_p)_{p \geq 0}$, but with initial state X_n . Although it may seem confusing, we can take the Markov chain $(Z_p)_p = (X_p)_p$, under probability \mathbb{P}_{X_n} . Indeed since we look at mean values, only the probability laws are important, not the random variables themselves.

The Markov property states the equality of the two members for all functions h , seen as test functions. Forgetting X_0, \dots, X_{n-1} and keeping only X_n , we actually do not lose information in the law of the future process (X_n, X_{n+1}, \dots) . This statement for constant times constitutes the Weak Markov property. The Strong property states that things go the same way if T is a stopping time, *i.e.* if T ranges over a large class of random times.

II-4.5 Formulation for Sequential Nets. Let $\mathcal{N} = (S, A, F, s_0)$ be the sequential net associated with a probabilistic transition system $(S, A, s_0, (\mu_s)_{s \in S})$ (Ch. 1, III-3). Let (\mathcal{U}, ρ) denote the unfolding of the marked net \mathcal{N} , and let \mathcal{T} be the event structure associated with \mathcal{U} . We recall that \mathcal{T} is a tree of events. Using the Markov chain theory, we equip the space (Ω, \mathcal{F}) with a probability \mathbb{P} and two Markov chains $(X_n)_{n \geq 0}$ and $(Y_n)_{n \geq 1}$, that respectively describe the successive markings and transitions of the net \mathcal{N} , and we have (Ch. 1, III-4.5):

$$\forall n \geq 1, \quad \mathcal{F}_n = \langle Y_1, \dots, Y_n \rangle = \langle X_0, \dots, X_n \rangle .$$

If T is a stopping time for $(X_n)_n$, since $\langle X_0 \rangle$ is the trivial σ -algebra $\{\emptyset, \Omega\}$, we have that $T \geq 1$ \mathbb{P} -*a.s.*, excepted for the trivial case $T = 0$, \mathbb{P} -*a.s.* Hence we assume without loss of generality that $T \geq 1$ \mathbb{P} -*a.s.*, and thus stopping times for $(X_n)_n$ and $(Y_n)_n$ coincide. We just call them stopping times $T : \Omega \rightarrow \overline{\mathbb{N}}$.

II-4.6 Stopping Operator Associated with a Stopping Time. Still in the case of a sequential net, let $T : \Omega \rightarrow \overline{\mathbb{N}}$ be a stopping time. We define the mapping $V^T : \Omega \rightarrow \mathcal{W}$ as follows. Each $\omega \in \Omega$ admits a unique decomposition through events $\omega = (Y_1, Y_2, \dots)$. We set:

$$V^T(\omega) = (Y_1, \dots, Y_{T(\omega)}) \quad \text{if } T(\omega) < \infty, \quad V^T(\omega) = \omega \quad \text{otherwise.}$$

Conversely, if $V : \Omega \rightarrow \mathcal{W}$ is a stopping operator (II-3.1), we define the integer random variable $T^V : \Omega \rightarrow \overline{\mathbb{N}}$ as follows. If $V(\omega)$ is finite, and with $\omega = (Y_1, \dots)$:

$$T^V(\omega) = n, \quad n \text{ such that } V(\omega) = (Y_1, \dots, Y_n) ,$$

and $T^V(\omega) = \infty$ if $V(\omega) = \omega$.

Hence we map the canonical total order $\{1, \dots, T(\omega)\}$ onto the “concrete” order $\{Y_1(\omega), \dots, Y_{T(\omega)}\}$ in the unfolding, and this for each $\omega \in \Omega$.

II-4.7 Proposition. *The mappings $V \rightarrow T^V$ and $T \rightarrow V^T$ are one-to-one between stopping operators and stopping times, and inverse one from the other.*

Proof – We denote by $\pi_n : \Omega \rightarrow \mathcal{W}$ the mapping $\pi_n = (Y_1, \dots, Y_n)$, and we have $\mathcal{F}_n = \langle \pi_n \rangle$. Let T be a stopping time. V^T is measurable since T is measurable, and $V^T(\omega) \subseteq \omega$ for all $\omega \in \Omega$. Let $\omega, \omega' \in \Omega$ such that $\omega \supseteq V^T(\omega')$, and let $n = T(\omega')$. Then $\pi_n(\omega) = \pi_n(\omega')$. The subset $\{T = n\}$ is \mathcal{F}_n -measurable and thus π_n -saturated. It follows that $T(\omega) = n$, and thus $V^T(\omega) = V^T(\omega')$. Hence V^T is a stopping operator.

Conversely, let V be a stopping operator. T^V is then an integer random variable. Let $n \geq 1$, we show that $\{T = n\}$ is \mathcal{F}_n -measurable, *i.e.* π_n -saturated. Let $\omega \in \Omega$ such that $T(\omega) = n$, and let $\omega' \in \Omega$ with $\pi_n(\omega) = \pi_n(\omega')$. Then $\omega \supseteq V(\omega')$, which implies $V(\omega') = V(\omega)$ and thus $T^V(\omega') = T^V(\omega)$. This shows that T^V is a stopping time.

The mappings $V \rightarrow T^V$ and $T \rightarrow V^T$ are obviously inverse one from the other. \square

II-5 The Strong Markov Property for Concurrent Homogeneous Systems.

We now return to unfoldings of general safe marked nets. We try to adapt the theory of Markov processes to our framework, as we left it at the end of II-3.

II-5.1 Shift Operator. Let $V : \Omega \rightarrow \mathcal{W}$ be a stopping operator. We define the random variable γ_V as follows:

$$\forall \omega \in \Omega \quad \begin{cases} \gamma_V(\omega) = \gamma(V(\omega)) & \text{if } V(\omega) \in \mathcal{W}_0, \\ \gamma_V(\omega) & \text{is undefined otherwise.} \end{cases}$$

To define a shift operator we use the cancellation operation \ominus , defined in Ch. 3, II-2.1. For V a stopping operator, we define the **shift operator** θ_V by:

$$\forall \omega \in \Omega, \quad \theta_V(\omega) = \omega \ominus V(\omega). \quad (5.9)$$

For each $\omega \in \Omega$ such that $V(\omega)$ is finite, we have $\theta_V(\omega) \in \mathcal{U}^{\gamma_V(\omega)}$, which is written:

$$\theta_V \in \mathcal{U}^{\gamma_V}.$$

We remark that the shift operator of (5.9) does not act on Ω . Recall that in the sequential case, we have defined a shift operator $\theta : \Omega \rightarrow \Omega$ (II-4.2). We have then defined for h a measurable function $\Omega \rightarrow \mathbb{R}$ and T a stopping time, the composite functions $h \circ \theta$, $h \circ \theta_2$, *etc.*, and $h \circ \theta_T$. We can relate this facility to the product form of the space Ω . The concurrent framework is different: $h \circ \theta$ is not defined, neither is $h \circ \theta_V$, since $\theta_V : \Omega \rightarrow \mathcal{U}^{\gamma_V}$ changes the space, ω by ω .

This problem is indeed one of the main difference that we encounter between concurrent and sequential systems. Unlike the classical framework of dynamic systems theory [42], there is no semi-group of operators $(\theta_n, n \geq 0)$. We propose thus to slightly extend the framework, introducing the following objects.

II-5.2 Definition. (Homogeneous functions) Denote by \mathcal{D} the set of equivalence classes of cuts in \mathcal{U} (which identify with reachable markings). Let g be a collection of measurable mappings $g = (g^c)_{c \in \mathcal{D}}$, with $g^c : \Omega^c \rightarrow \mathbb{R}$. We say that g is a **homogeneous function on Ω** .

Equivalently, a homogeneous function is given by a collection $g = (g^c)_c$ of mappings $g^c : \Omega^c \rightarrow \mathbb{R}$, where c ranges over the cuts of \mathcal{U} , satisfying:

$$c \cong c' \Rightarrow g^c = g^{c'} .$$

II-5.3 Non-negative and Integrable Homogeneous Functions. We say that g is non-negative if all g^c are non-negative. If \mathbb{P} is a probability on Ω , we say that g is integrable if for all cut c of \mathcal{U} , if $\mathbb{P}(c) > 0$ then g^c is integrable *w.r.t.* the probability \mathbb{P}^c .

II-5.4 The \mathcal{F}_V σ -Algebra. We adapt the σ -algebra \mathcal{F}_T for stopping times (II-4.1, Eq. (5.8)) as follows. Let V be a stopping operator. We define the σ -algebra \mathcal{F}_V as the collection of subsets $A \in \mathcal{F}$ such that:

$$\forall \omega, \omega' \in \Omega, \quad \omega \in A, \quad \omega' \supseteq V(\omega) \Rightarrow \omega' \in A . \tag{5.10}$$

It is readily checked that \mathcal{F}_V is indeed a σ -algebra. We see as in II-4.7, for \mathcal{N} a sequential net and T a stopping time, that $\mathcal{F}_T = \mathcal{F}_{V^T}$, where V^T is the stopping operator associated with T (II-4.6).

We now state a lemma before establishing the Markov property.

II-5.5 Lemma. *Let V be a stopping operator, and let:*

$$U = \{u \in \mathcal{W}_0 \mid \exists \omega \in \Omega, V(\omega) = u\} .$$

1. For all $u \in U$, we have: $\{V = u\} = \Omega(u)$.
2. Let ϕ be a \mathcal{F}_V -measurable real function. There is a function $\hat{\phi} : U \rightarrow \mathbb{R}$ such that:

$$\forall u \in U, \quad \phi|_{\Omega(u)} = \hat{\phi}(u) .$$

Proof— 1. Let $u \in U$. Every ω such that $V(\omega) = u$ contains u , hence $\{V = u\} \subseteq \Omega(u)$. Conversely, let $\omega \in \Omega(u)$. By definition there is an element $\omega' \in \Omega$ such that $V(\omega') = u$. Then $\omega \supseteq V(\omega')$, and it implies $V(\omega) = V(\omega') = u$ by point 2 in II-3.1. We have shown that $\{V = u\} = \Omega(u)$.

2. Let $\phi : \Omega \rightarrow \mathbb{R}$ be \mathcal{F}_V -measurable. Let $u \in U$ and $\omega_0 \in \Omega$ with $V(\omega_0) = u$. Let A be the subset $A = \{\phi = \phi(\omega_0)\}$, \mathcal{F}_V -measurable since ϕ is \mathcal{F}_V -measurable. Applying (5.10) to A , we get that $\Omega(u) \subseteq A$. It follows that ϕ is constant on $\Omega(u)$. Remark that we have: $p(u)\hat{\phi}(u) = \int_{\Omega(u)} \phi \, d\mathbb{P}$. \square

II-5.6 Theorem. (*Strong Markov property for concurrent systems*) Let \mathbb{P} be a homogeneous probability on (Ω, \mathcal{F}) . For every stopping operator V , for any positive homogeneous function $g = (g^c)_{c \in \mathcal{D}}$, we have:

$$\mathbb{E}(g \circ \theta_V | \mathcal{F}_V) = \mathbb{E}^{\gamma_V}(g^{\gamma_V}), \quad \mathbb{P}\text{-a.s.} \tag{5.11}$$

The two members vanish by convention on $\{V \notin \mathcal{W}_0\}$. The function $g \circ \theta_V$ vanishes on $\{V \notin \mathcal{W}_0\}$, and is defined on $\{V \in \mathcal{W}_0\}$ by:

$$g \circ \theta_V(\omega) = g^{\gamma_V}(\theta_V(\omega)) = g^{\gamma_V}(\omega \ominus V(\omega)). \tag{5.12}$$

The right hand member in (5.11) is the composite of $\omega \mapsto \gamma(\omega)$ and $c \mapsto \mathbb{E}^c(g^c)$. \mathbb{E}^c denotes the expectation in the probability space $(\Omega^c, \mathcal{F}^c, \mathbb{P}^c)$.

II-5.7 Remark. Observe the need for homogeneous functions to formulate the Markov property. If $g : \Omega \rightarrow \mathbb{R}$ is a usual function, then $g(\omega \ominus \theta_V(\omega))$ is not defined.

Proof— Let ϕ be a positive, \mathcal{F}_V -measurable function. We have to compute $\mathbb{E}(g \circ \theta_V \phi)$. We assume without loss of generality that ϕ vanishes on $\{V \notin \mathcal{W}_0\}$. Let U denote the set of finite values of V : $U = \mathcal{W}_0 \cap V(\Omega)$. Since U is at most countable, and $g, \phi \geq 0$, we have:

$$\mathbb{E}(g \circ \theta_V \phi) = \sum_{u \in U} \mathbb{E}(\mathbf{1}_{\{V=u\}} \phi g \circ \theta_V).$$

According to lemma II-5.5, $\{V = u\} = \Omega(u)$. We also denote by $\hat{\phi}$ the function such that $\phi|_{\Omega(u)} = \hat{\phi}(u)$ given by lemma II-5.5. Using (5.12), we get:

$$\begin{aligned} \mathbb{E}(g \circ \theta_V \phi) &= \sum_{u \in U} \hat{\phi}(u) \int_{\Omega(u)} g^u(\omega \ominus u) \, d\mathbb{P}(\omega) \\ &= \sum_{u \in U} p(u) \hat{\phi}(u) \mathbb{E}^u(g^u), \end{aligned}$$

where p denotes the likelihood of \mathbb{P} , defined by $p(v) = \mathbb{P}(\Omega(v))$.

Let $u, u' \in U$ such that $\gamma(u) \cong \gamma(u')$. We have $g^u = g^{u'}$ since g is a homogeneous function, and $\mathbb{P}^u = \mathbb{P}^{u'}$ since \mathbb{P} is homogeneous. It follows that $\mathbb{E}^u(g^u) = \mathbb{E}^{u'}(g^{u'})$. Recall the notation \mathcal{D} for the set of equivalence classes of cuts, and set $\bar{\gamma}_V$ the random variable in \mathcal{D} , the equivalence class of γ_V . We have:

$$\begin{aligned} \mathbb{E}(g \circ \theta_V \phi) &= \sum_{u \in U} \sum_{\substack{c \in \mathcal{D} \\ \bar{\gamma}(u)=c}} p(u) \hat{\phi}(u) \mathbb{E}^u(g^u) \\ &= \sum_{c \in \mathcal{D}} \mathbb{E}^c(g^c) \sum_{\substack{u \in U \\ \bar{\gamma}(u)=c}} p(u) \hat{\phi}(u). \end{aligned} \tag{5.13}$$

Since ϕ equals $\hat{\phi}(u)$ on $\Omega(u) = \{V = u\}$, we compute the right sum in (5.13) as follows:

$$\begin{aligned} \forall c \in \mathcal{D}, \quad \sum_{\substack{u \in U \\ \bar{\gamma}(u)=c}} p(u)\hat{\phi}(u) &= \sum_{\substack{u \in U \\ \bar{\gamma}(u)=c}} \int_{V=u} \phi \, d\mathbb{P} \\ &= \int_{\bar{\gamma}_V=c} \phi \, d\mathbb{P} \\ &= \mathbb{P}(\bar{\gamma}_V = c)\mathbb{E}(\phi | \bar{\gamma}_V = c). \end{aligned} \quad (5.14)$$

Using (5.13) and (5.14), we get:

$$\begin{aligned} \mathbb{E}(g \circ \theta_V \phi) &= \sum_{c \in \mathcal{D}} \mathbb{P}(\bar{\gamma}_V = c)\mathbb{E}(\phi \mathbb{E}^c(g^c) | \bar{\gamma}_V = c) \\ &= \sum_{c \in \mathcal{D}} \mathbb{P}(\bar{\gamma}_V = c)\mathbb{E}(\phi \mathbb{E}^{\gamma_V}(g^{\gamma_V}) | \bar{\gamma}_V = c) \\ &= \mathbb{E}(\phi \mathbb{E}^{\gamma_V}(g^{\gamma_V})). \end{aligned}$$

It follows that $\mathbb{E}(g \circ \theta_V | \mathcal{F}_V) = \mathbb{E}^{\gamma_V}(g^{\gamma_V})$. \square

II-5.8 Application to Sequential Systems. Assume that \mathcal{N} is a sequential net associated with a transition system as in I-1.2. \mathbb{P} is the probability that comes with the two Markov chains $(X_n)_n$ and $(Y_n)_n$ of states and transitions of the system. Assume that \mathbb{P} is a homogeneous probability in the sense of II-2.1 (this will be shown in IV-3). Then we show that the usual Strong Markov property II-4.3 holds for $(Y_n)_{n \geq 1}$. It implies also the Strong Markov property for $(X_n)_n$, since $\langle X_0, \dots, X_n \rangle = \langle Y_1, \dots, Y_n \rangle$.

Proof – Let h be a positive measurable function and let T be a stopping time. $h : \Omega \rightarrow \mathbb{R}$ is defined on the product space $S^{\mathbb{N}}$, where A is the set of actions of the system. For each $c \in \mathcal{D}$, there is an injection $\Omega^c \rightarrow A^{\mathbb{N}}$. We set $g^c = h|_{\Omega^c}$, and $g = (g^c)_{c \in \mathcal{D}}$ defines a homogeneous function (II-5.2). Let V denote the stopping operator V^T associated with T (II-4.6). Then we have $\theta_V = \theta_T$, where θ_T is the shift operator of II-4.2, and for all $\omega = (e_1, \dots)$ we have:

$$g \circ \theta_V(\omega) = g^{\gamma_V(\omega)}(e_{T+1}, \dots) = h(e_{T+1}, \dots) = h \circ \theta_T(\omega).$$

Hence we have: $g \circ \theta_V = h \circ \theta_T$. We have already observed that $\mathcal{F}_V = \mathcal{F}_T$. We apply the Markov property for concurrent systems (II-5.6) to get:

$$\begin{aligned} \mathbb{E}(h \circ \theta_T | \mathcal{F}_T) &= \mathbb{E}(g \circ \theta_V | \mathcal{F}_V) \\ &= \mathbb{E}^{\gamma_V}(g^{\gamma_V}) \\ &= \mathbb{E}^{X_T}(h), \end{aligned}$$

the later since $\gamma_V = X_T$. The two members vanish on $\{V \notin \mathcal{W}_0\} = \{T = \infty\}$. This is the Markov property II-4.3 for the stopping time T . \square

III— d -Homogeneous Nets.

The construction that we propose for probabilities, the distributed product, is based on well-stopped configurations. It appears that, in general, the distributed product that one defines is not homogeneous. Does it mean that the study of Section II is useless?

We take the study of Section II, and we replace all occurrences of the words “finite configuration” by “finite well-stopped configuration”. This defines in particular d -homogeneous probabilities. Then all results concerning homogeneous probabilities hold, with this weak form, and in particular the Markov property. A novelty is introduced with the *embedded Markov chain*: We show that the normal decomposition, defined in Chapter 3, defines a Markov chain that characterises the process. The study of the embedded Markov chain is our first application of the Markov property; others will be found in next chapter.

Although the construction of distributed product does not provide a homogeneous net in general, we will see in next section that it always gives a d -homogeneous net. This justifies the restriction of our study to d -homogeneous nets.

In III-1 and III-2 we state the above results, with the well-stopped configurations restriction. The embedded Markov chain is studied in III-3.

III-1 d -Homogeneous Nets.

By definition, homogeneous probabilities induce the same probability in the futures of any two finite configurations v, v' such that $\gamma(v) \cong \gamma(v')$ (II-2.1). For a notion adapted to the distributed probabilities, based on well-stopped configurations, we only require the invariance of the probabilities for v, v' finite and well-stopped.

III-1.1 Definition. (*d -homogeneous net, well-stopped cuts and markings*) Let $(\mathcal{N}, \mathbb{P})$ be a compact probabilistic net. Let \mathcal{U} be the unfolding of \mathcal{N} , and let p denote the likelihood of \mathbb{P} . We say that $(\mathcal{N}, \mathbb{P})$ is **d -homogeneous** if the following property holds:

$$\forall v, v' \in \mathcal{X}_0, p(v), p(v') > 0, \quad \gamma(v) \cong \gamma(v') \Rightarrow \mathbb{P}^v = \mathbb{P}^{v'}. \quad (5.15)$$

Say that a cut c is **well stopped**, if the configurations v such that $c = \gamma(v)$ is well-stopped. An equivalence class of cut c is said to be **well-stopped** if there is a $v \in \mathcal{X}_0$ with $\gamma(v)$ within class c . We denote by \mathcal{D}' the set of well-stopped classes of cuts.

In a similar way, we say that a marking M is (positive) **well-stopped** if there is a (positive) well-stopped configuration leading to M .

III-1.2 Remark. We insist that d -homogeneous nets are *compact*.

If \mathbb{P} is d -homogeneous, $c \rightarrow \mathbb{P}^c$ is a congruence *w.r.t.* the equivalence relation \cong restricted to well-stopped cuts. If M is a positive well-stopped marking, we write $(\Omega^M, \mathcal{F}^M, \mathbb{P}^M)$ to denote the probabilistic future $(\Omega^v, \mathcal{F}^v, \mathbb{P}^v)$, for any configuration v finite, positive and well-stopped, leading to M .

III-2 Well-Stopped Markov Property.

The stopping operators defined in II-3.1 can be well-stopped. A Markov property holds then *w.r.t.* d -homogeneous probabilities.

III-2.1 Definition. (*Well-stopping operator*) We say that $V : \Omega \rightarrow \mathcal{X}$ is a **well-stopping operator**, or that V is **well-stopped**, if $V : \Omega \rightarrow \mathcal{W}$ is a stopping operator (II-3.1).

Hence a well-stopping operator is a stopping operator, whose value $V(\omega)$ is maximal or is finite and well-stopped.

III-2.2 Shift Operator. We use the above notations for the shift operator θ_V (II-5.1) and the σ -algebra \mathcal{F}_V (II-5.4). We have $\theta_V(\omega) = \omega \ominus V(\omega)$.

III-2.3 Definition. (*d -homogeneous functions*) We define a **d -homogeneous function**, as a collection $(g^c)_{c \in \mathcal{D}'}$ of measurable mappings $g^c : \Omega \rightarrow \mathbb{R}$, with c ranging over \mathcal{D}' (by contrast with homogeneous functions where $c \in \mathcal{D}$, II-5.2). Equivalently, g is given by a collection $(g^v)_{v \in \mathcal{X}_0}$, such that:

$$\forall v, v' \in \mathcal{X}_0, \quad \gamma(v) \cong \gamma(v') \Rightarrow g^v = g^{v'} .$$

III-2.4 Theorem. (*Well-stopped Markov property*) Let $(\mathcal{N}, \mathbb{P})$ be a d -homogeneous probabilistic net. For all d -homogeneous positive function $g = (g^c)_{c \in \mathcal{D}'}$, and for all well stopping operator V , the following holds:

$$\mathbb{E}(g \circ \theta_V | \mathcal{F}_V) = \mathbb{E}^{\gamma^V}(g^{\gamma^V}), \quad \mathbb{P}\text{-a.s.}, \quad (5.16)$$

with the same conventions than in II-5.6. In particular the two members vanish on $\{V \notin \mathcal{X}_0\}$.

Proof – The proof is formally the same as the proof for homogeneous probabilities (II-5.6). We replace \mathcal{W}_0 by \mathcal{X}_0 . The invariance property of the d -homogeneous probability \mathbb{P} *w.r.t.* the equivalent elements of \mathcal{X}_0 (*i.e.*, leading to equivalent cuts) gives the result. \square

III-3 The Embedded Markov Chain.

We consider a distributed Markov net $(\mathcal{N}, \mathbb{P})$. The normal decomposition of maximal configurations, defined in Chapter 3, defines a sequence of random variables that converge to a maximal process. The random markings reached through the normal decomposition, together with the local random actions, constitute the *embedded Markov chain*, that is actually a finite Markov chain. *The embedded Markov chain characterises the probabilistic system $(\mathcal{N}, \mathbb{P})$.*

This is our first application of the Markov property. The embedded Markov chain has a huge space of states, and its manipulation should thus be avoided for computational reasons. However the embedded Markov chain appears as an auxiliary tool for the study of limit theorems in Chapter 7.

III-3.1 Random Variables from the Normal Decomposition. Let $(\mathcal{N}, \mathbb{P})$ be a d -homogeneous net, with \mathcal{U} the unfolding and \mathcal{E} the associated event structure. Since \mathcal{N} is compact, \mathcal{E} is locally finite and has finite concurrent width. In particular each maximal configuration ω has a unique normal decomposition (Ch. 3, III-3.1), given by a sequence $(v_n, z_n)_{n \geq 1}$ and $v_0 = \emptyset$, characterised by:

$$\forall n \geq 1, \quad v_n \subseteq \omega, \quad v_n = v_{n-1} \oplus z_n,$$

with z_n full-initial in $\mathcal{E}^{v_{n-1}}$, or equivalently: $z_n \in \Omega_{B^\perp(\mathcal{E}^{v_{n-1}})}$, for all $n \geq 1$.

We define the random variables $(V_n, Z_n)_{n \geq 0}$ and V_0 , by:

$$V_0(\omega) = \emptyset, \quad \forall n \geq 1, \quad V_n(\omega) = v_n, \quad Z_n(\omega) = z_n. \quad (5.17)$$

By construction, V_n takes its values in \mathcal{X}_0 . Z_n takes its values in a finite set \mathcal{Z} , given by:

$$\mathcal{Z} = \bigcup_{c \in \mathcal{D}} \Omega_{B^\perp(\mathcal{U}^c)},$$

where \mathcal{D} denotes the finite set of reachable markings of \mathcal{N} , which are in finite numbers. We define the random variable M_n as the marking in \mathcal{N} associated with the finite configuration V_n , for all $n \geq 0$. If $\rho : \mathcal{U} \rightarrow \mathcal{N}$ denotes the folding mapping, we have:

$$\forall n \geq 0, \quad M_n = \rho \circ \gamma(V_n).$$

M_n takes its values in the finite set of reachable markings.

III-3.2 The Embedded Markov Chain. Due to the following result, that states that $(M_n, Z_n)_{n \geq 0}$ is a Markov chain, we call $(M_n, Z_n)_{n \geq 0}$ the **embedded Markov chain** of $(\mathcal{N}, \mathbb{P})$.

It is intuitively clear that $(M_n, Z_n)_{n \geq 0}$ is a Markov chain: given (M_n, Z_n) , and especially the marking M_n , one computes the law of Z_{n+1} , that is the law of the full-initial configuration in the cone of future \mathcal{U}^{M_n} , under probability \mathbb{P}^{M_n} . This gives the joint law of (M_{n+1}, Z_{n+1}) conditional on (M_n, Z_n) .

III-3.3 Theorem. *Let $(\mathcal{N}, \mathbb{P})$ be a d -homogeneous net. Then V_n is a finite well-stopping operator for every $n \geq 0$, and $(M_n, Z_n)_{n \geq 0}$ is a homogeneous Markov chain in the finite set $\mathcal{D} \times \mathcal{Z}$.*

III-3.4 Corollary. *For every d -homogeneous net $(\mathcal{N}, \mathbb{P})$, there is a finite set S and a Markov chain $(X_n)_{n \geq 1}$ on S , defined on the canonical probability space $\mathcal{A} = S^{\mathbb{N}}$, and an isomorphism of probability spaces $\mathcal{A} \rightarrow \Omega$.*

Proof— X_n is given by the embedded Markov chain $X_n = (M_n, Z_n)$. The isomorphism of probability spaces $\Phi : \mathcal{A} \rightarrow \Omega$ is given by:

$$\Phi((X_n)_n) = Z_1 \oplus Z_2 \oplus \cdots ,$$

□

III-3.5 Proof of III-3.3. To show that V_n is a well-stopping operator, since $V_n \in \mathcal{X}_0$, we only have to show that V_n is a stopping operator. Let $\omega, \omega' \in \Omega$ such that $\omega' \supseteq V_n(\omega)$. Then an induction shows that the first n terms of the normal decomposition of ω' coincide with those of ω , hence $V_n(\omega') = V_n(\omega)$, what was to be shown.

To show that $(M_n, Z_n)_n$ is a Markov chain, we construct its probability transition Q , a square matrix indexed by $(\mathcal{D} \times \mathcal{Z}) \times (\mathcal{D} \times \mathcal{Z})$. We denote by $m(v)$ the marking reached by a configuration v .

Let $(m, z) \in \mathcal{D} \times \mathcal{Z}$. If there is no finite well-stopped configuration v in \mathcal{U} such that $m(v) = m$ and $p(v) > 0$, we fill the row (m, z) of Q with any probability vector. Otherwise, let $v \in \mathcal{X}_0$ with positive trace ($p(v) > 0$) and such that $m(v) = m$. Let μ_m be the probability on \mathcal{Z} given as follows. Let Z^v be the full-initial random configuration:

$$Z^v : \Omega^v \rightarrow \Omega_{B^\perp(\mathcal{E}^v)}, \quad \xi \rightarrow \xi \cap B^\perp(\mathcal{E}^v), \quad (5.18)$$

with Ω^v equipped with the probability $\mathbb{P}^v = \mathbb{P}(\cdot | \Omega(v))$. We define μ_m as the law of Z^v in $\Omega_{B^\perp(\mathcal{E}^v)}$, that we extend with zeros to \mathcal{Z} . Since \mathbb{P} is d -homogeneous, μ_m does not depend on the particular $v \in \mathcal{X}_0$ that has been chosen.

To fill the row (m, z) of matrix Q , we give probability:

$$\mu_m(z')$$

to the pairs $(m', z') \in \mathcal{D} \times \mathcal{Z}$ such that $m' = m(v \oplus z')$, where v is as above any $v \in \mathcal{X}_0$ leading to m , and probability 0 to the other pairs. This gives a stochastic matrix Q .

Now we show that $(M_n, Z_n)_{n \geq 0}$ is a Markov chain with transition matrix Q . Let $(m', z') \in \mathcal{D} \times \mathcal{Z}$. The conditional probability:

$$p(m', z') = \mathbb{P}((M_{n+1}, Z_{n+1}) = (m', z') | M_0, Z_0, \dots, M_n, Z_n)$$

is zero whenever $m' \neq M_n \oplus z'$. Otherwise we have:

$$p(m', z') = \mathbb{P}(Z_{n+1} = z' \mid M_0, Z_0, \dots, M_n, Z_n). \quad (5.19)$$

Since $V_n = M_0 \oplus Z_1 \oplus \dots \oplus Z_n$, we have the equality of σ -algebras:

$$\mathcal{F}_{V_n} = \langle Z_1, \dots, Z_n \rangle = \langle M_0, Z_0, \dots, M_n, Z_n \rangle. \quad (5.20)$$

For $v \in \mathcal{X}_0$ of positive trace, let $g^v : \Omega^v \rightarrow \mathbb{R}$ be the non-negative function defined by:

$$g^v(\xi) = \begin{cases} 1 & \text{if } \xi \cap B^\perp(\mathcal{E}^v) = z' \\ 0 & \text{otherwise.} \end{cases}$$

The mapping $v \rightarrow g^v$ is a congruence for the equivalence of cuts: if $v, v' \in \mathcal{X}_0$ satisfy $m(v) = m(v')$ then $g^v = g^{v'}$, modulo the identification $\Omega^v = \Omega^{v'}$ (II-1.3). The collection $(g^v)_v$ defines thus a non-negative d -homogeneous function $g = (g^c)_{c \in \mathcal{D}}$ (III-2.3). We have:

$$\mathbb{E}(g \circ \theta_{V_n} \mid \mathcal{F}_{V_n}) = \mathbb{P}(\xi \cap B^\perp(\mathcal{E}^{V_n}) = z' \mid \mathcal{F}_{V_n}) = \mathbb{P}(Z_{n+1} = z' \mid \mathcal{F}_{V_n}). \quad (5.21)$$

Using (5.19) and (5.20), we get:

$$\begin{aligned} p(m', z') &= \mathbb{P}(Z_{n+1} = z' \mid \mathcal{F}_{V_n}) \\ &= \mathbb{E}(g \circ \theta_{V_n} \mid \mathcal{F}_{V_n}), \quad \text{by (5.21),} \\ &= \mathbb{E}^{\gamma_{V_n}}(g^{\gamma_{V_n}}), \quad \text{by the well-stopped Markov property,} \\ &= \mathbb{P}^{\gamma_{V_n}}(Z^{V_n} = z'), \quad \text{where } Z^v \text{ is defined by (5.18),} \\ &= \mu_{M_n}(z') = Q_{(M_n, Z_n), (m', z')}. \end{aligned}$$

This shows that $(M_n, Z_n)_{n \geq 1}$ is a Markov chain with transition matrix Q . \square

IV—Distributed Markov Nets

We now have to construct d -homogeneous probabilities. In general, the distributed product does not give homogeneous probabilities: nothing is guaranteed for the invariance of the probabilities in the futures of finite configurations leading to the same marking. The invariance holds if the configurations are finite and well-stopped. We construct thus “almost” memory-less probabilities. It may happen though that the probabilities that we obtain are actually homogeneous. For instance, in sequential systems, the distinction has no effect since all finite configurations are well-stopped. This topic is discussed through examples in next section.

In IV-1, we give the definition of branching cells for occurrence nets—it was defined for event structures in Ch. 3. In particular for labelled occurrence nets,

an equivalence relation is inherited from the labelling mapping. We then define distributed Markov nets in IV-2, and we show that they define d -homogeneous nets. In IV-3, we show how our results apply to sequential systems. In particular, we complete the proof of the usual Strong Markov property for finite Markov chains with usual stopping times.

IV-1 Branching Cells in Occurrence Nets and in Unfoldings.

We adapt the notion of branching cell from event structures (Ch. 3, V-1.2) for occurrence nets. We introduce an equivalence class on branching cells if the occurrence net is given by the unfolding of a safe net.

IV-1.1 Branching Cells of an Occurrence Net. Let \mathcal{U} be an occurrence net, and let \mathcal{E} be the event structure associated with \mathcal{U} . Let $\overset{\circ}{\lambda}$ be a branching cell of \mathcal{E} (Ch. 3, V-1.2). We define the branching cell λ of \mathcal{U} as the smallest open set in \mathcal{U} that contains $\overset{\circ}{\lambda}$ (Ch. 1, II-3.5). λ is given by:

$$\lambda = \overset{\circ}{\lambda} + \sum_{e \in \overset{\circ}{\lambda}} (\bullet e + e \bullet).$$

If (\mathcal{U}, ρ) is the unfolding of a safe marked net (\mathcal{N}, M) , each branching cell $(\lambda, \rho|_{\lambda} : \lambda \rightarrow \mathcal{N})$ is a labelled occurrence net *w.r.t.* the net \mathcal{N} . We denote by \mathcal{W}_{λ} and by Ω_{λ} the set of maximal configurations of λ , which are isomorphic to $\mathcal{W}_{\overset{\circ}{\lambda}}$ and to $\Omega_{\overset{\circ}{\lambda}}$ respectively.

We denote by $\Lambda_{\mathcal{U}}$ the set of branching cells of \mathcal{U} , and $\Lambda_{\mathcal{U}}$ is one-to-one with $\Lambda_{\mathcal{E}}$.

IV-1.2 Equivalence of Branching Cells in Unfoldings. Let (\mathcal{U}, ρ) be the unfolding of a safe marked net \mathcal{N} . A branching cell λ , equipped with the restriction $\rho|_{\lambda}$, is a labelled occurrence net. We say that $\lambda, \lambda' \in \Lambda_{\mathcal{U}}$ are equivalent, denoted by:

$$\lambda \cong \lambda',$$

if λ and λ' are isomorphic as labelled occurrence nets (Ch. 1, II-4.2). The relation is indeed an equivalence relation, and we denote by $\mathcal{L}_{\mathcal{U}}$ the set of equivalence classes:

$$\mathcal{L}_{\mathcal{U}} = \Lambda_{\mathcal{U}} / \cong.$$

Let $\lambda \in \Lambda_{\mathcal{U}}$. By definition there is $v \in \mathcal{X}_0$ such that $\lambda \in \Delta^+(v)$. It follows that $c(\lambda) = \text{Min}_{\leq}(\lambda)$ is a $\|\text{-clique}$ in \mathcal{U} , since $c(\lambda) \subseteq \gamma(v)$. Let λ' be equivalent to λ . Then $c(\lambda) \cong c(\lambda')$ in the sense of II-1.1. Conversely, the following holds.

IV-1.3 Proposition. *Let $\lambda, \lambda' \in \Lambda_{\mathcal{U}}$. Assume that $\text{Min}_{\preceq}(\lambda) \cong \text{Min}_{\preceq}(\lambda')$. Then there is a unique isomorphism of labelled occurrence nets $\lambda \rightarrow \lambda'$.*

Proof— Consider $\lambda \in \Lambda_{\mathcal{U}}$, and fix $v \in \mathcal{X}_0$ such that $\lambda \in \Delta^+(v)$. We have that λ is a minimal non empty stopping prefix of \mathcal{U}^v . Let M be the marking associated with v , and we set: $N = \rho(\text{Min}_{\preceq}(\lambda))$, the projection of $\text{Min}_{\preceq}(\lambda)$ in the net \mathcal{N} . We have: $N \subseteq M$, and thus (\mathcal{N}, N) is a safe net.

For every configuration y of λ , $\gamma(v) \oplus y$ is a configuration of \mathcal{U}^v —recall that we have to add the initial cut to obtain configurations in occurrence nets. Since λ is a prefix of \mathcal{U}^v , the configuration $\gamma(v) \oplus y$ only consumes resources from N . It follows that $(\lambda, \rho|_{\lambda})$ is a folding of (\mathcal{N}, N) . Let \mathcal{V} be the unfolding of (M, N) . By the universal property of \mathcal{V} , there is a unique morphism of foldings $m : \lambda \rightarrow \mathcal{V}$. On the other hand, any playing sequence in (\mathcal{N}, N) only consumes resources from N , and thus induces a configuration in \mathcal{U}^v . We have thus a morphism of labelled occurrence nets $m' : \mathcal{V} \rightarrow \mathcal{U}^v$.

Since both m and m' are morphisms into unfoldings, and according to Lemma II-4.7 of Ch. 1, m and m' are injective. It follows that $m^{-1} : m(\lambda) \rightarrow \lambda$ satisfies $m^{-1} = m'|_{m(\lambda)}$.

We now show that $m(\lambda)$ is an initial branching cell of \mathcal{V} . $m(\lambda)$ is a prefix, and we check that:

$$\forall e, e' \in \mathcal{V}, \quad e \#_d e' \Rightarrow m'(e) \#_d^v m'(e').$$

Since λ is a stopping prefix in \mathcal{U}^v , it implies that $m(\lambda)$ is a stopping prefix in \mathcal{V} . Conversely, since m is a morphism $\lambda \rightarrow \mathcal{V}$, if B is a non empty stopping prefix of \mathcal{V} with $B \subseteq m(\lambda)$, then $m'(B)$ is a stopping prefix of \mathcal{U}^v , that encounters λ . It follows that $m'(B) = \lambda$ and then $B = m(\lambda)$. We have shown that $m(\lambda)$ is an initial branching cell of \mathcal{V} .

Now let $\lambda' \in \Lambda_{\mathcal{U}}$ with $\text{Min}_{\preceq}(\lambda) \cong \text{Min}_{\preceq}(\lambda')$. We have as above an injective morphism $q : \lambda' \rightarrow \mathcal{V}$. $m(\lambda)$ and $q(\lambda')$ have the same initial cut in \mathcal{V} . Two initial branching cells with a condition in common in their initial cuts coincide (it is a consequence of Ch. 3, IV-3.1 for event structures). It implies that $m(\lambda) = q(\lambda')$, whence the isomorphism $q^{-1} \circ m : \lambda \rightarrow \lambda'$. The uniqueness of $\lambda \rightarrow \lambda'$ follows from the uniqueness of m and q . \square

IV-1.4 Remark. If \mathcal{N} is a compact net, the set $\mathcal{L}_{\mathcal{U}}$ of equivalence classes of branching cells of the unfolding \mathcal{U} of \mathcal{N} is finite.

Indeed, a branching cell has the form:

$$\lambda \in \Delta^+(v) = \Delta^{\perp}(\mathcal{E}^v) = \Delta^{\perp}(\mathcal{E}^{\gamma(v)}).$$

There is a finite number of $\mathcal{E}^{\gamma(v)}$ modulo isomorphism of labelled occurrence nets. There are finitely many branching cells in $\Delta^+(v)$ for each v , hence there are finitely many classes of branching cells.

IV-2 Distributed Markov Nets.

IV-2.1 Collection of Spaces Ω_λ . Proposition IV-1.3 implies that, if $\lambda \cong \lambda'$, then the isomorphism $\lambda \rightarrow \lambda'$ is unique. This natural isomorphism induces through its actions on subsets an isomorphism of partial orders mapping \mathcal{W}_λ onto $\mathcal{W}_{\lambda'}$. In turn we get by restriction to maximal configurations a bijection between the finite sets $\Omega_\lambda \rightarrow \Omega_{\lambda'}$.

This allows to identify each class $l \in \mathcal{L}_U$ with a labelled occurrence net l , such that for each λ of the class l , there are isomorphisms:

$$\mathcal{W}_l \rightarrow \mathcal{W}_\lambda, \quad \Omega_l \rightarrow \Omega_\lambda,$$

induced by the action on sets of the unique isomorphism of labelled occurrence nets $\lambda \rightarrow l$. In particular, a probability law μ on Ω_λ is seen on Ω_l .

IV-2.2 Definition. (*Homogeneous family of branching probabilities*) Let \mathcal{U} be the unfolding of a safe compact marked net \mathcal{N} . A family of branching probabilities $(\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{U}}}$ is said to be **homogeneous** if the following holds:

$$\forall \lambda, \lambda' \in \Lambda_{\mathcal{U}}, \quad \lambda \cong \lambda' \Rightarrow \mu_\lambda = \mu_{\lambda'},$$

where we identify the probabilities μ_λ and $\mu_{\lambda'}$ and their images in Ω_l , with l the common class of λ and λ' . Equivalently, a homogeneous family of branching probabilities is given by a *finite* collection $(\mu_l)_{l \in \mathcal{L}_U}$, where μ_l is a probability measure on Ω_l for every class of branching cell $l \in \mathcal{L}_U$.

(*Distributed Markov net*) We define a **distributed Markov net** as a pair $(\mathcal{N}, (\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{U}}})$ as above: \mathcal{N} a compact safe marked net, and $(\mu_\lambda)_\lambda$ a homogeneous family of branching probabilities on its unfolding. We associate the probability \mathbb{P} on Ω given by the distributed product $\mathbb{P} = \bigotimes_{\lambda \in \Lambda_{\mathcal{U}}}^d \mu_\lambda$.

IV-2.3 Theorem. Let $(\mathcal{N}, (\mu_\lambda)_{\lambda \in \Lambda_{\mathcal{U}}})$ be a distributed Markov net. Let \mathbb{P} be the distributed product of the family $(\mu_\lambda)_\lambda$. Then $(\mathcal{N}, \mathbb{P})$ is d -homogeneous. In particular, the well-stopped Markov property holds.

Proof – Let $v, v' \in \mathcal{X}_0$ with $p(v), p(v') > 0$ and $\gamma(v) \cong \gamma(v')$. The isomorphism of foldings $f : \mathcal{U}^v \rightarrow \mathcal{U}^{v'}$ from II-1.2 induces a one-to-one mapping:

$$\hat{f} : \Lambda_{\mathcal{U}^v} \rightarrow \Lambda_{\mathcal{U}^{v'}}, \quad \hat{f}(\lambda) \cong \lambda \quad \forall \lambda \in \Lambda_{\mathcal{U}^v},$$

and for each $\lambda \in \Lambda_{\mathcal{U}}$, a bijection $f_\lambda : \Omega_\lambda \rightarrow \Omega_{\hat{f}(\lambda)}$ conjugated to the identity $\Omega_l \rightarrow \Omega_l$ where l is the common class of λ and of $\hat{f}(\lambda)$.

Since the family $(\mu_\lambda)_\lambda$ is homogeneous, we have $\mu_\lambda = \mu_{\hat{f}(\lambda)}$ for all $\lambda \in \Lambda_{\mathcal{U}}$. By the composition formula (4.18) in Ch. 4, Th. IV-2.2, we get:

$$\mathbb{P}^{v'} = \bigotimes_{\lambda' \subseteq \mathcal{U}^{v'}}^d \mu_{\lambda'} = \bigotimes_{\lambda \subseteq \mathcal{U}^v}^d \mu_{\hat{f}(\lambda)} = \bigotimes_{\lambda \subseteq \mathcal{U}^v}^d \mu_\lambda = \mathbb{P}^v. \quad (5.22)$$

□

IV-2.4 Remark. What we require for distributed Markov nets is stronger than a pair $(\mathcal{N}, \mathbb{P})$, where \mathbb{P} is distributed and d -homogeneous. Indeed, nothing in general prevents two induced branching probabilities \mathbb{P}^λ and $\mathbb{P}^{\lambda'}$ to be different, even if \mathbb{P} is d -homogeneous and $\lambda \cong \lambda'$, since λ and λ' may be activated by cuts that are not equivalent.

IV-3 Sequential Systems.

We continue the example I-1.2 of a sequential net \mathcal{N} associated with a probabilistic transition system $(S, A, x_0, (\mu_x)_{x \in S})$. \mathbb{P} denotes the probability on Ω that comes from the Markov chains $(X_n)_{n \geq 0}$ (the markings) and $(Y_n)_{n \geq 1}$ (the transitions). We show that we can obtain \mathbb{P} and the Markov property for \mathbb{P} from our construction.

Every branching cell $\lambda \in \Lambda_{\mathcal{T}}$ is one-to-one with its image in A , which is given by a set $\partial_-^{-1}(x)$, for x a certain state of S (Ch. 3, VII-2). The state x is unique and is the label of the initial cut of λ , which reduces to a unique condition. We set the branching probability over λ :

$$\nu_\lambda(t) = \mu_x(\partial_+(t)) . \quad (5.23)$$

IV-3.1 Theorem. *Let \mathcal{N} be the sequential net associated with a probabilistic transition system $(S, A, x_0, (\mu_x)_{x \in S})$. The canonical probability measure \mathbb{P} on Ω defined by the Markov chain $(X_n)_{n \geq 0}$ makes $(\mathcal{N}, \mathbb{P})$ a distributed Markov net, with branching probabilities given by (5.23). In particular, $(X_n)_n$ satisfies the usual Strong Markov property (II-4.3).*

Proof— The family $(\nu_\lambda)_{\lambda \in \Lambda_{\mathcal{U}}}$ is homogeneous, let \mathbb{Q} be the distributed product $\mathbb{Q} = \bigotimes_{\lambda \in \Lambda_{\mathcal{U}}}^d \nu_\lambda$. We compute the likelihood q of \mathbb{Q} on a chain of events $v = (e_1, \dots, e_n)$. Let $t_i = \rho(e_i)$ the label of e_i , with $(\lambda_1, \dots, \lambda_n)$ the sequence of encountered branching cells, and (x_0, \dots, x_n) the associated sequence of states. Using the form (5.23), we compute the likelihood of \mathbb{Q} according to Proposition III-3.4 in Ch. 3:

$$\begin{aligned} q(v) &= \nu_{\lambda_1}(e_1) \dots \nu_{\lambda_n}(e_n) \\ &= \mu_{x_0}(\partial_+(t_1)) \mu_{x_1}(\partial_+(t_2)) \dots \mu_{x_{n-1}}(\partial_+(t_n)) \\ &= \mu_{x_0}(x_1) \dots \mu_{x_{n-1}}(x_n) \\ &= \mathbb{P}((X_0, \dots, X_n) = (x_0, \dots, x_n)) = p(v) , \end{aligned}$$

where p denotes the likelihood of \mathbb{P} . Since \mathcal{T} is locally finite (Ch. 3, VII-2), the uniqueness in the extension theorem (Ch. 2, III-3.1) implies that $\mathbb{P} = \mathbb{Q}$. Now \mathbb{Q} , and thus \mathbb{P} , has the required properties: By IV-2.3, \mathbb{Q} is d -homogeneous as a distributed product. Since all finite configurations are well-stopped (Ch. 3, VII-1.1), \mathbb{Q} is actually homogeneous. By II-5.8, it follows that \mathbb{Q} satisfies the Strong Markov property for Markov chains II-4.3. \square

IV-3.2 Remark. The chains $(X_n, Y_n)_{n \geq 1}$ and the embedded Markov chain $(M_n, Z_n)_{n \geq 1}$ (defined in III-3.2) are image one from the other by $M_n = X_n$, $Y_n = \rho(Z_n)$, where $\rho : \mathcal{U} \rightarrow \mathcal{N}$ is the folding.

IV-3.3 Non homogeneous Markov Chains. For the sake of completeness, let us analyse the case of *non-homogeneous* Markov chains. Let E be a finite set, and let $(P^i)_{i \geq 0}$ be a sequence of stochastic matrixes on E . As for homogeneous Markov chains, the canonical probability space is $\mathcal{A} = E^{\mathbb{N}}$, and $(X_i)_i$ denotes the sequence of components. We say that $(X_i)_i$ is a (non homogeneous) Markov chain associated with the family $(P_i)_i$ if \mathcal{A} is equipped with the unique probability \mathbb{P} satisfying, for all i and for all tuple (x_0, \dots, x_i) :

$$\mathbb{P}(X_0 = x_0, \dots, X_i = x_i) = \nu(x_0)P_{x_0, x_1}^0 \dots P_{x_{i-1}, x_i}^{i-1},$$

with ν the starting probability on E . \mathbb{P} is a probability on the border at infinity of a tree, and we have shown that every probability on the border at infinity of a locally finite tree is a distributed probability, and can be expressed as a distributed product (Ch. 4, V-2.4).

The branching cells of the unfolding are given by the collection:

$$\lambda(b) = b + b^\bullet + (b^\bullet)^\bullet, \quad b = \text{Min}_{\leq}(\lambda(b)).$$

with b a condition of the unfolding. Remark that the equivalence \cong on cuts reduces to an equivalence on conditions. Let $h(b)$ denote the *height* of a condition b , *i.e.* the number of events in configuration $[b[$. A probability \mathbb{P} on Ω defines a (non homogeneous) Markov chain if and only if, for all conditions b, b' in \mathcal{U} :

$$h(b) = h(b'), \quad b \cong b' \Rightarrow \mathbb{P}^{\lambda(b)} = \mathbb{P}^{\lambda(b')}.$$

That is, on each row of the unfolding, the branching probability only depends on the initial condition. This is illustrated in Figure 5.2. Whereas the probability \mathbb{P} defines a *homogeneous Markov chain* if and only if, for all condition $b, b' \in \mathcal{U}$, and without the height constraint:

$$b \cong b' \Rightarrow \mathbb{P}^{\lambda(b)} = \mathbb{P}^{\lambda(b')}.$$




Figure 5.2: *Non homogeneous Markov chain.*

V—Examples of Distributed Markov Nets.

This Section examines some examples of distributed Markov nets. We show how parametric families of probabilities can be defined on compact nets. The lower precision bound of our study is the branching cell, and thus the branching probability. Which branching probability shall be chosen depends on what one needs to model. We present some simple examples of parametric families obtained by *renormalisation* of weights.

We also present the case of confusion-free nets, a model introduced by G. Winskel and Nielsen-Plotkin-Winskel ([47, 33]) in relation with concrete domains. Confusion-free nets are also studied as a particular class of “simple” safe nets. They are the model examined by H. Völzer and Varacca-Völzer-Winskel ([45, 44]) for probabilistic true-concurrent nets. They constitute a sub-class of compact nets, that share many properties with sequential systems. In particular, the material introduced in our theory (stopping prefixes, branching cells, well-stopped configurations) have simple expressions.

V-1 Additive and Multiplicative Probability.

Let \mathcal{N} be a safe compact marked net. To construct a distributed Markov net based on \mathcal{N} , it is enough to specify a homogeneous collection of branching probabilities. Although we do not know precisely the set $\mathcal{L}_{\mathcal{U}}$ of classes of branching cells, we can define a homogeneous collection of branching probabilities by following a renormalisation construction. We use this technique to define the additive and the multiplicative branching probabilities. Many others can be considered. We apply these examples to illustrate the difference between homogeneous and d -homogeneous probabilities on Ω .

V-1.1 Additive Probability. Let $\mathcal{N} = (P, T, M_0)$ be the safe compact marked net, and let $\phi > 0$ be a positive function $\phi : T \rightarrow \mathbb{R}$ defined on the transitions of \mathcal{N} . Denote by $\rho : \mathcal{U} \rightarrow \mathcal{N}$ the folding mapping, and let $l \in \mathcal{L}_{\mathcal{U}}$ be a class of branching cells. We define the additive branching probability over l by:

$$\forall z \in \Omega_l, \quad \mu_l(z) = \frac{1}{Y_l} \sum_{e \in z} \phi \circ \rho(e),$$

where e ranges over the events of z , and Y_l is a renormalisation constant. We sum the weights of transitions encountered, and then we renormalise. We define the d -product of the family $(\mu_l)_{l \in \mathcal{L}_{\mathcal{U}}}$ as the **additive** probability on Ω , with weight function ϕ .

V-1.2 Application: a Probability Non Strongly Homogeneous. Additive probabilities are d -homogeneous by construction. They can give examples of non homogeneous probabilities.

Consider for instance the net depicted in Figure 5.3 (a). Use the notation $\alpha = \phi(\alpha)$, $\beta = \phi(\beta)$, etc, for ϕ a positive weight function and transitions α, β , etc. Let κ be the cut of the unfolding, filled by the tokens of Figure 5.3 (b). Then κ is equivalent to the initial cut c_0 . Remark that κ is not well-stopped. We check that the probabilistic futures of κ and of c_0 are isomorphic if and only if (see [1] for details):

$$\alpha = \beta + \delta = \delta + \gamma. \tag{5.24}$$

It follows that \mathbb{P} is strongly homogeneous if and only if (5.24) holds. In general—*i.e.*, excepted on a thin set of parameters—the additive probability is not homogeneous on this example. This is related to the height of the branching cell λ .




Figure 5.3: Additive probability.

V-1.3 Multiplicative Probability. Using again the renormalisation technique, we define in a similar way the multiplicative probability. We consider a positive weight function $\phi : T \rightarrow \mathbb{R}$ on the transitions of the net. For $l \in \mathcal{L}_U$ a class of branching cell of the unfolding, let ν_l be the branching probability over l defined on Ω_l by:

$$\forall v \in \Omega_l, \quad \nu_l(v) = \frac{1}{A_l} \prod_{e \in v} \phi \circ \rho(e),$$

where e ranges over the events of v , Y_l is a renormalisation constant, and $\rho : \mathcal{U} \rightarrow \mathcal{N}$ is the folding mapping. ν_l is called the multiplicative branching probability over l , *w.r.t.* the weight function ϕ . The distributed product of multiplicative branching probabilities constitutes the multiplicative probability *w.r.t.* ϕ .

Modulo re-parametrisation, the multiplicative branching probabilities over a branching cell l form an exponential statistical model (see [1]).

V-1.4 Example. If \mathbb{Q} denotes a multiplicative probability defined for the net of Figure 5.3, then we check that \mathbb{Q} is strongly homogeneous. It follows that, on this example, additive and multiplicative probabilities coincide only on the set of probabilities described by (5.24).

V-2 A Case of Homogeneity: Confusion-Free Nets.

We follow [33] for the presentation of confusion-free nets and confusion free event structures. Confusion-free nets and event structures form a restricted class of systems where the concurrency behaviour is severely controlled. This class has been studied by H. Völzer, D. Varacca and G. Winskel for probabilistic applications from a domain theory point of view. Confusion-free nets are compact. This explains *a posteriori* that authors have studied this class for the construction of probability measures without the powerful tools of projective systems of probabilities.

We briefly present confusion-free nets and event structures from [33, 44], and we describe our familiar objects (stopping prefixes, branching cells) in this framework. We show that confusion free event structures share several properties with trees. In particular, we show that d -homogeneity and homogeneity of probabilities coincide. By comparison with confusion-free event structures, we can interpret locally finite event structures as event structures with *finite confusion*.

V-2.1 Definition. (*Confused nets*) Let (\mathcal{N}, M_0) be a safe marked net. We say that (\mathcal{N}, M_0) is **symmetrically confused** if there is a reachable marking M and transitions t, t', t'' that can play from M , and such that:

$$\left\{ \begin{array}{l} \bullet t \cap \bullet t' \neq \emptyset, \quad \bullet t' \cap \bullet t'' \neq \emptyset \\ \text{but: } \bullet t \cap \bullet t'' = \emptyset. \end{array} \right. \quad (5.25)$$

We say that (\mathcal{N}, M_0) is **asymmetrically confused** if there is a reachable marking M and transitions t, t', t'' such that:

$$\left\{ \begin{array}{l} t \text{ and } t'' \text{ can play concurrently from } M \\ t' \text{ cannot play from } M, \text{ but } t' \text{ can play from } M - \bullet t + t \\ \bullet t' \cap \bullet t'' \neq \emptyset \end{array} \right.$$

We say that (\mathcal{N}, M_0) is **confused** if it is either symmetrically or asymmetrically confused. Confused nets are illustrated in Figure 5.4.

Remark that confusion-freeness of a safe marked net is easily checked on the net, since reachable markings of safe nets are computable (Cf. Ch. 8 for more details).

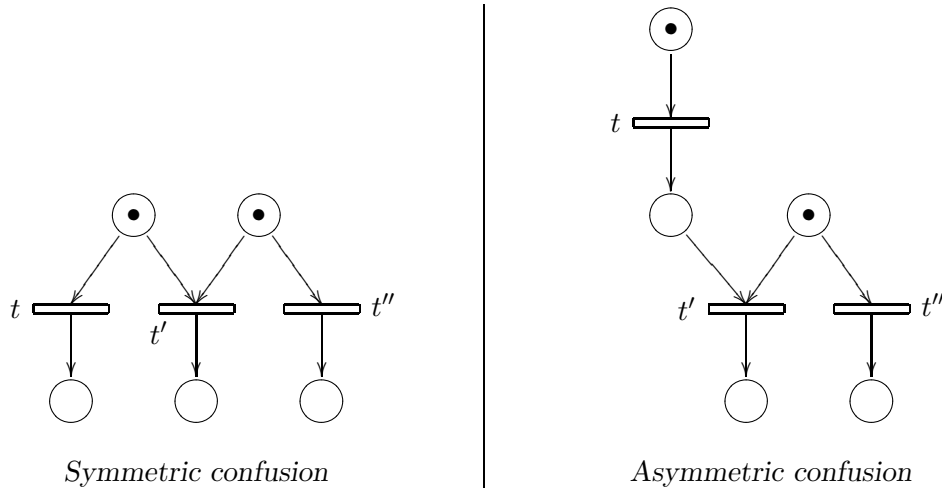


Figure 5.4: Confused nets.

V-2.2 Remark. The class of confusion-free nets contains in particular the well-known *extended free-choice* nets, defined for $\mathcal{N} = (P, T, M_0)$ by one of the equivalent properties ([39, 17]):

$$\forall t, t' \in T, \quad \bullet t \cap \bullet t' \neq \emptyset \Rightarrow \bullet t = \bullet t',$$

or:

$$\forall s, s' \in P, \quad s \bullet \cap s' \bullet \neq \emptyset \Rightarrow s \bullet = s' \bullet.$$

V-2.3 Confusion-Free Nets and Concrete Domains. Confusion-Free Event Structures. Confusion-free nets are introduced in [33] for their relation with concrete domains. Let \mathcal{E} be the event structure associated with the unfolding of a safe marked net (\mathcal{N}, M_0) . Then \mathcal{N} is confusion free if and only if the domain \mathcal{W} of configurations of \mathcal{E} satisfies the so-called Q axiom of concrete domains:

Axiom Q: For every $v, v' \in \mathcal{W}$ with $v \subseteq v'$ and for every $e \in \mathcal{E}$ such that:

$$e \in \text{Min}_{\preceq}(\mathcal{E}^v),$$

we have:

$$e \# v' \Rightarrow \exists! f \in \text{Min}_{\preceq}(\mathcal{E}^v) : f \in v', e \# f.$$

(Recall that $e \# v'$ means that e is in conflict with at least one event of v').

It is observed in [33] that the Q axiom has two components: existence and uniqueness. Existence corresponds to asymmetric confusion, and uniqueness corresponds to symmetric confusion.

We will say that an event structure \mathcal{E} is confusion-free if the domain \mathcal{W} of configurations of \mathcal{E} satisfies axiom Q . Clearly, any cone of future \mathcal{E}^v is then confusion-free.

V-2.4 Confusion-Free Event Structures and Stopping Prefixes. It can be shown, as stated in [44], that banning symmetric and asymmetric confusion in an event structure corresponds to the following properties of the dynamic conflict relation $\#_d$:

1. $e \#_d e' \Rightarrow [e[= [e'[$,
2. the reflexive closure of $\#_d$ is transitive.

For each event $x \in \mathcal{E}$, denote by $F(e)$ the following subset of \mathcal{E} —we call it the **flower** of event e :

$$F(e) = \{x \in \mathcal{E} \mid e \#_d x\},$$

and assume that \mathcal{E} is the unfolding of a safe net (confusion-free). Then $F(e)$ is finite, since the events of $F(e)$ are indexed by some of the transitions t that can play from the marking M , obtained from configuration $[e[$. Moreover consider the following subset, for e a given event:

$$B = \bigcup_{f \in [e[} F(f). \quad (5.26)$$

Then B is a $\#_d$ -closed prefix of \mathcal{E} , *i.e.* a stopping prefix, and B is finite. Actually $B = B(e)$, *i.e.* B is the smallest stopping prefix that contains e . It implies that \mathcal{E} is locally finite. Whence the following result.

V-2.5 Proposition. ([44]) *Any confusion-free net is compact.*

We re-obtain the result for sequential systems, since sequential nets are confusion-free. We also have the following result, which extends a property already shown for trees:

V-2.6 Proposition. *In a confusion-free event structure, every configuration is stopped.*

Proof— Let v be a configuration, and let ω be a maximal configuration containing v . From the expression (5.26) of $B(e)$ for every event e , we have:

$$v = \omega \cap B, \quad B = \bigcup_{e \in v} B(e),$$

which implies the statement. \square

V-2.7 Corollary. *Let (\mathcal{N}, M_0) be a confusion-free net. Then any distributed Markov net defined on (\mathcal{N}, M_0) is strongly homogeneous.*

V-2.8 Branching Cells for Confusion-Free Event Structures. To study the branching cells of a confusion-free event structure, and since the cones of future are confusion-free, it is enough to study the *initial* branching cells. Let $\lambda \in \Delta^\perp(\mathcal{E})$ be an initial branching cell, *i.e.* a minimal non void stopping prefix. For any pair (e, e') of events in λ , the transitivity of $\#_d$ implies that $e \neq e' \Rightarrow e \#_d e'$, and therefore $e \preceq e' \Rightarrow e = e'$. It follows that:

For $\lambda \in \Delta^\perp(\mathcal{E})$, with \mathcal{E} a confusion-free event structure, every event $e \in \lambda$ is minimal in \mathcal{E} and satisfies $\lambda = B(e)$. Each maximal configuration $v \in \Omega_\lambda$ consists of a unique event.

As noticed, this can be done with any branching cell. Remark that we consider branching cells in the unfolding, not the classes:

V-2.9 Proposition. *Any branching cell $\lambda \in \Lambda_{\mathcal{U}}$ of the unfolding of a confusion free net has the form:*

$$\lambda = \{x \in \mathcal{E} \mid x \#_d e\},$$

for any event $e \in \lambda$. Any event $x \in \lambda$ is minimal in λ , and any element $v \in \Omega_\lambda$ consists of a unique event.

As a consequence, branching cells are disjoint in the unfolding. This is an other property shared with trees. However, the projection of branching cells in the net might not be disjoint, as shown by the example of Ch. 6, Figure 6.2, p. 193. The computation of branching cells is simpler for confusion-free nets than for general compact nets; see Chapter 8.

VI—Conclusion

In this chapter we have proposed a definition for true-concurrent and memory-less random systems. Natural geometric isomorphisms come from the intrinsic memory-less properties of the dynamics of Petri nets. We add a probabilistic counterpart to this property to define the homogeneous probabilities, basis of a memory-less randomisation. The distributed product of branching probabilities defines “almost” homogeneous probabilities, for which a Strong Markov property holds in a true-concurrent framework. The concurrent Markov property implies the Strong Markov property for finite Markov chains.

Extending the definition of the transition matrix of a Markov chain, the branching probabilities form the finite set of parameters of a distributed Markov net. Each of these parameters is given by a finite probability vector, which gives the law of a *local process*.

With the embedded Markov chain, we interpret the random dynamics in the graph of markings of the net. For computational reasons, the embedded Markov chain is not intended to be of practical use excepted maybe for small examples, but constitutes a theoretical auxiliary tool. It shows that the global dynamics of any homogeneous net is isomorphic to the dynamics of a finite Markov chain on a large state space.

$\mathcal{N}, (\mathcal{U}, \rho), \mathcal{E}$	<ul style="list-style-type: none"> • a safe net or a safe marked net (according to the context), the unfolding of a safe marked net, the associated event structure.
$\gamma(v)$	<ul style="list-style-type: none"> • the cut $\gamma(v) = \text{Max}_{\preceq}(v)$ associated with a finite configuration v.
$\mathcal{U}^v, \mathcal{U}^c$	<ul style="list-style-type: none"> • the cone of future for occurrence nets associated with configuration v, or associated with cut c. $\mathcal{U}^v = \mathcal{U}^c$ if $c = \gamma(v)$.
\mathbb{P}, p	<ul style="list-style-type: none"> • a probability on the space Ω, the associated likelihood $p = \mathbb{P}(\Omega(\cdot))$
$c \cong c'$	<ul style="list-style-type: none"> • two equivalent cuts (II-1)
$\mathcal{D}, \mathcal{D}'$	<ul style="list-style-type: none"> • the set of equivalence classes of cuts, the set of equivalence classes of well-stopped cuts (III-1.1)
$(\mathcal{U}^c, \mathcal{F}^c, \mathbb{P}^c)$	<ul style="list-style-type: none"> • the probabilistic future associated with cut c, or with the class $c \in \mathcal{D}$ if \mathbb{P} is homogeneous, or with the class $c \in \mathcal{D}'$ if \mathbb{P} is d-homogeneous
V, γ_V	<ul style="list-style-type: none"> • a stopping operator $V : \Omega \rightarrow \mathcal{W}$ or a well-stopping operator $V : \Omega \rightarrow \mathcal{X}_0$. γ_V is the random cut $\gamma_V(\omega) = \gamma(V(\omega))$.
$g = (g^c)_{c \in \mathcal{D}}$	<ul style="list-style-type: none"> • a homogeneous function (II-5.2)
$g = (g_c)_{c \in \mathcal{D}'}$	<ul style="list-style-type: none"> • a d-homogeneous function (III-2.3)
$(V_n, Z_n)_{n \geq 1}$	<ul style="list-style-type: none"> • the sequence of random variables given by the normal decomposition (III-3.1)
$(M_n, Z_n)_{n \geq 1}$	<ul style="list-style-type: none"> • the embedded Markov chain in a probabilistic d-homogeneous net (III-3.1)
$\lambda \cong \lambda'$	<ul style="list-style-type: none"> • two equivalent branching cells in the unfolding (IV-1.2)
$\Lambda_{\mathcal{U}}, \mathcal{L}_{\mathcal{U}}$	<ul style="list-style-type: none"> • the set of branching cells of \mathcal{U}, the quotient set $\mathcal{L}_{\mathcal{U}} = \Lambda_{\mathcal{U}} / \cong$ (IV-1.2)

Table 5.1: Summary of notations for Markov nets.

Chapter 6

Recurrent Nets

In this chapter, we try to adapt some results from the theory of recurrent Markov chains to probabilistic nets. Recurrence will be a natural assumption in Chapter 7 for establishing limit theorems. Our main tool is the Markov property for concurrent systems.

Recurrent states of a Markov chain have a very simple and intuitive definition: those states where the chain comes back infinitely often. The definition is not so clear for nets: starting from a marking, several definitions can be proposed to say that the net is back in the initial marking. The definition that we propose requires that all the tokens have moved before returning back to the marking. We introduce a stopping operator that fits this definition. This tool allows to adapt some results from recurrent Markov chains theory. We show in particular the two following results for d -homogeneous nets:

- A marking has probability 0 or 1 to be recurrent.
- Markings reached from a recurrent marking are recurrent.

Since our definition considers the whole marking of a net, *i.e.* its global state, the stopping operator is called a *global renewal operator*. We use it the same way as we use the usual renewal operator (or *balayage operator*) in Markov chains theory, where it is defined as a stopping time.

We also introduce the notion of *local renewal*. For sequential systems, there is no difference between local and global renewal, since the global state of a sequential system consists in its unique local state. To obtain results concerning the local renewal, it is natural to consider *distributed* Markov nets, instead of general d -homogeneous nets. The sequence of germs defined by the successive arrivals of a same branching cell (modulo isomorphism) define an *i.i.d* sequence of random variables. This is a comfortable framework for statistical estimation, as shown in Chapter 8.

The recurrence is typically an *almost sure* property. When we say that a marking is recurrent, it means that it has probability one to return infinitely often. In other words, although there exists in general processes that never return to the initial marking, they are very rare since all together they have probability zero. Hence we

characterise properties of nets that cannot be stated without the probabilistic framework: properties that are not always true, but *almost always* true in the probabilistic sense.

In Section I, *Global renewal operator*, we introduce the global renewal operator. We also study the iterates of this operator. Globally recurrent nets and conservativity of the global recurrence are the topic of Section II, *Recurrent nets*. The local renewal is the subject of Section III, *Local renewal*, where we particularise the study to distributed probabilities.

I—Global Renewal Operator.

This section introduces some preliminary notions for studying the recurrence of markings. In particular, we introduce a stopping operator, called renewal operator, that stops any ω at its first return to the initial marking, in a sense to be precised. We insist on that all the tokens must have moved to bring the net back to its initial marking.

Using the renewal operator, we will follow the theory of Markov chains and adapt to our framework the well-known notion of recurrent state for a Markov chain. Hence we study probabilistic nets that have probability 1 to return infinitely often to the initial marking.

In I-1 we define the return to the initial marking, which brings us to the definition of the *renewal operator*. A recursive construction leads in I-2 to the recursive renewal operators.

I-1 The Renewal Operator.

I-1.1 Returning to the Initial Marking. Let (\mathcal{N}, M_0) be a safe compact marked net. Let \mathcal{U} be the unfolding of (\mathcal{N}, M_0) , which is locally finite, with c_0 the initial cut of \mathcal{U} . We will focus on topics related to well-stopped configurations, but the definitions can be relaxed to \mathcal{W}_0 instead of \mathcal{X}_0 , and then the corresponding results hold.

We recall that $\gamma(y)$ denotes the cut associated with a finite configuration y . For ω a maximal configuration of \mathcal{U} , we denote by:

$$X(\omega) = \{y \in \mathcal{X}_0 \mid y \subseteq \omega\}$$

the lattice of finite well-stopped sub-configurations of ω ($X(\omega)$ is a lattice according to Ch. 3, VI-2.3). We set the following subset of $X(\omega)$:

$$C(\omega) = \{y \in X(\omega) \mid \gamma(y) \cong c_0\}.$$

Hence $C(\omega)$ contains the finite well-stopped sub-configurations of ω that lead back to the initial marking. The two following results do not present any difficulty.

I-1.2 Lemma. *Let v, v' be two compatible finite configurations. Then we have:*

$$\gamma(v \cup v') = \text{Max}_{\leq}(\gamma(v) \cup \gamma(v')), \quad \gamma(v \cap v') = \text{Min}_{\leq}(\gamma(v) \cup \gamma(v')).$$

I-1.3 Lemma. *For every $\omega \in \Omega$, $C(\omega)$ is a sub-lattice of $X(\omega)$.*

l-1.4 Leaving the Initial Cut. The lattice $C(\omega)$ describes the many different ways ω has to come back to the initial marking. We want to identify configurations of $C(\omega)$ that have left the initial cut, or equivalently, that have moved all the tokens of the initial marking M_0 .

For this, we set the following subset of $C(\omega)$, for each condition b of the initial cut:

$$\forall b \in c_0, \quad C_b(\omega) = \{y \in C(\omega) \mid b \notin \gamma(y)\}.$$

It follows from Lemma l-1.2 that $C_b(\omega)$ is a sub-lattice of $C(\omega)$. We define thus the following element if $C_b(\omega)$ is non empty:

$$D_b(\omega) = \min(C_b(\omega)). \quad (6.1)$$

l-1.5 Definition. (Renewal operator) Let (\mathcal{N}, M_0) be a safe Petri net, and let \mathcal{U} denote the unfolding of (\mathcal{N}, M_0) . We define the **renewal operator** $R : \Omega \rightarrow \mathcal{W}$ by:

$$R(\omega) = \begin{cases} \omega, & \text{if } \exists b \in c_0 : C_b(\omega) = \emptyset, \\ \bigcup_{b \in c_0} D_b(\omega), & \text{if } C_b(\omega) \neq \emptyset \text{ for all } b \in c_0. \end{cases}$$

Hence the renewal operator applied to ω gives the smallest well-stopped sub-configuration of ω that leads back to M_0 , *without sharing any condition with the initial cut* c_0 . In the net, it means that all tokens have moved.

l-1.6 Lemma. *The renewal operator R is a well stopping operator.*

Proof – To show that R is measurable, remark first that for every finite stopping prefix B of \mathcal{U} , B and \mathcal{U} share the same initial cut. We denote by C^B , D_b^B and R^B the analogous of C , D_b and R in Ω_B . Well-stopped configurations of B are those of \mathcal{E} contained in B (Ch. 3, IV-2.1), hence we have:

$$\forall \omega \in \Omega, \quad C^B(\omega_B) = C(\omega) \cap B, \quad C_b^B(\omega_B) = C_b(\omega) \cap B,$$

with the usual notation $\omega_B = \omega \cap B$, and with the intersections $C(\omega) \cap B$ and $C_b(\omega) \cap B$ in the sense of restriction to configurations contained in B . It implies that we have:

$$\forall b \in c_0, \quad D_b(\omega) = \bigcup_{B \in \mathcal{S}_0} \uparrow D_b^B(\omega_B),$$

and thus: $R = \bigcup_{B \in \mathcal{S}_0} \uparrow R^B$. Since all R^B are \mathcal{F}_B -measurable, it follows that R is \mathcal{F} -measurable.

We show that R satisfies the property of stopping operators (Ch. 5, II-3.1): Let $\omega, \omega' \in \Omega$ such that $\omega' \supseteq R(\omega)$, we have to show that $R(\omega) = R(\omega')$. We assume without loss of generality that $R(\omega) \in \mathcal{X}_0$. Then for every initial condition $b \in c_0$, $D_b(\omega)$ is a finite well-stopped configuration contained in ω' , leading back to M_0 and such that $b \notin D_b(\omega)$, hence $D_b(\omega) \in C_b(\omega')$, and thus $D_b(\omega') \subseteq D_b(\omega) \subseteq R(\omega)$. Taking the union over $b \in c_0$, we get: $R(\omega') \subseteq R(\omega)$.

For the converse inclusion, let b be an initial condition, and let $v = D_b(\omega') \cap R(\omega)$. Since $D_b(\omega')$ and $R(\omega)$ are two compatible well-stopped configurations, v is well-stopped (Ch. 3, VI-2.3). It follows from I-1.3 that v leads back to M_0 , and it follows from I-1.2 that $b \notin \gamma(v)$. Since $v \subseteq R(\omega) \subseteq \omega$, we have $v \in C_b(\omega)$ and thus $v \supseteq D_b(\omega)$. It implies that $D_b(\omega') \supseteq D_b(\omega)$, from which follows that $R(\omega') \supseteq R(\omega)$, and finally: $R(\omega) = R(\omega')$. \square

I-1.7 Remark. The cut γ_R , if defined, is equivalent to the initial cut c_0 . The shift operator θ_R associated with R (Ch. 5, II-5.1) is defined by $\theta_R(\omega) = \omega \ominus R(\omega)$. Modulo the isomorphism $\Omega \rightarrow \Omega^{\gamma_R}$, θ_R satisfies:

$$\forall \omega \in \Omega, \quad R(\omega) \notin \Omega \Rightarrow \theta_R(\omega) \in \Omega .$$

The shift operator associated with the renewal operator acts “almost” on Ω , whereas in general, shift operators do not act at all on Ω .

I-2 Successive Renewal Operators.

We introduce a very natural construction, that gives a sequence of stopping operators. It consists in applying recursively the renewal operator.

I-2.1 Successive Renewal Operators. As noticed in I-1.7, the renewal shift operator θ_R is not far from being defined $\Omega \rightarrow \Omega$. We can thus define inductively the sequence $(S_n)_{n \geq 0}$:

$$S_0 = \emptyset, \quad S_{n+1} = S_n \oplus R \circ \theta_{S_n}, \quad \forall n \geq 0 . \tag{6.2}$$

More precisely, we set:

$$S_{n+1}(\omega) = \begin{cases} \omega & \text{if } S_n(\omega) \in \Omega, \\ S_n(\omega) \oplus R(\omega \ominus S_n(\omega)) & \text{if } S_n(\omega) \notin \Omega . \end{cases}$$

$S_1 = R$ is the first well-stopped sub-configuration of ω that leaves the initial cut and returns to the initial marking. Then $\xi = \omega \ominus S_1(\omega)$ is a maximal configuration of $\mathcal{U}^{\gamma_R(\omega)} \rightarrow \mathcal{U}$, so ξ is an element of Ω . We apply the renewal operator to ξ to get $S_2(\omega) = S_1(\omega) \oplus R(\xi)$, and we repeat the process: subtract $R(\xi)$ from ξ , get an element of Ω , *etc.* See an illustration of the process in Figure 6.1.

It is straightforward to see that all S_n are well-stopping operators satisfying $\gamma(S_n) \cong c_0$. The shift operators are defined $\theta_{S_n} : \Omega \rightarrow \Omega$, with the same restrictions than for R (Cf. Remark I-1.7). Remark that if $S_j(\omega) = \omega$, then $S_n(\omega) = \omega$ for all $n \geq j$.




Figure 6.1: The successive renewal operators S_n .

I-2.2 Remark. The construction of the successive renewal operators is general. Although we have used the particularity that the renewal operator is defined $\Omega \rightarrow \Omega$, what matters is that R is a *universal operator*: R acts on Ω^v for all finite configurations v of the unfolding \mathcal{U} , and the action on Ω^v is a well-stopping operator *w.r.t.* the cone \mathcal{U}^v , for every v . In this case, the recursive formulas:

$$S_1 = R, \quad S_n = S_{n-1} \oplus R \circ \theta_{S_n},$$

define as sequence S_n of well stopping operators. This has to be compared with the well-known result for stopping times in sequential systems: if S, T are two stopping times, then $S + T \circ \theta_S : \omega \rightarrow S(\omega) + T(\theta_S(\omega))$ is a stopping time.

I-2.3 Lemma. Let (\mathcal{N}, M_0) be a safe marked net. Then for every $\omega \in \Omega$, $(S_n(\omega))_{n \geq 0}$ is non-decreasing in \mathcal{X} , and converges to ω .

Proof – Let $v_n = S_n(\omega)$. That $(v_n)_n$ is non-decreasing follows from (6.2). To show the convergence to ω , and since $v_n \subseteq \omega$ for all $n \geq 0$, we assume without loss of generality that $S_n(\omega) \notin \Omega$ for all $n \geq 0$. Let $v = \bigcup_n v_n$, we show that v is maximal.

Assume that v is not maximal. Then there is an event $e \notin v$, compatible with v and such that $\bullet e \subseteq v$. Then every condition $b \in \bullet e$ is maximal in v (otherwise $e \in v$). Let $b \in \bullet e$, then there is an integer n with $b \in v_n$, and b is maximal in v_n , which implies: $b \in \gamma(v_n)$. The same holds for v_{n+1} , hence $\gamma(v_n)$ and $\gamma(v_{n+1})$ share condition b . It follows that $\gamma(v_{n+1} \ominus v_n)$ owns an initial condition of \mathcal{U}^{v_n} , which contradicts $v_{n+1} \ominus v_n = R(\omega \ominus v_n)$. This shows that v is maximal, and thus $\bigcup_n v_n = \omega$. \square

II—Recurrent Nets

Using the renewal operators introduced above, we adapt to probabilistic nets the notion of recurrent state of a Markov chain. This leads to define recurrent markings and recurrent nets. We show the counterpart of some well-known results for Markov chains; in particular, markings return infinitely often with probability 0 or 1.

In II-1, we recall the notion of recurrent states for Markov chains, and state the analogous for nets. We establish the 0-1 alternative for markings. In II-2, we show that markings reachable from a recurrent marking are recurrent, an other well-known result for Markov chains.

II-1 Recurrent Nets.

We first recall the notion of recurrent state for Markov chains.

II-1.1 Recurrent States of Markov Chains. We write below Proposition 1.2, p. 64, of [37]. We use the notations of Ch. 1, III-1.1, concerning a Markov chain $(X_n)_{n \geq 0}$ on a finite set E , defined on the canonical measurable space $(\mathcal{A}, \mathcal{F})$.

For every $x \in E$, there are only two possibilities:

- (i) $\mathbb{P}_x(\overline{\lim}_n \{X_n = x\}) = 0$,
- (ii) $\mathbb{P}_x(\overline{\lim}_n \{X_n = x\}) = 1$.

We recall that the limit sup of a sequence of subsets $A_n \subseteq \mathcal{A}$ is defined as follows (Cf. Ch. 4, I-5.1):

$$\begin{aligned} \overline{\lim}_n A_n &= \bigcap_{N \geq 0} \bigcup_{n \geq N} A_n \\ &= \{\omega \in \Omega \mid \omega \in A_n \text{ for infinitely many } n\} \\ &= \{A_n \text{ i.o.}\}, \end{aligned}$$

where “ A_n i.o.” is to be read: “ A_n infinitely often”.

In case (i) the state x is said to be *transient*, in case (ii) x is said to be *recurrent*. In the framework of safe nets, and with our definition of renewal, we are brought to the question of the infinite return of the initial marking, with an action on all tokens at each passage.

We state the counterpart of II-1.1 for nets as follows.

II-1.2 Proposition. Let $(\mathcal{N}, \mathbb{P})$ be a compact probabilistic net. Assume that \mathbb{P} is d -homogeneous, and let $(S_n)_{n \geq 1}$ denote the successive renewal operators of \mathcal{N} . Then there are only two possibilities:

- (i) $\mathbb{P}(\overline{\lim}_n \{S_n \notin \Omega\}) = 0$,
- (ii) $\mathbb{P}(\overline{\lim}_n \{S_n \notin \Omega\}) = 1$.

Case (ii) holds if and only if we have: $R \notin \Omega$, \mathbb{P} -a.s.

Proof— We follow the proof of [37]. Let $n \geq 0$. From (6.2) we have:

$$\{S_{n+1} \notin \Omega\} = \{S_n \notin \Omega\} \cap \{R \circ \theta_{S_n} \notin \Omega\}. \quad (6.3)$$

We use the following property of conditional expectations:

$$\mathbb{E}(h) = \mathbb{E}(\mathbb{E}(h | \mathcal{G})),$$

applied with $h = \mathbf{1}_{\{S_{n+1} \notin \Omega\}}$ and $\mathcal{G} = \mathcal{F}_{S_n}$ to get from (6.3):

$$\begin{aligned} \mathbb{P}(S_{n+1} \notin \Omega) &= \mathbb{E}(h) = \mathbb{E}(\mathbb{E}(\mathbf{1}_{\{S_n \notin \Omega\}} \mathbf{1}_{\{R \circ \theta_{S_n} \notin \Omega\}} | \mathcal{F}_{S_n})) \\ &= \mathbb{E}(\mathbf{1}_{\{S_n \notin \Omega\}} \mathbb{E}(\mathbf{1}_{\{R \circ \theta_{S_n} \notin \Omega\}} | \mathcal{F}_{S_n})), \end{aligned}$$

the later since $\{S_n \notin \Omega\}$ is \mathcal{F}_{S_n} -measurable, and by applying the property that \mathcal{F}_{S_n} -measurable functions pass through the conditional expectation ((4.5) in Ch. 4, I-3.4). By the well-stopped Markov property (Ch. 5, Th. III-2.4) we have:

$$\mathbb{E}(\mathbf{1}_{\{R \circ \theta_{S_n} \notin \Omega\}} | \mathcal{F}_{S_n}) = \mathbb{E}^{\gamma_{S_n}}(\mathbf{1}_{\{R \notin \Omega\}}) = \mathbb{P}(R \notin \Omega),$$

since $\gamma_{S_n} \cong c_0$. Let $a = \mathbb{P}(R \notin \Omega)$, we obtain:

$$\mathbb{P}(S_{n+1} \notin \Omega) = a \mathbb{P}(S_n \notin \Omega).$$

If $a = 1$ then $\mathbb{P}(S_n \notin \Omega) = 1$ for all $n \geq 0$ and thus $\mathbb{P}(S_n \notin \Omega \text{ i.o.}) = 1$, that is case (ii). If $a < 1$ then $\sum_{n \geq 0} \mathbb{P}(S_n \notin \Omega) < \infty$, and by the Borel-Cantelli Lemma (Ch. 4, I-5.2) it implies: $\mathbb{P}(S_n \notin \Omega \text{ i.o.}) = 0$, case (i). \square

II-1.3 Definition. (*Recurrent and transient nets*) A d -homogeneous probabilistic net $(\mathcal{N}, M_0, \mathbb{P})$ satisfying Proposition II-1.2 (ii) is said to be **recurrent**. In case (i), the net is said to be **transient**.

II-1.4 Example. Consider a d -homogeneous probabilistic net based on the net depicted in Figure 6.2 (notice that the net is confusion free), where the places with same labels (P and Q) are identified. Clearly, the net is recurrent if the branching probabilities are positive. Observe that there are maximal configurations ω satisfying $R(\omega) = \omega$, although they have probability zero. For instance:

$$\omega = \beta \oplus (a \oplus \alpha) \oplus (a \oplus \alpha) \oplus \cdots,$$

keeps a token frozen in place Q . Recurrence is thus an example of a natural property of nets, that needs to be formulated in the probabilistic framework.




Figure 6.2: A recurrent net. Places with same labels are identical.

II-2 Conservativity of Recurrence.

We continue our reading of Markov chains theory, with in mind the idea of applying it to probabilistic nets. A useful result is the conservativity of recurrence: a marking reached from a recurrent marking is recurrent itself. This is the topic of Proposition II-2.3.

Recall that a marking M is said to be positive well-stopped if there is a \mathbb{P} -positive configuration v , finite and well-stopped, leading to M .

For \mathbb{P} a d -homogeneous probability, and for M a positive well-stopped marking, \mathbb{P}^M denotes the probability on Ω^M given by any probabilistic future $\mathbb{P}^v = \mathbb{P}(\cdot \mid \Omega(v))$, for $v \in \mathcal{X}_0$, \mathbb{P} -positive and leading to M .

II-2.1 Lemma. *Let $(\mathcal{N}, \mathbb{P})$ be a d -homogeneous recurrent net, and let R and $(S_n)_{n \geq 0}$ denote respectively the renewal and the successive renewal operators. For each $k \geq 1$, we set up a random variable $\kappa_k : \Omega \rightarrow \mathcal{X}_0$, defined on $\overline{\lim}_n \{S_n \notin \Omega\}$ by:*

$$\kappa_k(\omega) = S_k(\omega) \ominus S_{k-1}(\omega).$$

Then $(\kappa_k)_{k \geq 1}$ is an i.i.d sequence of random variables with values in \mathcal{X}_0 , and with probability law the law of $\kappa_1 = R$.

Proof— Since the net is recurrent, $\mathbb{P}(\overline{\lim}_n \{S_n \notin \Omega\}) = 1$, hence $\kappa_k : \Omega \rightarrow \mathcal{X}_0$ is well defined. To show that $(\kappa_k)_{k \geq 1}$ is i.i.d, we show that for all $n \geq 2$, and for all tuples $(k_1, \dots, k_n) \in \mathcal{X}_0 \times \dots \times \mathcal{X}_0$, we have:

$$\mathbb{P}(\kappa_1 = k_1, \dots, \kappa_n = k_n) = \mathbb{P}(R = k_1) \dots \mathbb{P}(R = k_n). \tag{6.4}$$

Let $(k_1, \dots, k_n) \in \mathcal{X}_0 \times \dots \times \mathcal{X}_0$. We condition by $\mathcal{F}_{S_{n-1}}$ to get, using the properties of conditional expectations (4.4) and (4.5) of Ch. 4, I-3.4:

$$\begin{aligned} \mathbb{P}(\kappa_1 = k_1, \dots, \kappa_n = k_n) &= \mathbb{E}(\mathbf{1}_{\{\kappa_1=k_1\}} \dots \mathbf{1}_{\{\kappa_n=k_n\}}) \\ &= \mathbb{E}(\mathbf{1}_{\{\kappa_1=k_1\}} \dots \mathbf{1}_{\{\kappa_{n-1}=k_{n-1}\}} \mathbb{E}(\mathbf{1}_{\{\kappa_n=k_n\}} | \mathcal{F}_{S_{n-1}})) \end{aligned}$$

Let c_0 be the initial cut of the unfolding. We apply the well-stopped Markov property (Ch. 5, III-2.4) to get, since $\gamma(S_{n-1}) \cong c_0$, \mathbb{P} -a.s.:

$$\mathbb{E}(\mathbf{1}_{\{\kappa_n=k_n\}} | \mathcal{F}_{S_{n-1}}) = \mathbb{E}(\mathbf{1}_{\{R=k_n\}}) = \mathbb{P}(R = k_n) .$$

We obtain thus:

$$\mathbb{P}(\kappa_1 = k_1, \dots, \kappa_n = k_n) = \mathbb{E}(\mathbf{1}_{\{\kappa_1=k_1\}} \dots \mathbf{1}_{\{\kappa_{n-1}=k_{n-1}\}}) \mathbb{P}(R = k_n) .$$

By induction, (6.4) follows. \square

II-2.2 Lemma. *Let $(\mathcal{N}, \mathbb{P})$ be a d -homogeneous net. Let \mathcal{U} be the unfolding of \mathcal{N} , with initial cut c_0 , and let R be the renewal operator. Assume that an element $\omega \in \Omega$ contains a non-decreasing sequence of configurations $(v_n)_{n \geq 1}$, converging to ω and such that, for all $n \geq 1$:*

$$v_n \in \mathcal{X}_0, \quad \gamma(v_n) \cong c_0, \quad \gamma(v_n) \cap \gamma(v_{n-1}) = \emptyset .$$

Then $\omega \in \overline{\lim}\{S_n \notin \Omega\}$.

Proof – Assume that $\omega \notin \overline{\lim}\{S_n \notin \Omega\}$. There is then a maximal integer n such that $S_n \notin \Omega$. Since $v_p \uparrow_p \omega$, there is an integer p such that $v_p \supseteq S_n$. The existence of v_{p+1} implies $R(\omega \ominus S_n) \notin \Omega$, contradicting $S_{n+1}(\omega) = \omega$.

\square

II-2.3 Proposition. *Let $(\mathcal{N}, M_0, \mathbb{P})$ be a d -homogeneous recurrent net. Then for every positive well-stopped marking M , $(\mathcal{N}, M, \mathbb{P}^M)$ is recurrent.*

Proof – Let M be a positive well-stopped marking, and assume that $(\mathcal{N}, M, \mathbb{P}^M)$ is transient. Let v be a finite well-stopped configuration of \mathcal{U} , with positive likelihood and leading to M . According to Lemma I-2.3, we have:

$$\Omega(v) = \bigcup_{p \geq 1} \{S_p \supseteq v\} .$$

Since $\mathbb{P}(\Omega(v)) > 0$, it implies there is an integer $p \geq 1$ such that: $\mathbb{P}(S_p \supseteq v) > 0$. Referring to the notations of II-2.1, we set the sequence of random variables $(K_n)_{n \geq 1}$ by:

$$K_n = S_{np} \ominus \dots \ominus S_{(n-1)p} = \kappa_{(n-1)p+1} \oplus \dots \oplus \kappa_{np} .$$

It follows from II-2.1 that $(K_n)_{n \geq 1}$ is *i.i.d.*, with probability law the law of $K_1 = S_p$. Let $U_n = \{K_n \supseteq v\}$. Every $\omega \in \overline{\lim}_n U_n$ contains a sequence of sub-configurations $(v_q)_{q \geq 1}$ satisfying:

$$v_q \in \mathcal{X}_0, \quad v_q \uparrow_q \omega, \quad \gamma(v_q) \cong \gamma(v), \quad \gamma(v_q) \cap \gamma(v_{q-1}) = \emptyset. \quad (6.5)$$

The sequence of subsets $(U_n)_{n \geq 1}$ is independent and $\mathbb{P}(U_n) = \mathbb{P}(S_p \supseteq v) > 0$, a positive constant. The Borel-Cantelli lemma (I-5.2) implies: $\mathbb{P}(\overline{\lim}_n U_n) = 1$.

Let $\xi \in \Omega^M$. Then $\omega = v \oplus \xi$ is \mathbb{P}^M -*a.s.* in $\overline{\lim}_n U_n$. Using the sequence $(v_q)_{q \geq 1}$ associated with ω as in (6.5), we set $y_q = v_q \ominus v$, and we have:

$$y_q \in \mathcal{X}_0, \quad y_q \uparrow \xi, \quad \gamma(y_q) \cong \gamma(v), \quad \gamma(y_q) \cap \gamma(y_{q-1}) = \emptyset.$$

By Lemma II-2.2, it implies that the successive renewal operators S_n^M in Ω^M , satisfy:

$$\mathbb{P}^M(\overline{\lim}\{S_n^M \notin \Omega^M\}) = 1.$$

This shows that $(\mathcal{N}, M, \mathbb{P}^M)$ is recurrent. \square

III—Local Renewal.

In the framework that we have proposed in Chapter 3 for the analysis of processes, the concurrency of processes is reflected through the concurrency of germs and branching cells. Since germs are local processes, is it possible to define local states and a local renewal? We study this question for distributed probabilities, for which the concurrency of germs has a simple probabilistic expression.

The successive germs of a process in a same class of branching cell modulo isomorphism form a sequence of independent and identically distributed random variables. This is the main result of this section, only valid for distributed probabilities, and which has important consequences for the statistical estimation of parameters (See Ch. 8).

The techniques involved are similar to the techniques used in the study of the global renewal. In particular we define *local* and *successive local* renewal operators. We show that these stopping operators are adapted for the study of the local renewal of *distributed probabilities*. Again, the main tool is the Markov property for concurrent systems. Note that the distinction between global and local renewal is trivial for sequential systems.

III-1 defines the local renewal operator associated with a class of branching cell. The successive local renewal operators are defined in III-2, together with the successive *l*-branching cells of a maximal process ω . The successive *l*-germs are studied in III-3 for distributed probabilities.

III-1 Local Renewal Operator.

III-1.1 Hitting a Class of Branching Cells. Let (\mathcal{N}, M_0) be a safe marked net, with unfolding \mathcal{U} and event structure \mathcal{E} . Let l be an equivalence class of branching cell: $l \in \mathcal{L}_{\mathcal{U}}$, with the notation of Chapter 5. For λ a branching cell of \mathcal{U} , we write $\lambda \cong l$ to denote that λ is in the class of l . We have remarked that l actually identifies with a labelled occurrence net, up to a unique isomorphism, such that if $\lambda \cong l$, there is a unique isomorphism of labelled occurrence nets $l \rightarrow \lambda$ (Ch. 5, IV-2.1). The isomorphism induces through the action on sets a one-to-one mapping between the finite sets $\Omega_l \rightarrow \Omega_\lambda$.

Let $l \in \mathcal{L}_{\mathcal{U}}$. For ω an element of Ω , we want to identify the “first instant” where ω “activates” a branching cell $\lambda \cong l$. We set:

$$F^l(\omega) = \{v \in \mathcal{X}_0, v \subseteq \omega \mid \exists \lambda \in \Delta^+(v) : \lambda \cong l, \lambda \notin \Delta^\perp(\mathcal{E})\}. \quad (6.6)$$

F^l contains the sub-configurations of ω that activate a non-initial branching cell equivalent to l (Cf. Ch. 3, Table 3.1, p. 3.1, for the notations). We consider non initial branching cells in order to force the process to go forward.

III-1.2 Lemma. *For every $l \in \mathcal{L}_{\mathcal{U}}$ and for every $\omega \in \Omega$, $F^l(\omega)$ is a lattice.*

Proof— Let $\omega \in \Omega$, we note $F^l = F^l(\omega)$. We show that F^l is stable under pairwise intersections. Let $v, v' \in F^l$ and λ, λ' be the associated branching cells as in (6.6), and let $y = v \cap v'$. Then y is well-stopped (Ch. 3, VI-2.3). There are two cases.

If $\lambda = \lambda'$ then $\lambda \in \Delta^+(y)$ since $\text{Min}_{\preceq}(\lambda) \subseteq \gamma(y)$ by I-1.2, and thus we have: $y \in F^l$.

Otherwise $\overset{\circ}{\lambda} \cap \overset{\circ}{\lambda'} = \emptyset$, the branching cells have disjoint events by Ch. 3, VI-3.4. $\lambda \cap \omega$ and $\lambda' \cap \omega$ cannot be concurrent since their minimal elements have same labels, and the labelling is injective on \parallel -cliques. It follows that there is condition $b \in \text{Min}_{\preceq}(\lambda)$ and a condition $b' \in \text{Min}_{\preceq}(\lambda')$ such that, say: $b \preceq b'$. It implies that $v' \cap \overset{\circ}{\lambda} \neq \emptyset$ and thus $v' \cap \lambda$ is maximal in λ . In particular $\text{Min}_{\preceq}(\lambda) \subseteq v'$. Applying Lemma I-1.2, we get that $\text{Min}_{\preceq}(\lambda) \subseteq \gamma(y)$, and thus $\lambda \in \Delta^+(y)$. Since $\lambda \notin B^\perp(\mathcal{E})$, we have that $y \in F^l$.

The stability under union is shown in the same way. \square

The meaning of Lemma III-1.2 is clear: if a branching cell is activated at two compatible instants, it cannot disappear, neither through union nor through intersection. The above proof also shows that $\mathcal{L}^l(\omega)$ is totally ordered:

III-1.3 Lemma. *For $l \in \mathcal{L}_{\mathcal{U}}$, let $\mathcal{L}^l(\omega)$ be the set of branching cells of $\overline{\Lambda}(\omega)$ equivalent to l , that is:*

$$\mathcal{L}^l(\omega) = \{\lambda \in \overline{\Lambda}(\omega) \mid \lambda \cong l\}.$$

Then \mathcal{L}^l is totally ordered by the relation \preceq_ω , defined by for $\lambda, \lambda' \in \mathcal{L}^l(\omega)$ by $\lambda \preceq_\omega \lambda'$ if and only if:

$$\exists c \in \text{Min}_{\preceq}(\lambda), \exists c' \in \text{Min}_{\preceq}(\lambda') : c \preceq c'. \quad (6.7)$$

Since $F^l(\omega)$ may be empty, we define the local renewal operator associated with l as follows.

III-1.4 Definition. (*Local renewal operator*) Let \mathcal{U} be the unfolding of a safe marked net \mathcal{N} , and let $l \in \mathcal{L}_{\mathcal{U}}$, a class of branching cells. We define the local renewal operator of l by:

$$R^l(\omega) = \begin{cases} \min(F^l(\omega)), & \text{if } F^l(\omega) \neq \emptyset, \\ \omega, & \text{if } F^l(\omega) = \emptyset. \end{cases}$$

It is straightforward to check that R^l is a well stopping operator. Remark that, in general, $F^l(\omega)$ does not leave the initial cut, in the sense of I-1.4.

III-2 Successive Local Renewals.

The technique for defining successive local renewal operators is the same than for the global renewal.

III-2.1 Successive Local Renewals. For $l \in \mathcal{L}_{\mathcal{U}}$, we define the successive renewal $W_n^l : \Omega \rightarrow \mathcal{X}$ as in I-2.1 and I-2.2: the local renewal operator is indeed a universal well stopping operator. Recall that the shift operator θ_V associated with a stopping operator V is defined by $\theta_V(\omega) = \omega \ominus V(\omega)$, a configuration of $\mathcal{E}^{V(\omega)}$ (Ch. 5, II-4.2).

For a short definition, we set the **successive local renewal** by $W_1^l = R^l$, and for $n \geq 1$:

$$W_n^l = W_{n-1}^l \oplus R^l \circ \theta_{W_{n-1}^l}.$$

More precisely, we set for $n \geq 1$:

$$W_n^l(\omega) = \begin{cases} W_{n-1}^l(\omega) \oplus R^l(\theta_{W_{n-1}^l}(\omega)), & \text{if } W_{n-1}^l(\omega) \notin \Omega, \\ \omega, & \text{if } W_{n-1}^l(\omega) = \omega. \end{cases}$$

Remark that $R^l(\theta_{W_{n-1}^l}(\omega)) = R^l(\omega \ominus W_{n-1}^l(\omega))$ is the result of *the local renewal operator* R^l in $\mathcal{U}^{\gamma_{W_{n-1}^l}(\omega)}$, not in \mathcal{U} .

We formalise the “infinite return of local state” as the property that the above construction always remains in the case: $W_{n-1}^l(\omega) \notin \Omega$, for ω fixed, and for \mathbb{P} -a.s all ω . Equivalently, the set $\overline{\lim}_n \{W_n \notin \Omega\} = \{W_n \notin \Omega \text{ i.o.}\}$ has probability 1.

The following proposition says that, in a recurrent and d -homogeneous net, if a branching cell is activated once, it will be activated infinitely often, and in \mathbb{P} -a.s all maximal processes. Remark that we do not require \mathbb{P} to be a distributed probability.

III-2.2 Proposition. *Let $(\mathcal{N}, \mathbb{P})$ be a d -homogeneous probabilistic net, with \mathcal{U} the unfolding. If $(\mathcal{N}, \mathbb{P})$ is recurrent, then every positive class $l \in \mathcal{L}_{\mathcal{U}}$ (i.e.: $\exists v \in \mathcal{X}_0 : p(v) > 0, l \in \Delta^+(v)$) satisfies:*

$$\mathbb{P}(\overline{\lim}_n \{W_n^l \notin \Omega\}) = 1.$$

Proof – We proceed as in Proposition II-2.3. We use the sequence of *i.i.d* random variables $(\kappa_k)_{k \geq 1}$, defined in II-2.1 by $\kappa_k = S_k \ominus S_{k-1}$, with S_k the successive renewals of \mathcal{N} . Let $v \in \mathcal{X}_0$ with positive likelihood ($p(v) > 0$) and such that $l \in \Delta^+(v)$. Since \mathcal{N} is recurrent, we have $\Omega(v) = \bigcup_{k \geq 1} \{S_k \supseteq v\}$, as a consequence of I-2.3. As $\mathbb{P}(\Omega(v)) > 0$, there is thus an integer q such that $\mathbb{P}(S_q \supseteq v) > 0$. Then we consider the *i.i.d* sequence $K_n = S_{nq} \ominus S_{(n-1)q}$, with law S_q in \mathcal{X}_0 . Since $(K_n)_{n \geq 1}$ is independent, the Borel-Cantelli lemma implies that $\overline{\lim} \{K_n \supseteq v\}$ has probability 1.

We have shown for \mathbb{P} -*a.s* all ω , ω contains infinitely many branching cells $\lambda \cong l$. By a lemma analogous to Lemma II-2.2, it implies that \mathbb{P} -*a.s* all ω belong to $\overline{\lim} \{W_n^l \notin \Omega\}$. \square

III-2.3 The Successive l -Branching Cells. Assume that $(\mathcal{N}, \mathbb{P})$ is d -homogeneous recurrent, and let $l \in \mathcal{L}_{\mathcal{U}}$ be a positive equivalence class of branching cell. For every $n \geq 1$, and for \mathbb{P} -*a.s* all ω , $W_n^l(\omega) \subsetneq \omega$, and there is a branching cell $l_n \in \Delta^+(W_n^l(\omega))$ with $l_n \cong l$. For each n , there is a unique l_n that satisfies these two conditions. We say that the sequence of branching cells $(l_n)_{n \geq 1}$ is the sequence of **successive l -branching cells of ω** .

We have the following result, where $\mathcal{L}^l(\omega)$ denotes as in Lemma III-1.3 the branching cells of $\overline{\Lambda}(\omega)$ equivalent to l . Basically, the result is due to the safety assumption on the net.

III-2.4 Lemma. *Let $(\mathcal{N}, \mathbb{P})$ be a d -homogeneous recurrent net. Let $l \in \mathcal{L}_{\mathcal{U}}$ be a class of branching cell of the unfolding, and let $l_n(\omega)$ denote the successive occurrences of l of ω , defined for all $n \geq 1$ and for \mathbb{P} -*a.s* all ω . We have the equality of sets:*

$$\mathcal{L}^l(\omega) = \{l_n(\omega), n \geq 1\},$$

where we add $\{\lambda_0\}$ to the right hand member if there is a $\lambda_0 \in \Delta^\perp(\mathcal{E})$ such that $\lambda_0 \cong l$.

Proof – For every $n \geq 1$, the branching cell $l_{n+1}(\omega)$ is the successor of $l_n(\omega)$ in the total order $\mathcal{L}^l(\omega)$, *w.r.t.* the relation \preceq_ω defined in III-1.3. The result of the lemma follows. \square

III-3 The Case of Distributed Probabilities: Successive l -Germs.

With the successive l -branching cells of a maximal process ω , we want now to study the associated l -germs of ω . For this, it is natural to assume that the proba-

bility is distributed, *i.e.* that the net is a distributed Markov net in the sense seen in Chapter 5.

III-3.1 The Successive l -Germs. Let $(\mathcal{N}, \mathbb{P})$ be a d -homogeneous probabilistic net, and assume that $(\mathcal{N}, \mathbb{P})$ is recurrent. Fix a class of branching cells $l \in \mathcal{L}_{\mathcal{U}}$, and let $l_n(\omega)$ denote the successive l -branching cells of ω , defined for \mathbb{P} -a.s ω and for all $n \geq 1$ as seen in III-2.3. We have $l_n(\omega) \in \overline{\Lambda}(\omega)$, and since ω is maximal we have $\omega \cap l_n(\omega) \in \Omega_{l_n(\omega)} = \Omega_l$ for all n . We define thus a sequence of l -germs by setting:

$$\forall n \geq 1, \quad Z_n^l(\omega) = \omega \cap l_n(\omega) \in \Omega_l .$$

We can go further in the analysis if we consider distributed probabilities.

III-3.2 Theorem. *Let $(\mathcal{N}, (\mu_l)_{l \in \mathcal{L}_{\mathcal{U}}})$ be a distributed Markov net, with \mathcal{U} the unfolding, and let \mathbb{P} be the associated distributed probability. Assume that $(\mathcal{N}, \mathbb{P})$ is recurrent. Then for every positive class $l \in \mathcal{L}_{\mathcal{U}}$, the sequence $(Z_n^l)_{n \geq 1}$ of successive l -germs is *i.i.d* with law μ_l in Ω_l .*

We begin with a lemma.

III-3.3 Lemma. *With the notations of Theorem III-3.2, the probability law of Z_1^l in Ω_l is μ_l .*

Proof – We note $Z = Z_1^l$; we have to show that $\mathbb{P}(Z = z) = \mu_l(z)$ for all $z \in \Omega_l$. Let $z \in \Omega_l$, we will apply the well-stopped Markov property to the local renewal operator R^l . Let \mathcal{D} denote the class of equivalent cuts, and let $g = (g^c)_{c \in \mathcal{D}}$ be the d -homogeneous function defined by:

$$\forall c \in \mathcal{D}, \forall \xi \in \Omega^c, \quad g^c(\xi) = \begin{cases} 0, & \text{if: } \nexists \lambda \in \Delta^\perp(\mathcal{U}^c), \lambda \cong l, \\ \mathbf{1}_{\{\xi \cap \lambda = z\}}, & \text{if: } \exists \lambda \in \Delta^\perp(\mathcal{U}^c), \lambda \cong l. \end{cases}$$

Remark that, if there is a $\lambda \in \Delta^\perp(\mathcal{U}^c)$ with $\lambda \cong l$, then λ is unique, hence $\xi \cap \lambda \in \Omega_\lambda$ is well defined. We have the equality of sets:

$$\begin{aligned} \{Z = z\} &= \{\omega \in \Omega \mid g(\omega \ominus R^l(\omega)) = 1\} \\ &= \{g \circ \theta_{R^l} = 1\}, \end{aligned}$$

where θ_{R^l} denotes the shift operator associated with R^l (Ch. 5, II-5.1). We apply the well stopped Markov property (Ch. 5, III-2.4) to g and R^l to get:

$$\begin{aligned} \mathbb{P}(Z = z) &= \mathbb{E}(g \circ \theta_{R^l}) \\ &= \mathbb{E}(\mathbb{E}(g \circ \theta_{R^l} \mid \mathcal{F}_{R^l})), && \text{conditioning by } \mathcal{F}_{R^l}, \\ &= \mathbb{E}(\mathbb{E}^{\gamma_{R^l}}(g^{\gamma_{R^l}})), && \text{by the Markov property.} \end{aligned} \quad (6.8)$$

To compute this expectation, let c be a well-stopped cut, such that $\lambda \cong l$ for a $\lambda \in \Delta^\perp(\mathcal{U}^c)$. Then we have:

$$\mathbb{E}^c(g^c) = \mathbb{P}^c(\xi \cap \lambda = z),$$

where ξ denotes the variable in Ω^c . Since \mathbb{P} is distributed, and c is well-stopped, \mathbb{P}^c is distributed, and induces the same branching probabilities $(\mu_\lambda)_\lambda$ than \mathbb{P} . It follows that $\mathbb{P}^c(\xi \cap \lambda = z) = \mu_l(z)$. Hence $c \mapsto \mathbb{E}^c(g^c)$ is constant, equal to $\mu_l(z)$. From (6.8), it follows that: $\mathbb{P}(Z = z) = \mu_l(z)$, what was to be shown. \square

III-3.4 Proof of III-3.2. We fix $l \in \mathcal{L}_\mathcal{U}$ and $n \geq 1$, and we have to show that for every $n \geq 0$, and for every tuple $(z_1, \dots, z_n) \in \Omega_l \times \dots \times \Omega_l$, the following holds:

$$\mathbb{P}(Z_1^l = z_1, \dots, Z_n^l = z_n) = \mu_l(z_1) \dots \mu_l(z_n). \quad (6.9)$$

We will apply the well stopped Markov property (Ch. 5, III-2.4). Let g be the homogeneous function defined as in the proof of Lemma III-3.3, with $z = z_n$, that is for every cut c :

$$\forall \xi \in \Omega^c, \quad g^c(\xi) = \begin{cases} 1, & \text{if } l \in \Delta^\perp(\mathcal{U}^c) \text{ and if } \xi \cap l = z_n, \\ 0, & \text{otherwise.} \end{cases}$$

Assume for the moment that the following holds:

$$\mathbb{E}(g \circ \theta_{W_n} \mid \mathcal{F}_{W_{n-1}}) = \mu_l(z_n). \quad (6.10)$$

Let $q = \mathbb{P}(Z_1^l = z_1, \dots, Z_n^l = z_n)$. We have by the usual transformations:

$$\begin{aligned} q &= \mathbb{E}(\mathbf{1}_{\{Z_1^l = z_1\}} \dots \mathbf{1}_{\{Z_{n-1}^l = z_{n-1}\}} \mathbb{E}(g \circ \theta_{W_n} \mid \mathcal{F}_{W_{n-1}})) \\ &= \mathbb{P}(Z_1^l = z_1, \dots, Z_{n-1}^l = z_{n-1}) \mu_l(z_n), \end{aligned}$$

the later by (6.10). By induction, (6.9) follows.

It remains to show (6.10). We note $W_n = W_n^l$ and $R = R^l$. Then we have $W_n = W_{n-1} \oplus R \circ \theta_{W_{n-1}}$, from which follows:

$$\begin{aligned} \omega \ominus W_n &= (\omega \ominus W_{n-1}) \ominus R \circ \theta_{W_{n-1}} \\ \theta_{W_n} &= \theta_{W_{n-1}} \ominus R \circ \theta_{W_{n-1}}. \end{aligned} \quad (6.11)$$

We set the d -homogeneous function h , defined by $h(\xi) = g(\xi \ominus R(\xi))$. Applying the well stopped Markov property, and by (6.11), we get:

$$\begin{aligned} \mathbb{E}(g \circ \theta_{W_n} \mid \mathcal{F}_{W_{n-1}}) &= \mathbb{E}(g(\theta_{W_{n-1}} \ominus R \circ \theta_{W_{n-1}}) \mid \mathcal{F}_{W_{n-1}}) \\ &= \mathbb{E}(h \circ \theta_{W_{n-1}} \mid \mathcal{F}_{W_{n-1}}) \\ &= \mathbb{E}^{\gamma_{W_{n-1}}}(h^{\gamma_{W_{n-1}}}) \\ &= \mathbb{P}^{\gamma_{W_{n-1}}}(Z_1^l = z_n). \end{aligned} \quad (6.12)$$

Remark that the right hand member of (6.12) is the composite of $\omega \in \Omega \rightarrow \gamma_{W_{n-1}}$ and $c \rightarrow \mathbb{P}^c(Z_1^l = z_n)$, where Z_1^l represents the first l -germ in \mathcal{U}^c . For every well stopped cut c , we apply Lemma III-3.3 in the cone of future \mathcal{U}^c , since \mathbb{P}^c is distributed and consistent with the family $(\mu_l)_{l \in \mathcal{L}_U}$, to get that the probability law of Z_1^l is μ_l . Hence for every well-stopped cut c , $\mathbb{P}^c(Z_1^l = z_n) = \mu_l(z_n)$. With (6.12), it implies (6.10). \square

IV—Conclusion

We have shown that the Markov property for concurrent systems is a basis for a recurrence theory of probabilistic nets. We have established for compact d -homogeneous nets some results on recurrent markings, classical for finite Markov chains: alternative 0-1 between recurrent and transient markings, conservativity of recurrence.

We have also introduced the notion of local renewal. Defining a local renewal operator, we establish the character *i.i.d* of the sequence of successive l -germs, for l a class of branching cell.

The “next” result to be established, and that we miss, is the *positive recurrence of markings*. In a finite recurrent Markov chain, not only the states are recurrent, they also come back within a random time with *finite expectation*. An analogous can certainly be expected for concurrent systems. We will face this difficulty in next chapter.

Chapter 7

Ergodicity and Limit Theorems

Among the most powerful results from probability theory are certainly the so-called limit theorems. Limit theorems as the Strong law of large numbers and the Central Limit Theorem, and more generally *ergodic properties*, characterise the asymptotic behaviour of a dynamical system. Since we have adapted until a certain point the recurrence theory of finite Markov chains to concurrent systems, the question of limit theorems becomes natural.

The aim of this chapter is first to propose a formulation of the Strong law of large numbers for concurrent systems, in the model of compact probabilistic nets. In concurrent systems, unlike in sequential systems, the appropriate unit of time is not the event. For instance the prefixes of an event structure formed by chains of events with bounded length do not define interesting prefixes from the dynamics point of view, since in general they are not intrinsic. For distributed probabilities, branching cells represent the atomic probabilistic “actions”, related to the local resolution of conflicts. Therefore we adopt the branching cell as a unit of time, by simply counting the branching cells that the process passes through. The difference with sequential systems is that causally related events can be involved in a same atomic action. With this new unit of time, we propose a statement for the Strong law of large numbers.

For this we also introduce *distributed functions*. A distributed function can be seen as a local test function, that outputs for instance 1 each time the net satisfies a given local property. By summing the outputs of the distributed function along a finite process, we count the number of occurrences of the local property along the finite process. For instance the time elapsed along a process is obtained by summing the distributed function with constant value 1. The Strong law of large numbers for concurrent systems that we propose gives the limit of the summed outputs of any distributed function, normalised by the growth of concurrent time for the same process. The limit is almost sure, and is taken with the process growing to its maximal value.

The stationary measure of an ergodic finite Markov chain, that gives the asymptotic repartition of the chain through its set of states, becomes in the concurrent framework a measure on the finite set of branching cells (modulo isomorphism), that

is a *density coefficient* for each class of branching cell. The asymptotic ratios of distributed functions have then two components: a *local* component, concerned by the local branching probability, and a global component, given by the density coefficients of branching cells. This only holds for *distributed Markov nets* and not for general concurrent memory-less systems. For distributed probabilities, the interactions between all components reduce to the finite collection of density coefficients.

To establish the Strong law of large numbers, we use the embedded Markov chain of a safe net as an auxiliary tool. To derive an intrinsic result from the ergodic result for the embedded Markov chain, we need to analytically control the concurrency range of the system. For this we introduce the notion of concurrent height of a net, given by an integer random variable, and the definition of a net with *integrable concurrent height*. With the recurrence property, the integrable concurrent height is a second example of an analytical tool for studying concurrent systems. For instance a *bounded* concurrent height would be too much restricting. Although it is not shown, we can expect that every net decomposes through components with integrable concurrent height.

We show two ergodic results directly on nets. A first lemma translates for nets a useful lemma, often used in the theory of dynamical systems. A second result establishes, with the vocabulary of ergodic theory, that recurrent nets are ergodic. To obtain the Strong law of large numbers we use two other tools: a simplification derives from the properties of distributed probabilities, and the embedded Markov chain brings the last ingredient. Finally the state-of-work concerning the use of Martingales is the topic of a separated section. The aim is to obtain the Central Limit Theorem in a future work.

In Section I, *Ergodic means*, we recall the form of ergodic means for sequential systems, and the convergence result for finite ergodic Markov chains. We introduce the *distributed functions* and we define *concurrent ergodic means*. In II, *The Strong law of large numbers*, we introduce the remaining material needed to establish the Strong law of large numbers, in particular the notion of net with *integrable concurrent height*.

I—Ergodic Means

This section describes the form of ergodic means for sequential systems like Markov chains, and introduces an analogous notion for concurrent systems. Concurrent ergodic means are based on the notion of *distributed function*. Both forms of ergodic means are identical for a sequential net. We also introduce the “unit of time” that we use, obtained by counting branching cells along a process. We show that, due to concurrency, to expect the Strong law of large numbers to hold, it is natural to consider processes that grow to infinity in a “regular” way.

In I-1, we recall the form of sequential ergodic means and the convergence result for ergodic means of a finite ergodic Markov chain. We introduce in I-2 the form that we consider for concurrent ergodic means. Then we state in I-3 the Strong law of large numbers for concurrent systems, to be established later.

I-1 Sequential Ergodic Means.

I-1.1 The Strong Law of Large Numbers for Markov Chains. Let $X = (X_n)_{n \geq 1}$ be a Markov chain on a finite set E , starting from state x_0 . Let $f : E \rightarrow \mathbb{R}$ be a function. Assume that X is recurrent (all the states $x \in E$ are recurrent) and irreducible (there is a positive path from x to y for every pair $x, y \in E$). Then the ergodic theory of stationary processes ([10, 36]) states that the ergodic means:

$$\frac{1}{n} \sum_{k=1}^n f(X_k) \quad (7.1)$$

converge for \mathbb{P}_{x_0} -a.s all ω , and the limit is given by $\int f d\mu_{x_0}$, where μ_{x_0} is a probability measure on E . Moreover the measure does not depend on x_0 , and thus the limit does not depend on x_0 . The measure $\mu = \mu_{x_0}$ on E is called the stationary measure of X . This result constitutes the Strong Law of large numbers, strong referring to the almost sure convergence. If X is not recurrent and irreducible, the convergence still holds, but the stationary measure depends on x_0 .

I-1.2 Measure of Time. In the model of nets for concurrent systems, since the nature of time varies from an execution to an other, what should be the unit of time seems not clear. We propose to use the tools of well-stopped configurations and branching cells for the definition of a unit of time.

If v is a finite well-stopped configuration of a locally finite event structure, we have defined in Ch. 4, III-3.3, the number of branching cells associated with v :

$$\langle N, v \rangle = \text{Card}(\Lambda_{\mathcal{E}}(v)) .$$

We will consider $\langle N, v \rangle$ as the time elapsed during the process v . We will also consider that $\langle N, \cdot \rangle$ is a function defined for all configurations of all locally finite event structures.

If y is a finite well-stopped configuration of the cone \mathcal{E}^v , then we have the disjoint union: $\Lambda_{\mathcal{E}}(v \oplus y) = \Lambda_{\mathcal{E}}(v) \sqcup \Lambda_{\mathcal{E}^v}(y)$ (Ch. 3, VI-3.7). It implies:

$$\langle N, v \oplus v' \rangle = \langle N, v \rangle + \langle N, v' \rangle ,$$

an additivity relation expected for a measure of time.

I-2 Distributed Functions and Concurrent Ergodic Means.

Let (\mathcal{N}, M) be a safe compact net, and let \mathcal{U} be the unfolding. Let $\Lambda_{\mathcal{U}}$ denote the set of branching cells of \mathcal{U} , and let $\mathcal{L}_{\mathcal{U}}$ denote the quotient modulo equivalence of branching cells, as in Ch. 5, IV-1.2. We recall that $\mathcal{L}_{\mathcal{U}}$ is a finite set.

A *distributed function* is like a local test function.

I-2.1 Definition. We define a *distributed function* of \mathcal{N} as a function defined on the disjoint union:

$$f : \bigsqcup_{\lambda \in \mathcal{L}_{\mathcal{U}}} \Omega_{\lambda} \rightarrow \mathbb{R} .$$

Equivalently, f is given by a collection $(f^{\lambda})_{\lambda \in \Lambda_{\mathcal{E}}}$ of functions $f^{\lambda} : \Omega_{\lambda} \rightarrow \mathbb{R}$, such that:

$$\forall \lambda, \lambda' \in \Lambda_{\mathcal{E}}, \quad \lambda \cong \lambda' \Rightarrow f^{\lambda} = f^{\lambda'} ,$$

modulo the unique isomorphism $\Omega_{\lambda} \rightarrow \Omega_{\lambda'}$ (Ch. 5, IV-1.3). We identify the distributed function f , the collection $(f^{\lambda})_{\lambda}$ with λ ranging over $\mathcal{L}_{\mathcal{U}}$ and the collection $(f^l)_l$ with l ranging over $\mathcal{L}_{\mathcal{U}}$.

If $f = (f^{\lambda})_{\lambda}$ is a distributed function, f defines a mapping $\langle f, \cdot \rangle : \mathcal{X}_0 \rightarrow \mathbb{R}$ by:

$$\forall v \in \mathcal{X}_0, \quad \langle f, v \rangle = \sum_{\lambda \in \Lambda(v)} f^{\lambda}(v \cap \lambda) , \quad (7.2)$$

where $\Lambda(v)$ denotes the dynamic puzzle around v (Ch. 3, VI-3.1). Since $v \cap \lambda$ is maximal in λ for every $\lambda \in \Lambda(v)$ (Ch. 3, VI-3.3), $f^{\lambda}(v \cap \lambda)$ is well-defined.

I-2.2 Remark. Distributed functions form a \mathbb{R} -vector space of finite dimension, equal to $\text{Card}(\mathcal{L}_{\mathcal{U}})(\sum_{l \in \mathcal{U}} \text{Card}(\Omega_l))$. The scalar notation of (7.2) is justified since $\langle \cdot, \cdot \rangle$ is indeed bilinear, *w.r.t.* the usual addition for the left argument, and *w.r.t.* the concatenation \oplus for the right argument.

I-2.3 Example. The function $\langle N, v \rangle$ of I-1.2 is given by the distributed function $(1)_{\lambda}$, with constant functions $\Omega_{\lambda} \rightarrow \mathbb{R}$ with value 1 for all λ .

I-2.4 Example. Let $\lambda_0 \in \mathcal{L}_{\mathcal{U}}$. We define the distributed function N^{λ_0} by $N^{\lambda_0} = (f^l)_{l \in \mathcal{L}_{\mathcal{U}}}$, with:

$$\forall l \in \mathcal{L}_{\mathcal{U}}, \quad \forall z \in \Omega_l, \quad f^l(z) = \begin{cases} 1, & \text{if } l = \lambda_0, \\ 0, & \text{if } l \neq \lambda_0. \end{cases}$$

Hence $\langle N^{\lambda_0}, \cdot \rangle$ counts the number of occurrences of branching cells within the class λ_0 . We have:

$$\forall v \in \mathcal{X}_0, \quad \langle N, v \rangle = \sum_{\lambda \in \mathcal{L}_{\mathcal{U}}} \langle N^{\lambda}, v \rangle. \tag{7.3}$$

I-2.5 Ergodic Means for Nets. We interpret the sequential ergodic means (7.1) as a mean along the path drawn by X_n in the covering tree. For concurrent systems, we are prompted to define ergodic means as a function $Tf : \mathcal{X}_0 \rightarrow \mathbb{R}$:

$$\forall v \in \mathcal{X}_0, \quad Tf(v) = \frac{1}{\langle N, v \rangle} \langle f, v \rangle. \tag{7.4}$$

In particular for B a finite stopping prefix, and with the usual notation $\omega_B = \omega \cap B$, we define the \mathcal{F}_B -measurable random variable:

$$T_B f : \Omega \rightarrow \mathbb{R}, \quad T_B f(\omega) = Tf(\omega_B) = \frac{1}{\langle N, \omega_B \rangle} \langle f, \omega_B \rangle.$$

I-2.6 Ergodic Means Coincide for Sequential Nets. Assume that \mathcal{N} is a sequential net associated with a probabilistic transition system $(S, A, x_0, (\mu_x)_x)$ (Ch. 1, III-2). Let $X = (X_n)_{n \geq 0}$ and $Y = (Y_n)_{n \geq 1}$ be the canonical Markov chains associated with the transition system, X for the chain of states and Y for the chain of transitions. The unfolding is associated with a tree of events \mathcal{T} , all branching cells are disjoint in \mathcal{T} and in the set of transitions A (Ch. 3, VII-2). It follows that a distributed function identifies with a function $f : A \rightarrow \mathbb{R}$.

Let B be a stopping prefix of \mathcal{T} , and let $t_B : \Omega \rightarrow \mathbb{N}$ be the \mathcal{F}_B -measurable random variable defined by $t_B(\omega) = \langle N, \omega_B \rangle$ (we have seen in Ch. 5 that t_B is actually a stopping time in the usual sense). Then we have:

$$T_B f(\omega) = \frac{1}{t_B(\omega)} \sum_{k=1}^{t_B(\omega)} f(Y_k). \tag{7.5}$$

Assume that $(B_n)_n$ is a non decreasing sequence of finite stopping prefixes, with $B_n \rightarrow \mathcal{E}$. Then $(t_{B_n})_n$ converges \mathbb{P} -*a.s.* to $+\infty$. From the ergodic theory of Markov chains, the ergodic sums (7.5) with $B = B_n$ converge \mathbb{P} -*a.s.*, and the limit does not depend on the sequence $(B_n)_n$. Typically, this last point will not hold anymore with concurrency: see I-3.3.

I-3 Regular Sequences of Stopping Operators. Nets Satisfying the Strong Law of Large Numbers.

We give the statement that we consider for the Strong law of large numbers. The growth of time is formalised by a sequence $(V_n)_n$ of finite stopping operators, such that $V_n(\omega) \uparrow \omega$. We require that the growth of V_n is linear in n . This condition concerns the *observations*. A second condition, intrinsic to the net, is introduced in II-1.

I-3.1 Definition. (*Regular sequences of stopping operators*) Let $(V_n)_{n \geq 1}$ be a sequence of finite stopping operators of a locally finite probabilistic event structure. We say that the sequence is **regular** if:

1. $(V_n)_n$ is non decreasing: $\forall n, V_n \subseteq V_{n+1}$,
2. $\bigcup_n V_n(\omega) = \omega$, for \mathbb{P} -a.s ω of Ω ,
3. There are two constants K_1, K_2 such that:

$$\forall n \geq 1, \quad 0 < K_1 \leq \frac{\langle N, V_n \rangle}{n} \leq K_2 < \infty, \quad \mathbb{P}\text{-a.s.} \quad (7.6)$$

For example, assume that $\mathcal{E} = \mathcal{T}$ is a tree of events with no maximal event. If B_n is the prefix of height n , $(B_n)_n$ is regular with $K_1 = K_2 = 1$.

I-3.2 Definition. (*Nets with the Strong law of large numbers*) Let $(\mathcal{N}, \mathbb{P})$ be a (compact) probabilistic net. We say that the Strong law of large numbers holds for $(\mathcal{N}, \mathbb{P})$ if every distributed function f defined on the unfolding of \mathcal{N} satisfies the following property: For every regular sequence $(V_n)_n$ of stopping operators, the ergodic means:

$$Tf(V_n) = \frac{1}{\langle N, V_n \rangle} \langle f, V_n \rangle,$$

converge \mathbb{P} -a.s, and for every other regular sequence $(W_n)_n$, we have:

$$\lim_{n \rightarrow \infty} Tf(V_n) = \lim_{n \rightarrow \infty} Tf(W_n), \quad \mathbb{P}\text{-a.s.}$$

I-3.3 Remark. As shown by the following example, and for concurrent systems, the regularity of stopping operators is mandatory to expect the uniqueness in Definition I-3.1. Let \mathcal{N} be the safe marked net represented in Figure 7.1, consisting in two independent deterministic loops.

The unfolding \mathcal{U} is given by the disjoint union of two chains: (A, t_1, A, t_1, \dots) and (B, t_2, B, t_2, \dots) . Ω has a unique element ω , and \mathbb{P} is thus trivial: $\mathbb{P}(\omega) = 1$. \mathcal{U} has two classes of branching cells, say λ_1 and λ_2 . Let f be the distributed function $f = N^{\lambda_1}$ that




Figure 7.1: Two independent deterministic nets.

counts the occurrences of λ_1 (defined in I-2.4). Let $(C_n)_{n \geq 1}$ and $(D_n)_{n \geq 1}$ be the sequences of stopping prefixes defined by:

$$C_n = \underbrace{(A, t_1, \dots, t_1, A)}_{n \text{ occurrences of } A} \sqcup \underbrace{(B, t_2, \dots, t_2, B)}_{n^2 \text{ occurrences of } B},$$

and $(D_n)_{n \geq 1}$ defined similarly, with n^2 occurrences of A and n occurrences of B . We have $\langle N, \omega \rangle = n + n^2$, and thus:

$$T_{C_n} f(\omega) = \frac{1}{1+n}, \quad T_{D_n} f(\omega) = \frac{n}{1+n}.$$

Therefore the limit depends on the sequence.

The following lemma shows that the normal decomposition of maximal configurations (Ch. 5, III-3.1) leads to a regular sequence of stopping operators.

I-3.4 Lemma. *If the net is recurrent, the sequence $(V_n)_n$ of random variables given by the normal decomposition of maximal configurations is a regular sequence of stopping operators.*

Proof – We have shown in Ch. 5, III-3.3 that V_n is a finite well-stopping operator. Since $(V_n)_n$ is non-decreasing by construction, it remains only to show that $(V_n)_n$ satisfies point 3 of Definition I-3.1. Since the net is recurrent, $V_n(\omega)$ is not maximal for \mathbb{P} -a.s all ω . Therefore the cone of future \mathcal{U}^{V_n} has a non empty stopping prefix, and contains thus at least a branching cell. Hence: $\langle N, V_{n+1} \rangle \geq 1 + \langle N, V_n \rangle$, and thus:

$$\forall n \geq 1, \quad \frac{\langle N, V_n \rangle}{n} \geq 1.$$

For the converse equality, we have $\langle N, V_{n+1} \rangle \leq \langle N, V_n \rangle + k$, where k is for instance (not the optimal bound) the maximal number of concurrent transitions of the net, whence:

$$\forall n \geq 1, \quad \frac{\langle N, V_n \rangle}{n} \leq k.$$

□

II—The Strong Law of Large Numbers

The aim of this section is to show the Strong law of large numbers for concurrent systems, with the statement that we have given above in I-3.2. We introduce for this the class of probabilistic nets with *integrable concurrent height*. This analytical condition is intended to control the range of concurrency inside a probabilistic net. We underline the geometric interpretation of this condition in the unfolding. We establish in the framework of nets some preliminary results that can be found in the theory of mathematical dynamical systems. Then we show that, from the Strong law of large numbers applied to the embedded Markov chain, we can derive the Strong law of large numbers for a distributed Markov net.

In II-1 we introduce the definition of nets with *integrable concurrent height*. We show two preliminary results on d -homogeneous probabilistic nets in II-2. Then we show in II-3 the Strong law of large numbers for recurrent and distributed Markov nets with integrable concurrent height.

II-1 Integrable Concurrent Height.

Let \mathcal{U} be the unfolding of a probabilistic net $(\mathcal{N}, \mathbb{P})$. Let v be a finite well-stopped and positive configuration of \mathcal{U} , the probabilistic future (Ω^v, \mathbb{P}^v) is well defined. For each condition $b \in \gamma(v)$, we set the following random variable:

$$L_b^v : \Omega(v) \rightarrow \mathcal{W}, \quad L_b^v(\xi) = \text{Sup}\{w \in \mathcal{X}_0, w \subseteq v \oplus \xi : b \in \gamma(w)\}.$$

The interpretation in the net of L_b^v is that of a playing sequence that continues v without moving the token associated with b .

Recall that \mathbb{E}^v denotes the expectation *w.r.t.* probability \mathbb{P}^v . We define the **branching distance** $D(v, L_b^v(\xi))$ by counting the branching cells between v and L_b^v :

$$\forall \xi \in \Omega(v), \quad D(v, L_b^v(\xi)) = \langle N, L_b^v(\xi) \rangle - \langle N, v \rangle.$$

$D(v, L_b^v(\xi))$ may be infinite, and says how much $L_b^v(\xi)$ is larger than v . It says how far a process can go towards direction $\xi \in \Omega(v)$, keeping in place in the net the token associated with b (See Example II-1.3 and Figure 7.2). The mean value is given by:

$$\mathbb{E}^v(D(v, L_b^v)).$$

II-1.1 Definition. (*Nets with integrable concurrent height*) We say that a probabilistic net $(\mathcal{N}, \mathbb{P})$ has **integrable concurrent height** if for every $v \in \mathcal{X}_0$ with $\mathbb{P}(\Omega(v)) > 0$:

$$\mathbb{E}^v\left(\sup_{b \in \gamma(v)} D(v, L_b^v)\right) < \infty.$$




Figure 7.2: Integrable concurrent height.

II-1.2 Remark. For each $v \in \mathcal{X}_0$, set $R^v : \Omega^v \rightarrow \mathbb{R}$ defined by:

$$\forall \xi \in \Omega^v, \quad R^v(\xi) = \sup_{b \in \gamma(v)} D(v, L_b^v(\xi)).$$

Then $R = (R^v)_{v \in \mathcal{X}_0}$ is a d -homogeneous function, *i.e.* R^v only depends on the marking $m(v)$, such that v leads to $m(v)$ (Cf. Ch.5 III-2.3 for d -homogeneous functions). Saying that the net has integrable concurrent height is equivalent to say that R is integrable.

From $R^v \leq \sum_{b \in \gamma(v)} D(v, L_b^v)$, it follows that R is integrable if and only if the mean values $\mathbb{E}^v(D(v, L_b^v))$ are finite.

II-1.3 Example. Typically, the net of Ch. 6, Fig. 6.2, p. 193, has integrable concurrent height. Figure 7.2 geometrically illustrates the random variable L_b^v with finite mean $\mathbb{E}^v(L_b^v) < \infty$. The piece cut up from the cone of future \mathcal{U}^v along $L_b^v(\cdot)$ has *finite area*. The net has integrable concurrent height if all the pieces cut up this way have finite area.

In the unfolding of a net with integrable concurrent height, there cannot be branches that never interact. In other words, all branches must synchronise. Moreover, the delay between synchronisations has a finite mean.

II-2 Preliminary Asymptotic and Ergodic Results.

We establish two asymptotic results as a preliminary for the Strong law of large numbers. The analogous of these results for sequential systems are found in the

theory of dynamical systems ([42]). Theorem II-2.2 is formulated as an ergodicity result for recurrent nets: under some invariance condition, test functions are \mathbb{P} -a.s constant. Remark the use of a Martingale argument in the proof.

Proposition II-2.1 is a key for analytically controlling concurrent terms in the equations. It has to be compared with the following well known result, a lemma often used in the analysis of dynamical systems. Let (E, τ, μ) be a dynamical system, i.e. μ is a probability on the space E and $\tau : E \rightarrow E$ is a μ -invariant pointwise transformation: $\tau\mu = \mu$. If $f : E \rightarrow \mathbb{R}$ is integrable, then:

$$\lim_{n \rightarrow \infty} \frac{f \circ \tau^n}{n} = 0, \quad \mu\text{-a.s.}$$

We recall that a d -homogeneous function (Ch. 4, III-2.3) is a collection $H = (H^v)_{v \in \mathcal{X}_0}$, where $H^v : \Omega^v \rightarrow \mathbb{R}$ is a random variable for every $v \in \mathcal{X}_0$, such that $H^v = H^{v'}$ if $\gamma(v) \cong \gamma(v')$, that is if v and v' lead to equivalent cuts, or equivalently if v and v' lead to the same marking. H is said to be integrable if every H^v is \mathbb{P}^v -integrable.

II-2.1 Proposition. *Let $(\mathcal{N}, \mathbb{P})$ be a d -homogeneous probabilistic net, and let $(V_n)_n$ be a regular sequence of stopping operators defined in the unfolding \mathcal{U} of \mathcal{N} . For each n , θ_n denotes the shift operator θ_{V_n} , defined by $\theta_n(\omega) = \omega \ominus V_n(\omega)$, with $\theta_n(\omega) \in \Omega^{V_n(\omega)}$.*

Let $H = (H^v)_{v \in \mathcal{X}_0}$ be a d -homogeneous integrable function defined on \mathcal{U} . Then we have:

$$\lim_{n \rightarrow \infty} \frac{H^{V_n}(\theta_n)}{\langle N, V_n \rangle} = 0, \quad \mathbb{P}\text{-a.s.} \quad (7.7)$$

II-2.2 Theorem. (Ergodicity of recurrent nets) *Let $(\mathcal{N}, \mathbb{P})$ be a d -homogeneous and recurrent net. Let $H = (H^v)_{v \in \mathcal{X}_0}$ be an integrable d -homogeneous function. Assume that, for \mathbb{P} -a.s all ω , we have:*

$$\forall v \in \mathcal{X}_0, v \subseteq \omega, \quad H^v(\omega \ominus v) = H^\emptyset(\omega).$$

Then H^\emptyset is \mathbb{P} -a.s constant.

II-2.3 Proof of Prop. II-2.1. Since $(V_n)_n$ is a regular sequence of stopping operators, there is a constant $K_1 > 0$ such that $\langle N, V_n \rangle \geq K_1 n$, therefore it is enough to show:

$$\lim_{n \rightarrow \infty} \frac{H^{V_n}(\theta_n)}{n} = 0, \quad \mathbb{P}\text{-a.s.}$$

For this, denoting by X_n the random variable $X_n = \frac{H^{V_n}(\theta_n)}{n}$, we use the classical criterion that implies the convergence \mathbb{P} -a.s of X_n to zero:

$$\forall \epsilon > 0, \quad \sum_{n \geq 1} \mathbb{P}(X_n > \epsilon) < \infty. \quad (7.8)$$

Let \mathcal{D} denote the set of well-stopped markings of the net: those markings associated with finite well-stopped configurations. We denote by $m(v)$ the marking associated with a finite configuration v . Since \mathbb{P} is d -homogeneous, we write $\mathbb{P}^m = \mathbb{P}^v$ for any $m \in \mathcal{D}$ with $m = m(v)$. The d -homogeneous function H identifies with the finite collection $H = (H^m)_{m \in \mathcal{D}}$.

For each $n \geq 0$, we denote by E_n the set of values of V_n . By assumption, $V_n(\omega)$ is finite for \mathbb{P} -a.s ω , hence E_n is at most countable. As V_n is a stopping operator, we have already observed this simple property:

$$\forall u \in E_n, \quad \{V_n = u\} = \Omega(u),$$

which implies: $\mathbb{P}(\cdot | V_n = u) = \mathbb{P}^u(\cdot)$. From this we get:

$$\begin{aligned} \mathbb{P}(X_n > \epsilon) &= \sum_{u \in E_n} \mathbb{P}(V_n = u) \mathbb{P}\left(\frac{H^{V_n}(\omega \ominus V_n)}{n} > \epsilon \mid V_n = u\right) \\ &= \sum_{u \in E_n} \mathbb{P}(V_n = u) \mathbb{P}^u\left(\frac{H^u}{n} > \epsilon\right) \\ &= \sum_{m \in \mathcal{D}} \left(\sum_{\substack{u \in E_n \\ m(u)=m}} \mathbb{P}(V_n = u) \right) \mathbb{P}^m\left(\frac{H^m}{n} > \epsilon\right) \\ &\leq \sum_{m \in \mathcal{D}} \mathbb{P}^m\left(\frac{H^m}{n} > \epsilon\right). \end{aligned} \tag{7.9}$$

To show (7.8), and from (7.9), it is enough to show that for every $m \in \mathcal{D}$, the following sum is finite: $\sum_n \mathbb{P}^m(H^m > n\epsilon) < \infty$. Recall the usual transformation $\mathbb{E}(\sum_{n \geq 0} \mathbf{1}_{\{f \geq n\}}) \leq \mathbb{E}(f) + 1$ for f a non negative random variable. Using this transformation, we get for each $m \in \mathcal{D}$:

$$\begin{aligned} \sum_{n \geq 1} \mathbb{P}^m(H^m > n\epsilon) &= \mathbb{E}^m\left(\sum_{n \geq 1} \mathbf{1}_{\{\frac{1}{\epsilon} H^m > n\}}\right) \\ &\leq 1 + \frac{1}{\epsilon} \mathbb{E}^m(H^m) < \infty, \end{aligned}$$

since H^m is \mathbb{P}^m -integrable for each $m \in \mathcal{D}$. This completes the proof. \square

II-2.4 Proof of Th. II-2.2. Let \mathcal{U} denote the unfolding of \mathcal{N} . We begin with the following observation:

\diamond If V is a finite well-stopping operator of \mathcal{U} , then for every integrable function $f : \Omega \rightarrow \mathbb{R}$, we have:

$$\mathbb{E}(f | V) = \int_{\Omega^{\gamma_V}} f(V \oplus \xi) d\mathbb{P}^{\gamma_V}(\xi).$$

This has been shown in Ch. 5, I-3.3, for $V(\cdot) = \cdot \cap B$ and B a finite stopping prefix. This is shown in general in the same way, using again that $\{V = u\} = \Omega(u)$ for every $u \in \mathcal{X}_0$ such that $\mathbb{P}(V = u) > 0$.

Now let $(R_n)_n$ be the sequence of successive renewal operators. All R_n are \mathbb{P} -a.s finite since we assume that the net is recurrent. Let \mathcal{F}_n denote the σ -algebra $\mathcal{F}_n = \langle R_n \rangle$ (since R_n is finite, we also have that $\mathcal{F}_n = \mathcal{F}_{R_n}$ in the sense seen in Chapter 6, but we will not use this fact). It is easily checked that we have $\omega = \bigcup_n R_n(\omega)$ for \mathbb{P} -a.s all ω , hence: $\langle \mathcal{F}_n, n \geq 1 \rangle = \mathcal{F}$. And since H^θ is \mathbb{P} -integrable, the Martingale convergence theorem implies the convergence:

$$H^\theta = \lim_{n \rightarrow \infty} \mathbb{E}(H^\theta | \mathcal{F}_n), \quad \mathbb{P}\text{-a.s.} \quad (7.10)$$

Using \diamond with $f = H^\theta$, we get the following expression for the conditional expectation $\mathbb{E}(H^\theta | \mathcal{F}_n)$:

$$\begin{aligned} \mathbb{E}(H^\theta | \mathcal{F}_n) &= \int_{\Omega^{\gamma_{R_n}}} H^\theta(R_n \oplus \xi) \, d\mathbb{P}^{\gamma_{R_n}}(\xi) \\ &= \int_{\Omega} H^\theta(R_n \oplus \xi) \, d\mathbb{P}(\xi), \end{aligned} \quad (7.11)$$

the later since \mathbb{P} is d -homogeneous and since $\gamma_{R_n} \cong c_0$ by construction of the successive renewal operators. Applying the assumption $H^\theta(\omega) = H^v(\omega \ominus v)$ with $v = R_n$ we get:

$$\forall \xi \in \Omega, \quad H^\theta(R_n \oplus \xi) = H^{R_n}(\xi) = H^\theta(\xi),$$

since H is also a d -homogeneous function. Using together (7.10) and (7.11), we obtain:

$$H^\theta(\omega) = \lim_{n \rightarrow \infty} \int_{\Omega} H^\theta(\xi) \, d\mathbb{P}(\xi) = \mathbb{E}(H^\theta), \quad \mathbb{P}\text{-a.s.}$$

This completes the proof. \square

II-3 The Strong Law of Large Numbers.

Our goal in II-3 is to prove the following result.

II-3.1 Theorem. (*Strong law of large numbers*) Let $(N, M_0, (\mu_l)_l)$ be a distributed positive Markov net, with \mathbb{P} the probability associated. We assume that the net satisfies the following assumptions:

1. \mathcal{N} is recurrent,
2. \mathcal{N} has integrable concurrent height,

Then $(\mathcal{N}, \mathbb{P})$ satisfies the Strong law of large numbers. There are non negative numbers $\alpha(l)$ for $l \in \mathcal{L}_{\mathcal{U}}$, such that for every distributed function $f = (f^l)_{l \in \mathcal{L}_{\mathcal{U}}}$ defined on the unfolding, and for every regular sequence $(V_n)_n$ of stopping operators of \mathcal{U} , we have:

$$\lim_{n \rightarrow \infty} Tf(V_n) = \sum_{l \in \mathcal{L}_{\mathcal{U}}} \alpha(l) \mu_l(f^l), \quad \mathbb{P}\text{-a.s.}$$

The Strong law of large numbers holds for any marking M reachable from M_0 , with the same coefficients $\alpha(l)$.

We have already shown some intermediate results in II-2. The rest of the proof involves three steps.

II-3.2 First Step: Reduction to Densities of Branching Cells. We establish a result that reduces the Strong law of large numbers to the existence of limits of the ergodic means TN^l , with l ranging over $\mathcal{L}_{\mathcal{U}}$. This result cannot be expected in general for non distributed probabilities.

Assume that $(\mathcal{N}, \mathbb{P})$ is a probabilistic net satisfying Th. II-3.1. Applying the Strong law to the distributed function N^l that counts the occurrences of a class of branching cells $l \in \mathcal{L}_{\mathcal{U}}$ (I-2.4), we find a non-negative random variable $\alpha(l)$ given by the limit \mathbb{P} -a.s:

$$\alpha(l) = \lim_{n \rightarrow \infty} TN^l(V_n), \tag{7.12}$$

for $(V_n)_n$ a regular sequence of stopping operators. The coefficient $\alpha(l)$ represents the asymptotic ratio of the class l , among the other classes. Therefore we call it the **density** of l .

The natural question is then to recover the Strong law of large numbers from the coefficients $\alpha(l)$, if they exist. This holds for distributed probabilities. But it cannot be expected for general probabilities since we miss the correlation information between local processes.

Recall that a distributed Markov net $(\mathcal{N}, (\mu_l)_{l \in \mathcal{L}_{\mathcal{U}}})$ is said to be positive if all μ_l are positive. In this case, the associated probability gives positive probability to every finite shadow $\Omega(v)$.

II-3.3 Lemma. *Let $(\mathcal{N}, (\mu_l)_{l \in \mathcal{L}_{\mathcal{U}}})$ be a distributed and recurrent positive Markov net, with \mathbb{P} the associated probability on Ω . Assume that for every class of branching cell $l \in \mathcal{L}_{\mathcal{U}}$, and for every regular sequence $(V_n)_n$ of stopping operators, the ergodic means converge:*

$$\lim_{n \rightarrow \infty} TN^l(V_n) = \alpha(l), \quad \mathbb{P}\text{-a.s.},$$

with a limit independent of the regular sequence $(V_n)_n$. Then the Strong law of large numbers holds for $(\mathcal{N}, \mathbb{P})$, and we have for every distributed function $f = (f^l)_{l \in \mathcal{L}_{\mathcal{U}}}$ and every regular sequence $(V_n)_n$:

$$\lim_{n \rightarrow \infty} Tf(V_n) = \sum_{l \in \mathcal{L}_{\mathcal{U}}} \alpha(l) \mu_l(f^l), \tag{7.13}$$

where $\mu_l(f^l)$ denotes the expectation of f^l under μ_l .

II-3.4 Second Step: Using the Embedded Markov Chain. Let $(\mathcal{N}, (\mu_l)_l)$ be a distributed positive Markov net. We fix a class l of branching cells in the unfolding \mathcal{U} of \mathcal{N} . We set f the distributed function $f = N^l$.

We consider the random variables $(V_n, Z_n)_{n \geq 1}$ given by the normal decomposition of maximal configurations (Cf. Ch. 5, III-3.1). The embedded Markov chain (Ch. 5, III-3.2) is given by $(M_n, Z_n)_{n \geq 1}$, with $M_n = m(V_n)$ the marking associated with V_n . We have:

$$\begin{aligned} \frac{\langle f, V_n \rangle}{\langle N, V_n \rangle} &= \frac{\langle f, V_n \rangle}{n} \frac{n}{\langle N, V_n \rangle} \\ &= \frac{\sum_{k=1}^{n-1} \mathbf{1}_{\{l \in \Delta^\perp(\mathcal{U}^{M_k})\}}}{n} \frac{n}{\sum_{k=1}^{n-1} \text{Card}(\Delta^\perp(\mathcal{U}^{M_k}))}. \end{aligned}$$

ergodic means
ergodic means

The ergodic theory of finite Markov chains, applied to the embedded Markov chain, implies that the above expression converges \mathbb{P} -a.s. to a random variable G . Since $\langle f, V_n \rangle = \langle N^l, V_n \rangle \leq \langle N, V_n \rangle$, G is bounded by 1. The idea is to study G without studying the embedded Markov chain, which has a set of states that we do not want to manipulate. From the embedded Markov chain, we only derive the existence of G .

For each $v \in \mathcal{X}_0$, the same construction applies to the probabilistic future $(\mathcal{U}^v, \mathbb{P}^v)$. This defines a collection of measurable mappings $H^v : \Omega^v \rightarrow \mathbb{R}$, with $H^\emptyset = G$. The collection $H = (H^v)_{v \in \mathcal{X}_0}$ is d -homogeneous by construction.

II-3.5 Lemma. Assume that \mathcal{N} satisfies the assumptions of Theorem II-3.1. Fix l a class of branching cells. Using the above notation $f = N^l$, $H = (H^v)_{v \in \mathcal{X}_0}$ and $G = H^\emptyset$, we have:

1. For every regular sequence $(W_n)_n$ of stopping operators, the following convergence holds:

$$\lim_{n \rightarrow \infty} Tf(W_n) = G, \quad \mathbb{P}\text{-a.s.}$$

2. For each $v \in \mathcal{X}_0$, and for \mathbb{P} -a.s. all $\omega \in \Omega(v)$:

$$H^v(\omega \ominus v) = H^\emptyset(\omega).$$

II-3.6 Last Step: Putting All Together. Proof of Theorem II-3.1.

Let $(\mathcal{N}, M_0, (\mu_l)_{l \in \mathcal{L}_\mathcal{U}})$ be a distributed Markov net as in Th. II-3.1. We fix a class of branching cells $l \in \mathcal{L}_\mathcal{U}$. According to the second step II-3.4, we define the d -homogeneous function $H = (H^v)_{v \in \mathcal{X}_0}$. According to Lemma II-3.5, H satisfies:

$$\forall v \in \mathcal{X}_0, \quad H^v(\cdot \ominus v) = H^\emptyset(\cdot).$$

H is bounded by 1, and is thus integrable. Theorem II-2.2 implies that H^θ is constant, and it also follows that H^v is constant with the same value for all $v \in \mathcal{X}_0$. Let $\alpha(l)$ denote this constant. Then we have for every regular sequence $(V_n)_n$ of stopping operators:

$$\lim_{n \rightarrow \infty} TN^l(V_n) = \alpha(l), \quad \mathbb{P}\text{-a.s.}$$

This holds for all classes $l \in \mathcal{L}_U$. By Lemma II-3.3, it implies the Strong law of large numbers. Since H^v has the same value than $H^\theta = \alpha(l)$, the Strong law of large numbers holds in the cones of future with the same density coefficients $\alpha(l)$.

II-3.7 Proof of Lemma II-3.3. By linearity of $f \rightarrow Tf$, we assume without loss of generality that there is a class $l \in \mathcal{L}_U$ such that $f^\lambda = 0$ if $\lambda \neq l$.

Let $(V_n)_n$ be a regular sequence of stopping operators. We denote by $(W_n)_{n \geq 1}$ the successive local l -renewals, and by $(l_n(\omega))_{n \geq 1}$ the successive l -branching cells of ω (Ch. 6, III-2.3). Since $(V_n)_n$ is non decreasing, the sequence:

$$I_n(\omega) = \{ \lambda \in \mathcal{L}_U \mid \lambda \cong l, \lambda \in \Lambda(V_n) \}$$

is a non decreasing sequence of intervals *w.r.t.* the relation \preceq_ω defined in Ch. 6, III-1.3. According to Lemma III-2.4 of Ch. 6, for every branching cell $\lambda \cong l$ satisfying $\lambda \in \overline{\Lambda}(\omega)$, there is an integer j such that $\lambda = l_j(\omega)$. It follows that for every $n \geq 1$, there is an integer $J(n) \geq 1$ such that:

$$I_n = \{ l_p(\omega), 1 \leq p \leq J(n) \}.$$

Since we assume that f vanishes out of l , we have:

$$\begin{aligned} Tf(V_n) &= \frac{1}{\langle N, V_n \rangle} \langle f, V_n \rangle \\ &= \frac{1}{\langle N, V_n \rangle} \langle f, W_{J(n)} \rangle \\ &= \frac{\langle N^l, V_n \rangle}{\langle N, V_n \rangle} \frac{1}{\langle N^l, V_n \rangle} \sum_{k=1}^{J(n)} f^l(Z_k^l), \end{aligned} \tag{7.14}$$

where $(Z_n^l)_{n \geq 1}$ denote the successive l -germs (Ch. 6, III-3.1). Since the net is recurrent and all μ_l are positive, it follows from Proposition III-2.2 of Chapter 6 that $\lim_{n \rightarrow \infty} J(n) = +\infty$. Since the probability \mathbb{P} is distributed, the sequence $(Z_n^l)_{n \geq 1}$ is *i.i.d* with law μ_l in Ω_l , according to Theorem III-3.2 of Chapter 6. It follows from the Strong law of large numbers for *i.i.d* sequences that we have:

$$\begin{aligned} \frac{1}{\langle N^l, V_n \rangle} \sum_{k=1}^{J(n)} f^l(Z_k^l) &= \frac{1}{\langle N^l, W_{J(n)} \rangle} \sum_{k=1}^{J(n)} f^l(Z_k^l) \\ &= \frac{1}{J(n)} \sum_{k=1}^{J(n)} f^l(Z_k^l) \xrightarrow{n \rightarrow \infty} \mu_l(f^l), \quad \mathbb{P}\text{-a.s.} \end{aligned}$$




Figure 7.3: Proof of Lemma II-3.5.

We also have by hypothesis:

$$\lim_{n \rightarrow \infty} \frac{\langle N^l, V_n \rangle}{\langle N, V_n \rangle} = \alpha(l), \quad \mathbb{P}\text{-a.s.}$$

We get thus, from (7.14):

$$\lim_{n \rightarrow \infty} Tf(v_n) = \alpha(l) \mu_l(f^l), \quad \mathbb{P}\text{-a.s.}$$

which is (7.13) for f . This completes the proof. \square

II-3.8 Proof of Lemma II-3.5. 1. Let $(W_n)_n$ be a regular sequence of stopping operators. $(V_j, Z_j)_j$ denotes as above the sequence of random variables given by the normal decomposition of maximal configurations. For each $n \geq 1$, we set $J_n : \Omega \rightarrow \mathbb{N}$ the integer random variable defined by:

$$J_n = \inf\{p \geq 1 : V_p(\omega) \supseteq W_n(\omega)\} < \infty.$$

V_{J_n} is seen as an approximation of W_n . The normal decomposition does not go too fast: we have by Th. VI-3.5 of Ch. 3:

$$\bar{\Lambda}(\omega) = \bigcup_{n \geq 0} \Delta^+(V_n).$$

It follows that the cuts $\gamma(W_n)$ and $\gamma(V_{J_n})$ intersect, since we have (see an illustration in Figure 7.3):

$$J_n = \sup\{p \geq 1 \mid \Delta^+(V_p) \cap \Delta^+(W_n) \neq \emptyset\}.$$

We note: $R^v = \sup_{b \in \gamma(v)} D(v, L_b^v)$, where L_b^v denotes the concurrent height random variable defined in II-1. Since $\gamma(W_n) \cap \gamma(V_{J_n}) \neq \emptyset$ and since $W_n \subseteq V_{J_n}$, we have, with D denoting the branching distance defined in II-1:

$$D(W_n, V_{J_n}) = \langle N, V_{J_n} \ominus W_n \rangle \leq R^{W_n}.$$

Now we show that $\delta_n = Tf(W_n) - \frac{\langle f, V_{J_n} \rangle}{\langle N, V_{J_n} \rangle}$ converges \mathbb{P} -*a.s* to zero.

$$\begin{aligned} \delta_n &= \frac{\langle f, W_n \rangle}{\langle N, W_n \rangle} - \frac{\langle f, V_{J_n} \rangle}{\langle N, V_{J_n} \rangle} \\ &= \frac{\langle f, W_n \rangle - \langle f, V_{J_n} \rangle}{\langle N, W_n \rangle} + \langle f, V_{J_n} \rangle \left(\frac{1}{\langle N, W_n \rangle} - \frac{1}{\langle N, V_{J_n} \rangle} \right) \\ &= \frac{\langle f, V_{J_n} \rangle}{\langle N, V_{J_n} \rangle} \frac{\langle N, V_{J_n} \ominus W_n \rangle}{\langle N, W_n \rangle} - \frac{\langle f, V_{J_n} \ominus W_n \rangle}{\langle N, W_n \rangle}. \end{aligned}$$

We use that $\langle f, x \rangle = \langle N^l, x \rangle \leq \langle N, x \rangle$ for all $x \in \mathcal{X}_0$ to get:

$$|\delta_n| \leq 2 \frac{\langle N, V_{J_n} \ominus W_n \rangle}{\langle N, W_n \rangle} \leq 2 \frac{R^{W_n}(\omega \ominus W_n)}{\langle N, W_n \rangle}. \tag{7.15}$$

The collection $R = (R^v)_{v \in \mathcal{X}_0}$ is d -homogeneous, and defines thus a d -homogeneous function. R is integrable since we assume that \mathcal{N} has integrable concurrent height. It follows from Proposition II-2.1 that δ_n converges \mathbb{P} -*a.s* to zero, what was to be shown.

2. The second part of the Lemma follows from an analogous computation. \square

III—Using Martingales

In this section we show that, for distributed Markov nets, a distributed function naturally defines a partially ordered Martingale. It seems to be a convenient framework to establish a Central Limit Theorem.

III-1 Definition of a Martingale.

III-1.1 Definition. (*Mean and covariance distributed function*) Let $f = (f^l)_{l \in \mathcal{L}\mathcal{U}}$ be a distributed function, defined on the unfolding \mathcal{U} of a distributed Markov net $(\mathcal{N}, (\mu^l)_{l \in \mathcal{L}\mathcal{U}})$. We define the **mean** of f as the distributed function denoted by $\mathbb{E}(f)$, and given by $\mathbb{E}(f) = (g^l)_{l \in \mathcal{L}\mathcal{U}}$, with $g^l : \Omega_l \rightarrow \mathbb{R}$ the constant:

$$\forall z \in \Omega_l, \quad g^l(z) = \mu^l(f^l),$$

the expectation of f^l under μ^l . We say that f has **zero mean** if $\mathbb{E}(f)$ is identically null. The **covariance** of f is the distributed function denoted by $\sigma^2(f)$ and given by $\sigma^2(f) = (h^l)_{l \in \mathcal{L}\mathcal{U}}$, with h^l constant for each $l \in \mathcal{L}\mathcal{U}$:

$$\forall z \in \Omega_l, \quad h^l(z) = \sigma^2(f^l) = \mu^l((f^l)^2 - \mu^l(f^l)),$$

the covariance of f^l under μ^l .

Then we have the following.

III-1.2 Proposition. *Let f be a distributed function with zero mean defined on the unfolding of a distributed Markov net $(\mathcal{N}, (\mu^l)_{l \in \mathcal{L}_U})$. For every $B \in \mathcal{S}_0$, let $F_B : \Omega \rightarrow \mathbb{R}$ be the \mathcal{F}_B -measurable function induced by f , defined by:*

$$F_B(\omega) = \langle f, \omega_B \rangle .$$

Then $(F_B)_{B \in \mathcal{S}_0}$ is a martingale, that is:

$$\forall B, B' \in \mathcal{S}_0, \quad B \subseteq B' \Rightarrow \mathbb{E}(F_{B'} | \mathcal{F}_B) = F_B .$$

The covariance of the random variable F_B is given by:

$$\sigma^2(F_B) = \mathbb{E}(\langle \sigma^2(f), \omega_B \rangle) , \tag{7.16}$$

where $\sigma^2(f)$ is the covariance distributed function, defined in III-1.1.

III-1.3 Remark. For sequential systems, the martingale is defined as follows. Let $(X_n)_{n \geq 0}$ be a Markov chain on a finite set E . For each $x \in E$, let $f_x : E \rightarrow \mathbb{R}$ such that, if μ_x denotes the x^{th} row of the transition matrix, we have: $\mu_x(f_x) = 0$. Then the sequential martingale that we consider is:

$$F_n = \sum_{k=1}^n f_{X_{k-1}}(X_k) .$$

Let $Y_k = f_{X_{k-1}}(X_k)$. Equation (7.16) is analogous to the addition formula for square means:

$$\sigma^2(F_n) = \mathbb{E}(F_n^2) = \mathbb{E}(Y_1^2) + \dots + \mathbb{E}(Y_n^2) ,$$

and takes into account the horizontal independence—due to concurrency—that comes with distributed probabilities.

III-1.4 Remark. Let f be a distributed function, and let $\mathbb{E}(f)$ denote the mean of f as in III-1.1. Then $g = f - \mathbb{E}(f)$ has zero mean. Let $F_B, G_B : \Omega \rightarrow \mathbb{R}$ be the \mathcal{F}_B -measurable functions defined by $F_B(\omega) = \langle f, \omega_B \rangle$, $G_B(\omega) = \langle g, \omega_B \rangle$. Then we have:

$$F_B = \langle \mathbb{E}(f), \omega_B \rangle + G_B ,$$

with $\mathbb{E}(f)$ a constant distributed function and $(G_B)_{B \in \mathcal{S}_0}$ a martingale.

III-2 Towards a Central Limit Theorem.

III-2.1 Background: Convergence in Law. The convergence in law is the convergence stated by the Central Limit Theorem. Let $(X_n)_{n \geq 1}$ be a sequence of real random variables defined on a probability space $(\mathcal{A}, \mathcal{F}, \mathbb{P})$. Let F_n denote the probability law of X_n , F_n is a probability measure on \mathbb{R} . Let μ be a probability measure on \mathbb{R} . We say that the sequence $(X_n)_{n \geq 1}$ converges in law to the law μ if, for all continuous and bounded function $f : \mathbb{R} \rightarrow \mathbb{R}$, we have:

$$\lim_{n \rightarrow \infty} \int_{\mathbb{R}} f(x) dF_n(x) = \int_{\mathbb{R}} f(x) d\mu(x).$$

We will denote this convergence by: $X_n \xrightarrow{\mathcal{L}} \mu$

We recall below the Central Limit Theorem for martingale differences.

III-2.2 Theorem. (*Central Limit Theorem for martingale differences, [35]*) Let $(\mathcal{F}_n)_{n \geq 0}$ be a non decreasing sequences of σ -algebras of a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and let $(Y_n)_{n \geq 1}$ be a sequence of martingale differences, that is:

$$\forall n \geq 1, \quad \mathbb{E}(Y_n | \mathcal{F}_{n-1}) = 0.$$

We assume that Y_n is square integrable for all n , and we set:

$$\sigma_n^2 = \sum_{k=1}^n \mathbb{E}(Y_k^2), \quad R_n = \sum_{k=1}^n \mathbb{E}(Y_k^2 | \mathcal{F}_{k-1}).$$

We assume that two following conditions are satisfied:

1. The convergence: $\frac{R_n}{\sigma_n^2} \rightarrow 1$ holds in probability,
2. (Lindberg condition):

$$\forall \epsilon > 0, \quad \lim_{n \rightarrow \infty} \frac{1}{\sigma_n^2} \sum_{k=1}^n \mathbb{E}(Y_k^2 \mathbf{1}_{\{|Y_k| > \epsilon \sigma_n\}}) = 0.$$

Then we have the convergence in law:

$$\frac{Y_1 + \cdots + Y_n}{\sigma_n} \xrightarrow{\mathcal{L}} \mathcal{N}(0, 1),$$

where $\mathcal{N}(0, 1)$ is the normal law.

III-2.3 Conjecture. (*Central Limit Theorem for concurrent systems*) Let $(\mathcal{N}, (\mu^l)_{l \in \mathcal{L}_U})$ be a distributed and recurrent Markov net, that we assume of integrable concurrent height. Let f be a distributed function with zero mean defined on the unfolding \mathcal{U} . Let $\sigma^2(f)$ denote the covariance form of f (defined in III-1.1), and for each $B \in \mathcal{S}_0$, set:

$$\sigma_B^2 = \mathbb{E}(\langle \sigma^2(f), \omega_B \rangle).$$

Let $(B_n)_{n \geq 1}$ be a non-decreasing sequence of stopping prefixes, such that $B_n \rightarrow \mathcal{E}$. If $(B_n)_n$ is regular (in the sense of regular stopping operators), we expect the convergence in law to the normal law $\mathcal{N}(0, 1)$:

$$\frac{\langle f, \omega_{B_n} \rangle}{\sigma_{B_n}} \xrightarrow{\mathcal{L}} \mathcal{N}(0, 1).$$

IV—Conclusion

We have given a framework to state the Strong law of large numbers for concurrent systems. This framework is based on the notion of *distributed function*, that is like a *local test function*. Distributed functions also provide a natural unit of time for concurrent systems, that is not in general a single event, unlike in sequential systems. The Strong law of large numbers establishes the limit of the asymptotic ratio of the integrated value of a distributed function along a process, normalised with the concurrent time elapsed along the process.

The *density coefficients* of branching cells replace the stationary measure for Markov chains. We miss in this study the *positivity* of the density coefficients, related to the finite expectation of the renewal.

Chapter 8

Algorithms for Statistics

In this chapter, we make a short insight into statistics for our model. We address the problem of global parametric estimation for distributed Markov nets. Local estimation is the crucial point for applications in management of networks, and is the topic of a discussion in Chapter 9. A first case of interest consists however in the simpler *global* estimation problem.

Unlike the local estimation problem, the global statistical estimation does not present theoretical difficulties. Due to the treatment of concurrency by distributed probabilities, a full observation of a recurrent and distributed Markov net leads to the analysis of *i.i.d* sequence of random variables, the most well-know situation for statistical estimation. Hence the statistical estimation itself does not present difficulty.

The more interesting points of the estimation problem are twofolds. First, we need to compute the geometric characteristics of a net, basis of the probabilistic model. A first question is to decide if a net is compact—*i.e.*, if its unfolding is locally finite—, a second question is to compute the finite collection of branching cells of compact nets. We propose procedures that always end for these two tasks, but:

1. The procedure that decides if a net is compact may find that a net is compact, whereas the net is not.
2. Computing branching cells involves several sub-procedures, one of them is missing.

Although the study is not complete, we can however expect that compactness of a net is decidable and that branching cells of compact nets are computable.

The second question addressed by the estimation problem examines the complications introduced by a sequential treatment of concurrent data, since we consider that an operational treatment of data always involves a phase with sequential computations. The risk is to wait for an infinite time before receiving the data that we are interested in. We formalise the situation with a non-deterministic variable¹ that

¹In Computer science, unlike in Probability, a non-deterministic variable is not a random variable, but is a variable with undetermined value.

gives a sequentialisation—the observed data—of a maximal true-concurrent process. We show how this variable can be bounded by a purely random variable with finite mean. As a consequence, the observation delay due to concurrency is asymptotically negligible, and we obtain thus an operational estimation procedure.

Section I, *Computation of branching cells*, states and studies the computation of geometric characteristics of nets: compactness, branching cells of the unfolding. Section II, *Operational statistical estimation*, establishes the statistical estimation procedure. We introduce and study the observation delay due to concurrency.

I—Computation of Branching Cells

For efficient probabilistic procedures, we certainly need to compute the basic theoretical objects of this study: stopping prefixes and branching cells. We state the computational questions that one encounters for the statistical procedure proposed in II, and we bring some (partial) answers. We show that some questions reduce to a reachability problem in a sub-net of the original net.

I-1 Computational Problems.

I-1.1 Computing the Unfolding of a Safe Marked Net. It is well known that prefixes of arbitrary size of the unfolding of a safe net can be recursively computed. The *model-checking* of concurrent systems provides tools for constructing a *complete prefix* of the unfolding, that is a finite prefix where all the reachable markings of the net are present as cuts of the prefix [32, 19].

I-1.2 Computing Stopping Prefixes. However, we must be careful when using the term “computed”, for a prefix of the unfolding. Consider for instance the unfolding of the very simple net depicted in Figure 8.1. The condition labelled by B has infinite branching, but this is not an information directly delivered by the complete prefix. Whence the following questions.

1. Is there an algorithm that decides if a safe net is compact, *i.e.* if the unfolding of the net is locally finite?
2. Is there an algorithm that computes the classes of branching cells of the unfoldings of compact nets?

I-1.3 Flower of an Event. Both questions reduce more or less to the problem of computing $B(e)$, the smallest stopping prefix that contains an event e of the unfolding. We recall that a stopping prefix is a prefix closed under the dynamic conflict. The dynamic conflict relation is defined on events $x, y \in \mathcal{E}$ by:

$$x \#_d y \Leftrightarrow x \# y, \quad [x] \cup [y] \in \mathcal{W}, \quad [x] \cup [y] \in \mathcal{W}, \quad (8.1)$$

where \mathcal{W} is the partial order of configurations. An intermediate object of interest for computing $B(e)$ is certainly the set of events:

$$F(e) = \{x \in \mathcal{E} : x \#_d e\}.$$

We define $F(e)$ as the **flower** of event e , and we formulate these questions:




Figure 8.1: A single loop and a prefix of the unfolding. The infinite conflict is not dynamic.

1. *a.* Can the finiteness of a flower be decided? *b.* Does the finiteness of all flowers imply the compactness of the net?
2. If \mathcal{N} is a compact net, the flowers of the unfolding of \mathcal{N} are finite.
 - a.* Can we compute the flowers? *b.* Can we compute the branching cells from the flowers?

Question 2*b*, the computation of branching cells from flowers, is treated in I-2. For question 2*a*, we reduce the computation of flowers to a reachability problem (I-3) combined with an other related question. For question 1*a*, we give in I-3.6 a computable sufficient criterion for non compactness. The criterion concerns the existence of a cycle in a finite and computable graph. The computation always ends, and decides the validity of a condition that implies the non compactness of the net. It must be noted that all the examples of this document behave as if the criterion was necessary and sufficient.

I-2 Computing Branching Cells from Flowers.

If one can compute the flower $F(e)$ of any given event in the unfolding of a net, then one can compute the finite collection of branching cells of the unfolding of a compact net. This is the result of I-2.

I-2.1 From Flowers to Stopping Prefixes. Assume that there is an algorithm that computes the flower $F(x)$ of any event of the unfolding of a safe marked

and compact net. Then we can compute the stopping prefix $B(e)$, smallest stopping prefix that contains e , for every event e .

Indeed fix \mathcal{N} a safe and compact marked net, and e an event of the unfolding of \mathcal{N} . The following increasing sequence of subsets of \mathcal{U} is then computable, where $\downarrow A$ denotes the downward closure of a subset $A \subseteq \mathcal{U}$:

$$B_0 = \emptyset, B_1 = [e], \quad B_{n+1} = B_n \cup \bigcup_{x \in B_n \setminus B_{n-1}} \downarrow F(x). \quad (8.2)$$

I-2.2 Lemma. *The sequence $(B_n)_{n \geq 1}$ defined by (8.2) is a non decreasing sequence of prefixes, subsets of $B(e)$. There is an integer j such that $B_j = B_{j+1}$, and then $B_n = B_j = B(e)$ for all $n \geq j$.*

Proof – It is obvious that for every j , we have:

$$B_j = B_{j+1} \Rightarrow \forall n \geq j, \quad B_n = B_j.$$

An induction shows the following properties of $(B_n)_{n \geq 1}$:

1. B_n is a prefix subset of $B(e)$.
2. For all $x \in B_n$, and for all event $y \in \mathcal{U}$, $x \#_d y \Rightarrow y \in B_{n+1}$.

Since $(B_n)_n$ is non decreasing, and since the net is compact, it follows from 1 that there is an integer j such that $B_j = B_{j+1}$. The point 2 implies that B_j is $\#_d$ -closed. Hence B_j is a stopping prefix that contains e , and is a subset of $B(e)$: so $B_j = B(e)$. \square

I-2.3 From Stopping Prefixes to Initial Branching Cells. Assume that we know how to compute $B(e)$ for any event e in the unfolding of a safe and compact marked net \mathcal{N} . Then we compute the initial branching cells $\lambda \in \Delta^\perp(\mathcal{U})$ as follows. Observe that the minimal events in the unfolding represent the transitions that can play from the initial marking of the net; they are thus in finite number.

Procedure: Compute $B(e)$ for each minimal event $e \in \mathcal{U}$. The collection $\Delta^\perp(\mathcal{U})$ of initial branching cells of \mathcal{U} is formed by the minimal elements of the finite family $\{B(e), e \text{ minimal in } \mathcal{U}\}$.

Proof – Let $B(e)$ be a minimal element of the family $\{B(e)\}$ as stated by the above procedure. We have to show that $B(e)$ is an initial branching cell, *i.e.* that $B(e)$ is a minimal non void stopping prefix. Since \mathcal{U} is an unfolding, the event structure \mathcal{E} associated with \mathcal{U} has finite concurrent width. By Lemma III-1.2 of Ch. 3, $B(e)$ contains an initial branching cell λ . (We could use instead that \mathcal{U} is locally finite to obtain the same result). Then λ admits a minimal event f , that is also minimal in \mathcal{U} , and $B(f) \subseteq \lambda \subseteq B(e)$. Since $B(e)$ is minimal it follows that $B(e) = B(f) = \lambda$. Hence $B(e) \in \Delta^\perp(\mathcal{U})$.

Conversely, let λ be an initial branching cell of \mathcal{U} . Then for any event e , minimal in λ , e is minimal in \mathcal{U} and satisfies $\lambda = B(e)$. \square

I-2.4 Conclusion: From Flowers to Branching Cells. From the two above procedures, if (\mathcal{N}, M_0) is a safe compact marked net, and if the flowers are computable, we can compute the initial branching cells of (\mathcal{N}, M_0) . With the complete prefix technique we can compute all the reachable markings M , from which we can compute the initial branching cells of (\mathcal{N}, M) . The result of these finite computations contain all the branching cells of the unfolding of (\mathcal{N}, M_0) . We obtain more than the branching cells, since the set of reachable markings contains markings that are not well-stopped.

I-3 Finiteness of Flowers.

We analyse the flowers of events. We give a sufficient and computable condition for non-finiteness of a flower, which implies the non local finiteness of the unfolding.

I-3.1 Labels of the Flower $F(e)$. We say as in [33] that two events x, y of the unfolding are in **immediate conflict** if they share a precondition: $x \neq y$ and $\bullet x \cap \bullet y \neq \emptyset$. We have the following observation:

If x, y are two events in dynamic conflict, then x and y are in immediate conflict.

Proof – Let x, y with $x \#_d y$. Since in particular x and y are in conflict, there is a condition b and two distinct events $e, f \in b^\bullet$ such that $e \preceq x$ and $f \preceq y$. Since the conflict between x and y is dynamic, $[x \cup [y]$ is conflict-free. It implies that $e \notin [x[$, and then $e = x$. Symmetrically, we get that $f = y$. We have thus $\bullet x \cap \bullet y = \bullet e \cap \bullet f \ni b$, so x and y are in immediate conflict. \square

As a consequence, if e is an event labelled by the transition $\tau = \rho(e)$ of the net, the events in the flower $F(e)$ are labelled by transitions t that satisfy:

$$\bullet t \cap \bullet \tau \neq \emptyset.$$

Hence we have reduced the labels of events in $F(e)$ to a subset of transitions. An other reduction arises if we consider a minimal event e . Remark that the minimal events are easily computable, since they are one-to-one with the transitions of the net that are enabled from the initial marking.

I-3.2 Lemma. *Let \mathcal{N} be a safe marked net, with unfolding $\rho : \mathcal{U} \rightarrow \mathcal{N}$. Let e be a minimal event of the unfolding, labelled by the transition $\tau = \rho(e)$, and let $L = \bullet \tau$ denote the preset of transition τ in \mathcal{N} .*

Let $x \in F(e)$, and let $v = [x[$. Then for every condition b of v , we have:

$$\rho(b) \in L \Rightarrow b \in \bullet e.$$

Proof— Let b be a condition of $v = [x[$. As $x \in F(e)$, e is compatible with v , so we have one of the three possibilities:

$$b \parallel e, \quad b \preceq e, \quad b \succeq e.$$

If $b \succeq e$ then $x \succeq e$, which contradicts that $x \# e$, so this possibility is discarded. If $b \preceq e$, as e is a minimal event, then $b \in \bullet e$. Thus the possibilities reduce to:

$$b \parallel e, \quad b \in \bullet e.$$

Remark that if $b \parallel e$, then $b \parallel \bullet e$. The mapping ρ is injective on \parallel -cliques. Therefore: $\rho(b) \in L$ and $b \parallel e$ imply $b \in \bullet e$. Hence we always have: $\rho(b) \in L \Rightarrow b \in \bullet e$. \square

I-3.3 Reduction to a Reachability Problem. Let v denote as in Lemma I-3.2 the configuration $v = [x[$, with x an event in the flower of a minimal event e , and with $\tau = \rho(e)$. Let $s = (t_1 \dots, t_k)$ be a playing sequence of the net compatible with v . Then Lemma I-3.2 implies that $\tau \neq t_i$ for all i . Denote by M_0 the initial marking of the net. Since $\tau \neq t_i$ for all i , s is a playing sequence of the net (\mathcal{N}, M) , with $M = M_0 \setminus \bullet \tau$. We have, in the net (\mathcal{N}, M) :

$$(\mathcal{N}, M) : M \xrightarrow{s} m,$$

where m is a marking of \mathcal{N} that satisfies: $m \cap \bullet \tau = \emptyset$. In the net (\mathcal{N}, M_0) , the playing sequence s , compatible with $v = [x[$, enables the transition $t = \rho(x)$ since $e \#_d x$. The marking reached by s in (\mathcal{N}, M_0) is $m \cup \bullet \tau$. Therefore:

$$\bullet t \subseteq m \cup \bullet \tau.$$

We have seen in I-3.1 that t satisfies: $\bullet t \cap \bullet \tau \neq \emptyset$.

Conclusion: If x is in the flower $F(e)$ of a minimal event e with $\rho(e) = \tau$, the transition $t = \rho(x)$ satisfies:

1. $t \neq \tau, \quad \bullet t \cap \bullet \tau \neq \emptyset$.
2. Let $M = M_0 \setminus \bullet \tau$. There is a marking m reachable from (\mathcal{N}, M) and such that: $\bullet t \subseteq m \cup \bullet \tau$.

The converse is the topic of the following result, which proof is straightforward.

I-3.4 Proposition. Let e be a minimal event in the unfolding (\mathcal{U}, ρ) of a safe marked net (\mathcal{N}, M_0) , labelled by transition $\tau = \rho(e)$. Let $M = M_0 \setminus \bullet \tau$. Let s be a finite playing sequence of the safe net (\mathcal{N}, M) , leading from M to m . Let v be the configuration of \mathcal{U} that lifts the sequence s , seen as a playing sequence of (\mathcal{N}, M_0) . Assume that there is a transition t of \mathcal{N} that satisfies:

1. $t \neq \tau, \quad \bullet t \cap \bullet \tau \neq \emptyset$.
2. $\bullet t \subseteq m \cup \bullet \tau$.

Then there is an event x in the flower $F(e)$, with $\rho(x) = t$ and enabled by v .

From this characterisation we derive a computable sufficient condition for the non finiteness of flowers, which implies the non-local finiteness of the unfolding.

1-3.5 A Relation on the Sub-Markings. We still consider a minimal event e of the unfolding (\mathcal{U}, ρ) of (\mathcal{N}, M_0) , with $\tau = \rho(e)$, and we set:

$$M = M_0 \setminus \bullet\tau .$$

Using the complete prefix algorithm applied with (\mathcal{N}, M) , the following set is computable:

$$\mathcal{M} = \{m \text{ reachable from } (\mathcal{N}, M), \text{ and such that there is} \\ \text{a transition } t \text{ satisfying points 1 and 2 of 1-3.4} \} .$$

Remark that, due to the safety hypothesis on (\mathcal{N}, M_0) , we have that $m \cap \bullet\tau = \emptyset$ for every $m \in \mathcal{M}$.

We define the binary transitive relation \leq on \mathcal{M} as follows. Let $m, m' \in \mathcal{M}$. Let t denote a generic transition that comes with the definition of $m \in \mathcal{M}$. By Proposition 1-3.4 applied to m and t , there is an event $x \in F(e)$ labelled by t and enabled by v . We set $m \leq m'$ if and only if there is a sequence a transitions s such that:

$$m \xrightarrow{s} m' ,$$

and containing a transition q with $\bullet q \cap \bullet t \neq \emptyset$, for at least one such transition t .

1-3.6 Theorem. *With the notations of 1-3.5, if (\mathcal{M}, \leq) admits a non trivial cycle, then the flower $F(e)$ of the minimal event e is infinite.*

Proof— (\Leftarrow) Set $<$ the relation $\leq \setminus \text{Id}_{\mathcal{M}}$ on \mathcal{M} , and assume that \mathcal{M} admits a chain $m_1 < \dots < m_n$ with $m_1 = m_n$. Let r and s be the playing sequences of \mathcal{N} such that we have:

$$M \xrightarrow{r} m_1 \xrightarrow{s} m_1 .$$

Then we consider the concatenation $z_k = rs^k$ for every integer $k \geq 0$. Let v_k be the lifted configuration of z_k in \mathcal{U} . There is an event x_0 , labelled by a transition t and enabled by v_0 , and such that: $x_0 \in F(e)$, and there is a transition q of s , with $\bullet q \cap \bullet t \neq \emptyset$.

For each $k \geq 1$, there is an event $x_k \in F(e)$ labelled by the same transition t , and enabled by v_k . Since transition $q \in s$ satisfies $\bullet q \cap \bullet t \neq \emptyset$, all x_k are disjoint, and thus $F(e)$ is infinite. \square

1-4 Computing Flowers.

The work presented here to compute the flowers is not complete. We indicate the result that can be expected, in our point of view.

I-4.1 Reduction to Flowers of Initial Events. Let \mathcal{U} be the unfolding of a safe marked net (\mathcal{N}, M_0) . We recall from Ch. 3, II-4, that the dynamic conflict relation $\#_d^{\mathcal{U}^v}$ in the cone of future of a configuration v satisfies:

$$\#_d^{\mathcal{U}^v} = \#_d \cap (\mathcal{U}^v \times \mathcal{U}^v).$$

In particular if e is an event of \mathcal{U} , the flower $F(e)$ coincides with the flower of e in the cone of future $\mathcal{U}^{[e]}$ of configuration $[e]$. Since e is minimal in $\mathcal{U}^{[e]}$, this shows that computing flowers is equivalent to computing flowers of *minimal events*.

I-4.2 Computing the Flower of a Minimal Event. We keep the notations of Proposition I-3.4, in particular the sub-marking $M = M_0 \setminus \bullet\tau$. We assume that the net \mathcal{N} is *compact*. We fix t a transition satisfying points 1 and 2 of I-3.4, and we want to determine the events of the flower $F(e)$ labelled with transition t . Repeating the operation for every t leads to the flower $F(e)$.

Let Ω' denote the set of maximal configurations of the unfolding of (\mathcal{N}, M) , we have an injection $\Omega' \hookrightarrow \mathcal{W}$, with \mathcal{W} the configurations of \mathcal{U} . For each $\omega' \in \Omega'$, there is a minimal $v \subseteq \omega'$ such that the cut $\gamma(v)$ contains all conditions labelled by $\bullet t$. Each cut enables an event of the flower $F(e)$ labelled by t . It is clear that the set of all these cuts is finite, otherwise the flower $F(e)$ would not be finite, and the unfolding would not be locally finite.

Although this set is finite, I am not sure of the way to compute it. Assume that it has been computed. Then we repeat the operation in the cone of future \mathcal{U}^c for each of these cuts c : computing the minimal cuts that contain the conditions of $\bullet t$, that enable as above an event of $F(e)$ labelled with t . We continue the computation as much as we can, and the computation ends since the net is compact.

I-4.3 Conclusion. The computation of flowers reduces to the computation of flowers of *minimal events*. If some finite set in the unfolding of a *sub-marking* of the initial marking can be computed, then one can compute the flower of a minimal event.

II—Operational Statistical Estimation

In this section, we address the statistical estimation problem from a very pragmatic point of view. A recurrent net is given, equipped with a distributed and homogeneous probability. How can the parameters of the net be retrieved from observations? It is important to precise what we mean by “observations”.

Unlike what can be called a “logical” observation, which is formally represented by a partially ordered set of events, we can assume that a “physical” observation is

given by a sequence of events. Indeed, if the events correspond for instance to alarms in a network, the centralised observer will receive the events through a channel, that delivers totally ordered sequences. The point is that the sequentialisation could make the computational process wait during an infinite time: but this does not happen, in probability.

The goal of this section is thus to provide an estimation algorithm, with accepts as input a totally ordered sequence of transitions of a net, compatible with a maximal trace of the net. The output at physical time n , the instant where the n^{th} event is received, is an approximation of the probabilistic parameters of the net. We show that, under conditions on the concurrent height of the net, the approximations converge to the true value, and we characterise the convergence rate.

II-1 Sequential Observations.

II-1.1 Computing the Unfolding. We assume that the geometric characteristics of the unfolding of a safe and compact Petri net have been computed. Section I has given some indications on this topic.

II-1.2 Receiving Sequences. Our theoretical results deal with stopped and well-stopped configurations. But what we really observe are sequences of transitions. We recall that any playing sequence of transitions $(t_n)_n$ in a net \mathcal{N} admits a unique lifted sequence of events in the unfolding. If $\rho : \mathcal{U} \rightarrow \mathcal{N}$ denotes the folding mapping, the lifted sequence $(e_n)_n$ is the unique sequence of events of the unfolding such that:

1. $\{e_1, \dots, e_k\}$ is a configuration of the event structure of the unfolding,
2. $\rho(e_k) = t_k$ for all k .

Drawing the lifted configuration from the sequence is straightforward, since we know the structure of the net \mathcal{N} . Receiving a sequence of transitions, we want to build the dynamic puzzle of branching cells around the lifted configuration, to isolate the germs “on the fly”.

The remaining point is that nothing insures that the observations will entirely fill in a germ before beginning an other one, which is needed for statistical computations.

II-2 Observation Delay.

II-2.1 Floating Events and Floating Front. Let w be a finite configuration of the unfolding \mathcal{U} of a compact net \mathcal{N} . Since compatible well-stopped configurations form a lattice, there is in particular an upper bound for the $v \in \mathcal{X}_0$ included in w . We set:

$$\forall w \in \mathcal{W}_0, \quad C(w) = \sup\{v \in \mathcal{X}_0 \mid v \subseteq w\}.$$

We say that the events contained in $w \setminus C(w)$ are the **floating events** of w . The collection of floating events is called the **floating front** of w .

II-2.2 Integrable Floating Front. We denote by $|v|$ the number of events of a finite configuration v of the unfolding. Let v be a finite well-stopped configuration. For each $\xi \in \Omega(v)$, we set:

$$f^v(\xi) = \sup\{|w| - |v| : w \subseteq \xi, C(w) = v\} .$$

The integer $f^v(\xi)$ is the maximal number of floating events of a configuration $w \subseteq \xi$, satisfying $C(w) = v$. Clearly, $v \rightarrow f^v$ is a congruence for $v \in \mathcal{X}_0$. In other words, $(f^v)_{v \in \mathcal{X}_0}$ is a d -homogeneous function (Ch. 5, II-2.3).

We say that the net \mathcal{N} has an **integrable floating front** if $(f^v)_v$ is integrable, *i.e.* if there is a constant K such that:

$$\forall v \in \mathcal{X}_0, \quad \mathbb{E}^v(f^v) \leq K < \infty .$$

II-2.3 Consequences for the Receiving Sequences. Assume that we observe a net \mathcal{N} with integrable floating front and integrable concurrent height. Let ω be the maximal execution of the net from which we get an observation sequence $(t_n)_n$. Let $v \in \mathcal{X}_0$ be a well-stopped sub-configuration of ω . We want to estimate the instant n such that the lifted v_n of $(t_k)_{1 \leq k \leq n}$ in the unfolding contains v .

Let n be the minimal integer such that $v_n \supseteq v$. We have the chain:

$$v \subseteq C(v_n) \subseteq v_n .$$

There is a condition b common to both $\gamma(v)$ and $\gamma(v_n)$. Hence b belongs also to $\gamma(C(v_n))$. We have the decomposition:

$$v_n = v \oplus (C(v_n) \ominus v) \oplus (v_n \ominus C(v_n)) ,$$

which implies the following:

$$n = |v_n| = |v| + p_n + q_n ,$$

where p_n and q_n are variables with a random part and a non-deterministic part, due to the ordering of the sequence $(t_n)_n$. Both variables p_n and q_n are controlled by purely random variables. One recalls the definitions of the random variable L_b^v and of the branching distance D (Cf. Ch. 7, II-1.1). If we denote by k the maximal number of events contained in a branching cell, we have:

$$p_n \leq k \sup_{b \in \gamma(v)} D(v, L_b^v), \quad q_n \leq f^v ,$$

where f^v is the front variable (II-2.2). Setting $Q^v = k \sup_{b \in \gamma(v)} D(v, L_b^v) + f^v$, we get the following.

Conclusion: Let v be a well-stopped sub-configuration of the maximal configuration to be observed through an observation sequence $(t_n)_n$. Assume that the net has integrable floating front and integrable concurrent height. The first instant n at which we observe a lifted configuration v_n that contains v , satisfies:

$$|v| \leq n \leq |v| + Q^v,$$

where Q^v is an integrable homogeneous function.

II-3 The Sequential Estimation Algorithm.

We consider a compact marked net (\mathcal{N}, M_0) . A statistical model associated to (\mathcal{N}, M_0) is the family of distributed Markov nets $(\mathcal{N}, M_0, (\mu_\theta^l)_{\theta \in \Theta})$, where θ is a parameter. The natural parameter is the finite collection of branching probabilities $\theta = (\mu_\theta^l)_{l \in \mathcal{L}_U}$. Hence we assume that Θ is a subset of the \mathbb{R} -vector space of finite dimension that contains the probability vectors $(\mu_\theta^l)_l$. For simplicity we assume that μ_θ^l is a positive branching probability for each $\theta \in \Theta$.

We observe a process through a sequential observation, compatible with a maximal process under an unknown true value θ^0 of the parameter. We want to estimate θ^0 .

II-3.1 An Algorithm for Computing the Observed Well-Stopped Configurations. Let (t_1, t_2, \dots) be the sequential observation of the net (\mathcal{N}, M_0) . We consider first an algorithm that computes by successive increments the lifted configuration v_n of (t_1, \dots, t_n) and the well-stopped configuration $C(v_n)$. We recall (Cf. II-2.1) that $C(v_n)$ is the upper bound of the well-stopped configurations contained in v_n . Computing v_n from v_{n-1} and t_n is straightforward: add a new event e_n labelled by t_n —the preconditions labelled by $\bullet t_n$ are present in $\gamma(v_{n-1})$ —and add the postconditions labelled by t_n^\bullet to e_n . We denote it shortly by: $v_n := v_{n-1} \oplus e_n$.

The following algorithm computes the lifted v_n together with $C_n = C(v_n)$, provided that the branching cells can be computed.

Initialisation: $v_0 := \emptyset, C_0 := \emptyset$.

Step n :

1. Receive t_n , set $v_n := v_{n-1} \oplus e_n$ with e_n labelled by t_n .
2. Compute the branching cells in $\Delta^+(C_{n-1})$.
3. **For:** Each branching cell $\lambda \in \Delta^+(C_{n-1})$,
Do: Compute $v_n \cap \lambda$ and set:

$$C_n := \begin{cases} C_{n-1} \cup (v_n \cap \lambda), & \text{if } v_n \cap \lambda \in \Omega_\lambda, \\ C_{n-1}, & \text{otherwise.} \end{cases}$$

Done.

Step $n + 1$: ...

Remark that the result of the loop 3 (b) does not depend on the order in which the loop is processed.

II-3.2 The Successive Branching Cells Algorithm. It is then straightforward to obtain an algorithm that computes the successive branching cells λ within a fixed class $l \in \mathcal{L}_U$, evaluated on the observed trace ω . In Chapter 6, we have called them the *successive l -branching cells*.

From the above algorithm, we extract the sequence of integers $(n_i)_{i \geq 1}$ where C_n is actually incremented:

$$n_{i+1} = \inf\{n > n_i \mid C_n \neq C_{n-1}\}.$$

A procedure to obtain the successive l -branching cells λ_j and the associated germs Z_j is then for instance the following.

Initialisation: $W_0^l = \emptyset, k_0 = 0$.

Step j :

1. Compute: $k_j := \inf\{i > k_{j-1} \mid \exists \lambda \in \Delta^+(C_{n_i}) : \lambda \cong l\}$.
2. Set λ_j the unique branching cell $\lambda \in \Delta^+(C_{n_{k_j}})$ such that $\lambda \cong l$.
3. Set $Z_{j-1}^l := C_{n_{k_j}} \cap \lambda_{j-1}$.

Step $j + 1$: ...

II-3.3 The Sequential Empirical Estimator. The successive l -germs Z_j^l are independent *i.i.d* with the law μ_θ^l in Ω_l (Ch. 6, Th. III-3.2). Hence the empirical

estimator of μ_θ^l is computed, for the j^{th} l -branching cell, by:

$$\forall v \in \Omega_l, \quad \hat{\mu}_j^l(v) = \frac{1}{j} \text{Card}\{i \leq j : Z_i^l = v\}. \quad (8.3)$$

We call the estimator (8.3) the **sequential empirical l -estimator**.

The consistency and the convergence rate of the estimator (8.3) are well known, provided that the sample $(Z_j^l)_j$ is *infinite*. We recall the result below. Then we will have to consider the growth of index j *w.r.t.* the arrivals of events. Unlike the convergence *w.r.t.* index j , this last growth let appear the concurrency properties of the system.

II-3.4 Proposition. *Let $(\mathcal{N}, (\mu_\theta^l)_{l,\theta})$ be a statistical model associated to a compact marked net \mathcal{N} . We assume that for the true value θ_0 of the parameter, with probability \mathbb{P}_{θ_0} associated, the net $(\mathcal{N}, \mathbb{P}_{\theta_0})$ is recurrent. Then the sequential empirical estimator (8.3) is strongly consistent, that is we have the convergence \mathbb{P}_{θ_0} -a.s.:*

$$\forall v \in \Omega_l, \quad \lim_{j \rightarrow \infty} \hat{\mu}_j^l(v) = \mu_{\theta_0}^l(v).$$

If $\Omega_l = \{v_1, \dots, v_r\}$, we have the convergence in law:

$$\sqrt{j} \left\{ \frac{1}{\sqrt{\mu_{\theta_0}^l(v_k)}} \left(\hat{\mu}_j^l(v_k) - \mu_{\theta_0}^l(v_k) \right) \right\}_{1 \leq k \leq r} \xrightarrow{\mathcal{L}}_{j \rightarrow \infty} \mathcal{N}_r(0, \Gamma),$$

with $\mathcal{N}_r(0, \Gamma)$ the normal law in \mathbb{R}^r with zero mean and covariance matrix $\Gamma_{i,j} = \delta_i^j - \sqrt{\mu_{\theta_0}^l(v_i)} \sqrt{\mu_{\theta_0}^l(v_j)}$.

Proof – As noticed above, it is enough to show that the Z_j^l are in infinite number. But this is stated by Prop. III-2.2, Chapter 6, since the net is recurrent and d -homogeneous. The convergence in law comes from the Central Limit Theorem applied to the sequence $(Z_j^l)_j$ of *i.i.d* variables ([15], p. 92). \square

II-4 Convergence Rate of the Sequential Empirical Estimator.

Estimating the rate of convergence of the sequential empirical estimator involves two steps. We have to compute first the rate of convergence in j of the estimator (8.3)—this comes from classical statistical theory, and is stated in Proposition II-3.4. Second, we have to estimate the number of observed events needed to obtain the j^{th} l -germ Z_j^l .

This second term is typically a concurrency delay. We show below that this delay is related to the concurrent height and to the size of the floating front.

II-4.1 Lemma. *Let $(\mathcal{N}, (\mu^l)_{l \in \mathcal{L}_U})$ be a distributed and recurrent Markov net. Assume that l is a positive branching cell (i.e., there is a positive $v \in \mathcal{X}_0$ with $l \in \Delta^+(v)$), and that the net has integrable concurrent height. Let $(W_n^l)_{n \geq 1}$ be the successive local renewal operators of l (Cf. Ch. 6, III). Then we have:*

$$W_n^l(\omega) \uparrow_{n \rightarrow \infty} \omega, \quad \mathbb{P}\text{-a.s.}$$

Proof – Since the net is recurrent, and since l is positive, Prop. III-2.2 of Ch. 6 states that the successive renewal operators W_n^l are defined for all integers $n \geq 1$. Let ω such that $\bigcup_n W_n^l(\omega) \neq \omega$. Then there is a branching cell $\lambda \in \Delta^+(W_n^l(\omega))$ for infinitely many n . Let $v = W_n^l(\omega)$ for such an integer n . For b a condition $b \in \text{Min}_{\leq}(\lambda)$, we have (Cf. Ch.7, II-1):

$$\begin{aligned} L_b^v(\omega) &= \sup\{w \in \mathcal{X}_0 : w \subseteq \omega, b \in \gamma(b)\}, \\ &\supseteq W_j^l, \quad \text{for infinitely many } j. \end{aligned}$$

It follows that we have an infinite branching distance: $D(v, L_b^v(\omega)) = \infty$. Since the net has integrable concurrent height, we have the finite expectation: $\mathbb{E}^v(D(v, L_b^v)) < \infty$. It implies that the set of ω with $\bigcup_n W_n^l(\omega) \neq \omega$ has probability 0. This is equivalent to the statement of the lemma. \square

Lemma II-4.1 is intended to justify our application of the Strong law of large numbers, along the sequence W_l^n . However, we have not shown the Strong law of large numbers for nets with integrable concurrent height. We admit the following result.

II-4.2 The Positive Recurrence of Branching Cells. We now encounter the problem that we have already remarked: the positivity of the density coefficients of branching cells. We miss here this result, so we have to admit the following statement:

Let $(\mathcal{N}, (\mu^l)_{l \in \mathcal{L}_U})$ be a recurrent and distributed Markov net, with integrable concurrent height. Assume that every branching probability μ^l is positive (i.e.: $\mu^l(v) > 0$ for all $v \in \Omega_l$). Then the density coefficients $\alpha(l)$ are positive.

II-4.3 Consequence on the Arrival Rate of the l -Germs. The positivity of branching cells also implies that the local renewal operators are regular. Applying the Strong law of large numbers to the distributed function N^l we get:

$$\frac{\langle N^l, W_j^l \rangle}{\langle N, W_j^l \rangle} = \frac{j}{\langle N, W_j^l \rangle} \xrightarrow{j \rightarrow \infty} \alpha(l).$$

Since $\alpha(l) > 0$, it implies the following equivalent:

$$\langle N, W_j^l \rangle \sim_{j \rightarrow \infty} \frac{j}{\alpha(l)}. \tag{8.4}$$

To observe the j^{th} l -germ Z_j^l , we have to observe the configuration $W_j^l \oplus Z_j^l$. Using the conclusion of II-2.3, this requires a number n of events estimated by:

$$|W_j^l \oplus Z_j^l| \leq n \leq |W_j^l \oplus Z_j^l| + Q^{W_j^l \oplus Z_j^l},$$

where Q is an integrable homogeneous function. Denoting by k the maximal number of events contained in a branching cell, we get:

$$n \leq k(1 + \langle N, W_j^l \rangle) + Q^{W_j^l \oplus Z_j^l}.$$

By Proposition II-2.1 of Ch. 7 we have:

$$\lim_{j \rightarrow \infty} \frac{Q^{W_j^l \oplus Z_j^l}}{j} = 0, \quad \mathbb{P}\text{-a.s.}$$

Using (8.4), we obtain the following result.

II-4.4 Theorem. *(Under positive recurrence of branching cells) Let $(\mathcal{N}, (\mu_l)_l)$ be a distributed and recurrent Markov net with all branching probabilities positive, with integrable concurrent height and with integrable floating front. Let $l \in \mathcal{L}_U$ be a class of branching cells. Let n be the number of transitions received from an observation sequence $(t_n)_n$ before Algorithm II-3.2 computes the j^{th} l -germ from an observation sequence $(t_n)_n$. Then we have for a constant K :*

$$\frac{n}{j} \leq K + \frac{Q^j}{j},$$

where $(Q^j)_j$ is a sequence of random variables defined on $\Omega^{W_j^l \oplus Z_j^l}$, satisfying $\lim_{j \rightarrow \infty} \frac{1}{j} Q^j = 0$.

III—Conclusion

We have made a short analysis of the computational questions that arise with the theoretical objects introduced in this document. In particular we have brought some answers for deciding the compactness of a safe net and for computing the geometric characteristics of its unfolding.

These elements are set up to be applied to a statistical procedure that addresses the problem of global parametric estimation in the probabilistic model of distributed Markov nets.

Conclusions and Perspectives

I—Contributions

In this document, we have presented a mathematical analysis of models from Concurrency theory: prime event structures and safe Petri nets. The probabilistic model is applied to a statistical estimation procedure.

The contributions of this document cover the following topics:

- Projective formalism for topological event structures. Compactness of the border at infinity for a class of concurrent systems.
- Geometric tools for event structures: cone of future, well-stopped configurations, germs and branching cells. Algorithms for computing germs and branching cells.
- Definition and construction of distributed probabilities and of memory-less concurrent systems: distributed Markov nets.
- Formalism for a Strong Markov property for concurrent systems: stopping and shift operators, homogeneous functions. The Markov property. Application to recurrent nets. Global and local renewal. Existence of the embedded Markov chain.
- A unit of time for concurrent systems. Ergodic nets. The Strong law of large numbers. Definition of a martingale.
- Statistical estimation. The problem of sequentialisation.

We have introduced a new formalism adapted to concurrent systems in order to generalise the techniques and results from finite Markov chains theory (Chapters 2 and 5). The new formalism allows to express for concurrent systems several notions from dynamical systems theory, without reference to a global totally ordered time. In several aspects, the behaviour of concurrent systems can be compared with that of Markov chains': for instance global recurrence of both models are similar, and both models satisfy the Strong law of large numbers (Chapter 7).

The concurrency properties of models introduce specificities that have been treated through the use of new mathematical objects. The class of *distributed probabilities*, leading to the definition of *distributed Markov nets* (Chapters 4 and 5), properly expresses local properties of concurrent systems, as illustrated by the study of *local renewal* for concurrent systems (Chapter 6). As shown by the statistical study of Chapter 8, the model that we propose improves the study of probabilistic distributed systems.

Many questions remain open. We discuss some of them below.

II—Open Questions

II-1 Topological and Computational Questions

We have studied in Chapter 2 the topology of the space of maximal configuration of event structures, showing its compactness if the event structure is locally finite. We expect this condition to be necessary, which remains to be proved. The examples that we have studied encourage us to expect the equivalence. This would fully justify the name of “compact nets” for nets with a locally finite unfolding.

From a computational point of view, there is an open question concerning the decidability of compactness of safe nets. We expect that the procedure that we have given in Chapter 8 is actually sound and complete for deciding the compactness of nets.

II-2 Limit Theorems

We have shown in Chapter 7 the Strong law of large numbers for nets. This important ergodic result justifies the introduction of a new unit of time, adapted to concurrent systems. In general this unit of time differs from the single event, that is the natural unit of time for concurrent systems, although it actually reduces to the single event for systems without concurrency. The stationary measure involved in the formulation of the Strong law of large numbers for concurrent systems such as finite Markov chains is replaced for concurrent systems by a finite collection of *density coefficients*.

Our ergodic study misses the result according to which the density coefficients are positive. We related this topic to the finite mean, counting with the concurrent unit of time, of the renewal operator.

The conditions that we have introduced, recurrence and integrable concurrent height of concurrent systems, seem well adapted to the treatment of ergodic proper-

ties. We can expect decompositions of nets into components with these properties, analogous to the recurrence classes of Markov chains.

The introduction of the density coefficients of a net, associated with the finite collection of branching cells of its unfolding, let us consider this collection as a new space of states, adapted to the treatment of concurrent systems. In particular, the density coefficients sum to 1 over this space of states. The advantage is that in most cases, this space of states will be much smaller than the set of reachable markings. A more extensive study from this point of view could be initiated. In particular, can we impose a “net structure” on this space of states?

III—Perspectives

Our study is based on *locally finite* event structures, what about non locally finite event structures? For nets, it appears that the non-compactness is related to compositionality problems. We also use this point of view to consider the problem of local statistical estimation.

III-1 Product of Nets.

G. Winskel has introduced a categorical framework for Petri nets, defining in particular categories of safe Petri nets and of event structures [48, 50] in which a product exists. Combined with a restriction operation, and with a labelling defined on two nets, the product leads to a synchronous product of nets defined *w.r.t.* the labelling, generalising for instance the product “à la Arnold-Nivat” often used by the *model-checking* school. Several classical applications of Probability in Computer science, such as simulation and bisimulation [29, 22], are based on a close framework.

The product of nets is the essential ingredient for defining a language of communicating processes closed to the family of CCS, SCCS from Milner, and semantically interpreted in the category of safe nets. In the following grammar of processes, X is a set of elementary processes and L a set of labels¹, with x and λ variables in X and in L [48, 50]:

$$p ::= \text{NIL} \mid x \mid \lambda p \mid p + p \mid p|_Q \mid p[\Xi] \mid p \text{SYN}_L p \mid \text{REC } x.p$$

Q is a subset of labels $Q \subseteq L$, and $\Xi : L \rightarrow L$ is a relabelling mapping.

The semantics in nets is informally described as follows. NIL represents a dead process that cannot perform any event. λp stands for a *guarding label* λ , that must

¹More precisely, L is a *synchronisation algebra*.

occur for process p to hold. The sum $p + q$ of two processes consists in p or in q . If Q is a set of labels, an execution of process $p|_Q$ is any execution of p where only Q -events occur. $p[\Xi]$ stands for the action of the relabelling function Ξ on p . $p \text{ SYN}_L q$ is the synchronous product defined *w.r.t.* L . Finally, $\text{REC } x.b$ stands for a recursive evaluation of the expression b , closed *w.r.t.* the variable x .

III-2 Probabilistic Composition. Non Compact Probabilistic Nets.

The important observation is that the composition operations in the above grammar of processes, in particular the sum and the product, *do not preserve local finiteness*. Consider for instance the two nets \mathcal{N}_0 and \mathcal{N}_1 depicted in Figure 9.1. Taking their synchronous product by synchronising on transitions labelled by the common label c , we obtain our usual non compact example, depicted in Figure 9.2.




Figure 9.1: Two safe nets \mathcal{N}_0 and \mathcal{N}_1 .

III-2.1 Meaning of the Space Ω . Let Ω denote the set of maximal processes of the product net. An element ω contains three informations: ω determines first a process v_0 in \mathcal{N}_0 , obtained by forgetting simply the actions in \mathcal{N}_1 and looking only at \mathcal{N}_0 . In the same way, ω determines a process v_1 in \mathcal{N}_1 . Finally, ω determines also the *precise interleaving* of v_0 and v_1 . Whence the following observations:

1. A simplification is brought if we assume that v_0 and v_1 are *maximal processes* of \mathcal{N}_0 and of \mathcal{N}_1 .
2. Do we really need to keep the whole information about the interleaving of processes?

Condition 1 means that no deadlock is introduced by the synchronisation product. Figure 9.3 depicts a simple example of product without this property. It seems that this deadlock problem has not been treated in the literature. For point 2, I think that there is a better space than Ω to consider, and I give some details below.

I propose to consider two compact nets $\mathcal{N}_0, \mathcal{N}_1$ such that their product is **max-synchronous**: maximal processes of the product project into maximal processes of each net. Given two probabilities $\mathbb{P}_0, \mathbb{P}_1$ respectively on Ω_0, Ω_1 (with obvious notations), consider the **synchronisation set**: the set of pairs $(\omega_0, \omega_1) \in \Omega_0 \times \Omega_1$




Figure 9.2: Non compact net \mathcal{N} obtained by synchronisation of \mathcal{N}_0 and \mathcal{N}_1 from Figure 9.1 (synchronisation on c). We recognise in \mathcal{U} the event structure of Ch. 2, Fig. 2.1.




Figure 9.3: Synchronisation on transition a : introduction of a dead-lock. Configuration \emptyset is maximal in the product but not in the left component.

that can synchronise. Remark that we do not pay attention to the many ways in which two maximal processes can synchronise in general (although in the examples depicted here, this consideration is meaningless since there is at most 1 way for a pair ω_0, ω_1 to synchronise). If this set has $\mathbb{P}_0 \times \mathbb{P}_1$ probability non zero, consider the conditional probability:

$$\mathbb{Q} = \mathbb{P}_0 \times \mathbb{P}_1(\cdot \mid \mathcal{D}) ,$$

where \mathcal{D} is the synchronisation set.

Observing one of the net under the product dynamics is like observing the net under an **observation probability** different from the original probability in general. For instance for the net depicted in Figure 9.2, the observation set is, with obvious notations:

$$\mathcal{D} = \{a^n \oplus c, n \geq 0\} \times \{c\} \cup \{(a^\infty, b)\}.$$

Since a^∞ has \mathbb{P}_0 probability 0, \mathcal{D} differs with a $\mathbb{P}_0 \times \mathbb{P}_1$ -zero set from:

$$\{a^n \oplus c, n \geq 0\} \times \{c\}.$$

Therefore, the observation probability in \mathcal{N}_0 matches the initial probability, whereas the observation probability in \mathcal{N}_1 is the degenerated probability $\mathbb{P}_1(c) = 1$. For the live double loop depicted in Figure 9.5, obtained from the product of nets of Figure 9.4, the observation probabilities match the original probabilities. In this example the global space Ω can be equipped with a probability \mathbb{Q} that makes (Ω, \mathbb{Q}) isomorphic as a probability space to the product space $(\Omega_0 \times \Omega_1, \mathbb{P}_0 \times \mathbb{P}_1)$. For compactness reasons, the isomorphism is not topological.




Figure 9.4: Two safe nets for the live double loop. Synchronisation on c .

III-2.2 Distributed Estimation and Systems Management. In this compositional context, a distributed estimation can be addressed as an estimation of the intrinsic probabilistic parameters from local observations under the observation probability. Estimating the observation probability is like the global estimation problem treated in Chapter 8. Retrieving the original probability from the observation probability is an other problem. As shown by the example of Figure 9.1, this is not always possible: in this example, the observation probability is always trivial, and contains thus no information on the original probability.

Distributed estimation of parameters is an example of application of the probabilistic model to the management of networked systems. The labelling of systems also




Figure 9.5: *The live double loop.*

introduces many ways to modelise hidden informations. The HMM (*Hidden Markov Models*) are common tools in classical system management. Analogous theory and tools for distributed Markov nets are expected, with applications for instance to the distributed diagnosis of systems [6].

III-2.3 Conclusion. Compositionality seems to be a rich topic for probabilistic and statistical applications, as well as for probabilistic bisimulations of true-concurrent systems. They provide examples of randomisation of non compact nets. A challenging work with applications to the management of systems is the introduction of HMM techniques.

Bibliography

- [1] S. Abbes. A probabilistic model for true concurrency, second version. Technical Report PI-1591, IRISA, 2004.
- [2] F. Baccelli. Ergodic theory of stochastic petri networks. *The Annals of Prob.*, 20:375–396, 1992.
- [3] F. Baccelli, G. Bohen, G. Olsder, and J.P. Quadrat. *Synchronisation and linearity*. Wiley, 1992.
- [4] F. Baccelli, S. Foss, and B. Gaujal. Free choice nets, an algebraic approach. *IEEE Transaction on Automatic Control*, 41(12):1751–1778, 1996.
- [5] C. Baier, H. Hermanns, J.-P. Katoen, and V. Wolf. Comparative branching-time semantics for markov chains. In *CONCUR 03*, volume 2761 of *LNCS*, pages 482–497. Sp. V., 2003.
- [6] A. Benveniste, E. Fabre, S. Haar, and C. Jard. Diagnosis of asynchronous discrete event systems, a net unfolding approach. *IEEE Trans. on Aut. Control*, 48(5):714–727, 2003.
- [7] A. Benveniste, S. Haar, and E. Fabre. Markov nets: probabilistic models for distributed and concurrent systems. *IEEE Trans. on Aut. Control*, 48(11):1936–1950, Nov. 2003.
- [8] A. Benveniste, S. Haar, E. Fabre, and C. Jard. Distributed monitoring of concurrent and asynchronous systems. In *Proc. of CONCUR 2003*, volume 2761 of *LNCS*, pages 1–26, 2003.
- [9] Nicolas Bourbaki. *Intégration, Chapitre IX*. Éléments de mathématique, fascicule xxxv. Hermann, 1969.
- [10] L. Breiman. *Probability*. SIAM, 1992.
- [11] E. Brinksma, J.-P. Katoen, R. Langerak, and D. Latella. A stochastic causality-based process algebra. *The Computer Journal*, 38(7), 1995.
- [12] G. Choquet. *Cours d’analyse. Tome 2, Topologie*. Masson, 1964.

-
- [13] K. Ciesielski. *Set theory for the working mathematician*. Cambridge University Press, 1997.
- [14] D. Dacunha-Castelle and M. Duflo. *Probabilités et statistiques, Tome 1*. Masson, 1990.
- [15] D. Dacunha-Castelle and M. Duflo. *Probabilités et statistiques, Tome 2*. Masson, 1990.
- [16] B.A. Davey and H.A. Priestley. *Introduction to lattices and orders*. Cambridge Mathematical Textbooks. Cambridge University Press, 1990.
- [17] J. Desel and R. Esparza. *Free choice Petri nets*. Cambridge University Press, 1995.
- [18] J. Engelfriet. Branching processes of Petri nets. *Acta Informatica*, 28:575–591, 1991.
- [19] J. Esparza and K. Heljanko. A new unfolding approach to LTL model-checking. In *ICALP 2000*, volume 1853 of *LNCS*, pages 475–486, 2000.
- [20] S. Genc and S. Lafortune. Distributed diagnosis of discrete-event systems using petri nets. In *Applications and Theory of Petri Nets 2003, 24th International Conference, ICATPN 2003*, volume 2679 of *LNCS*. Sp. V., 2003.
- [21] S. Haar. Clusters, confusion and unfoldings. *Fundamenta Informaticae*, 47(3-4):259–270, 2001.
- [22] H. Hermanns, U. Herzog, and J.-P. Katoen. Process algebra for performance evaluation. *T.C.S.*, 274(1):43–88, 2002.
- [23] P.W. Hoogers, H.C.M. Kleijn, and P.S. Thiagarajan. A trace semantics for petri nets. *Information and Computation*, 117:98–114, 1995.
- [24] J.-P. Katoen, C. Baier, and Letella D. Metric semantics for true-concurrent real time. *Theoretical Computer Science*, 254(1):501–542, 2001.
- [25] J.-P. Katoen and P.R. D’Argenio. *Lectures on formal methods and performance analysis: first EEF/Euro summer school on trends in computer science*, chapter General distributions in process algebra. Sp. V., 2002.
- [26] H.C.M. Kleijn, R. Morin, and B. Rozoy. Event structures for local traces. *Electronic Notes in Computer Science*, 16(2), 1998.
- [27] R. Kummetz and D. Kuske. The topology of mazurkiewicz traces. *T.C.S.*, 305:237–258, 2003.
- [28] M. Z. Kwiatowska. A metric for traces. *Inf. Processing Let.*, 35:129–135, 1990.

-
- [29] K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1), 1991.
- [30] J. Mairesse and S. Gaubert. Modeling and analysis of timed petri nets using heaps of pieces. *IEEE Trans. Automatic Control*, 44(4):683–397, 1999.
- [31] G. Markowsky and B. Rosen. Bases for chain-complete posets. *IBM J. res. Develop.*, 20:138–147, 1976.
- [32] K.L. McMillan. A technique of state space search based on unfolding. *Formal Methods in System Design*, 6(1):45–65, 1995.
- [33] M. Nielsen, G. Plotkin, and G. Winskel. Petri nets, event structures and domains, part 1. *T. C. S.*, 1981.
- [34] M. O. Rabin. Probabilistic automata. *Inf. and Control*, 6(3):230–245, 1963.
- [35] A. Raugi and J.-P. Conze. Cours de dea. théorie des martingales et systèmes dynamiques. IRMAR, Univ. de Rennes 1, 2000.
- [36] Pál Révész. *The laws of large numbers*. Ac. Press, 1968.
- [37] D. Revuz. *Markov chains*. Number 11 in N.-H. Math. Library. North-Holland, 1975.
- [38] H.L. Royden. *Real Analysis*. Macmillan, 1988.
- [39] G. Rozenberg and P.S. Thiagarajan. *Petri nets: Basic notions, structure, behaviour*, volume 424 of *LNCS*. Sp. V., 1986.
- [40] R. Segala. A compositional trace-based semantics for probabilistic automata. In *CONCUR'95*, volume 962 of *LNCS*, pages 234–248. Sp. V., 1995.
- [41] R. Segala and N. Lynch. Decision algorithms for probabilistic bisimulation. In *Proc. of CONCUR 2002*, volume 2421 of *LNCS*, pages 371–396, 2002.
- [42] Y.G. Sinai, editor. *Dynamical systems, ergodic theory and applications*. Sp. V, 2000.
- [43] D. Varacca. *Probability, nondeterminism and concurrency: two denotational models for probabilistic computation*. PhD thesis, BRICS, University of Aarhus, 2003.
- [44] D. Varacca, H. Völzer, and G. Winskel. Probabilistic event structures and domains. In *Proc. of CONCURR 04*, LNCS, 2004.
- [45] H. Völzer. Randomized non-sequential processes. In *Proc. of CONCUR 2001*, volume 2154 of *LNCS*, pages 184–201, 2001.
- [46] F. Wallner. Model-checking ltl using net unfoldings. In *CAV' 95*, volume 1427 of *LNCS*, pages 207–218, 1998.

- [47] G. Winskel. *Events in computation*. PhD thesis, University of Edinburgh, 1980.
- [48] G. Winskel. *Event structure semantics in CCS and related languages*. Number 140 in LNCS. Sp. V., 1982.
- [49] G. Winskel. Categories of models for concurrency. In *Seminar on Concurrency, Carnegie-Mellon University, July 1984*, number 197 in LNCS, pages 246–267. Sp. V., 1985.
- [50] G. Winskel. Petri nets, algebras, morphisms, and compositionality. *Information and Computation*, 72(3):197–238, 1987.

Index

- σ -algebra, 64
 - Borel, 64
 - of a stopping operator, 163
 - operational, 76
- a.s.*, 65
- additive probability, 176
- algorithm
 - branching cell, 228
 - compactness, 232
 - well-stopped configurations, 236
- Borel-Cantelli Lemma, 129
- bound
 - least upper, 39
- branching cell, 108
 - computing, 228
 - confusion-free, 181
 - counting, 138
 - equivalence, 171
 - equivalence class, 171
 - initial, 98
 - occurrence net, 171
 - positive, 134
 - successive, 198
 - algorithm, 237
 - tree, 118
- branching probability, 133
 - homogeneous family, 173
 - induced, 133, 135
 - positive, 133
- cancellation
 - \ominus , 92
- causality
 - in event structure, 40
 - in occurrence net, 47
- chain, 39
- closed
 - downwards, 41
 - dynamic conflict, 86
- coefficient, density, *see* density
- cofinal, 67
 - in event structure, 70
- compatible
 - configurations, 42
 - events, 42
- concatenation
 - \oplus , 92
 - cones, 92
 - well-stopped configuration, 103
- concurrency
 - in event structure, 40
 - in occurrence net, 48
- concurrent
 - subsets, 41
 - width, 48
 - event structure, 96
 - unfolding, 52
- condition, 47
- cone of future
 - in event structure, 90
 - in occurrence net, 155
 - probabilistic, 131
 - tree, 118
- configuration
 - compatibles, 42
 - cut, 49
 - domain, 42
 - full-initial, 98
 - in a cone, 90
 - in a prefix, 42
 - in event structure, 41

- in occurrence net, 49
 - marking, 51
 - maximal, 42
 - stopped, B -stopped, 89
 - stopped-initial, 105
 - sub-configuration, 41
- conflict
 - dynamic, 86
 - in a cone, 94
 - immediate, 47, 230
 - in event structure, 40
 - in occurrence net, 47
 - inherited, 40
 - subsets in, 41
- confusion-free
 - event structure, 179
 - net, 178
- cut, 48
 - configuration, 49
 - equivalence, 157
 - initial, 48
 - well-stopped, 166
- d -homogeneous
 - function, 167
 - net, 166
- decomposition
 - germ-decomposition, 112
 - initial stop. prefix, 104
 - normal, 99
 - induced, 100
 - random variables, 168
 - step-by-step, 107
- decreasing
 - sequence, 39
- delay, *see* observation
- density coefficient, 215
- distributed function, 206
 - martingale, 220
 - mean, covariance, 219
- distributed Markov net, 173
 - example, 176
- domain
 - concrete, 179
 - of configurations, 42
- equivalence
 - branching cell, 171
 - cut, 157
- ergodic
 - mean
 - nets, 207
 - sequential, 205
 - net, 212
- estimator, 237
- event
 - in event structure, 40
 - in occurrence net, 47
- event structure, 40
 - confusion-free, 179
 - empty, 40
 - locally finite, *see* locally finite
 - prime, 38
 - probabilistic, 76
 - induced, 76
 - sub-structure, 41
- expectation, 125
 - conditional
 - w.r.t.* a stopping prefix, 127
 - on a σ -algebra, 126
 - on a subset, 125
- extension
 - compact event structures, 78
 - locally finite event structures, 80
 - Prokhorov's Theorem, 68
- floating event, front, 234
- flow relation, 44
- flower, 227
 - computing, 232
 - finiteness, 230
- folding, 50
 - morphism, 51
- function
 - d -homogeneous, 167
 - distributed, 206
 - homogeneous, *see* homogeneous
- future, *see* cone of future

- germ, 112
 - absolute random, 134
 - conditional random, 136
 - initial, 105
 - successive l -germs, 199
- Hausdorff space, 63
- height
 - integrable concurrent, 210
- homogeneous
 - collection of branching probabilities, 173
 - function, 163
 - integrable, 163
 - probability, 158
- i.o.*, 128
- increasing
 - sequence, 39
- independent
 - σ -algebra, 124
 - i.i.d. sequence, 128
 - random variables, 124
 - subsets, 124
- integrable
 - concurrent height, 210
 - floating front, 235
- interleaving relation, 46
- lattice, 39
 - complete, 39
 - semi upper, 39
- Law of large numbers, 208, 214
- likelihood, 130
 - future, 131
- limit sup, 128
- locally finite
 - event structure, 87
 - tree, 87
- marking, 44
 - initial, 44
 - reachable, 45
 - trace, 46
- Markov chain, 53
 - dual, 54
 - embedded, 168, 216
 - non homogeneous, 175
 - probabilistic transition system, 55
 - recurrent state, 191
- Markov property
 - concurrent, 164
 - concurrent and well-stopped, 167, 173
 - sequential, 160, 174
- martingale, 220
- max-synchronous, 246
- maximal
 - configuration, 42
 - element, 39
- minimal
 - element, 39
 - non void stopping prefix, 98
- morphism
 - foldings, 51
 - labelled occurrence net, 51
 - partial orders, 39
- multiplicative probability, 177
- multiset, 44
- net, 43
 - d -homogeneous, 166
 - compact, 154
 - confusion-free, 178
 - distributed Markov, 173
 - integrable concurrent height, 210
 - marked, Petri, 44
 - node, 44
 - occurrence, *see* occurrence net
 - probabilistic, 154
 - recurrent, 192
 - safe, 45
 - sequential, 56
- node, 44
- non decreasing
 - sequence, 39
- non increasing
 - sequence, 39
- observation

- delay, 234
- probability, 248
- sequential, 234
- occurrence net, 47
 - labelled, 50
 - morphism, 51
 - unfolding, 51
- open subset, 48
- operational
 - σ -algebra, 76
 - topology, 72
- operator
 - local renewal, 197
 - successive, 197
 - renewal, 188
 - successive, 189
 - shift
 - concurrent, 162
 - sequential, 160
 - stopping, 159
 - for sequential nets, 161
 - from an intrinsic prefix, 159
 - regular sequence, 208
 - universal, 190
 - well-stopped, 167
- order
 - partial, 39
- place, 43
- playing
 - sequence, 45
 - transition, 45
- positive
 - net, 154
- prefix
 - event structure, 41
 - in occurrence net, 49
 - intrinsic, 76
 - stopping, *see* stopping prefix
- prime
 - algebraic
 - coherent partial order, 43
 - complete lattice, 42
 - event structure, 38
- probability
 - \star -regular, 131
 - branching, *see* branching probability
 - conditional, 124
 - consistent w.r.t. a family of branching probabilities, 141
 - distributed, 137, 141
 - homogeneous, 158
 - image, 64
 - law, 64
 - observation, 248
 - product, 125
 - space, 64
 - finite, 65
 - starting, initial, 53
 - trace of, 134
- process, 42, *see* configuration
- product
 - initial stop. prefix, 105
 - nets, 245
 - probability, 125
 - topological, 66
- projection
 - global π_P , 69
 - relative $\pi_{P,P'}$, 69
- projective limit, 67
- projective system, 66
 - complete, of configurations, 70
 - limit of, 67
 - of probabilities, 67
 - operational, of maximal configuration, 73
 - topological, measurable, 67
- projective topology, 67
 - on configurations, 70
- puzzle, 114
 - concatenation, 117
- recurrent
 - net, 192
 - positive recurrent branching cell, 239
 - state of a Markov chain, 191

- renewal
 - global, 188
 - successive, 189
 - local, 197
 - successive, 197
- separable space, 63
- shadow, 90
 - finite, 72
 - thick, 133
- stopping
 - operator, *see* operator
 - time, 160
- stopping prefix, 86
 - $B(e)$, $B(A)$, 87
 - confusion-free, 180
 - initial, 98
 - decomposition, 104
 - lattices \mathcal{S} , \mathcal{S}_0 , 86
 - non void \mathcal{S}^* , 96
 - tree, 118
- sub-structure, 41
- time
 - concurrent unit, 205
 - stopping, 160
- trace
 - equivalence, 46
 - of a probability, 134
 - partial order, 47
- transition, 43
 - matrix, 53
 - playing, 45
- transition system, 55
 - covering, 57
 - dual, 55
 - probabilistic, 55
 - sequential net, 56
- tree
 - graph, 57
 - of events, 57
- unfolding, 51
 - concurrent width, 52
 - sequential net, 58
- well founded
 - partial order, 39
- well-stopped
 - configuration, 102
 - in a stop. prefix, 103
 - lattice of, 113
 - cut, 166
 - marking, 166
 - sequence, 102
- Zorn's lemma, 39