



HAL
open science

Verification of Stochastic Timed Automata

Pierre Carlier

► **To cite this version:**

Pierre Carlier. Verification of Stochastic Timed Automata. Modeling and Simulation. Université Paris Saclay (COMUE); Université de Mons, 2017. English. NNT : 2017SACLN058 . tel-01696130

HAL Id: tel-01696130

<https://theses.hal.science/tel-01696130>

Submitted on 30 Jan 2018

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.

NNT: 2017SACLN058

Verification of Stochastic Timed Automata

Thèse de doctorat de l'Université Paris-Saclay
préparée à l'ENS Paris-Saclay

École doctorale n°580
Sciences et Technologies de l'Information et de la Communication
Spécialité de doctorat: Informatique

Thèse présentée et soutenue à Mons (Belgique), le 8 décembre 2017, par

Pierre Carlier

Composition du Jury :

Véronique Bruyère Professeur, Université de Mons	Présidente
Parosh Aziz Abdulla Professeur, Uppsala Universitet	Rapporteur
Christel Baier Professeur, Technische Universität Dresden	Rapporteuse
Claudine Picaronny Maître de Conférences, ENS Paris-Saclay	Examinatrice
Mickaël Randour Chercheur Qualifié F.R.S – FNRS, Université de Mons	Examineur
Patricia Bouyer Directrice de recherche, ENS Paris-Saclay	Directrice de thèse
Thomas Brihaye Professeur, Université de Mons	Co-Directeur de thèse

Labaratoire Spécification et Vérification
Ecole Normale Supérieure Paris-Saclay, UMR 8643 du CNRS
61 avenue du Président Wilson, 94235 Cachan Cedex, France

Thèse de Doctorat issue d'une cotutelle entre
l'Université de Paris-Saclay préparée à
l'Ecole Normale Supérieure de Cachan (ENS Paris-Saclay)
et l'Université de Mons

université
PARIS-SACLAY

LSU

école
normale
supérieure
paris-saclay

UMONS
Université de Mons

Faculté
des Sciences

Verification of Stochastic Timed Automata

présentée par
Pierre Carlier

pour l'obtention du grade de Docteur en Sciences.

Thèse soutenue le 8 décembre 2017.

Jury

M. PAROSH AZIZ ABDULLA (Rapporteur)
Professeur, Uppsala Universitet, Suède
Mme. CHRISTEL BAIER (Rapporteuse)
Professeur, Technische Universität Dresden, Allemagne
Mme. PATRICIA BOUYER (Co-Directrice)
Directrice de recherche, ENS Paris-Saclay, France
M. THOMAS BRIHAYE (Co-Directeur)
Professeur, Université de Mons, Belgique
Mme. VÉRONIQUE BRUYÈRE (Présidente)
Professeur, Université de Mons, Belgique
Mme. CLAUDINE PICARONNY (Examinatrice)
Maître de conférences, ENS Paris-Saclay, France
M. MICKAËL RANDOUR (Examinateur)
Chercheur Qualifié F.R.S. - FNRS, Université de Mons, Belgique

Acknowledgements

Tout d'abord, je souhaite remercier mes deux co-directeurs de thèse Patricia Bouyer et Thomas Brihaye. Leurs nombreux conseils et leur aide furent précieux dans l'élaboration de cette thèse. Je suis reconnaissant envers Thomas pour m'avoir ouvert au monde de la recherche au travers de mon mémoire et envers Patricia qui m'a fait confiance en m'acceptant comme doctorant et qui m'a donc permis de travailler dans un nouvel environnement. Je la remercie également particulièrement pour son accueil à chacune de mes visites à Cachan. Je leur suis aussi très redevable pour le temps qu'ils m'ont consacré.

Mes prochains remerciements vont à tous les autres membres de mon jury, pour avoir accepté d'en faire partie et pour avoir pris le temps de lire ma thèse: pour commencer, mes rapporteurs Parosh A. Abdulla et Christel Baier pour leurs retours, mais aussi Véronique Bruyère, Claudine Picaronny et Mickaël Randour.

Je tiens maintenant à remercier mes co-auteurs: Nathalie Bertrand et Quentin Menet. Travailler avec eux fut très agréable, leur aide et expérience furent très utiles.

Je veux aussi remercier spécialement Virginie Guenard, pour son aide dans les formalités administratives à Cachan, et pour sa patience malgré mes nombreux retards...

Je remercie également le projet ERC EQualIS pour avoir permis mon financement.

Ensuite, je remercie tous mes professeurs de l'UMONS pour avoir confirmé mon goût pour les maths et m'avoir donné envie d'aller plus loin. Je remercie aussi particulièrement Mme Dubrulle pour m'avoir donné le goût des maths en secondaire.

Mes prochaines pensées vont aux nombreux collègues (ex-)doctorants qui ont croisé ma route à Mons et à Cachan. A Mons, je commence par remercier Quentin H., compagnon de route depuis la Bac1, il n'y pas la place pour rappeler tous les

bons moments de ces 8 dernières années. Je remercie également particulièrement Marion, collègue de bureau de ma dernière année, pour l'ambiance toujours excellente (et rarement calme) qui régnait au bureau. Je remercie aussi Aline Goe. qui a contribué à cette ambiance lors de ces 3 derniers mois. Tous les rires partagés dans ce bureau ont indéniablement grandement servi à l'élaboration de cette thèse. Et je n'oublie évidemment pas les nombreux autres occupants du 2D11!

Un grand merci à tous les autres (ex-)doctorants des départements de math et d'info: Céline, Maximilien, Adrien, Aline Gou., David H., Gauvain, Horacio, Jérémy, Monia, Nathanaël, Pierre H., Quentin B., Quentin L, merci pour les discussions (toujours de haut niveau!) sur les temps de midi, pour les pauses "thé", les soirées jeux... J'ai également une pensée pour les anciens, notamment Dany, Fabien (et Lucie), Gwendolyn, Quentin M. pour les nombreuses parties de belote le midi, mais aussi David S., Mathieu et Noémie. Je n'oublie pas non plus, mes collègues de Cachan qui m'ont accueilli à chacune de mes visites! Je pense à Daniel, Jérémy, Marie, Nathan, Patrick, Samy, Simon,... Je remercie particulièrement Engel, collègue de bureau d'un temps, pour les nombreuses discussions cinéma/littérature/mangas. Je termine par remercier Mickaël pour les temps passés en collocation à Cachan lors de mes visites la première année. Je n'oublie également pas tous les bons moments passés en conférences avec bon nombre d'entre vous!

Je dis merci aussi à tous mes amis, ils ont tous contribué d'une façon ou d'une autre à l'aboutissement de ce travail. Merci à Charline S. pour les rires et les encouragements, merci à Alexandre et à Damien, mes deux autres compagnons de route de Bac1, pour les bons souvenirs que je garde de mes années d'étude (et pour les futurs souvenirs!), merci à Johann pour les (trop) nombreux verres et merci à Chahrazade, Charline C., Julien, Naomi, Pauline, ma soeur Margot et mon frère Thomas, je vous remercie tous pour tous les excellents moments passés ensemble!

Je finis par remercier ma famille: ma soeur, mon frère, mes parents et mes grands-parents, pour leur soutien depuis toujours.

Abstract

Verification is now a well-known branch in computer science. It is crucial when dealing with computer programs in automatic systems: we want to check if a given system is correct and satisfies some specifications that should be met. One way to analyse those systems is to model them mathematically. The question is then: can we check if the model satisfies the required specifications? This is called the *model-checking problem*.

Several models have been studied in the literature. We have an interest for models that can mix both timing and randomized aspects. In this thesis we thus study the *stochastic timed automaton* model (STA). The contributions of this document are twofold.

First, we study the *qualitative* and *quantitative* model-checking problems of STA. STA are, in particular, general probabilistic systems and with such model, one is thus interested in questions like “Is a property satisfied, within a given model, with probability 1?” (*qualitative*) or “Can we compute an approximation of the probability that the model satisfies a given property?” (*quantitative*).

We study those questions for general stochastic systems using, amongst other, the notion of *decisiveness* used in infinite Markov chains in order to get strong qualitative and quantitative results, and that we extend here in or more general context. We prove several results for the qualitative and quantitative model-checking problems of those probabilistic systems, some of them being extensions of previous work on Markov chains, others being new, and we show how it can be applied to subclasses of STA.

Then we study the *compositional verification* in STA. In general, a system is the result of several smaller systems working together. Compositional verification allows then one to reduce the analysis of a big system to the analyses of the smaller systems which compose it. It is then crucial to have a good compositional framework in mathematical models, and this lacks in STA.

In this thesis, we define an operator of composition for STA. We first make the assumption that the STA composed run completely independently from each other, *i.e.* they do not communicate between them. We prove that our definition satisfies indeed this independence assumption. Such an operator of composition is not very interesting as in general, systems do communicate. But it is a necessary first step. We then introduce the new model of *interactive STA* (ISTA) that will allow for interactions between the systems. We define an operator of composition in ISTA that will make synchronisations possible between the systems and that is built on the previous composition in STA.

We end this thesis with the identification of a subclass of ISTA in which all the qualitative and quantitative results provided in this thesis can be applied, and which thus comes with the nice compositional framework defined in the model.

Contents

1	Introduction and Motivations	1
1.1	Model-checking problem	1
1.1.1	Background models	2
1.1.2	Stochastic timed automata	4
1.2	Compositional verification	5
1.3	Contributions	6
1.3.1	Qualitative and quantitative analysis of STSs	6
1.3.2	Parallel composition in STA	7
1.4	Other related works	8
1.5	Plan of the thesis	12
2	Background	15
2.1	Timed automata	15
2.1.1	Region graph	24
2.1.2	Time-converging aspects	27
2.1.3	Composition of timed automata	30
2.2	Denumerable Markov chains	34
2.2.1	Attractors	38
2.2.2	Decisiveness	39
2.3	Continuous-time Markov chains	43
2.4	Composition and interactive Markov chains	47
2.4.1	Composition of general transitions systems and application to DMCs and CTMCs	47
2.4.2	Interactive Markov chains	49

3	Stochastic Timed Automata	55
3.1	Definition and illustration of the model	55
3.2	Thick region graph and Markov chain	65
3.3	Fairness and classes of STA	69
I	Qualitative and Quantitative Analysis of Stochastic Transition Systems and Application to Stochastic Timed Automata	73
4	Stochastic Transition Systems	75
4.1	Definition and illustration of the model	76
4.1.1	Formulas for STSs	84
4.1.2	Labelled STSs and their properties	85
4.1.3	Qualitative and quantitative model-checking problem	90
4.2	Properties of STSs	91
4.2.1	Several decisiveness notions	91
4.2.2	Attractors	96
4.2.3	Fairness	98
4.2.4	Relationships between the various properties	99
4.3	Concluding remarks	103
5	Abstraction Between STSs	105
5.1	Abstraction	106
5.1.1	Properties of abstractions	107
5.1.2	Soundness and completeness of abstractions	110
5.2	Transfer of properties through abstractions	112
5.2.1	The case of sound abstractions	112
5.2.2	Trickier transfers of properties	114
5.3	Conditions for completeness and soundness	123
6	Qualitative and Quantitative Analysis	127
6.1	Qualitative analysis	128
6.1.1	Reachability and repeated reachability properties	129
6.1.2	Properties given by a DMA in DMCs	134
6.1.3	Properties given by a DMA in general STSs via denumerable abstractions	139
6.2	Quantitative analysis	148
6.2.1	Quantitative reachability analysis	149
6.2.2	Quantitative repeated reachability analysis	150
6.2.3	Properties given by a DMA in DMCs	151

6.2.4	Properties given by a DMA in general STSs via denumerable abstractions	152
6.3	Summary of the results on STSs	152
7	Application to STA	157
7.1	From STA to STS	158
7.1.1	The thick region graph abstraction	159
7.2	Reactive STA	161
7.3	Single-clock STA	164
8	Conclusion and Future Work	169
II	Composition of Stochastic Timed Automata	173
9	Interleaving Parallel Composition in STA	175
9.1	Definition of the parallel composition	176
9.2	Properties of the parallel composition	187
9.3	Bisimulation and congruence	209
9.3.1	Bisimulation	210
9.3.2	Congruence	214
10	Interactive Stochastic Timed Automata and Handshaking Composition	221
10.1	Syntax of ISTA	222
10.2	Semantics of ISTA	226
10.2.1	Parallel composition and hiding operator	226
10.2.2	Semantics through STA	242
10.3	Decidability of ISTA	243
11	Conclusion and Future Work	255
	Bibliography	259
	Résumé en français	267

List of Figures

2.1	A timed automaton \mathcal{A}	17
2.2	A timed automaton modelling a railroad crossing, $\mathcal{A}_{\text{railroad}}$	21
2.3	A timed automaton modelling a mouse	23
2.4	Partition on \mathbb{R}_+^2 induced by $\approx_{\mathcal{A}}$ and zoom on $[0, 1]^2$	25
2.5	The run ρ and the regions r_0, \dots, r_5	26
2.6	A two-clock timed automaton \mathcal{A}_{cvg} with a time-convergence phenomena	29
2.7	The composition of two simple timed automata \mathcal{A}_1 and \mathcal{A}_2	31
2.8	Timed automata $\mathcal{A}_{\text{train}}$ modelling the train (on the left), and $\mathcal{A}_{\text{controller}}$ modelling the controller (on the right).	33
2.9	A finite Markov chain \mathcal{M}_1	35
2.10	Random walk over \mathbb{N}	36
2.11	A queuing system with a maximum of n tasks.	45
2.12	Two simple IMCs \mathcal{M}_1 (on the left) and \mathcal{M}_2 (on the right).	50
2.13	IMC $\mathcal{M}_1 \parallel_1 \mathcal{M}_2$	52
3.1	The IPv4 Zeroconf STA for $N = 3$	62
3.2	A STA modelling a G/G/1/k-queue.	64
3.3	The first steps of $R(\mathcal{A})$	66
4.1	A Muller automaton M and the product $\mathcal{T}_2 \times M$	88
4.2	A DMC \mathcal{T}_3 that is not strongly fair.	99
5.1	Scheme for the proof of Proposition 5.2.6.	116
6.1	Left, \mathcal{T}_1 a random walk over \mathbb{N} and right, its sound finite abstraction \mathcal{T}_2	141

9.1	The product of two STA modelling the IPv4 Zeroconf	184
9.2	$\mathcal{A}_1 \notin \text{CSTA}$	186
9.3	\mathcal{A}_2 is Zeno	192
9.4	A simple example for bisimulation.	213
9.5	\mathcal{B} is bisimilar to \mathcal{A}	215
10.1	An ISTA $\mathcal{A}_{\text{cool}}$ describing a cooling system	223
10.2	A slight variant of the cooling system; ISTA $\mathcal{A}_{\text{cool}_i}$	229
10.3	The handshaking composition $\mathcal{A}_{\text{cool}_1} \parallel_{\{\text{burn}\}} \mathcal{A}_{\text{cool}_2}$	230
10.4	An ISTA $\mathcal{A}_{\text{worker}}$ representing a worker in a power plant, fixing cooling systems.	231
10.5	The handshaking composition $(\mathcal{A}_{\text{cool}_1} \parallel_{\{\text{burn}\}} \mathcal{A}_{\text{cool}_2}) \parallel_{\Gamma'} \mathcal{A}_{\text{worker}}$.	232
10.6	Bisimulation in ISTA	237

Introduction and Motivations

1.1. Model-checking problem

It is a known fact that real-life systems are now often controlled by computer programs. When dealing with planes or nuclear plants (for instance) those systems can be critical. It is thus important that we can verify if the computer programs governing these systems are correct: we do not want any bug to occur in such systems since consequences could be disastrous.

One way to analyse those programs is made through mathematical models. Those models give us tools allowing us to decide if a system is correct. The real-life system is described via some mathematical model that we call Mod and the purpose is then to check that Mod satisfies some desired property (or does not satisfy some undesired property) φ for the real-life system. The question whether Mod satisfies φ (written $\text{Mod} \models \varphi$) is called the *model-checking problem*. We are then interested in algorithms solving this problem.

The model-checking problem finds its roots in [CE81] and [QS82], but one can go to [BK08] for an advanced picture of the question. Several mathematical models for Mod have been proposed like finite automata [RS59], Petri nets [Pet81], Markov chains [Var85] or the one that will be of interest to us, timed automata [AD94]. The properties that we want our models to meet or not, *i.e.* φ , have also to be defined in this mathematical environment. A large amount of different logics have been proposed, and our interest will go to temporal logics or more precisely, to the logic LTL [Pnu77]. Each model has its own motivation and its own power of expressiveness, leading to various decidability results depending on the class of property we want to check. In the sequel, we

will be interested in probabilistic and timed systems.

1.1.1 Background models

Timed automata. When dealing with real-life systems, it is important to allow for timing constraints: in computer systems, it is often the case that an event has to occur within some time interval. For instance a gate at railroad crossing has to be lowered fast enough so that the train can safely cross the road. Another example is a traffic light in which time is obviously important as it is set to switch light at each fixed time. Roughly speaking, time constraints need to allow requirements like: a given action can be performed only in a given time interval

Timed automata (see [AD90] and [AD94]) answer to this purpose. They can be seen as finite automata enriched with a set of real-valued variables called clocks and with constraints over the clocks (called *guards*) on the transitions. As time elapses and as actions are taken, clock values change. This allows us to express timing restrictions. The model-checking problem in timed automata can be expressed as follows: given a timed automaton \mathcal{A} and a property φ , we would like to know whether \mathcal{A} satisfies φ . The timed aspect allows to express richer properties with bounds (for instance "does formula φ holds true in the interval $[1, 2]$?").

The interest for this model comes from a whole class of decidability results. It should be noted that the semantics of a timed automaton comes with an infinite set of states that is not denumerable, leading to difficulties when tackling the model-checking problem. However each timed automaton importantly comes with a finite abstraction, known as the region graph, which is equivalent in some sense to the timed automaton. This allows for several decidability results. For instance in [AD94], the authors showed that the model-checking problem for *reachability properties* is decidable in timed automata. In [ACD93], the authors also showed the decidability of the model-checking problem of TCTL properties, another temporal logic that allows to express time bounded properties. Moreover, tools have been developed, like Uppaal (see [BDL⁺06] and [Upp]), in order to verify real-life systems modelled by timed automata and which has numerous applications.

Probabilistic systems. As a lot of uncertainties can occur in real-life systems, probabilistic models are often useful as well. We consider the model of Markov chains (see [Var85]). A finite Markov chain can be seen as a finite automaton in which each state is equipped with a distribution over the set of outgoing edges. If the set of states can be infinite and at most denumerable, then we

talk of denumerable Markov chains (DMCs for short). We will have a particular interest for the work of [ABM07]. In any case, a probabilistic model Mod gives rise to a probability measure Prob over the set of all possible behaviours. Given a measurable property φ , we are then interested in the value $\text{Prob}(\text{Mod} \models \varphi)$. This gives rise to two kinds of problems:

- the *qualitative* model-checking problem: does it hold true that $\text{Prob}(\text{Mod} \models \varphi) = 1$ or not (or $\text{Prob}(\text{Mod} \models \varphi) = 0$ or not)? This is also known as the *almost-sure* model-checking problem (or the *0-almost-sure* model-checking problem);
- the *quantitative* model-checking problem: can we compute or approximate the value of $\text{Prob}(\text{Mod} \models \varphi)$?

If the probabilistic model is a finite Markov chain, considering LTL properties, those problems are easily solved: the qualitative problem is reduced to structural properties of the underlying graph (see [Var85]), while the exact value of $\text{Prob}(\text{Mod} \models \varphi)$ can be computed for the quantitative problem (see [CY88]).

In [ABM07], the authors are interested in DMCs. They define the notion of *decisiveness*. Roughly speaking, a Markov chain is *decisive w.r.t.* a set of states B if it reaches B or a state from which B can never be reached with probability 1. This notion allows one to transfer good properties from finite Markov chains to denumerable ones. The authors show the qualitative and quantitative problems for reachability and repeated reachability properties to be decidable for decisive Markov chains.

Finally still in this setting of probabilistic systems, we briefly mention the model of continuous-time Markov chains (CTMCs for short). This model is a first step for a mix between probabilistic and timed systems. It has a finite set of states like in finite Markov chains, but time progresses continuously along the behaviour. It is assumed that time delays between events are chosen randomly according to exponential distributions. Like in timed automata, this allows to express richer properties with interval bounds. In [BHHK03], the authors proved CTMCs to be decidable for the temporal logic CSL (which expresses those richer properties with intervals).

Those models come with several tools, including the well-known model-checker Prism [Pri]. It allows to construct and model systems and to verify them for a wide range of temporal logics (including PCTL which adds probabilistic aspects, and CSL for CTMCs) and has numerous applications [KNP11].

1.1.2 Stochastic timed automata

As said earlier we are interested in systems which have both probabilistic and timed features. If CTMCs enter in this setting, they are not entirely satisfying as they do not allow for timing constraints on transitions between states and as time between events only delays following exponential distributions. We are interested here in the recent model of stochastic timed automata (STA for short) as introduced in [BBB⁺07]. We give a quick idea of the model. STA are timed automata equipped with distributions over the delays and the edges and this, in each state of the automaton. In particular, distributions over the delays are continuous and can, roughly speaking, be of any form (respecting some reasonable constraints as we will see later) like for instance exponential distributions (as in CTMCs) or also uniform distributions for bounded intervals.

As it will be defined in the sequel, the STA model gives rise to a probability distribution over all possible behaviours and it thus enters the setting of probabilistic systems. Therefore the qualitative and quantitative model-checking problems have also a meaning in STA and this will be one of the focus of this thesis. Given a STA \mathcal{A} and a measurable property φ can we

- decide if $\text{Prob}(\mathcal{A} \models \varphi) = 1$ or $\text{Prob}(\mathcal{A} \models \varphi) = 0$?
- approximate the value $\text{Prob}(\mathcal{A} \models \varphi)$?

Several decidability results are already known for the almost-sure model-checking problem. In [BBB⁺08], the authors showed that the qualitative model-checking problem is decidable for LTL properties in one-clock STA (under some technical but easily satisfied conditions) while in [BBJM12] the authors identified another class of STA, called *reactive* STA (roughly speaking STA with only exponential distributions), for the same decidability results again on LTL properties. This has been published with full proofs in [BBB⁺14]. As in timed automata, it relies on the finite abstraction of the region graph that can be here considered as a finite Markov chains and is equivalent in some sense to the STA. The identified notion for this equivalence to hold true was *fairness* (if an edge with non-null probability is enabled infinitely many times then it is chosen infinitely often with probability 1) and the authors showed that one-clock and reactive STA are *almost-surely* fair (*i.e.* fair with probability 1). Those proofs required *ad hoc* methods. We end this brief discussion on STA by noting that at this point, the only known quantitative model-checking problem decidability result for STA comes from [BBBM08] where the authors consider a restricted class of one-clock STA (basically one-clock STA with only exponential distributions) and deploy *ad hoc* methods in order to provide an approximation scheme for *e.g.* LTL properties.

1.2. Compositional verification

With the expansion of technology, real-life systems become bigger and bigger. The verification of $\text{Mod} \models \varphi$ can thus be very complicated if Mod describes a big system. This is where compositional verification has its importance: the objective is to divide Mod into n smaller systems $\text{Mod}_1, \dots, \text{Mod}_n$ and analyse those smaller systems separately in order to deduce properties on the big system Mod . This comes also from the fact that in general, a real-life system is the result of several smaller systems. It is then simpler to model separately the small systems and to define an operator of composition with adapted interactions between the different models. The product result after parallel composition must then describe the behaviour of the initial system we wanted to verify.

Given n systems $\text{Mod}_1, \dots, \text{Mod}_n$, we write $\text{Mod}_1 \parallel \dots \parallel \text{Mod}_n$ for the parallel composition of the n systems. Then compositional verification can be expressed as follows: given a property φ , can we build properties $\varphi_1, \dots, \varphi_n$ such that $\text{Mod}_i \models \varphi_i$ for each $i \in \{1, \dots, n\}$ implies $\text{Mod}_1 \parallel \dots \parallel \text{Mod}_n \models \varphi$? The need of a composition operator is thus crucial.

Compositionality has been first studied in the context of process algebras (see [HBR84] or more recently [DK05b] in a probabilistic and timed context). Since we are interested in (probabilistic and timed) transition systems, we will rather be interested in the approach of [HZ11]. It studies composition in the context of discrete- and continuous-time Markov chains. The authors propose three kinds of composition: *fully synchronised*, *interleaving* and *handshaking* compositions. Fully synchronised composition and interleaving composition are used respectively for discrete-time Markov chains and CTMCs. Fully synchronised composition only considers that the systems composed always interact between them while interleaving composition assumes that the systems run completely independently from each other. Hence the interest for handshaking composition which mixes both previous types.

The approach leads to interactive Markov chains (IMCs for short) [Her02], which are convenient for a compositional setting with handshaking. IMCs will be a source of inspiration for our work. They are studied as process algebras in [Her02] we will thus rather consider the work of [HK09] where the semantics is given as a probabilistic and non-determinist transition system. Briefly, an IMC can be seen as a CTMC in which we add non-determinism through non-probabilistic transitions. It comes with a very nice compositional framework for verification: for instance in [HK09], the authors provide techniques in order to approximate time-bounded reachability properties in this setting with non-determinism.

Bisimulation. For a nice compositional framework, an important notion comes with *bisimulation* (see for instance [LS91] and [LY93]). A bisimulation is a relation between the states of a model that identifies states with similar behaviours. It is extended to a relation between systems, identifying thus systems with similar behaviours. When dealing with composition, an important property that must be satisfied is the following: bisimulation is a congruence w.r.t. parallel composition. This means that you can replace in the product a component by another component that has a similar behaviour, *i.e.* by a bisimilar component. For instance with IMCs, in [Her02] several definitions of bisimulation are given and are shown to be a congruence w.r.t. parallel composition.

1.3. Contributions

The contributions of this thesis can be divided in two parts: firstly the qualitative and quantitative model-checking problem for general stochastic transition systems (STSs for short), *i.e.* with a continuous state-space, with an application to STA. Secondly, we are interested in parallel composition for STA which, at this point, lacks in the model and is crucial in order to simplify verification.

1.3.1 Qualitative and quantitative analysis of STSs

Inspired from [ABM07], we extend the notion of *decisiveness* in discrete-time Markov chains to more general probabilistic systems: STSs. Those represent systems with a continuous state-space and a Markov kernel or in other words, Markov chains with an infinite and non-denumerable set of states and thus with continuous probability distributions between states. We identify hypotheses that lead to decidability results and we show that STA can be seen as STSs allowing us to transfer the decidability results for STSs to STA. More precisely, we establish the following points.

- We define different notions for STSs (*decisiveness, fairness, finite attractor*) and we show the different links between those notions.
- We define a notion of abstraction for STSs (as in: the region graph is an abstraction of a timed automaton). We have a particular interest when this abstraction is a DMC. We identify conditions under which the abstraction is said *sound* and *complete*, that will allow to reduce the qualitative and quantitative analysis of the initial STS to its abstraction.
- We consider the qualitative and quantitative model-checking problems of LTL properties through the product of STSs with deterministic Muller au-

tomata. Using similar techniques as on lossy probabilistic channel systems (see [ABRS05] and [Ber06]), we show that when the abstraction is a Markov chain and is sound and complete, then

- ▷ the qualitative model-checking problem of LTL properties in STSs is reduced to the qualitative model-checking problem of LTL properties in the abstraction;
 - ▷ there is an approximation scheme for the quantitative model-checking problem of reachability properties given by some graph of the abstraction and that can be used for the quantitative model-checking of all LTL properties.
- We finally identify classes of STA in which the previous results can be applied, leading to new approximation schemes for STA!

These contributions were firstly described in [BBBC16] under a different formalism, but are the subject of [BBBC17] which is currently submitted for publication.

1.3.2 Parallel composition in STA

Now inspired from the approach of [HZ11], we are interested in the definition of an operator of composition for STA. We first consider the simple case where the composed automata run completely independently. This yields to an interleaving semantics. We define such a parallel composition and we identify a class of STA

- in which parallel composition is well-defined and internal, and
- such that parallel composition corresponds to the interleaving semantics for STA.

We then define a notion of bisimulation in STA and importantly show that bisimulation is a congruence w.r.t. parallel composition. These contributions are the subject of [BBCM16].

As explained in [HZ11], parallel composition has more interest when it comes with synchronisations. Therefore, inspired from the IMC model, we define the interactive stochastic timed automaton model (ISTA for short). Based on [HK09] we define a parallel composition with handshaking. We also identify a class of ISTA in which parallel composition is well-defined and internal, and we define a notion of bisimulation that is importantly a congruence w.r.t. parallel composition.

We end the thesis with a result that links both parts of the document. We identify a class of ISTA in which parallel composition is well-defined and internal, and in which all the previous results on the qualitative and quantitative model-checking problems can be applied.

These last contributions are new and are not yet published nor submitted.

1.4. Other related works

We would like to briefly mention some models and papers that somehow relate to our work (whether for a similar model to the STA model, or for interesting similar results in other settings). We do not have the pretension that the following list is exhaustive.

Probabilistic timed automata. We briefly mention the probabilistic timed automaton model [KNSS02] (PTA for short) since it is another probabilistic extension of timed automata. PTA extend timed automata only with distributions over the edges (distributions over the delays are not present in this model, to the contrary of our STA model). It has been widely studied in the literature with also numerous case studies. A list of related papers can be find on [Pri]. It comes with nice decidability results, *e.g.* in [KNSS02], the authors showed the decidability of PCTL properties allowing thus the approximate probabilities of a rich class of properties. Note also that the model-checker Prism [KNP11] can verify PTA. Finally, observe that PTA cannot be viewed as STSs due to the presence of non-determinism via time passage (while STSs are purely stochastic).

A continuous variant of PTA has also been studied in [KNSS00]. Like the STA model, the delays are randomized with continuous probabilities. Few results are known on the model, however in [KNSS00] the authors provide a method for the model-checking of formulas expressed in the logic PTCTL.

Generalised semi Markov processes. Generalised semi-Markov processes (GSMPs for short) are probabilistic systems with continous state-space [Gly89]. The model very much compares to our STA model. We briefly describe the model as in [BKKR11]. A GSMP is given with a finite set of states, a finite set of transitions and a finite set of events. Each event is associated with a time interval corresponding to its firing time and it is equipped with a density function on this interval. Each state and each transition is equipped with a set of events. Roughly speaking when entering a state, the events of the state are active and a random value is chosen for each event accordingly to its density function. Then time elapses just like in timed automata, and as soon as the values of all events

of a transition are reached, the transition is chosen. Like in STA, the semantics of GSMP can be given as a STS. In [AB06], the authors provide an algorithm that can approach the probability that a GSMP satisfies some kind of “Until”-formulas happening before time T and within k discrete events. In [BA07], the authors provide another algorithm for a restricted class of GSMPs and this time for “Until”-formulas. Finally in [BKKŘ11], the authors identify a class of GSMPs allowing two technical lemmas on a finite abstraction (which is just like in timed automata, the region graph). In [BBBC17], we showed that these two technical lemmas imply soundness of the abstraction leading to the whole decidability and approximation results on STSs for this class of GSMPs.

Stochastic time Petri nets. We now have an interest for the recent paper [PHV16] that considers the model-checking of time-bounded “Until”-formulas for stochastic time Petri nets (STPNs for short). Petri nets [Pet81] can be seen as finite automata in which transitions can have several sources and several targets and in which each node has a marking which is a natural number. A state of a Petri net is a marking function. A transition t between two markings occurs as follows: first, it must be the case that the transition is enabled, *i.e.* for each source of t , its marking is at least 1, then you remove 1 from each source of t and finally, you add 1 to each target of t . A time Petri net equips each transition with an interval firing times. Finally, a STPN equips each transition with a continuous distribution over its interval firing times. Now a state is a marking function along with a tuple of positive numbers representing firing times for each transition that are enabled in the current marking. Here an enabled transition can occur from a given state only if it has the minimum firing time. Similarly as in the GSMP model, time elapses and as soon as the minimum firing time is reached, one enabled transition is chosen. Then accordingly to the new marking, some enabled transitions inherit their previous firing time while newly enabled transition get a new firing time according to the corresponding distribution.

In general, the semantics of STPNs are given as GSMPs. As we have seen it before, few results for the quantitative model-checking are provided for GSMPs at our knowledge, and those results are limited to restrictive class of GSMPs. In [PHV16], the authors consider the model-checking of time-bounded “Until” properties thanks to an abstraction of the system, the transient stochastic tree, and through the visit of *regeneration classes*. A regeneration class is a state in the abstraction in which every previous history is not required to be memorised *i.e.* roughly speaking classes of states in which all enabled transitions are either newly enabled or exponentially distributed. In [HPRV12], the authors provide conditions for the stochastic transient tree to be finite. In [PHV16], the authors

consider STPNs in which regeneration is encountered with probability 1 within a bounded number of discrete events. Under this condition, the authors provide algorithms that approximate the probability of time-bounded “Until” formulas from an initial regeneration point.

Stochastic hybrid systems. We consider the discrete-time version of stochastic hybrid systems (SHSs for short) [AKLP10] and [SA13]. In general, SHSs are defined with a dynamic over time that evolves accordingly to differential equations (see [Aba07], [FHH⁺11] and [HLS00]). Hybrid systems [Hen96] can be seen, like timed automata, as finite automata enriched of a finite set of real-valued variables. Whereas those variables are only needed to measure time in timed automata, in hybrid systems those variables allow to measure a great amount of things like ambient temperature. They can express dynamics much more complicated than time whose evolution is linear. In hybrid systems, the evolution of the variables is described by means of differential equations. This is the main difference with timed automata. Then we still have edges between a finite set of locations and those transitions allow to reset the different variables. SHSs are thus a stochastic extension of hybrid systems, with probabilities over the edges and the evolution of variables are defined through stochastic differential equations.

We are here interested in discrete-time SHSs (DTSHSs for short) that relies better on our STA model. DTSHSs are defined as STSs in [AKLP10] and [SA13]. In this model, the dynamics of variables through differential equations, are not present. A DTSHS is given with a finite set of locations, and each location is equipped with a continuous component (whose dimension can depend on the location). States are thus couples composed of the location and the current value of the continuous component. The transition kernel is quite elaborated as it is defined by the means of three probability distributions. Entering a new state depends first on a discrete distribution over the set of locations. Then it distinguishes the cases where the location has changed or not: if the location does not change, the continuous component evolves according to a first distribution; if the location changes, the continuous component is reset according to a second distribution.

In [AKLP10], the authors are interested in the evaluation of the probability of invariance properties on a finite horizon (the system always stays in a safe set of states within the N first steps) in DTSHSs. It is based on the generation of finite abstractions, which are finite Markov chains, that approximate the value of such probabilities. The construction is quite technical and is based on a decomposition of the safe set of states into a partition of sets of diameter at most

δ . The invariance property is then studied on those Markov chains for which algorithms exist, and the authors provide conditions under which the sequence of probabilities of the invariance property in the Markov chains converges towards the probability of the invariance property in the corresponding DTSHS, when δ converges towards 0. In [SA13], the authors are interested in similar constructions but for STSs and apply the results to DTSHSs.

Composition of STSs. We would like now to mention the following recent paper [GBK16] which defines a parallel composition operator for stochastic transition systems (as defined in [CSKN05]¹). This parallel composition makes no assumption on the probability distributions of the initial systems which differs from our contributions! In our work, when dealing with composition in (I)STA, we will assume that the (I)STA which are composed are stochastically independent (even though they could interact on transitions in ISTA), *i.e.* we assume that the distributions of a system are independent of the distributions of another system. In [GBK16], the authors define thus a parallel composition operator in a more general setting and based on couplings of probability measures.

Another STA model. Finally, we briefly discuss another model that is similar to our STA model and that is equipped with a compositional framework, but whose approach is based on process algebras. In [DK05a, DK05b], the authors are concerned with the stochastic process algebra \diamond , whose semantics is given as \diamond -stochastic timed automata (we write \diamond -STA). Our model very much compares to the latter as it is a mix between timed automata and GSMPs. We will briefly describe it. In such a system, when a clock variable is activated, it is sampled according to a predefined distribution (like events in GSMPs), and then it acts as a countdown timer: when time elapses, the clock variables decrease down to 0². Transitions can be fired once all clocks specified on the transition have reached value 0. First notice that both STA and \diamond -STA allow to express timing constraints to be satisfied by the system (which is not the case of CTMCs or IMCs). Then \diamond -STA comes with a compositional framework and several notions of bisimulations with nice congruence properties. It is interesting to mention as well that \diamond -STA allow for infinitely many states and clock variables, whereas STA do not (they have been defined on top of timed automata, with desirable decidability properties in mind). Then it is worth noting that \diamond -STA model is

¹This model relates very much to our STS model defined in Chapter 4: our model can be seen as the model of [CSKN05] in which there is a single label.

²Observe that it is not how we described GSMPs, however the more classical way to define the semantics of GSMPs is through this decrease of time.

at the basis of the modelling language Modest [BDHK06]. The semantics of this language is in fact given as a very general notion of stochastic timed automata (we call them Modest-STA) and which encompasses all the models we have mentioned (STA, \diamond -STA, GSMPs, CTMCs,...). STA in general can be viewed as a fragment of Modest-STA. Modest is a description language and hence does not come with algorithms. However it has a nice tool suite³ for verification of several probabilistic and timed models (including \diamond -STA but also PTA) [HH14] and [Mod].

1.5. Plan of the thesis

The thesis is organized as follows. Chapters 2 and 3 are here to identify the background notions needed. In Chapter 2 we introduce timed automata, denumerable Markov chains along with the work of [ABM07], CTMCs and IMCs. In Chapter 3 we introduce the notion of STA as well as the results of [BBB⁺14] that are relevant with this document. Then, the report is divided in two parts that can be read independently from each other (except for Section 10.3 that establishes a link between the two parts).

In Part I, we investigate the qualitative and quantitative model-checking problem of STS and apply the results to classes of STA. In Chapter 4 we introduce the notion of STS and define several useful notions. In Chapter 5, we define a notion of abstraction of STS and show how this may help to simplify the analysis of STS. In Chapter 6, we analyse STSs through their abstractions: we study the qualitative and quantitative model-checking problem of STS for LTL properties, and use abstractions in order to get simpler results. Finally in Chapter 7, we show that STA can be seen as STSs, and we identify two classes of STA for which the decidability results of Chapter 6 can be applied. This yields new approximation schemes for STA!

In Part II, we are interested in a notion of composition in STA. In Chapter 9, we define a parallel composition of STA under the assumption that the STA composed run completely independently. We show that it corresponds to the interleaving semantics of STA and importantly define a bisimulation that is a congruence w.r.t. composition. In Chapter 10 we define the new model of ISTA and define a handshaking composition operator for the model. We again define a bisimulation that is a congruence w.r.t. composition, and we identify conditions under which the semantics of an ISTA can be given as a STA. Finally, we identify a class of ISTA whose corresponding STA enter the framework of

³The language is too expressive for a single tool to analyse all Modest-STA.

Chapter 6 allowing us to infer all decidability results! Which leads to a class of ISTA that are decidable and in which there is a compositional framework.

CHAPTER 2

Background

In this chapter, we introduce several models of interest for our work and define several notions that we will need in the sequel.

In Section 2.1, we define the *timed automaton* model [AD90] and [AD94]. We then introduce the classical notion of *region graph* and we give some of the decidability results it allows to get. We also exhibit some converging aspects that will be problematic in some cases and we end the section with notions of interleaving and handshaking parallel composition in timed automata.

In Section 2.2, we define the notion of *denumerable Markov chain* (DMC for short). We then present the work of [ABM07] in which the notion of *decisiveness* is introduced, leading to results for the qualitative and quantitative model-checking problems of reachability and repeated reachability properties.

We then take some space in Section 2.3, to briefly introduce the model of continuous-times Markov chains (CTMCs for short), which is a first step for a system that mixes probability and timed aspects. We end the chapter with considering a compositional approach in DMCs and CTMCs following [HZ11], in Section 2.4. This leads to the notion of interactive Markov chain (IMC for short) [Her02] which is an extension of CTMCs with non-probabilistic transitions and with possibly non-determinism and that is suitable for a compositional framework.

2.1. Timed automata

In this section, we define and illustrate the notion of timed automaton ([AD90], [AD94]). We first introduce some notations.

Let $X = \{x_1, \dots, x_n\}$ be a finite set of real-valued variables called *clocks*. A *clock valuation* over X is a map $\nu : X \rightarrow \mathbb{R}_+$ where \mathbb{R}_+ is the set of nonnegative real numbers. We write \mathbb{R}_+^X for the set of clock valuations over X . If $\nu \in \mathbb{R}_+^X$, we write ν_i for $\nu(x_i)$ and we then denote ν by (ν_1, \dots, ν_n) . We can interpret ν_i as the value of clock x_i . If $t \in \mathbb{R}_+$, we write $\nu + t$ for the clock valuation defined by $(\nu_1 + t, \dots, \nu_n + t)$. If $Y \in 2^X$ (the power set of X), $[Y \leftarrow 0]\nu$ is the valuation that assigns to each clock x , 0 if $x \in Y$ and $\nu(x)$ otherwise. A *guard* over X is a finite conjunction of expressions of the form $x_i \sim c$ where $x_i \in X$, $c \in \mathbb{N}$ and $\sim \in \{<, \leq, =, \geq, >\}$. We denote by $\mathcal{G}(X)$ the set of guards over X . Given a clock valuation $\nu \in \mathbb{R}_+^X$ and a guard $g \in \mathcal{G}(X)$, we write $\nu \models g$ if ν satisfies g and we say that ν satisfies a guard of the form $x_i \sim c$ whenever $\nu_i \sim c$. We can now define the notion of timed automaton.

Definition 2.1.1. A *timed automaton* is a tuple $\mathcal{A} = (L, X, Act, E, Inv, AP, \mathcal{L})$ where


- (i) L is a finite set of locations,
- (ii) X is a finite set of clocks,
- (iii) Act is a finite set of actions,
- (iv) $E \subseteq L \times Act \times \mathcal{G}(X) \times 2^X \times L$ is a finite set of edges,
- (v) $Inv : L \rightarrow \mathcal{G}(X)$ is an invariant function,
- (vi) AP is a finite set of atomic propositions and $\mathcal{L} : L \rightarrow 2^{AP}$ is a labelling function.

We will describe later the semantics of a timed automaton. We first give a simple example in order to illustrate Definition 2.1.1.

Example 2.1.2. We consider the timed automaton depicted in Figure 2.1. In this example, the set of locations L is $\{l_0, l_1\}$; the set of clocks X is $\{x_1, x_2\}$; the set of actions Act is $\{a, b\}$; the set of edges (described by the arrows; the notation $x := 0$ means that clock x is reset to zero) is defined as follows: $E = \{e_1, e_2, e_3, e_4\}$ where

$$\begin{aligned} e_1 &= (l_0, a, 0 < x_1 \leq 1, \{x_1\}, l_0); & e_2 &= (l_0, b, (x_1 \leq 2) \wedge (x_2 \geq 2), \{x_2\}, l_1); \\ e_3 &= (l_1, a, x_2 \leq 2, \{x_2\}, l_1); & e_4 &= (l_1, b, x_2 = 1, \{x_1\}, l_0); \end{aligned}$$

the invariants are given by $Inv(l_0) = (x_1 \leq 2) \wedge (x_2 \leq 4)$ and $Inv(l_1) = x_2 \leq 2$; and the set of atomic propositions is $AP = \{p_1, p_2\}$ and the labelling function is defined by $\mathcal{L}(l_0) = \{p_1\}$ and $\mathcal{L}(l_1) = \{p_2\}$.

Figure 2.1: A timed automaton \mathcal{A}

Remark 2.1.3. Note that if a guard or an invariant is **true** (*i.e.* corresponding to $x \geq 0$ for some clock x), then it won't be written on the edge or on the location. Similarly, we will not write the label of a location if it is the empty set.

Note also that the labelling of the edges with the set of actions Act , the labelling of the locations over AP and the invariant function will not always be needed. In some cases we will thus omit them (maybe not the three of them at the same time). In those cases, it should be understood that

- all edges are labelled with the same action,
- all locations are labelled with the same atomic proposition, and
- all locations have the invariant **true**.

We will always make clear in which case we are.

Given an edge $e = (l, a, g, Y, l') \in E$, we define $\text{source}(e) = l$ for the location source of e and $\text{target}(e) = l'$ for the location target of e . We will sometimes write $e = l \xrightarrow{a, g, Y} l'$. Now, before explaining precisely how we interpret timed automata, let us define the notion of *transition system* which will be useful when defining the semantics of timed automata.

Definition 2.1.4. A *transition system* T is a tuple (Q, Γ, \rightarrow) where

- (i) Q is a set of states,
- (ii) Γ is an alphabet of actions, and
- (iii) $\rightarrow \subseteq Q \times \Gamma \times Q$ is a transition relation.

T is called *finite* if Q and Γ are finite.

The semantics of a timed automaton \mathcal{A} is given as a transition system. We first define the states of \mathcal{A} and the transitions between two states of \mathcal{A} with the aim of defining this transition system. Fix a timed automaton $\mathcal{A} = (L, X, E, \text{Inv})$ (here, actions and atomic propositions are not needed and we omit them in order to avoid heavy notations).

Definition 2.1.5. A *state* of \mathcal{A} is a pair $q = (l, \nu)$ where $l \in L$ and $\nu \in \mathbb{R}_+^X$ are such that $\nu \models \text{Inv}(l)$. We denote by $Q \subseteq L \times \mathbb{R}_+^X$ the set of all states in \mathcal{A} .

We then distinguish two types of transitions: *time-transitions* and *switch-transitions*.

Definition 2.1.6. Given $q = (l, \nu)$ and $q' = (l', \nu')$ two states of \mathcal{A} , there is a *time-transition* in \mathcal{A} between q and q' if $l = l'$ and there exists $t \geq 0$ such that $\nu' = \nu + t$. We denote this transition by $q \xrightarrow{t} q'$ and we write $q' = q + t$.

Definition 2.1.7. Given $q = (l, \nu)$ and $q' = (l', \nu')$ two states of \mathcal{A} , there is a *switch-transition* in \mathcal{A} between q and q' if there exists $g \in \mathcal{G}(X)$ and there exists $Y \in 2^X$ such that $e = (l, g, Y, l') \in E$, $\nu \models g$ and $\nu' = [Y \leftarrow 0]\nu$. We denote this switch-transition by $q \xrightarrow{e} q'$.

Finally in the sequel, we will be interested in a mix of those two types of transitions.

Definition 2.1.8. Given $q = (l, \nu)$ and $q' = (l', \nu')$ two states of \mathcal{A} , there is a *mixed-transition* in \mathcal{A} between q and q' if there exist $t \in \mathbb{R}_+$ and $e \in E$ such that $q \xrightarrow{t} q + t \xrightarrow{e} q'$. We denote this mixed-transition by $q \xrightarrow{t,e} q'$.

Remark 2.1.9. A mixed-transition corresponds thus to the succession of two transitions: a time-transition followed by a switch-transition. Then, given states $q = (l, \nu)$ and $q' = (l', \nu')$, given $t \in \mathbb{R}_+$ and $e \in E$, we have a mixed-transition $q \xrightarrow{t,e} q'$ whenever $\text{source}(e) = l$, $\text{target}(e) = l'$, $\nu + t \models g$ and $\nu' = [Y \leftarrow 0](\nu + t)$, where $e = (l, g, Y, l')$.

When it is not relevant, we will sometimes omit the labels on the mixed-transitions, *i.e.* we will sometimes write $q \rightarrow q'$ instead of $q \xrightarrow{t,e} q'$. We illustrate the previous notions on Example 2.1.2.

Example 2.1.10. If we consider the timed automaton \mathcal{A} of Figure 2.1, we can see that

- the set of states

$$\begin{aligned} Q &= \{(l, \nu) \in L \times \mathbb{R}_+^2 \mid (l = l_0 \wedge (\nu_1 \leq 2 \wedge \nu_2 \leq 4)) \vee (l = l_1 \wedge \nu_2 \leq 2)\} \\ &= \{l_0\} \times [0, 2]^2 \cup \{l_1\} \times (\mathbb{R}_+ \times [0, 2]); \end{aligned}$$

- $(l_0, (0, 0)) \xrightarrow{0.5} (l_0, (0.5, 0.5))$ is an instance of a time-transition in \mathcal{A} ;
- $(l_0, (0.5, 0.5)) \xrightarrow{e_1} (l_0, (0, 0.5))$ is an instance of a switch-transition in \mathcal{A} , where e_1 is defined in Example 2.1.2, since $(0.5, 0.5)$ satisfies the guard of e_1 , $0 < x_1 \leq 1$, and e_1 resets the clock x_1 to zero;
- $(l_0, (0, 0)) \xrightarrow{0.5, e_1} (l_0, (0, 0.5))$ is an instance of a mixed-transition in \mathcal{A} , corresponding to the succession of the two previous transitions.

We can now define the transition system associated with a timed automaton.

Definition 2.1.11. Let $\mathcal{A} = (L, X, E, \text{Inv})$ be a timed automaton. The *transition system associated with \mathcal{A}* is given by $T_{\mathcal{A}} = (Q, \mathbb{R}_+ \times E, \rightarrow)$ where Q is the set of states as defined in Definition 2.1.5, and the transition relation \rightarrow corresponds to the set of mixed-transitions in \mathcal{A} as defined in Definition 2.1.8.

We now define the notion of *run* in a timed automaton. Let $\mathcal{A} = (L, X, E, \text{Inv})$ denote a timed automaton and $T_{\mathcal{A}} = (Q, \mathbb{R}_+ \times E, \rightarrow)$ denote the transition system associated with \mathcal{A} .

Definition 2.1.12. Let k be a positive integer. A *finite run* (also called a *finite path*) of \mathcal{A} is a finite sequence of states $(q_i)_{i \in \{0, \dots, k\}} \subseteq Q$ such that for all $i \in \{0, \dots, k-1\}$ there is $t_i \in \mathbb{R}_+$ and there is an edge $e_i \in E$ such that $q_i \xrightarrow{t_i, e_i} q_{i+1}$ is a mixed-transition in \mathcal{A} . We denote it as follows:

$$\rho = q_0 \xrightarrow{t_0, e_0} q_1 \xrightarrow{t_1, e_1} \dots \xrightarrow{t_{k-1}, e_{k-1}} q_k.$$

Definition 2.1.13. An *infinite run* (or *infinite path*) of \mathcal{A} is an infinite sequence of states $(q_i)_{i \in \mathbb{N}}$ such that for all $i \in \mathbb{N}$ there is a mixed-transition $q_i \rightarrow q_{i+1}$ between q_i and q_{i+1} in \mathcal{A} . We denote it as follows:

$$\rho = q_0 \xrightarrow{t_0, e_0} q_1 \xrightarrow{t_1, e_1} q_2 \xrightarrow{t_2, e_2} \dots$$

where for each $i \in \mathbb{N}$, $q_i \xrightarrow{t_i, e_i} q_{i+1}$ denotes a mixed-transition.

We now illustrate these notions on Example 2.1.2.

Example 2.1.14. Let \mathcal{A} be the timed automaton of Figure 2.1. Let us consider the runs ρ_1 and ρ_2 depicted below.

$$\begin{aligned} \rho_1 &= (l_0, (0, 0)) \xrightarrow{0.5, e_1} (l_0, (0, 0.5)) \xrightarrow{1.6, e_2} (l_1, (1.6, 0)) \xrightarrow{0, e_3} (l_1, (1.6, 0)); \\ \rho_2 &= (l_0, (0, 0)) \xrightarrow{0.8, e_1} (l_0, (0, 0.8)) \xrightarrow{2, e_2} (l_1, (2, 0)) \xrightarrow{0.5, e_3} (l_1, (2.5, 0)) \\ &\quad \xrightarrow{0.5, e_3} (l_1, (3, 0)) \xrightarrow{0.5, e_3} (l_1, (3.5, 0)) \xrightarrow{0.5, e_3} \dots \end{aligned}$$

It holds that ρ_1 is a finite run of \mathcal{A} and ρ_2 is an infinite run of \mathcal{A} .

Remark 2.1.15. In the sequel, we are only interested with infinite behaviours, and thus with infinite runs. Infinite runs require that it is always-possible in the future to perform a mixed-transition, and thus a time- and a switch-transition. Observe however that it could be possible to reach a state from which there are no enabled edges in the future and thus from which it is not possible to perform any mixed-transition. Consider the timed automaton \mathcal{A} of Example 2.1.2, but this time assume that $\text{Inv}(l_0) = (x_1 \leq 3) \wedge (x_2 \leq 4)$. Then from state $(l_0, (2.5, 3))$ for instance, no edges are enabled in the future (the value of clock x_1 has overpassed the bounds 1 and 2 of edges e_1 and e_2). Such states are called *blocking states*. We thus have to prohibit such states.

We fix a timed automaton \mathcal{A} and $T_{\mathcal{A}}$ its transition system. We also consider a state $q = (l, \nu) \in Q$ and an edge $e \in E$, and we define the following set of delays $I(q, e) = \{t \in \mathbb{R}_+ \mid \nu + t \models \text{Inv}(l) \text{ and } \exists q' \in Q \text{ s.t. } q \xrightarrow{t, e} q'\}$ corresponding to the times after which, starting from q , edge e is enabled. Observe that if $\text{source}(e) \neq l$ then $I(q, e) = \emptyset$. We then define $I(q) = \bigcup_{e \in E} I(q, e)$, *i.e.* the set of delays after which, starting from q , an edge is enabled. We can now define the notion of *blocking state*.

Definition 2.1.16. State q is a *blocking state* whenever $I(q) = \emptyset$. The timed automaton \mathcal{A} is *non-blocking* whenever it has no blocking state, *i.e.* $I(q) \neq \emptyset$ for each $q \in Q$.

Remark 2.1.17. Accordingly to Remark 2.1.15, we thus make the assumption that in the sequel, all considered timed automata are non-blocking.

We give a brief example of the previous notions.

Example 2.1.18. We consider again the timed automaton \mathcal{A} of Example 2.1.2. Fix $q_0 = (l_0, (0, 0))$. We easily compute that $I(q_0, e_1) =]0, 1]$ and $I(q_0, e_2) = \{2\}$. It follows that $I(q_0) =]0, 1] \cup \{2\}$. It can also be shown that for each state q , $I(q) \neq \emptyset$, and thus \mathcal{A} is non-blocking.

However, if we consider the variant of \mathcal{A} described in Remark 2.1.15, where $\text{Inv}(l_0) = (x_1 \leq 3) \wedge (x_2 \leq 4)$, it can easily be checked that, for instance, $I((l_0, (2.5, 3))) = \emptyset$ and thus $(l_0, (2.5, 3))$ is blocking state.

Given a timed automaton \mathcal{A} and a state q we write $\text{Runs}(\mathcal{A}, q)$ for the set of infinite runs in \mathcal{A} starting from q . Thanks to Remark 2.1.17, it has sense since we will never end up in a blocking state. We write $\text{Runs}_f(\mathcal{A}, q)$ for the set of finite runs in \mathcal{A} starting from q .

We now give two other examples in order to illustrate the notion of timed automaton. The first example depicts a railroad crossing. It is inspired from [BK08].

Example 2.1.19 (A railroad crossing). When a train is approaching a railroad crossing, a controller has to lower the gate in order to stop the road traffic, and to raise it once the train has left the railroad crossing. A requirement that should be met in the model of such a situation is that the gate is always closed when a train is crossing the road. Timed aspect is thus important in the model: the controller must have enough time to lower the gate before the train is at the crossing. An example of a timed automaton modelling such a situation is depicted on Figure 2.2 and can be described as follows.

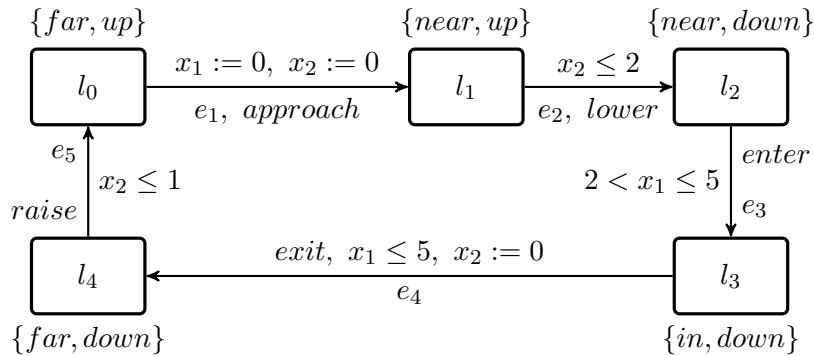


Figure 2.2: A timed automaton modelling a railroad crossing, $\mathcal{A}_{\text{railroad}}$

Two clocks are needed: the first one for the train and the second one for the controller (*i.e.* the time required to lower and to raise the gate). The set of actions is given by $Act = \{approach, enter, exit, lower, raise\}$ and the set of atomic propositions is given by $AP = \{far, near, in, up, down\}$. We assume that the invariant function is **true** in each location. The actions on the edges and the labels on the locations are quite self-explanatory, however we give a brief descriptions of the automaton.

At location l_0 , there is no train near the crossing and the gate is up (given by $\mathcal{L}(l_0) = \{far, up\}$). When a train is approaching, a signal is sent to the controller through the edge e_1 : the action *approach* depicts this signal, clock x_1 and x_2 are reset to 0 and we arrive in location l_1 where the controller is informed that a train is approaching: $\mathcal{L}(l_1) = \{near, up\}$. With the aim of simplifying the model we suppose that, once a train has sent a signal, no other train can approach the railroad crossing until the first train has crossed the road and the controller has raised the gate.

Now the two next steps describe: first the lowering of the gate by the controller and then the crossing of the road by the train. Here the timed aspect is needed in order to make sure that the execution of the different steps is well

ordered and that the controller and the train are well synchronised. Edge e_2 with action *lower* ensures that the controller does not take more than 2 time units to lower the gate, and edge e_3 with action *enter* makes sure that the train takes at least 2 time units and at most 5 in order to cross the road since it has sent the signal. This prevents the situation in which the train crosses the road while the gate is still open. Location l_2 is thus the situation in which the gate has been closed and the train has not yet entered the crossing ($\mathcal{L}(l_2) = \{near, down\}$), while location l_3 represents the state in which the train is crossing the road while the gate is closed ($\mathcal{L}(l_3) = \{in, down\}$).

Now, the train has to leave the railroad crossing before the controller raises the gate which is done by edge e_4 with action *exit*: the train leaves the crossing after at most 5 time units (since the sending of the signal) and location l_4 depicts the situation in which the train has crossed the road but the gate is still closed ($\mathcal{L}(l_4) = \{far, down\}$). We reset clock x_2 in order to control the time needed to raise the gate in edge e_5 : at most 1 time unit. We then return to location l_0 .

As an example,

$$\begin{aligned} \rho = (l_0, (0, 0)) &\xrightarrow[\textit{approach}]{1, e_1} (l_1, (0, 0)) \xrightarrow[\textit{lower}]{1.3, e_2} (l_2, (1.3, 1.3)) \xrightarrow[\textit{enter}]{1.1, e_3} (l_3, (2.4, 2.4)) \\ &\xrightarrow[\textit{exit}]{0.8, e_4} (l_0, (0, 3.2)) \xrightarrow[\textit{raise}]{0.6, e_5} (l_0, (0.6, 3.8)) \end{aligned}$$

is a finite run which represents the crossing of the road by one train.

Observe that this timed automaton has a blocking state. Indeed for instance, $(l_1, (3, 3))$ is a state of $\mathcal{A}_{\text{railroad}}$ (the invariant is **true**) and it is blocking as from state $(l_1, (3, 3))$ edge e_2 is no longer enabled. This configuration represents the situation where, after reaching state $(l_1, (0, 0))$ (*i.e.* after the train sends a signal), the controller waits too much time before lowering the gate, *i.e.* a timed-transition of 3 time units is performed as follows: $(l_1, (0, 0)) \xrightarrow{3} (l_1, (3, 3))$. It represents a flaw in the model. Observe however that since we only consider mixed-transition (see Definition 2.1.11), this situation cannot happen. Observe also that the situation could be avoided with invariants that would prevent the system from reaching such a state.

The second example is a one-clock timed automaton which depicts a mouse producing a simple or double click. This example can be found in [Bri06] and is inspired from [KT05], we give here a brief description.

Example 2.1.20. A mouse produces a double click when the button is pressed twice quickly enough, otherwise it just produces two simple clicks. The timed aspect is thus clearly essential to model such a system since an action has to

be executed twice within a short time interval. Let us assume that a mouse produces a double click when the button is pressed twice within one time unit. A timed automaton for this system is depicted on Figure 2.3.




Figure 2.3: A timed automaton modelling a mouse

There are three locations and one clock in this timed automaton. The set of actions is given by $Act = \{click\}$ and the set of atomic propositions is given by $AP = \{Idle, Simple, Double\}$. In location l_0 , associated with the label *Idle*, the mouse is waiting for the button to be pressed and in location l_1 the button has been pressed once and the mouse produces at least a simple click (which explains the label *Simple*). The system switches from l_0 to l_1 by transition $e_1 = (l_0, click, \mathbf{true}, \{x\}, l_1)$, clock x is reset to zero. Action *click* is then performed, *i.e.* the button has been pressed once. In location l_1 , if the button is pressed within one time unit, then the system goes to location l_2 (by transition $e_2 = (l_1, click, x < 1, \emptyset, l_2)$ where the action *click* is performed again). Otherwise the mouse has produced a simple click and returns to l_0 (transition $e_3 = (l_1, \emptyset, x = 1, \emptyset, l_0)$, no action is performed). Once in l_2 , the button has been pressed twice within one time unit (hence the proposition *Double* on l_2) and the mouse produces a double click. From l_2 we return to l_0 as soon as the double-click is produced (transition $e_4 = (l_2, \emptyset, x \leq 1, \emptyset, l_0)$, no action is performed). As an example,

$$\begin{aligned} \rho = & (l_0, 0) \xrightarrow[click]{1.3, e_1} (l_1, 0) \xrightarrow[click]{0.6, e_2} (l_2, 0.6) \xrightarrow{0.3, e_4} (l_0, 0.9) \\ & \xrightarrow[click]{4.3, e_1} (l_1, 0) \xrightarrow{1, e_3} (l_0, 1) \end{aligned}$$

is a finite run in which the mouse has first produced a double click and then a simple click.

2.1.1 Region graph

In this section, we present a version of the region graph of a timed automaton, see [AD90] and [AD94]. The region graph of a timed automaton is a finite abstraction of the initial model which is in some sense equivalent to the timed automaton. It has helped to prove decidability results on timed automata: *e.g.* in [AD94] it has been proved that the model-checking problem of reachability properties is decidable for timed automata, and in [ACD93] the authors showed the decidability of the model-checking problem of TCTL properties for timed automata.

For this section, we fix a timed automaton $\mathcal{A} = (L, X, E, \text{Inv})$ and its transition system $T_{\mathcal{A}} = (Q, \mathbb{R}_+ \times E, \rightarrow)$. We first define an equivalence relation between clock valuations and we then extend this relation to states of \mathcal{A} . We give some notations. We write $M_{\mathcal{A}}$ for the maximal constant appearing in guards of \mathcal{A} . Given $t \in \mathbb{R}_+$, we write $\lfloor t \rfloor$ for the integer part of t and $\{t\}$ for its fractional part.

Definition 2.1.21. Let $\nu, \nu' \in \mathbb{R}_+^X$. We say that ν and ν' are region-equivalent, and we write $\nu \approx_{\mathcal{A}} \nu'$, whenever the following conditions hold:

1. for every $x \in X$, $\lfloor \nu(x) \rfloor = \lfloor \nu'(x) \rfloor$ or $\nu(x), \nu'(x) > M_{\mathcal{A}}$,
2. for every $x \in X$ such that $\nu(x) \leq M_{\mathcal{A}}$, $\{\nu(x)\} = 0$ if and only if $\{\nu'(x)\} = 0$,
3. for every $x, y \in X$ such that $\nu(x), \nu(y) \leq M_{\mathcal{A}}$, $\{\nu(x)\} \leq \{\nu(y)\}$ if and only if $\{\nu'(x)\} \leq \{\nu'(y)\}$.

This equivalence relation extends to states of $T_{\mathcal{A}}$ with the following condition: given $q = (l, \nu)$ and $q' = (l', \nu') \in Q$, $q \approx_{\mathcal{A}} q'$ if and only if $l = l'$ and $\nu \approx_{\mathcal{A}} \nu'$.

Given $\nu \in \mathbb{R}_+^X$, we write $[\nu]_{\mathcal{A}}$ for the equivalence class of ν under $\approx_{\mathcal{A}}$. The equivalence classes of \mathbb{R}_+^X under $\approx_{\mathcal{A}}$ are called *regions*. We write $R_{\mathcal{A}}$ for the set of regions. Given $q = (l, \nu) \in Q$, it then holds that the equivalence class of q under $\approx_{\mathcal{A}}$, also written $[q]_{\mathcal{A}}$, corresponds to $[q]_{\mathcal{A}} = \{l\} \times [\nu]_{\mathcal{A}}$. We will note $[q]_{\mathcal{A}} = (l, [\nu]_{\mathcal{A}})$. A region r is said *memoryless* if for each clock $x \in X$, either $\nu(x) = 0$ for each $\nu \in r$, or $\nu(x) > M_{\mathcal{A}}$ for each $\nu \in r$. This special kind of regions will be of particular interest to us in the sequel. We illustrate the notion of region on the timed automaton of Example 2.1.2.

Example 2.1.22. Let us consider the timed automaton \mathcal{A} of Example 2.1.2 depicted in Figure 2.1. In this example, it holds that the maximal constant appearing in guards is $M_{\mathcal{A}} = 2$. The regions induced by $\approx_{\mathcal{A}}$ are represented on the left side of Figure 2.4 with a zoom on the partitions of the square $[0, 1]^2$ on

the right side. Observe that this square is partitioned into 11 regions: 4 points, 5 open segments and 2 open triangles.




Figure 2.4: Partition on \mathbb{R}_+^2 induced by $\approx_{\mathcal{A}}$ and zoom on $[0, 1]^2$

We can now define the region graph of a timed automaton \mathcal{A} .

Definition 2.1.23. Let $\mathcal{A} = (L, X, E, \text{Inv})$ be a timed automaton. The region graph $\mathcal{R}_{\mathcal{A}} = (V, F)$ of \mathcal{A} is an oriented graph where:

- $V = L \times R_{\mathcal{A}}$ is the set of vertices;
- $F \subseteq V \times V$ is the set of edges such that $(l, r) \rightarrow (l', r')$, with (l, r) and $(l', r') \in R$, if the following condition is met: there exist $\nu \in r$, $\nu' \in r'$, $t \geq 0$ and $e \in E$ such that $(l, \nu) \xrightarrow{t, e} (l', \nu')$ is a transition of $T_{\mathcal{A}}$.

Given a run $\rho = (q_k)_{k \in K}$ of $T_{\mathcal{A}}$, with $K = \{0, \dots, n\}$ or $K = \mathbb{N}$, we write $\iota(\rho)$ for the corresponding path of ρ in $\mathcal{R}_{\mathcal{A}}$: $\iota(\rho) = ([q_k]_{\mathcal{A}})_{k \in K}$. We now give an example of run in the region graph.

Example 2.1.24. We consider again the timed automaton \mathcal{A} of Example 2.1.2. We have given the regions of this timed automaton in Example 2.1.22 and due to the high number of regions (already 11 regions on the square $[0, 1]^2$), we do not give the entire graph region. However we illustrate here a run in $\mathcal{R}_{\mathcal{A}}$. Consider the following run of $T_{\mathcal{A}}$:

$$\rho = (l_0, (0.7, 0.4)) \xrightarrow{0.8, e_1} (l_0, (0, 1.2)) \xrightarrow{1.1, e_2} (l_1, (1.1, 0))$$

we then have that

$$\iota(\rho) = (l_0, r_0) \rightarrow (l_0, r_2) \rightarrow (l_1, r_4)$$

where

- $r_0 = \{(\nu_1, \nu_2) \in \mathbb{R}_+^2 \mid 0 < \nu_2 < \nu_1 < 1\}$,
- $r_2 = \{(\nu_1, \nu_2) \in \mathbb{R}_+^2 \mid (\nu_1 = 0) \wedge (1 < \nu_2 < 2)\}$, and
- $r_4 = \{(\nu_1, \nu_2) \in \mathbb{R}_+^2 \mid (1 < \nu_1 < 2) \wedge (\nu_2 = 0)\}$.

The run ρ and the regions are depicted on Figure 2.5. In order to make it clearer, we decomposed each mixed-transition into the corresponding time- and switch-transitions; *i.e.*:

- $(l_0, (0.7, 0.4)) \xrightarrow{0.8} (l_0, (1.5, 1.2)) \xrightarrow{e_1} (l_0, (0, 1.2))$, and
- $(l_0, (0, 1.2)) \xrightarrow{1.1} (l_0, (1.1, 2.3)) \xrightarrow{e_2} (l_1, (1.1, 0))$.

Region $r_1 = \{(\nu_1, \nu_2) \in \mathbb{R}_+^2 \mid 1 < \nu_2 < \nu_1 < 2\}$ and region $r_3 = \{(\nu_1, \nu_2) \in \mathbb{R}_+^2 \mid (1 < \nu_1 < 2) \wedge (\nu_2 > 2)\}$ are the regions reached after the two time-transitions.




Figure 2.5: The run ρ and the regions r_0, \dots, r_5

As said before, the region graph has the particular interest that it is finite (whereas the semantics of a timed automaton is given as an infinite (and non-denumerable) transition system) and it abstracts the corresponding automaton leading to, in some sense, an equivalent system. We now explain what we mean by “equivalent”. It uses the notion of *timed-abstract bisimulation* as introduced in [LY93]. A timed-abstract bisimulation is an equivalence relation over the set of states of a timed automaton, that identifies states that have, roughly speaking, the same one-step behaviour for mixed-transitions. This is formalised in the next definition. Given a relation $\mathcal{R} \subseteq Q \times Q$, we write $q\mathcal{R}q'$ instead of $(q, q') \in \mathcal{R}$.

Definition 2.1.25. Given a timed automaton $\mathcal{A} = (L, X, Act, E, Inv, AP, \mathcal{L})$ and $T_{\mathcal{A}} = (Q, \mathbb{R}_+ \times E, \rightarrow)$ its transition system, we say that a symmetric relation $\mathcal{R} \subseteq Q \times Q$ is a *timed-abstract bisimulation* for \mathcal{A} if for each $q\mathcal{R}q'$,

- $\mathcal{L}(q) = \mathcal{L}(q')$, and
- if there is $t \geq 0$, $e \in E$ and $q_1 \in Q$ such that $q \xrightarrow{t,e} q_1$, then there are $t \geq 0$ and $q'_1 \in Q$ such that $q' \xrightarrow{t,e} q'_1$ and $q_1 \mathcal{R} q'_1$.

Observe that since \mathcal{R} is a symmetric relation, the second item holds also in the other sense. We say that two states q and q' are *bisimilar* if there exists a bisimulation \mathcal{R} for \mathcal{A} such that $q \mathcal{R} q'$. In [AD94], the authors showed that the region-equivalence is a timed-abstract bisimulation for timed automata. This result allows for several decidability results, *e.g.* the ones already mentioned earlier: the decidability of reachability properties [AD94] and of TCTL properties [ACD93].

Bisimulation will also have a key role in the compositional framework.

2.1.2 Time-converging aspects

In this section, we present two time-converging aspects that one can observe in TA and that are problematic. Those will be of a peculiar interest in the sequel.

Zenoness. Runs in a given timed automaton can be seen as possible behaviours of the system that it depicts. However not all of them are realistic. We have already seen flaws in the model with blocking states as quickly observed in Example 2.1.19. In this section, we will see that some infinite runs could also represent unrealistic behaviours of the system depicted and should be avoided. We aim here at introducing the particular case of *zeno runs*. We give here a brief note on the subject, but it is detailed in [BK08].

Fix a timed automaton $\mathcal{A} = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L})$ and its transition system $T_{\mathcal{A}} = (Q, \mathbb{R}_+ \times E, \rightarrow)$. We give immediately the definition of this type of runs.

Definition 2.1.26. An infinite run $\rho = q_0 \xrightarrow{t_0, e_0} q_1 \xrightarrow{t_1, e_1} q_2 \xrightarrow{t_2, e_2} \dots$ is said *zeno* whenever $\sum_{i \geq 0} t_i$ is finite.

A zeno run depicts thus a situation where infinitely many actions are performed (the edges e_i) in a finite amount of time ($\sum_{i \geq 0} t_i < \infty$). This is obviously unrealistic. We illustrate this behaviour on Example 2.1.20 in order to make it plain.

Example 2.1.27. Let us consider the mouse system of Example 2.1.20. We

define ρ as follows:

$$\begin{aligned}
 \rho = (l_0, 0) & \xrightarrow[\text{click}]{\frac{1}{2}, e_1} (l_1, 0) \xrightarrow[\text{click}]{\frac{1}{4}, e_2} (l_2, 0.25) \xrightarrow{e_4} (l_0, 0.25) \\
 & \xrightarrow[\text{click}]{\frac{1}{8}, e_1} (l_1, 0) \xrightarrow[\text{click}]{\frac{1}{16}, e_2} \left(l_2, \frac{1}{16}\right) \xrightarrow{e_4} \left(l_0, \frac{1}{16}\right) \\
 & \quad \vdots \\
 & \xrightarrow[\text{click}]{\frac{1}{2^{2n-1}}, e_1} (l_1, 0) \xrightarrow[\text{click}]{\frac{1}{2^{2n}}, e_2} \left(l_2, \frac{1}{2^{2n}}\right) \xrightarrow{e_4} \left(l_0, \frac{1}{2^{2n}}\right) \\
 & \quad \vdots
 \end{aligned} \tag{2.1}$$

Observe that since $\sum_{i \geq 1} 1/2^i = 1 < \infty$, it holds that ρ is a zeno run. In this run, infinitely many actions are executed in one time unit: the button of the mouse is pressed faster and faster so that infinitely many double clicks are produced in one time unit. This obviously describes an unrealistic behaviour.

Zeno runs can be considered as flaws of the model and should be avoided. A timed automaton \mathcal{A} is said *non-zeno* if there does not exist a run that is zeno. One would thus be interested only in non-zeno timed automata

While this seems reasonable, it could in fact be too restrictive. Let \mathcal{A} be a timed automaton. To make sure that there does not exist a zeno run, *i.e.* a run where time only elapses finitely while infinitely many switch-transitions are chosen, we should require that all edges cannot be enabled before a certain constant time. In other words, one should ask that all switch-transitions (or all actions) cost a minimal constant time unit or, in a less restrictive way, we should ask that a sufficient number of the switch-transitions costs a minimal constant time unit (since some actions could be executed instantaneously). It is illustrated in [BK08, pp. 694-695] but not described here as it is not convenient with our work. The authors also provide a sufficient criterion for non-zenoness.

Here, in the probabilistic context of Chapter 3, we will be able to relax the condition of non-zenoness: it will allow us to measure the set of zeno runs and to ask that the set of zeno runs has measure null.

Another time-converging aspect. In this short paragraph, we will mention an example of a timed automaton that will be of a particular interest to us in the sequel. This timed automaton exhibits a time-converging problem different of the zenoness. In this example, we will encounter some kind of convergence on a clock which will lead (as we will see in Chapters 3 and 7) to undesired behaviours. This example can be found in [BBB⁺14].

Example 2.1.28. We consider the two-clock timed automaton \mathcal{A}_{cvg} described in Figure 2.6. Writing $q_0 = (\ell_0, (0, 0))$, we are interested in infinite runs starting from q_0 , *i.e.* $\text{Runs}(\mathcal{A}_{\text{cvg}}, q_0)$. We can make the following statement: each time we come back to ℓ_0 through edges e_3 or e_6 , the value of clock y increases but always stays under the bound of 1. Let us illustrate it with an example of a run:

$$\begin{aligned} \rho = (\ell_0, (0, \mathbf{0})) &\xrightarrow{0.7, e_1} (\ell_1, (0.7, 0.7)) \xrightarrow{0.3, e_2} (\ell_2, (1, 0)) \xrightarrow{0.2, e_3} (\ell_0, (0, \mathbf{0.2})) \\ &\xrightarrow{1, e_4} (\ell_3, (1, 1.2)) \xrightarrow{0.8, e_5} (\ell_4, (1.8, 0)) \xrightarrow{0.5, e_5} (\ell_0, (0, \mathbf{0.5})) \dots \end{aligned}$$

It can be shown that the values of clock y in location ℓ_0 lead to an increasing




Figure 2.6: A two-clock timed automaton \mathcal{A}_{cvg} with a time-convergence phenomena

sequence $(y_n)_{n \geq 0} \subseteq \mathbb{R}_+$ which is strictly bounded by 1 (*i.e.* for each $n \geq 0$, $y_n < 1$) and converges towards some y^* such that $0 < y^* \leq 1$.

Observe that this behaviour is very different of zeno behaviours and actually, it prevents them: it can be shown that each time locations ℓ_2 or ℓ_4 are reached, a strictly larger amount of time has to be waited before returning in ℓ_0 than the last time locations ℓ_2 or ℓ_4 were visited. Finally, observing that the first time location ℓ_2 (resp. ℓ_4) is reached, clock values are at $(1, 0)$ (resp. $(2, 0)$), a delay $0 < t < 1$ time units is needed before taking edge e_3 (resp. e_6). Yielding to the fact that for

each run $\rho \in \text{Runs}(\mathcal{A}_{\text{cvg}}, q_0)$, time diverges (*i.e.* if $\rho = q_0 \xrightarrow{t_0, e_0} q_1 \xrightarrow{t_1, e_1} \dots$, then $\sum_{i \geq 0} t_i = \infty$). However, this time-convergent aspect will lead to some unfair behaviours which will be problematic as we will see in Chapters 3 and 7.

2.1.3 Composition of timed automata

In this section, we briefly explain how composition in timed automata can be defined. Studying an approach that will be useful when addressing the problem in stochastic timed automata, we first define an independent operator of composition in timed automata (*i.e.* that corresponds to an interleaving semantics), and then we will define an operator of composition with synchronisations, also called *handshaking*. It is based on [BK08]. We will briefly discuss in Section 2.4.1 three different types of parallel composition and explain why the handshaking one is the most interesting; the discussion comes from [HZ11]. Finally, we will briefly mention the essential role of bisimulations in a compositional result and mention that the timed-abstract bisimulation defined in Section 2.1.1 ([LY93]) is a congruence w.r.t. parallel composition.

We assume that we have two timed automata $\mathcal{A}_1 = (L_1, X_1, E_1, \text{Inv}_1, \text{AP}_1, \mathcal{L}_1)$ and $\mathcal{A}_2 = (L_2, X_2, E_2, \text{Inv}_2, \text{AP}_2, \mathcal{L}_2)$ such that $X_1 \cap X_2 = \emptyset$. For this first definition, we do not need actions on the edges. We recall the standard (interleaving) parallel composition operator.

Definition 2.1.29. The parallel composition of \mathcal{A}_1 and \mathcal{A}_2 is the timed automaton $\mathcal{A}_1 \parallel \mathcal{A}_2 = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L})$ where

- $L = L_1 \times L_2$, $X = X_1 \cup X_2$, $\text{AP} = \text{AP}_1 \cup \text{AP}_2$,
- for each location $l = (l_1, l_2) \in L$, we define $\text{Inv}(l) = \text{Inv}_1(l_1) \wedge \text{Inv}_2(l_2)$ and $\mathcal{L}(l) = \mathcal{L}_1(l_1) \cup \mathcal{L}_2(l_2)$, and
- $E = E_{1,\bullet} \cup E_{\bullet,2}$ with
 - ▷ $E_{1,\bullet} = \{((l_1, l_2), g, Y, (l'_1, l_2)) \mid (l_1, g, Y, l'_1) \in E_1, l_2 \in L_2\}$, and
 - ▷ $E_{\bullet,2} = \{((l_1, l_2), g, Y, (l_1, l'_2)) \mid (l_2, g, Y, l'_2) \in E_2, l_1 \in L_1\}$.

We illustrate this notion on a very simple example.

Example 2.1.30. Consider the timed automata $\mathcal{A}_i = (L_i, X_i, E_i, \text{Inv}_i, \text{AP}_i, \mathcal{L}_i)$ for each $i \in \{1, 2\}$, depicted in Figure 2.7, defined for each $i \in \{1, 2\}$ with the following components: $L_i = \{l_i, l'_i\}$, $E_i = \{(l_i, g_i, Y_i, l'_i)\}$ for some $g_i \in \mathcal{G}(X_i)$ and $Y_i \subseteq X_i$, $\text{Inv}_i(l_i) = \text{Inv}_i(l'_i) = g'_i \in \mathcal{G}(X_i)$, $\text{AP}_i = \{p_i\}$ and $\mathcal{L}_i(l_i) = \mathcal{L}_i(l'_i) = \{p_i\}$. Then, $\mathcal{A}_1 \parallel \mathcal{A}_2 = (L, X_1 \cup X_2, E, \text{Inv}, \text{AP}_1 \cup \text{AP}_2, \mathcal{L})$ where

$L = \{(l_1, l_2), (l'_1, l_2), (l_1, l'_2), (l'_1, l'_2)\}$, for each $l \in L$, $\text{Inv}(l) = g'_1 \wedge g'_2$ and $\mathcal{L}(l) = \{p_1, p_2\}$, and

$$E = \left\{ \begin{aligned} &((l_1, l_2), g_1, Y_1, (l'_1, l_2)), ((l_1, l_2), g_2, Y_2, (l_1, l'_2)), \\ &((l'_1, l_2), g_2, Y_2, (l'_1, l'_2)), ((l_1, l'_2), g_1, Y_1, (l'_1, l'_2)) \end{aligned} \right\}.$$

It is depicted in Figure 2.7.

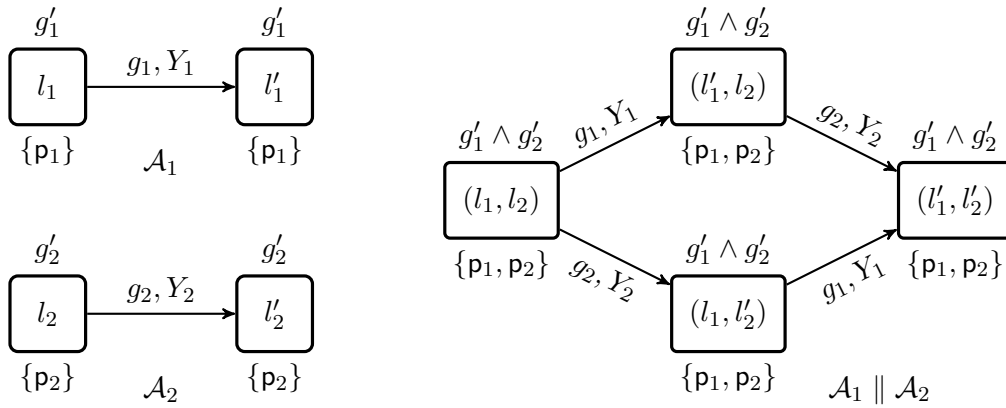


Figure 2.7: The composition of two simple timed automata \mathcal{A}_1 and \mathcal{A}_2

Remark 2.1.31. The set $E_{1,\bullet}$ (resp. $E_{\bullet,2}$) can be seen as the set of switch-transitions (edges) in $\mathcal{A}_1 \parallel \mathcal{A}_2$ that only perform a change of location in \mathcal{A}_1 (resp. \mathcal{A}_2). Hence, we abusively denote $E_{1,\bullet}$ by E_1 and $E_{\bullet,2}$ by E_2 . Let us observe that there are no edges in $\mathcal{A}_1 \parallel \mathcal{A}_2$ that depict switch-transitions in both timed automata \mathcal{A}_1 and \mathcal{A}_2 . More precisely, we have that elements of the form $((l_1, l_2), g_1 \wedge g_2, Y_1 \cup Y_2, (l'_1, l'_2))$ where $(l_1, g_1, Y_1, l'_1) \in E_1$ and $(l_2, g_2, Y_2, l'_2) \in E_2$ are not considered as edges of $\mathcal{A}_1 \parallel \mathcal{A}_2$. This is due to the fact that it is very unlikely (or, in a context with probabilities like in stochastic timed automata in Chapter 3, it has probability null), that both automata perform an action at the exact same time. However, we will see how this can be enforced using synchronisations in Definition 2.1.33.

Remark 2.1.32. We can extend Definition 2.1.29 to the composition of n timed automata $\mathcal{A}_1, \dots, \mathcal{A}_n$ with $\bigcap_{i=1}^n X_i = \emptyset$, where $\mathcal{A}_i = (L_i, X_i, E_i, \text{Inv}_i, \text{AP}_i, \mathcal{L}_i)$ for any $i \in \{1, \dots, n\}$, as follows. We define

$$\mathcal{A}_1 \parallel \dots \parallel \mathcal{A}_n = \left(L_1 \times \dots \times L_n, \bigcup_{i=1}^n X_i, E, \text{Inv}, \bigcup_{i=1}^n \text{AP}_i, \mathcal{L} \right),$$

where

$$E = \bigcup_{i=1}^n E_{\bullet, i, \bullet},$$

$$E_{\bullet, i, \bullet} = \left\{ \left((l_1, \dots, l_{i-1}, l_i, l_{i+1}, \dots, l_n), g, Y, (l_1, \dots, l_{i-1}, l'_i, l_{i+1}, \dots, l_n) \right) \mid \right. \\ \left. (l_i, g, Y, l'_i) \in E_i \text{ and } \forall j \neq i, l_j \in L_j \right\}$$

for any $i \in \{1, \dots, n\}$, and Inv and \mathcal{L} are obviously extended from Definition 2.1.29. We then abusively write E_i for $E_{\bullet, i, \bullet}$.

We are now interested in a notion of composition with synchronisations, also called *handshaking* parallel composition. This time, we consider two timed automata with the labelling over the edges $\mathcal{A}_1 = (L_1, X_1, \text{Act}_1, E_1, \text{Inv}_1, \text{AP}_1, \mathcal{L}_1)$ and $\mathcal{A}_2 = (L_2, X_2, \text{Act}_2, E_2, \text{Inv}_2, \text{AP}_2, \mathcal{L}_2)$ with $X_1 \cap X_2 = \emptyset$ and $\text{Act}_1 \cap \text{Act}_2 \neq \emptyset$. We fix $A \subseteq \text{Act}_1 \cap \text{Act}_2$.

Definition 2.1.33. The parallel composition of \mathcal{A}_1 and \mathcal{A}_2 on A is the timed automaton $\mathcal{A}_1 \parallel_A \mathcal{A}_2 = (L, X, \text{Act}, E, \text{Inv}, \text{AP}, \mathcal{L})$ where L , X , Inv , AP and \mathcal{L} are defined as in Definition 2.1.29 and: $\text{Act} = \text{Act}_1 \cup \text{Act}_2$ and E is defined as follows:

- for each $a \in A$, if $l_1 \xrightarrow{a, g_1, Y_1} l'_1 \in E_1$ and if $l_2 \xrightarrow{a, g_2, Y_2} l'_2 \in E_2$, then $(l_1, l_2) \xrightarrow{a, g_1 \wedge g_2, Y_1 \cup Y_2} (l'_1, l'_2) \in E$,
- for each $a \notin A$, if $l_1 \xrightarrow{a, g_1, Y_1} l'_1 \in E_1$ then for each $l_2 \in L_2$, $(l_1, l_2) \xrightarrow{a, g_1, Y_1} (l'_1, l_2) \in E$, and
- for each $a \notin A$, if $l_2 \xrightarrow{a, g_2, Y_2} l'_2 \in E_2$ then for each $l_1 \in L_1$, $(l_1, l_2) \xrightarrow{a, g_2, Y_2} (l_1, l'_2) \in E$.

We immediately illustrate this definition on an example. It is the railroad crossing of Example 2.1.19 that can be found in details in [BK08].

Example 2.1.34. We consider again the modelisation of a railroad crossing like in Example 2.1.19 depicted in Figure 2.2. However this time, we assume that the system is composed of two components: the train that is approaching the railroad crossing and the controller that has to lower (resp. raise) the gate when the train is approaching (resp. leaving) the railroad crossing. Those components are represented on Figure 2.8 as timed automata $\mathcal{A}_{\text{train}}$ and $\mathcal{A}_{\text{controller}}$. Then if we synchronise $\mathcal{A}_{\text{train}}$ and $\mathcal{A}_{\text{controller}}$ on $A = \{\text{approach}, \text{exit}\}$ the product $\mathcal{A}_{\text{train}} \parallel_A \mathcal{A}_{\text{controller}}$ is the timed automaton $\mathcal{A}_{\text{railroad}}$ of Example 2.1.19 depicted

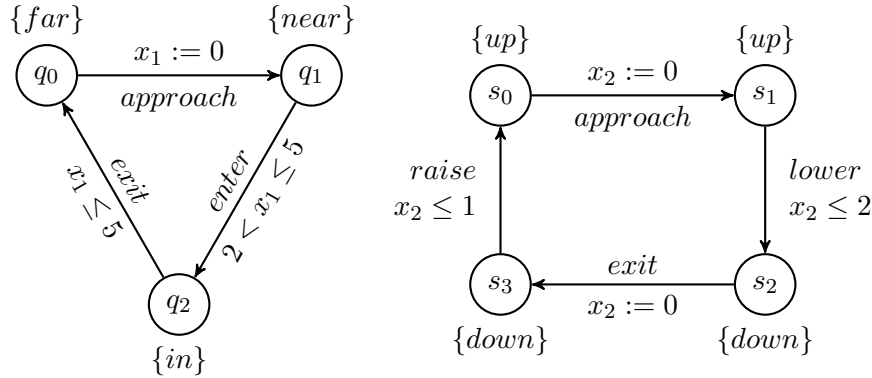


Figure 2.8: Timed automata \mathcal{A}_{train} modelling the train (on the left), and $\mathcal{A}_{controller}$ modelling the controller (on the right).

on Figure 2.2, where $l_0 = (q_0, s_0)$, $l_1 = (q_1, s_1)$, $l_2 = (q_1, s_2)$, $l_3 = (q_2, s_2)$ and $l_4 = (q_0, s_3)$.

Observe that there should also be the edge $(q_1, s_1) \xrightarrow{\text{enter}} (q_1, s_2)$. This represents the case where the train has entered the railroad crossing while the gate was still open. This corresponds to a faulty behaviour. However, this assumes that clock x_2 has waited longer than 2 time units in location s_1 *i.e.* the controller has waited too long before lowering the gate. This leads to a blocking state as observed in Example 2.1.19. As said previously in Remark 2.1.17, we do not consider blocking states, and we hence do not consider such behaviours. We can thus assume that this edge can never happen.

We now come back the notion of bisimulation on timed automata or more precisely on the timed-abstract bisimulation (see Definition 2.1.25). Roughly speaking, a bisimulation is an equivalence relation between states that have similar behaviours. Definition 2.1.25 introduces bisimulation as a relation between states, however it can be extended to a relation between timed automata in a standard way. It can be done by constructing a disjoint union of two timed automata and by defining a bisimulation on this disjoint union timed automaton. We will then say that \mathcal{A}_1 and \mathcal{A}_2 are *bisimilar* if there exists such a bisimulation. We write it $\mathcal{A}_1 \sim \mathcal{A}_2$. We will be more precise on these techniques when abording the definition in stochastic timed automata in Chapter 9.

Bisimulation is important when dealing with composition: if two systems have similar behaviours (*i.e.* are bisimilar), then when you compose them with another system, you should still have similar systems (*i.e.* bisimilar systems). This is formalised with the next result.

Theorem 2.1.35. *Bisimilarity is a congruence w.r.t. composition. More precisely: we consider three timed automata $\mathcal{A}_1 = (L_1, X_1, Act_1, E_1, Inv_1, AP_1, \mathcal{L}_1)$, $\mathcal{A}_2 = (L_2, X_2, Act_2, E_2, Inv_2, AP_2, \mathcal{L}_2)$ and $\mathcal{A} = (L, X, Act, E, Inv, AP, \mathcal{L})$ such that $X_1 \cap X_2 = X_2 \cap X = \emptyset$ and $Act_1 \cap Act_2 \cap Act \neq \emptyset$, and we consider $A \subseteq Act_1 \cap Act_2 \cap Act$. Then we have that*

- $\mathcal{A}_1 \sim \mathcal{A}_2$ implies $\mathcal{A}_1 \parallel_A \mathcal{A} \sim \mathcal{A}_2 \parallel_A \mathcal{A}$, and
- $\mathcal{A}_1 \sim \mathcal{A}_2$ implies $\mathcal{A} \parallel_A \mathcal{A}_1 \sim \mathcal{A} \parallel_A \mathcal{A}_2$.

It should be noted that timed-abstraction bisimulation is not the unique existing notion of bisimulation. In particular in [LY93], the authors provide several other types of bisimulation, including *strong timed bisimulation* and *weak timed bisimulation*. But we only focus here on timed-abstract bisimulation.

2.2. Denumerable Markov chains

In this section, we introduce the denumerable Markov chain model (DMC for short) [KSK76] and recall some notions of [ABM07]. We will not prove any results here as we will provide them later in Chapter 6 in a more general context.

Markov chains are probabilistic transition systems, *i.e.* transition systems (see Definition 2.1.4) in which each state is equipped with a distribution over the set of outgoing edges. In the case of DMCs, the set of states can be infinite but is, at most, denumerable.

Finite Markov chains (*i.e.* DMCs with a finite set of states) enjoy a lot of nice decidability results. Amongst others, we cite [Var85] and [CY88]: in the first paper, it is shown that the almost-sure model-checking problem of LTL properties is reduced to structural properties of the underlying graph while in the latter, it has been shown that the exact probability of LTL properties can be computed.

In [ABM07], the authors are interested in DMCs and allow thus for a denumerable set of states. They define the notion of *decisiveness*. Roughly speaking, a DMC is *decisive w.r.t.* a set of states B if it reaches B or a state from which B can never be reached with probability 1. This notion allows one to lift good properties from finite Markov chains to denumerable ones. The authors show the qualitative and quantitative problems for reachability and repeated reachability properties to be decidable for decisive Markov chains.

The aim of this section is thus to define the notions and present the results of [ABM07]. We first define the model of DMC and introduce some classical notions. We assume that the reader is familiar with basic notions of probability theory (see for instance [Fel66] and [Fel69]).

Definition 2.2.1. A *denumerable Markov chain* (DMC for short) is a couple $\mathcal{M} = (S, P)$ where S is an at most denumerable set of states and $P : S \times S \rightarrow [0, 1]$ is the probability transition such that for each $s \in S$, $\sum_{s' \in S} P(s, s') = 1$.

A DMC induces a transition system (see Definition 2.1.4) where the set of states is S , the alphabet Γ is given by $\{P(s, s') \mid s, s' \in S\}$ and there is a transition $s \rightarrow s'$ if $P(s, s') > 0$, which is then labelled with $P(s, s')$. Note that if $P(s, s') = 1$ we may omit the label on the transition. We can then always represent a DMC by its transition system. We give now two examples. The first example is a simple finite Markov chain.

Example 2.2.2. Consider the Markov chain $\mathcal{M}_1 = (S_1, P_1)$ depicted in Figure 2.9. In this example $S_1 = \{s_0, s_1, s_2, s_3, s_4, s_5\}$ and P_1 is trivially defined as represented on the figure: for instance, $P_1(s_0, s_1) = 0.3$ and $P_1(s_4, s_5) = 1$.




Figure 2.9: A finite Markov chain \mathcal{M}_1


The second example is the classical random-walk over the natural number.

Example 2.2.3. The second example is the Markov chain depicted in Figure 2.10. We consider here $\mathcal{M}_2 = (S_2, P_2)$ where

- $S_2 = \mathbb{N}$,
- for each $i \geq 1$, $P_2(i, i + 1) = p$ and $P_2(i, i - 1) = 1 - p$ with $p \in]0, 1[$, and
- $P_2(0, 1) = 1$.

This represents a random walk over the natural numbers.

Like in timed automata, we can define a run in a DMC $\mathcal{M} = (S, P)$: a finite run is a sequence of states $\rho = s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_n$ for some $n \geq 0$ such that for each $0 \leq i \leq n - 1$, $P(s_i, s_{i+1}) > 0$. The definition extends to

Figure 2.10: Random walk over \mathbb{N} .

infinite runs trivially. We uniformly write a run as $(s_k)_{k \in K}$ where either $K = \mathbb{N}$ or $K = \{0, \dots, n\}$ for some $n \geq 0$. Given a DMC \mathcal{M} and a state s , we write $\text{Runs}(\mathcal{M}, s)$ for the set of infinite runs starting from s .

A probability measure over the set of infinite runs of a DMC $\mathcal{M} = (S, P)$ is naturally defined. More precisely, fix an initial state $s_0 \in S$ and given a finite run starting in s_0 , $\rho = (s_k)_{k \in \{0, \dots, n\}}$ for some $n \geq 0$, we define the cylinder generated by ρ as

$$\text{Cyl}(\rho) = \{\rho' = (s'_k)_{k \in \mathbb{N}} \in \text{Runs}(\mathcal{M}, s_0) \mid \forall 0 \leq i \leq n, s'_i = s_i\},$$

i.e. the set of infinite runs that have ρ as a prefix. We then define the probability space $(\text{Runs}(\mathcal{M}, s_0), \Omega_{s_0}^{\mathcal{M}}, \text{Prob}_{s_0}^{\mathcal{M}})$ where

- $\Omega_{s_0}^{\mathcal{M}}$ is the σ -algebra generated by the cylinders starting in s_0 ,
- $\text{Prob}_{s_0}^{\mathcal{M}}$ is defined as follows: for each finite run $\rho = (s_k)_{k \in \{0, \dots, n\}}$,

$$\text{Prob}_{s_0}^{\mathcal{M}}(\text{Cyl}(\rho)) = \prod_{0 \leq i < n} P(s_i, s_{i+1})$$

and if $n = 0$, $\text{Prob}_{s_0}^{\mathcal{M}}(\text{Cyl}(\rho)) = 1$; thanks to Caratheodory's extension theorem, it extends to all measurable sets.

Observe that given $\rho = s_0 s_1 s_2 \dots s_n s_{n+1}$ and $\rho' = s_1 s_2 \dots s_n s_{n+1}$, it holds that

$$\text{Prob}_{s_0}^{\mathcal{M}}(\text{Cyl}(\rho)) = P(s_0, s_1) \cdot \text{Prob}_{s_1}^{\mathcal{M}}(\text{Cyl}(\rho')).$$

Remark 2.2.4. If we write $\text{Runs}(\mathcal{M})$ for the set of infinite runs of \mathcal{M} (here without a fixed initial state) and $\Omega^{\mathcal{M}}$ for the σ -algebra generated by all cylinders, observe that $\text{Prob}_{s_0}^{\mathcal{M}}$ defines also a probability distribution over $(\text{Runs}(\mathcal{M}), \Omega^{\mathcal{M}})$. Now if we fix some initial distribution μ over the set of states S instead of an initial state, we can define a probability distribution $\text{Prob}_{\mu}^{\mathcal{M}}$ over $\text{Runs}(\mathcal{M})$ as follows: given $\rho = s_0 s_1 \dots s_n$,

$$\text{Prob}_{\mu}^{\mathcal{M}}(\text{Cyl}(\rho)) = \mu(s_0) \cdot \text{Prob}_{s_0}^{\mathcal{M}}(\text{Cyl}(\rho)).$$

We illustrate those notions on the previous examples.

Example 2.2.5. In Example 2.2.2 an example of a finite run in the DMC \mathcal{M}_1 is $\rho_1 = s_0 s_4 s_5 s_4$ and an example of infinite run is $\rho_2 = s_0 s_3 s_3 s_3 \dots$. If we fix s_0 as the initial state, we can measure the probability of the cylinder generated by ρ_1 as follows:

$$\begin{aligned} \text{Prob}_{s_0}^{\mathcal{M}_1}(\text{Cyl}(\rho_1)) &= P_1(s_0, s_4) \cdot P_1(s_4, s_5) \cdot P_1(s_5, s_4) \\ &= 0.5 \cdot 1 \cdot 0.8 = 0.4. \end{aligned}$$

In DMC \mathcal{M}_2 of Example 2.2.3, some examples of finite runs are given by $\rho'_n = 0 \rightarrow 1 \rightarrow \dots \rightarrow n \rightarrow n+1$ for each $n \geq 0$ and an example of infinite run is $\rho' = 0 \rightarrow 1 \rightarrow \dots \rightarrow n \rightarrow n+1 \rightarrow \dots$. Again, if we consider 0 as the initial state, we can compute:

$$\begin{aligned} \text{Prob}_0^{\mathcal{M}_2}(\text{Cyl}(\rho'_n)) &= P_2(0, 1) \cdot P_2(1, 2) \cdot \dots \cdot P_2(n, n+1) \\ &= p^n. \end{aligned}$$

The σ -algebra generated by the cylinders allows one to express a rich variety of sets of runs, including sets of runs satisfying a given property expressed in LTL (see [Var85]). For the purpose of the section, we consider LTL formulas over the set of states S :

$$\varphi ::= B \mid \varphi_1 \wedge \varphi_2 \mid \neg\varphi \mid \varphi_1 \mathbf{U} \varphi_2 \ ,$$

where $B \subseteq S$, φ , φ_1 and φ_2 are LTL formulas. The satisfaction relation is then defined as follows. Given an infinite run $\rho = s_0 s_1 s_2 \dots \in \text{Runs}(\mathcal{M})$, writing $\rho_{\geq i} = s_i s_{i+1} s_{i+2} \dots \in \text{Runs}(\mathcal{M})$ for each $i \geq 0$, it holds that

$$\begin{aligned} \rho \models B &\iff s_0 \in B \\ \rho \models \varphi_1 \wedge \varphi_2 &\iff \rho \models \varphi_1 \text{ and } \rho \models \varphi_2 \\ \rho \models \neg\varphi &\iff \rho \not\models \varphi \\ \rho \models \varphi_1 \mathbf{U} \varphi_2 &\iff \exists i \geq 0 \text{ s.t. } \rho_{\geq i} \models \varphi_2 \text{ and } \forall 0 \leq j < i, \rho_{\geq j} \models \varphi_1. \end{aligned}$$

It holds that those formulas are measurable ([Var85]). Given a formula φ and given an initial state s_0 , we will simply write $\text{Prob}_{s_0}^{\mathcal{M}}(\varphi)$ for the probability of the set of infinite runs starting in s_0 and satisfying φ . We will also use classical notations like: $\top = S$; $\perp = \emptyset$; $\varphi_1 \vee \varphi_2 = \neg(\neg\varphi_1 \wedge \neg\varphi_2)$; $\mathbf{F}\varphi = \top \mathbf{U} \varphi$; $\mathbf{G}\varphi = \neg\mathbf{F}(\neg\varphi)$.

We will also consider CTL-like notations in order to define some set of states. Given an LTL formula φ , we define state formulas $\exists\varphi$ and $\forall\varphi$. Their semantics is given as follows:

$$\begin{aligned}
s \models \exists\varphi &\iff \exists\rho \in \text{Runs}(\mathcal{M}, s) \text{ s.t. } \rho \models \varphi \\
s \models \forall\varphi &\iff \forall\rho \in \text{Runs}(\mathcal{M}, s), \rho \models \varphi
\end{aligned}$$

With such a probabilistic model, one is then interested with the *qualitative and quantitative model-checking problems*. These problems are not only interesting for DMCs. All probabilistic systems on which a distribution over the set of runs can be defined, is concerned with those. We define it here for DMCs and LTL formulas, but we will see in Chapters 3 and 4 how this is straightforwardly extended to other probabilistic models.

Definition 2.2.6. Given a DMC \mathcal{M} , an initial state s_0 and an LTL formula φ , the *qualitative model-checking problem* aims at verifying whether $\text{Prob}_{s_0}^{\mathcal{M}}(\varphi) = 1$ or not, or $\text{Prob}_{s_0}^{\mathcal{M}}(\varphi) = 0$ or not.

This is also called the *almost-sure* model-checking problem (resp. the 0-model-checking problem).

Definition 2.2.7. Given a DMC \mathcal{M} , an initial state s_0 and an LTL formula φ , the *quantitative model-checking problem* aims at computing an approximation of $\text{Prob}_{s_0}^{\mathcal{M}}(\varphi)$.

As said earlier, those problems are shown to be decidable for finite Markov chains in [Var85] and [CY88]. For DMCs, those problems are analysed for reachability and repeated reachability properties in [ABM07] and are shown to be decidable under some hypotheses. In order to present these results, we need first to introduce some notions.

2.2.1 Attractors

In this section, we recall the notion of attractor in DMCs.

Definition 2.2.8. Let $\mathcal{M} = (S, P)$ be a DMC. A set $A \subseteq S$ is an *attractor* if for each $s \in S$, $\text{Prob}_s^{\mathcal{M}}(\mathbf{F} A) = 1$.

In other words, A is an attractor if it is reachable with probability 1 from each state of the DMC. Note that all DMCs have an attractor by taking the trivial one: $A = S$. We have a first important result on attractors.

Proposition 2.2.9. *If A is an attractor of DMC $\mathcal{M} = (S, P)$, it holds that for each $s \in S$, $\text{Prob}_s^{\mathcal{M}}(\mathbf{G} \mathbf{F} A) = 1$.*

We will have a particular interest for finite attractors, when considering a DMC with an infinite set of states. An attractor of interest for a finite Markov chain is the union of all *bottom strongly connected components* (written BSCC for short).

Let us give a short reminder on BSCCs. A finite Markov chain can be associated with an oriented graph (V, E) , in which the set of vertices V is the set of states and the edges of E are given by the transitions of the induced transition system. Then we say that $B \subseteq S$ is *strongly connected* if for each pair of states $s, s' \in B$ there exists a finite run starting in s and ending in s' and vice-versa. A *strongly connected component* (SCC for short) is a maximal strongly connected set of vertices. Finally, a *bottom strongly connected component* is a SCC that you cannot leave once reached: C is a BSCC if it is a SCC and moreover, for each $s \in C$ and $s' \in S$, $(s, s') \in E$ implies $s' \in C$.

Given a finite Markov chain \mathcal{M} , we write $\text{BSCC}(\mathcal{M})$ for the set of BSCCs of the oriented graph associated with \mathcal{M} . We then have this standard result:

Proposition 2.2.10. *Let $\mathcal{M} = (S, P)$ be a finite Markov chain. It holds that the union of all BSCCs of \mathcal{M} is an attractor: $\{s \in S \mid \exists B \in \text{BSCC}(\mathcal{M}), s \in B\}$ is an attractor.*

We illustrate the previous notions and results on Examples 2.2.2 and 2.2.3.

Example 2.2.11. Consider the finite Markov chain \mathcal{M}_1 of Example 2.2.2. This Markov chain has two BSCCs: $\{s_1, s_2\}$ and $\{s_3\}$; and observe that $\{s_0, s_4, s_5\}$ forms a SCC. From Proposition 2.2.10, we thus get that $\{s_1, s_2, s_3\}$ is an attractor for \mathcal{M}_1 .

Consider now the random-walk of Example 2.2.3. Here, the set of states is infinite, we therefore cannot speak of the BSCCs of \mathcal{M}_2 . If $p > 1/2$, it can be shown that \mathcal{M}_2 has no finite attractor. Indeed, classical results on random-walks state that for each $n \geq 0$, $\text{Prob}_{n+1}^{\mathcal{M}_2}(\mathbf{F}\{n\}) < 1$. Now if $p \leq 1/2$, it holds that $\{0\}$ is an attractor. And we can infer that any bounded subset $A \subseteq \mathbb{N}$ is an attractor.

2.2.2 Decisiveness

We are now able to present the notions and results of [ABM07]. It relies on a notion of decisiveness. Roughly speaking, a DMC \mathcal{M} is decisive w.r.t. a set of states B if for each state s , the probability to reach B or a state from which B can never be visited equals 1. This notion of decisiveness allows to lift good properties of finite Markov chains to DMCs.

Given a set of states $B \subseteq S$ we define $\widetilde{B} = \{s \in S \mid s \not\in \exists \mathbf{F} B\}$ i.e. \widetilde{B} denotes the set of states from which B can never be reached.

Definition 2.2.12. Let $\mathcal{M} = (S, P)$ be a DMC and $B \subseteq S$ be a set of states. We say that \mathcal{M} is *decisive w.r.t. B* if $\text{Prob}_s^{\mathcal{M}}(\mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$ for each state $s \in S$. Similarly, we say that \mathcal{M} is *strongly decisive w.r.t. B* if for each $s \in S$, $\text{Prob}_s^{\mathcal{M}}(\mathbf{G} \mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$.

It is proven in [ABM07, Lemma 3.2] that the two notions are equivalent for DMCs. But both notions give interesting behaviours: the notion of decisiveness asks that the set of runs that can always reach B but never reach it is almost-surely empty (*i.e.* it has a null probability), while the notion of strong decisiveness requires that the set of runs that can always reach B but only reach it finitely many times is almost-surely empty. Observe that it is obvious that all DMCs are decisive (and thus strongly decisive) w.r.t. to the set of all their states, but this is often not very interesting.

There is a link between finite attractors and decisiveness ([ABM07, Lemma 3.4]).

Proposition 2.2.13 ([ABM07]). *Let $\mathcal{M} = (S, P)$ be a DMC and $A \subseteq S$ a set of states. If A is a finite attractor of \mathcal{M} then \mathcal{M} is decisive w.r.t. any set of states B .*

This implies that all finite Markov chains are decisive w.r.t. any set of states. Let us come back to Examples 2.2.2 and 2.2.3.

Example 2.2.14. Observe the DMC \mathcal{M}_1 of Example 2.2.2. Consider the set of states $B = \{s_0, s_4, s_5\}$. Since $B_1 = \{s_1, s_2\}$ and $B_2 = \{s_3\}$ are BSCCs, it is trivial to get that $\widetilde{B} = B_1 \cup B_2$. Notice also that $\widetilde{B}_1 = B_2$ and $\widetilde{B}_2 = B_1$. Finally, since \mathcal{M}_1 is finite, we get that \mathcal{M}_1 is decisive (and thus strongly decisive) w.r.t. any set of states.

Now let us study the random-walk \mathcal{M}_2 of Example 2.2.3. First of all, since the chain is strongly connected, we get that for each $B \subseteq \mathbb{N}$, $\widetilde{B} = \emptyset$. Now let us recall the observations of Example 2.2.11. If $p > 1/2$, \mathcal{M}_2 has no finite attractor and thus we cannot infer anything from Proposition 2.2.13. As mentioned before, in this case we get that for any bounded set of states B , there exists a state $s \in S$ such that $\text{Prob}_s^{\mathcal{M}_2}(\mathbf{F} B) < 1$. And since $\widetilde{B} = \emptyset$, we get that \mathcal{M}_2 is not decisive w.r.t. B . Now if $p \leq 1/2$, \mathcal{M}_2 has a finite attractor. Proposition 2.2.13 allows hence to establish that in this case, \mathcal{M}_2 is decisive w.r.t any set of states.

As said before, decisiveness allows to lift good properties of finite Markov chains to denumerable ones. In [ABM07], the authors are invested in qualitative and quantitative model-checking problems (Definitions 2.2.6 and 2.2.7) for reachability and repeated reachability problems. They have the following results:

- under some decisiveness assumptions, the qualitative model-checking problems for reachability and repeated reachability are reduced to structural properties on the underlying graph;
- under some decisiveness and effectiveness assumptions, there exist algorithms which are correct and terminate, that approximate the probability of reachability and repeated reachability properties, *i.e.* the quantitative model-checking problems for reachability and repeated reachability problems are decidable.

We will now list precisely those results. We fix a DMC $\mathcal{M} = (S, P)$ and a set of states $B \subseteq S$. The first result concerns the qualitative problem of reachability properties. It corresponds to Lemmas 5.1 and 5.2 of [ABM07].

Proposition 2.2.15 ([ABM07]). *For each initial state $s_0 \in S$, it holds that:*

- $\text{Prob}_{s_0}^{\mathcal{M}}(\mathbf{F} B) = 1$ implies $s_0 \not\models \exists(\neg B \mathbf{U} \widetilde{B})$;
- if \mathcal{M} is decisive w.r.t. B , then $s_0 \not\models \exists(\neg B \mathbf{U} \widetilde{B})$ implies $\text{Prob}_{s_0}^{\mathcal{M}}(\mathbf{F} B) = 1$.

It states that, under a decisiveness assumption, it cannot be the case that a given set of states B is reached with probability 1 but that \widetilde{B} is reached through a run that never visits B beforehand. This reduces the almost-sure model-checking problem of reachability properties to the non-satisfaction of an “Until” formula on the underlying graph. There is also a result in [ABM07] about the 0-model-checking problem (does a formula have probability 0 or not?) of reachability properties: it is reduced to the non-satisfaction of a reachability property. It corresponds to Lemma 5.8.

Proposition 2.2.16 ([ABM07]). *For each initial state $s_0 \in S$, $\text{Prob}_{s_0}^{\mathcal{M}}(\mathbf{F} B) = 0$ if and only if $s_0 \not\models \exists \mathbf{F} B$ (*i.e.* $s_0 \in \widetilde{B}$).*

The next result is about the repeated reachability qualitative model-checking problem. Observe that by definition, it holds that $\widetilde{B} = \{s \in S \mid s \not\models \exists \mathbf{F} \widetilde{B}\}$, *i.e.* it is the set of states from which \widetilde{B} can never be reached. The following result corresponds to Lemmas 6.1 and 6.2 of [ABM07].

Proposition 2.2.17 ([ABM07]). *For each initial state $s_0 \in S$, it holds that:*

- $\text{Prob}_{s_0}^{\mathcal{M}}(\mathbf{G} \mathbf{F} B) = 1$ implies $s_0 \models \forall \mathbf{G} \exists \mathbf{F} B$;
- if moreover \mathcal{M} is strongly decisive w.r.t. B , then $s_0 \models \forall \mathbf{G} \exists \mathbf{F} B$ implies $\text{Prob}_{s_0}^{\mathcal{M}}(\mathbf{G} \mathbf{F} B) = 1$.

It states that, under a strong decisiveness assumption, a set of states B is visited infinitely often with probability 1 if and only if B can always be reached or, equivalently, \widetilde{B} can never be reached: indeed observe that $s_0 \models \forall \mathbf{G} \exists \mathbf{F} B$ is equivalent to $s_0 \in \widetilde{\widetilde{B}}$. This reduces the almost-sure model-checking problem of repeated reachability properties to the non-satisfaction of a certain reachability formula (to reach \widetilde{B}) on the underlying graph. There is also a result that reduces the positive model-checking problem (*i.e.* similar to Definition 2.2.6, except here the question is: is $\text{Prob}_{s_0}^{\mathcal{M}}(\varphi) > 0$ or not?) of repeated reachability properties to the satisfaction of a reachability property in the underlying graph; it corresponds to Theorem 6.11 of [ABM07].

Theorem 2.2.18 ([ABM07]). *For each initial state $s_0 \in S$, it holds that*

- *if \mathcal{M} is decisive w.r.t. B , then $s_0 \models \exists \mathbf{F} \widetilde{\widetilde{B}}$ implies $\text{Prob}_{s_0}^{\mathcal{M}}(\mathbf{G} \mathbf{F} B) > 0$;*
- *if moreover \mathcal{M} is decisive w.r.t. B and w.r.t. \widetilde{B} , then $\text{Prob}_{s_0}^{\mathcal{M}}(\mathbf{G} \mathbf{F} B) > 0$ implies $s_0 \models \exists \mathbf{F} \widetilde{\widetilde{B}}$.*

Remark 2.2.19. Observe that decisiveness w.r.t. a set of states B does not imply decisiveness w.r.t. \widetilde{B} .

All those previous results show thus that qualitative model-checking problems of reachability and repeated reachability problems can be reduced to structural properties on the underlying graph. It can thus be decided thanks to algorithms, but some effectiveness assumptions are required in order to make the computations possible. We do not give the details here however in [ABM07], the authors give classes in which the algorithms are solvable.

We consider now the quantitative model-checking problems for reachability and repeated reachability properties. Given a set of states B and an initial state s_0 , the objective is to approximate the exact values of $\text{Prob}_{s_0}^{\mathcal{M}}(\mathbf{F} B)$ and $\text{Prob}_{s_0}^{\mathcal{M}}(\mathbf{G} \mathbf{F} B)$. In [ABM07], the authors provide algorithms that approximate the probability of reachability and repeated reachability properties, and prove that those are correct and terminate under some decisiveness assumptions. We do not recall them here, however we will give the intuition of those algorithms.

For reachability properties, the algorithm takes as input the DMC \mathcal{M} , an initial state s_0 , a set of states B and a rational $\varepsilon > 0$ and computes a value θ such that $\theta \leq \text{Prob}_{s_0}^{\mathcal{M}}(\mathbf{F} B) < \theta + \varepsilon$. The algorithm runs as follows: the DMC is unfolded from the initial state and the algorithm check whether in a given state s : (a) $s \in B$, or (b) $s \in \widetilde{B}$, or (c) $s \notin B$ and $s \notin \widetilde{B}$. In case (a) the objective is reached. The probability of the run leading to s from s_0 is computed and added to some variable Yes , the unfolding is stopped there. In case (b) it is known

that B cannot be reached anymore and so the probability of the run leading to s from s_0 is computed and added to some other variable No , and the unfolding is stopped. Finally in case (c), the unfolding is continues. Observe that along the algorithm, it is always the case that $Yes \leq \text{Prob}_{s_0}^M(\mathbf{F} B) \leq 1 - No$. The algorithm runs until the difference between Yes and $1 - No$ is less than ε and it returns Yes .

The algorithm for repeated reachability properties is very similar except that Yes is incremented in the case where $s \in \widetilde{B}$: Proposition 2.2.17 states that under a strong decisiveness assumption, if $s \in \widetilde{B}$ then $\text{Prob}_s^M(\mathbf{G} \mathbf{F} B) = 1$. The condition to increment No does not change. And the condition to stop the algorithm is the same, the algorithm returns again Yes .

We can then summarize the results of [ABM07] as follows.

Proposition 2.2.20 (Lemmas 7.1 and 7.2 of [ABM07]). *Given an initial state $s_0 \in S$ and $\varepsilon > 0$ there exists an algorithm that computes θ such that $\theta \leq \text{Prob}_{s_0}^M(\mathbf{F} B) \leq \theta + \varepsilon$. Moreover if \mathcal{M} is decisive w.r.t. B , then the algorithm terminates.*

Proposition 2.2.21 (Lemmas 8.1 and 8.2 of [ABM07]). *Given an initial state $s_0 \in S$ and $\varepsilon > 0$, if \mathcal{M} is strongly decisive w.r.t. B , there exists an algorithm that computes θ such that $\theta \leq \text{Prob}_{s_0}^M(\mathbf{G} \mathbf{F} B) \leq \theta + \varepsilon$. Moreover if \mathcal{M} is decisive w.r.t. \widetilde{B} , then the algorithm terminates.*

In order to be able to apply those algorithms, it has to be able to check whether a state is in \widetilde{B} or in $\widetilde{\widetilde{B}}$. The authors of [ABM07] thus need some effectiveness assumptions. They provide also several classes in which the algorithms can be applied.

The objectives of Chapters 4, 5 and 6 will be to extend those notions and results to a richer class of probabilistic transition systems and to richer properties. And we will then see in Chapter 7 how it can be applied to our model of interest, that mixes probabilistic and timed aspects.

2.3. Continuous-time Markov chains

In this section, we briefly introduce the continuous-time Markov chain model (CTMC for short) [Fel69], which can be seen as a first step in a mix between timed systems and probabilistic behaviours. We will see in Chapter 3 how this is done for the STA model.

As defined in Section 2.2, a DMC is a probabilistic transition system over a denumerable set of states. This gives rise to a set of runs defined by the set of

infinite sequences of states such that there is a non null probability between two consecutive states. Note that here, we observe the state of the execution only at discrete time: if $\rho = (s_n)_{n \in \mathbb{N}}$ is a run of DMC \mathcal{M} , s_n corresponds to the state of the system at step n . Hence the time is somehow discrete.

In CTMCs, we assume that time evolves continuously and that the elapse of time follows exponential distributions. Hence here, we will as well be interested in the states that are visited along the executions, but also in the time it takes to move from one state to another one. The model enjoys also nice decidability results. We cite for instance [BHHK03] in which the authors proved that CTMCs are decidable for the model-checking problem of CSL properties.

We define here the syntax of the model, we will then briefly explain the semantics of this continuous model and then we will illustrate it on a simple example.

Definition 2.3.1. A *continuous-time Markov chain* (CTMC for short) is a couple $\mathcal{M} = (S, R)$ where S is a finite set of states and $R : S \times S \rightarrow \mathbb{R}_+$ is the rate function transition that satisfies: for each $s \in S$, there is $s' \in S$ such that $R(s, s') > 0$.

We follow here the definition of [BHHK03] where self-loops are possible: for each state s , we allow $R(s, s) > 0$ while sometimes, you will find $R(s, s) = -\sum_{s' \neq s} R(s, s')$.


Like in DMCs (see Section 2.2), a CTMC induces a transition system (see Definition 2.1.4) where the set of states is S , the alphabet Γ is given by $\{R(s, s') \mid s, s' \in S\}$ and there is a transition $s \rightarrow s'$ if $R(s, s') > 0$, which is then labelled with $R(s, s')$. A CTMC can thus also be represented by its transition system. We give now a simple example that we will use in the rest of the section in order to describe the semantics of CTMCs.

Example 2.3.2. We consider the CTMC $\mathcal{M} = (S, R)$ described on Figure 2.11, where $S = \{0, \dots, n\}$ and where the rate function R is defined as follows:

- $R(0, 1) = \lambda$, $R(n, n - 1) = \nu$, and
- for each $0 < i < n$, $R(i, i + 1) = \lambda$ and $R(i, i - 1) = \nu$,

where $\lambda, \nu > 0$. This represents a queuing system with a single queue, parameters λ for arrivals and ν for serving times, and a maximum of n tasks in the queue. We will describe the behaviour of this system in Example 2.3.3.

We now explain the semantics of CTMCs. We fix a CTMC $\mathcal{M} = (S, R)$. For each $s \in S$, we write $R(s) = \sum_{s' \in S} R(s, s') > 0$ from Definition 2.3.1. As

Figure 2.11: A queuing system with a maximum of n tasks.

explained previously, we allow time to evolve continuously. Then starting from a given state s , we are interested in the amount of time the system stays in s before performing some transition. Intuitively, $R(s, s')$ corresponds to the rate at which transition $s \rightarrow s'$ is taken, and thus $R(s)$ is the rate at which a transition is taken from s : this will describe how we can measure the time spent in state s .

Given states s and s' such that $R(s, s') > 0$, we assume that the time needed before performing transition $s \rightarrow s'$ follows an exponential distribution of parameter $R(s, s')$ (written $\text{Exp}(R(s, s'))^1$). More precisely, the probability to perform transition $s \rightarrow s'$ within t time units is given by $1 - e^{-R(s, s')t}$.

Then given a state s , if there are several states s' such that $R(s, s') > 0$, we assume that there is a competition between the outgoing transitions from s . It is known as the *race condition*. This is performed as follows: for each state s' with $R(s, s') > 0$, pick a time $t_{s'}$ according to the exponential distribution $\text{Exp}(R(s, s'))$. Then, the time after which we leave state s , *i.e.* we perform a transition, is given by the minimum of those times: $\min(\{t_{s'} \mid R(s, s') > 0\})$. It follows that the time of sojourn in state s follows the distribution defined as the minimum of all exponential distributions $\text{Exp}(R(s, s'))$ with $R(s, s') > 0$. This corresponds to the distribution² $\text{Exp}(R(s))$. Finally, the probability that transition $s \rightarrow s'$ wins the race within t time units, corresponds to the value:

$$P(s, s', t) = \frac{R(s, s')}{R(s)} \cdot (1 - e^{-R(s)t}).$$

We will abusively write $P(s, s', \infty)$ for the probability that transition $s \rightarrow s'$ wins the race at some point in the future. It is equal to $R(s, s')/R(s)$. We briefly illustrate this on Example 2.3.2.

Example 2.3.3. We consider again the CTMC \mathcal{M} of Example 2.3.2 represented on Figure 2.11. As stated before, this corresponds to a queuing system. We start

¹Recall that distribution $\text{Exp}(\mu)$ is given by the density function $f(t) = \mu e^{-\mu t}$, *i.e.* for each Borel set A , $(\text{Exp}(\mu))(A) = \int_A \mu e^{-\mu t} dt$.

²The minimum of two exponential distributions of parameters $\mu_1, \mu_2 > 0$ is the exponential distribution of parameter $\mu_1 + \mu_2$.

at state 0 with 0 task in the queue. Then, the probability that the first task arrives in the queue within t time units is given by $1 - e^{-\lambda t}$. Once there is one task in the queue, two actions are possible: either a new task can arrive in the queue ($1 \rightarrow 2$) or the task can be done (and thus be removed of the queue: $1 \rightarrow 0$). In any case, the probability that something happens within t time units is given by

$$1 - e^{-(\lambda+\nu)t}.$$

It should be observed that with probability 1, something will eventually happen. Then the condition race described above allows to establish that the probability that a new task arrives within t time units, is given by

$$P(1, 2, t) = \frac{\lambda}{\lambda + \nu} \cdot (1 - e^{-(\lambda+\nu)t})$$

and the probability that the task is done within t time units, is given by

$$P(1, 2, t) = \frac{\nu}{\lambda + \nu} \cdot (1 - e^{-(\lambda+\nu)t}).$$

This behaviour is repeated in each state of the CTMC, until state n where it is only possible to remove a task from the queue.

We terminate this section with a remark on the link between CTMCs and DMCs.

As already said previously, this model allows to measure the time spent in each state. However, one can abstract time in order to only consider the jumps in the system, *i.e.* the moments at which a transition is performed. This leads to the construction of a finite Markov chain which is, in some sense, equivalent to a given CTMC. This will be central in Part I when we will consider more general stochastic processes (Chapter 4) where it will be easier to check some properties on simpler models (Chapters 5 and 6) and it is comparable to the region graph for timed automata (see Section 2.1.1).

Given a CTMC $\mathcal{M} = (S, R)$, one can build the DMC $\mathcal{M}^* = (S, P)$ where for each $s, s' \in S$,

$$P(s, s') = \frac{R(s, s')}{R(s)}.$$

We say that \mathcal{M} and \mathcal{M}^* are equivalent in the sense that for each transition $s \rightarrow s'$,

$$P(s, s', \infty) = P(s, s').$$

A consequence of this equivalence is that the probability of LTL properties is the same in \mathcal{M} and \mathcal{M}^* . Therefore, the results on the qualitative and quantitative

model-checking problems of LTL properties for finite Markov chains of [Var85] and [CY88] can also be applied to CTMCs. Observe thus that LTL properties cannot express bounds on the time of sojourn in the different states. The logic CSL however, allows the expression of time-bounded properties. In [BHHK03], the authors showed CTMCs to be decidable for the qualitative and quantitative model-checking problem of CSL properties.

2.4. Composition and interactive Markov chains

In this section, we briefly introduce the discussion on composition in [HZ11]. We will see that concurrent systems can be composed with full synchronisations (DMCs), interleaving (CTMCs) or handshaking. The latter case will reveal to be the most interesting and will lead to the interactive Markov chains model [Her02]. This will be the subject of Section 2.4.2.

2.4.1 Composition of general transitions systems and application to DMCs and CTMCs

In this section, we follow [HZ11] in order to give an overview on composition in transition systems.

We fix for the section two transitions systems $T_1 = (Q_1, \Gamma_1, \rightarrow_1)$ and $T_2 = (Q_2, \Gamma_2, \rightarrow_2)$. The way we want to compose T_1 and T_2 , depends on how the real systems they depict, behave. We give here the three types of composition that are discussed in [HZ11]. Still following [HZ11], we explain which variants suit better for DMCs and CTMCs.

Firstly, requiring that $\Gamma_1 = \Gamma_2 = \{a\}$, composition can be done with *full synchronisation*. This gives rise to the product $T_1 \parallel_{fs} T_2 = (Q_1 \times Q_2, \{a\}, \rightarrow)$ where \rightarrow is defined as follows: for any states $q_1 \in Q_1$ and $q_2 \in Q_2$, if $q_1 \xrightarrow{a}_1 q'_1$ and $q_2 \xrightarrow{a}_2 q'_2$, then $(q_1, q_2) \xrightarrow{a} (q'_1, q'_2)$. Here it is assumed that the two transitions are executed at the same time. The unique label here, is to make the synchronisation plain. However, one could assume that $\Gamma_1 \neq \Gamma_2$ and $T_1 \parallel_{fs} T_2 = (Q_1 \times Q_2, \Gamma, \rightarrow)$ would be defined as follows: Γ is an alphabet that suits the best for $T_1 \parallel_{fs} T_2$ and for any states $q_1 \in Q_1$ and $q_2 \in Q_2$, $q_1 \xrightarrow{a}_1 q'_1$ and $q_2 \xrightarrow{b}_2 q'_2$ imply $(q_1, q_2) \xrightarrow{c} (q'_1, q'_2)$ for some $c \in \Gamma$.

Secondly, one could assume that T_1 and T_2 run independently from each other, hence the product should describe all possible transitions: $T_1 \parallel T_2 = (Q_1 \times Q_2, \Gamma_1 \cup \Gamma_2, \rightarrow)$ where \rightarrow is defined as follows: if $q_1 \xrightarrow{a}_1 q'_1$ then $(q_1, q_2) \xrightarrow{a} (q'_1, q_2)$ and if $q_2 \xrightarrow{a}_2 q'_2$ then $(q_1, q_2) \xrightarrow{a} (q_1, q'_2)$. Here, each transition of the product corresponds thus to one transition of one of the two systems. We do not

consider the case where the two systems could, on a coincidence, evolve towards a new state at the same time. This is the *interleaving* product. In Section 2.1.3, we chose this type in order to define a first composition operator in timed automata (see Definition 2.1.29).

Lastly, we consider the *handshaking* composition. Both previous cases are limited as they do not allow a mix between independence and communication: two systems could be running independently for a while, but exchanging information at some points. The handshaking composition treats this case. Here we assume to have a set of labels $A \subseteq \Gamma_1 \cap \Gamma_2 \neq \emptyset$, on which the systems will be communicating. This leads to the parallel composition $T_1 \parallel_A T_2 = (Q_1 \times Q_2, \Gamma_1 \cup \Gamma_2, \rightarrow)$ where \rightarrow is defined as follows: if $a \in A$, the composition is synchronised as in the first case; if $a \notin A$, then the composition is interleaving as in the second case. Formally, for each $a \in \Gamma_1 \cup \Gamma_2$, for each $q_1 \in Q_1$ and $q_2 \in Q_2$,

- if $a \in A$, $q_1 \xrightarrow{a}_1 q'_1$ and $q_2 \xrightarrow{a}_2 q'_2$, then $(q_1, q_2) \xrightarrow{a} (q'_1, q'_2)$,
- if $a \notin A$ and $q_1 \xrightarrow{a}_1 q'_1$, then $(q_1, q_2) \xrightarrow{a} (q'_1, q_2)$, and
- if $a \notin A$ and $q_2 \xrightarrow{a}_2 q'_2$, then $(q_1, q_2) \xrightarrow{a} (q_1, q'_2)$.

This handshaking composition encompass all three types: if $A = \Gamma_1 = \Gamma_2$, then we get the composition with full synchronisation; if $A = \emptyset$, we get the interleaving composition! We also call the handshaking composition by (parallel) composition *with synchronisations*. In timed automata, the composition with synchronisations of Definition 2.1.33 corresponds to this case.

We now briefly explain which types of composition suit better for DMCs and CTMCs (see [HZ11]).

DMCs. We fix two DMCs $\mathcal{M}_1 = (S_1, P_1)$ and $\mathcal{M}_2 = (S_2, P_2)$ and we recall that each DMC induces a transition system in which $s_1 \xrightarrow{p_1}_1 s'_1$ if and only if $P_1(s_1, s'_1) = p_1 > 0$ and similarly, $s_2 \xrightarrow{p_2}_2 s'_2$ if and only if $P_2(s_2, s'_2) = p_2 > 0$. It has been shown in [HZ11] that here, composition with fully synchronisation makes sense: $\mathcal{M}_1 \parallel_{fs} \mathcal{M}_2 = (S_1 \times S_2, P)$ where $P((s_1, s_2), (s'_1, s'_2)) = P_1(s_1, s'_1) \cdot P_2(s_2, s'_2)$, *i.e.* if $s_1 \xrightarrow{p}_1 s'_1$ and $s_2 \xrightarrow{q}_2 s'_2$, then $(s_1, s_2) \xrightarrow{pq} (s'_1, s'_2)$. This can be explained through the product of two independent stochastic processes.

CTMCs. We now consider two CTMCs $\mathcal{M}_1 = (S_1, R_1)$ and $\mathcal{M}_2 = (S_2, R_2)$ and we recall that each CTMC induces again a transition system in which $s_1 \xrightarrow{\lambda_1}_1 s'_1$ if and only if $R(s_1, s'_1) = \lambda_1 > 0$, and similarly $(s_2 \xrightarrow{\lambda_2}_2 s'_2$ if and only if $R(s_2, s'_2) =$

$\lambda_2 > 0$. It has again been shown in [HZ11] that here, interleaving composition makes sense: $\mathcal{M}_1 \parallel \mathcal{M}_2 = (Q_1 \times Q_2, R)$ where: $R_1(s_1, s'_1) > 0$ implies $R((s_1, s_2)(s'_1, s_2)) = R_1(s_1, s'_1)$, and $R_2(s_2, s'_2) > 0$ implies $R((s_1, s_2), (s_1, s'_2)) = R_2(s_2, s'_2)$. This again can be observed through the product of two independent stochastic processes and comes from the nice memoryless properties of exponential distributions.

2.4.2 Interactive Markov chains

In this section, we present the Interactive Markov chain model (IMC for short) [Her02]. They were first introduced as process algebras. As it is not the purpose of this thesis, we will introduce the model as defined in [HK09] which presents the model through transition systems.

IMCs can, roughly speaking, be seen as CTMCs in which we add non-determinism through some non-probabilistic transitions, that we call *interactive transitions*. As explained in [HZ11], those interactive transitions will allow to add some communication between the systems when composing them. In other words this will allow to extend the interleaving composition of CTMCs into a handshaking composition which is more interesting (see Section 2.4.1). The compositional framework for verification that comes with IMCs is very nice: for instance in [HK09], the authors provide techniques in order to approximate time-bounded reachability properties in this setting with non-determinism.

Here, we will define the model and define the notion of composition. We will then explain the semantics of the model. Finally, we will present a notion of bisimulation that will be a congruence w.r.t. composition, which is important when dealing with composition (as already discussed in Section 2.1.3).

Definition 2.4.1. An *interactive Markov chain* is a tuple $\mathcal{M} = (S, \Gamma^\tau, R, \twoheadrightarrow)$ where

- S is a finite set of states,
- Γ is a set of actions, τ is an internal action and $\Gamma^\tau = \Gamma \cup \{\tau\}$,
- $R : S \times S \rightarrow \mathbb{R}_+$ is rate function transition, and
- $\twoheadrightarrow \subseteq S \times \Gamma^\tau \times S$ is a set of interactive transitions.

This again induces a transition system (see Definition 2.1.4), where the set of states is S , the alphabet is given by $\Gamma^\tau \cup \{R(s, s') \mid s, s' \in S\}$ and there is a transition $s \rightarrow s'$ if

- $R(s, s') > 0$, in which case the transition is labelled with $R(s, s')$, or

- there exists $a \in \Gamma^\tau$ such that $(s, a, s') \in \rightarrow$, in which case it is labelled with a .

In the sequel, we will write $s \xrightarrow{a} s'$ instead of $(s, a, s') \in \rightarrow$. The transitions generated by the rate function are called *Markovian transition* while the transitions generated by \rightarrow are called *interactive transitions*. The internal action τ is here in order to have a distinction with the actions of Γ . Internal actions cannot be seen by other systems, hence one should not be able to synchronise on them. We will come back to that later. We will call the actions of Γ external.

We illustrate this on two simple examples.

Example 2.4.2. We consider two simple IMCs $\mathcal{M}_1 = (S_1, \Gamma^\tau, R_1, \rightarrow_1)$ and $\mathcal{M}_2 = (S_2, \Gamma^\tau, R_2, \rightarrow_2)$ where $\Gamma = \{a\}$, $S_1 = \{s_1, s_2\}$, $S_2 = \{s'_1, s'_2, s'_3\}$ and the rate functions and the interactive transitions are described on Figure 2.12 (λ , μ and ν are all positive and non null rates).




Figure 2.12: Two simple IMCs \mathcal{M}_1 (on the left) and \mathcal{M}_2 (on the right).

The internal action τ has a special role. Interactive transitions labelled with τ are not subject to interaction as they depict internal actions which can thus not be seen by other systems. Following [HK09], we will make an essential assumption: we suppose that in any IMC, internal interactive transitions take precedence over Markovian transitions. This is called the *maximal progress assumption*. It in fact assumes that τ -labelled transitions take place immediately. This assumption will be essential for the rest of this section.

We are not yet ready to give the semantics of an IMC. Indeed, as said before, interactive transitions with external actions wait for some other system(s) to synchronise on them. Hence, while the synchronisation cannot be made, the system should not be able to run (or otherwise should be blocked once it waits for a synchronisation that will never happen). Hence we need first to define an operator of composition. Then, once actions do not need to be synchronised anymore, we will abstract those actions into the internal action τ . This will be done with a hiding operator. Finally, we will say that a system is complete, when all external actions have been hidden, *i.e.* when all synchronisations needed can be done.

Definition 2.4.3. Let $\mathcal{M}_1 = (S_1, \Gamma_1^\tau, R_1, \rightarrow_1)$ and $\mathcal{M}_2 = (S_2, \Gamma_2^\tau, R_2, \rightarrow_2)$ be two IMCs and fix $A \subseteq \Gamma_1 \cap \Gamma_2$. The parallel composition of \mathcal{M}_1 and \mathcal{M}_2 on A is the IMC $\mathcal{M}_1 \parallel_A \mathcal{M}_2 = (S, \Gamma^\tau, R, \rightarrow)$ where:

- $S = S_1 \times S_2$ and $\Gamma = \Gamma_1 \cup \Gamma_2$,
- the rate function R is defined as follows: for each states $s = (s_1, s_2)$ and $s' = (s'_1, s'_2) \in S$,
 - ▷ $R(s, s') = R_1(s_1, s'_1) + R_2(s_2, s'_2)$ if $s = s'$,
 - ▷ $R(s, s') = R_1(s_1, s'_1)$ if $s \neq s'$ and $s_2 = s'_2$,
 - ▷ $R(s, s') = R_2(s_2, s'_2)$ if $s \neq s'$ and $s_1 = s'_1$,
 - ▷ $R(s, s') = 0$ otherwise;
- \rightarrow is defined as follows:
 - ▷ if $a \in A$, $s_1 \xrightarrow{a}_1 s'_1$ and $s_2 \xrightarrow{a}_2 s'_2$, then $(s_1, s_2) \xrightarrow{a} (s'_1, s'_2)$,
 - ▷ if $a \notin A$ and $s_1 \xrightarrow{a}_1 s'_1$, then for each $s_2 \in S_2$, $(s_1, s_2) \xrightarrow{a} (s'_1, s_2)$,
 - ▷ if $a \notin A$ and $s_2 \xrightarrow{a}_2 s'_2$, then for each $s_1 \in S_1$, $(s_1, s_2) \xrightarrow{a} (s_1, s'_2)$.


One can see from the definition that the parallel composition is interleaving on the rates just like in CTMCs. The definition of the new rate function R comes from the memoryless property of exponential distributions, again like in CTMCs. We will be more precise in Chapter 9, when tackling the definition of a composition operator in stochastic timed automata. Observe that the first item of the definition of R deals with self-loops: if $R_1(s_1, s_1) = \mu_1$ and $R_2(s_2, s_2) = \mu_2$, then we have a single self-loop in $\mathcal{M}_1 \parallel_A \mathcal{M}_2$ corresponding to $R((s_1, s_2), (s_1, s_2)) = \mu_1 + \mu_2$, instead of two distinct self-loops. Observe also that it is not allowed to synchronise on τ . We illustrate this definition on Example 2.4.2.

Example 2.4.4. We consider IMCs \mathcal{M}_1 and \mathcal{M}_2 of Example 2.4.2 depicted on Figure 2.12. Then the parallel composition of \mathcal{M}_1 and \mathcal{M}_2 w.r.t. $\{a\}$, i.e. $\mathcal{M}_1 \parallel_{\{a\}} \mathcal{M}_2$, is the IMC described on Figure 2.13.

We now define the hiding operator of a set of external actions on an IMC.

Definition 2.4.5. Let $\mathcal{M} = (S, \Gamma^\tau, R, \rightarrow)$ be an IMC and fix $A \subseteq \Gamma$. The hiding of \mathcal{M} w.r.t. A is the IMC $\mathcal{M} \setminus A = (S, \Gamma^\tau \setminus A, R, \rightarrow_*)$ where \rightarrow_* is defined as follows:

- if $a \in A$ and $s \xrightarrow{a} s'$, then $s \xrightarrow{\tau}_* s'$, and

Figure 2.13: IMC $\mathcal{M}_1 \parallel_1 \mathcal{M}_2$.

- if $a \notin A$ and $s \xrightarrow{a} s'$, then $s \xrightarrow{a}_* s'$.

Example 2.4.6. If we consider the IMC $\mathcal{M}_1 \parallel_{\{a\}} \mathcal{M}_2$ of Figure 2.13 in Example 2.4.4, the hiding of this IMC w.r.t. $\{a\}$ is the IMC $(\mathcal{M}_1 \parallel_{\{a\}} \mathcal{M}_2) \setminus \{a\}$ which is exactly the same as the initial one described in Figure 2.13, except that the interactive transition $(s_2, s'_1) \xrightarrow{a} (s_1, s'_2)$ becomes $(s_2, s'_1) \xrightarrow{\tau}_* (s_1, s'_2)$.

Before going further on the semantics that we can induce from the previous notions, we define a notion of bisimulation which will be an inspiration when defining such a notion in stochastic timed automata. We refer to Section 2.1.3 for a short discussion on bisimulations. Given a relation $\mathcal{R} \subseteq S \times S$, we write $s\mathcal{R}s'$ for $(s, s') \in \mathcal{R}$. Given an IMC \mathcal{M} , $s \in S$ and $C \subseteq S$, we write $R(s, C) = \sum_{s' \in C} R(s, C)$.

Definition 2.4.7. Given an IMC $\mathcal{M} = (S, \Gamma^\tau, R, \rightarrow)$, \mathcal{R} is a *bisimulation* for \mathcal{M} if it is an equivalence relation on S and if for each couple of states $s\mathcal{R}s'$ and for each equivalence class $C \in \mathcal{R}/S$,

- for each $a \in \Gamma^\tau$, if there is $s_1 \in S$ such that $s \xrightarrow{a} s_1$, then there is $s'_1 \in S$ such that $s' \xrightarrow{a} s'_1$ and $s_1\mathcal{R}s'_1$;
- if for all $s_1 \in S$, we do not have $s \xrightarrow{\tau} s_1$, then $R(s, C) = R(s', C)$.

Observe that since \mathcal{R} is an equivalence relation, the first item holds in both senses and that there are no interactive edges labelled with τ from s if and only if there are no such edges from s' . The fact that the equality between the rate functions holds true only when no internal interactive transitions are enabled, comes from the maximal progress assumption. We say that two states s and s' are *bisimilar* if there exists a bisimulation \mathcal{R} for \mathcal{M} such that $s\mathcal{R}s'$.

As explained in Section 2.1.3, bisimulation can be extended to a relation between IMCs in a standard way. We say that \mathcal{M}_1 and \mathcal{M}_2 are *bisimilar* if there exists a bisimulation between \mathcal{M}_1 and \mathcal{M}_2 . We write it $\mathcal{M}_1 \sim \mathcal{M}_2$. Again as said previously, we will be more precise on the subject when defining composition in stochastic timed automata (see Chapters 9 and 10).

We get then get the next result (see Section 2.1.3 for a short explanation on why congruence is important in composition) [Her02].

Theorem 2.4.8. *Bisimilarity is a congruence w.r.t. composition and hiding. More precisely: fix $\mathcal{M}_1 = (S_1, \Gamma_1^\tau, R_1, \rightarrow_1)$, $\mathcal{M}_2 = (S_2, \Gamma_2^\tau, R_2, \rightarrow_2)$ and $\mathcal{M} = (S, \Gamma^\tau, R, \rightarrow)$ three IMCs, fix $A \subseteq \Gamma_1 \cap \Gamma_2 \cap \Gamma$ and fix $A' \subseteq \Gamma_1 \cap \Gamma_2$, then we have that*

- $\mathcal{M}_1 \sim \mathcal{M}_2$ implies $\mathcal{M}_1 \parallel_A \mathcal{M} \sim \mathcal{M}_2 \parallel_A \mathcal{M}$,
- $\mathcal{M}_1 \sim \mathcal{M}_2$ implies $\mathcal{M} \parallel_A \mathcal{M}_1 \sim \mathcal{M} \parallel_A \mathcal{M}_2$, and
- $\mathcal{M}_1 \sim \mathcal{M}_2$ implies $\mathcal{M}_1 \setminus A' \sim \mathcal{M}_2 \setminus A'$.

An IMC \mathcal{M} is said *complete* whenever all interactive transitions are internal. In particular, once we have composed multiple IMCs leading to an IMC that do not require any more interactions, then we can hide all actions leading to a complete IMC. Concretely, assume that we have $\mathcal{M}_i = (S_i, \Gamma_i^\tau, R_i, \rightarrow_i)$ for $i \in \{1, \dots, n\}$ and that the system

$$\mathcal{M}_1 \parallel_{A_1} \mathcal{M}_2 \parallel_{A_2} \dots \parallel_{A_{n-1}} \mathcal{M}_n$$

does not require anymore interactions, where for each $i \in \{1, \dots, n-1\}$, $A_i \subseteq \Gamma_i \cap \Gamma_{i+1}$. We then consider the IMC

$$\left(\mathcal{M}_1 \parallel_{A_1} \mathcal{M}_2 \parallel_{A_2} \dots \parallel_{A_{n-1}} \mathcal{M}_n \right) \setminus A,$$

where $A = \bigcup_{i=1}^n \Gamma_i$ which is thus a complete IMC. We can give a semantics on complete IMCs. In state s , if there are outgoing interactive transitions (which are thus internal), one is immediately chosen non-deterministically: it corresponds to an interaction, and the assumption is made that once an interaction is possible, it is performed immediately (the maximal progress assumption). Note that non-determinism can be handled with a scheduler. Otherwise, if no interactive transition starts from s , then in this state, the behaviour is the same as in a CTMC (see Section 2.3, the race condition).

The resulting complete IMC can thus be reduced to a simpler system. In each state s that has outgoing interactive transitions, all Markovian transitions

can be removed. Moreover in [HK09], the authors are also involved in a notion of *weak bisimulation*³ that allows to collapse sequences of τ -transitions into one. In the case where the complete IMC is determinist, then it corresponds to a CTMC and all known techniques on the model can be applied (see Section 2.3 for a short word on the matter). If non-determinism remains, then new techniques have to be developed. For instance in [HK09], the authors provide techniques in order to approximate the probability of time-bounded reachability properties.

³Weak bisimulation corresponds to Definition 2.4.7 except that, roughly speaking, it allows for several successive τ -interactive transitions to occur regardless the exact number of such transitions.

CHAPTER 3

Stochastic Timed Automata

In this chapter, we introduce the notion of stochastic timed automaton (STA for short) as defined in [BBB⁺14] and also describe some results of interest to us.

In Section 3.1, we define and illustrate the STA model as in [BBB⁺07] and [BBB⁺14]. STA are a probabilistic extension of timed automata: we equip each state of a timed automaton with distributions over the delays and over the edges. It can also be seen as an extension of the CTMC model: here, it is allowed to have other kind of distributions than exponential distributions (including distributions on bounded intervals like uniform distributions), and like in the timed automaton model, time constraints can be put on the edges.

In Section 3.2, we refine the notion of region graph of Section 2.1.1 into the thick region graph [BBB⁺14]. We then explain how to construct a finite Markov chain from this thick region graph. The objective is to have a finite abstraction on which the analysis of LTL properties can be done, instead of the STA. However we exhibit an example with a bad behaviour.

Finally in Section 3.3, we define the notion of fairness which was identified in [BBB⁺14] as a condition to ensure the finite Markov chain abstraction to be in some sense, equivalent to the initial STA. We then present the decidability results of [BBB⁺14]: they provide to classes of STA for which the qualitative model-checking problem of LTL properties is decidable.

3.1. Definition and illustration of the model

In this section, we define the notion of STA [BBB⁺14] and illustrate it.

Before defining the model, we recall some notations of Section 2.1. The STA

model is a probabilistic extension of timed automata: it equips timed automata with probability distributions over the edges and the delays and this, in each state of the timed automaton. In order to do this, we need to define the set of delays in a given state.

Fix a timed automaton $\mathcal{A} = (L, X, Act, E, Inv, AP, \mathcal{L})$ and recall that it has an associated transition system $T_{\mathcal{A}} = (Q, \mathbb{R}_+ \times E, \rightarrow)$ (Definition 2.1.11 in Section 2.1). Given a state $q = (l, \nu) \in Q$ and an edge $e \in E$, recall that we write $I(q, e) = \{t \in \mathbb{R}_+ \mid \nu + t \models Inv(l) \text{ and } \exists q' \in Q \text{ s.t. } q \xrightarrow{t, e} q'\}$ (corresponding to the set of delays after which, starting from q , edge e is enabled) and that we write $I(q) = \bigcup_{e \in E} I(q, e)$, *i.e.* the set of delays after which, starting from q , an edge is enabled. This is the set of delays needed in order to define the STA model.

Remark 3.1.1. Observe that it is easily shown that for each state q and for each edge e , $I(q, e)$ is either the empty set, or a single point, or an interval of real positive numbers (which could be unbounded from above). This is due to the form of the guards on the edges and of the invariant function on the locations. Then, it holds that $I(q)$ is either the empty set or a finite union of single points or intervals.

We can now define the notion of STA.

Definition 3.1.2. A *stochastic timed automaton* (STA for short) is a tuple $\mathcal{A} = (L, X, Act, E, Inv, AP, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X})$ where

- (i) $(L, X, Act, E, Inv, AP, \mathcal{L})$ is a timed automaton,
- (ii) for each $q \in Q$, μ_q is a probability distribution over \mathbb{R}_+ such that $\mu_q(I(q)) = 1$ and $\mu_q(I(q)^c) = 0$, and
- (iii) for each $q = (l, \nu) \in Q$, p_q is a probability distribution over the set of edges that are enabled in q , *i.e.* $\{e = (l', a, g, Y, l'') \in E \mid l' = l \text{ and } \nu \models g\}$; it assigns 0 to each edge if there are none enabled in q .

Remark 3.1.3. We recall that from Remark 2.1.17, we assume that all timed automata are non-blocking, *i.e.* $I(q) \neq \emptyset$ for each state q . This ensures us that the probability μ_q and p_q are well-defined in each state q .

Let us illustrate the STA model on Example 2.1.2.

Example 3.1.4. Let us consider the timed automaton \mathcal{A} of Figure 2.1. We recall that it is non-blocking (see Example 2.1.18). In order to extend it into a STA, we need to equip it with probability distributions over delays and edges in each state of the timed automaton. We do not define μ_q and p_q for each state

q , but we give some examples. If $q = (l_0, (0, 0))$ then observe that no edge is enabled. Hence p_q assigns 0 to each edge. Now recall that: $I(q, e_1) =]0, 1]$, $I(q, e_2) = \{2\}$ and thus $I(q) =]0, 1] \cup \{2\}$ (see again Example 2.1.18). Hence examples of distribution on $I(q)$ are the uniform distribution on the interval $]0, 1]$ written $\mathcal{U}(]0, 1])$, or a combination between $\mathcal{U}(]0, 1])$ and the Dirac distribution on $\{2\}$ written δ_2 , for instance $\frac{1}{2}\mathcal{U}(]0, 1]) + \frac{1}{2}\delta_2$ meaning that the single point $\{2\}$ has probability $\frac{1}{2}$ and the interval $]0, 1]$ has probability $\frac{1}{2}$ and follows a uniform distribution. Consider now $q' = (l_0, (1.5, 3.5))$ then e_2 is the only enabled edge. The probability distribution $p_{q'}$ should thus assign probability 1 to e_2 and 0 to each other edge. We can then easily show that $I(q') = [0, 0.5]$, as the invariant is violated once the elapse of 0.5 time units is exceeded. A classical distribution on $I(q')$ is thus $\mu_{q'} = \mathcal{U}([0, 0.5])$. A final example is $q'' = (l_0, (0.5, 3))$ where this time, both edges e_1 and e_2 are enabled. An example of distribution is

$$p_{q''}(e) = \begin{cases} \frac{1}{2} & \text{if } e \in \{e_1, e_2\} \\ 0 & \text{otherwise.} \end{cases}$$

One can easily show that $I(q'') = [0, 1]$ and thus again, $\mu_{q''} = \mathcal{U}([0, 1])$ is an example of distribution for the delays.

Remark 3.1.5. It should be noted that each CTMC $\mathcal{M} = (S, R)$ (see Section 2.3) can be seen as a single-clock STA where the set of locations is S , each edge is guarded with `true` and resets the unique clock to 0, the distribution over the delays in location s is given by the exponential distribution $\text{Exp}(R(s))$ and the probability of the edge $e = (s, \text{true}, X, s')$ is given by $p_{(s, \nu) + t}(e) = R(s, s')/R(s)$ for each $\nu \in \mathbb{R}_+$ and $t \in \mathbb{R}_+$.

In the rest of this chapter, actions over the edges and the labelling function on the locations over AP will not be needed, we thus omit them (except for some general definitions).

Like in Section 2.2, the purpose of such a probabilistic model is to define a probability distribution over the set of runs. Given a STA \mathcal{A} , we will abusively write \mathcal{A} for its underlying timed automaton. Then observe that all notions seen in Section 2.1 have also sense for STA. Fix a STA $\mathcal{A} = (L, X, E, \text{Inv}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X})$. Recall that as stated in Section 2.1, we are only interested in infinite runs and that given an initial state q_0 , we write $\text{Runs}(\mathcal{A}, q_0)$ for the set of infinite runs starting in q_0 and $\text{Runs}_f(\mathcal{A}, q_0)$ for the set of finite runs starting in q_0 (see Definitions 2.1.12 and 2.1.13 for the notion of runs). We write $\text{Runs}(\mathcal{A})$ for the set of all infinite runs.

The aim is thus to define a probability measure over the set of infinite runs $\text{Runs}(\mathcal{A}, q)$ for every state q . Fix a state $q \in Q$, one should hence equip

$\text{Runs}(\mathcal{A}, q)$ with a σ -algebra. In order to define this σ -algebra, we first introduce the notions of *symbolic paths*, *constrained symbolic paths* and *cylinder* generated by a (constrained) symbolic path. Note that we consider that \mathbb{R}_+ is equipped with the Borel σ -algebra, written $\mathcal{B}(\mathbb{R}_+)$. Similarly for each $n \geq 2$, we consider \mathbb{R}_+^n equipped with the Borel σ -algebra, written $\mathcal{B}(\mathbb{R}_+^n)$.

Given q_0 a state in Q and $(e_i)_{1 \leq i \leq n}$ a finite sequence of edges in E , the *symbolic path* starting from q_0 and determined by $(e_i)_{1 \leq i \leq n}$ is the following set of finite runs:

$$\pi(q_0, e_1, \dots, e_n) = \{\rho = q_0 \xrightarrow{t_1, e_1} q_1 \cdots \xrightarrow{t_n, e_n} q_n \mid t_1, \dots, t_n \in \mathbb{R}_+\}.$$

If \mathcal{C} is a Borel subset of \mathbb{R}_+^n (i.e. $\mathcal{C} \in \mathcal{B}(\mathbb{R}_+^n)$), one defines the *constrained symbolic path* starting from q_0 , determined by $(e_i)_{1 \leq i \leq n}$ and satisfying \mathcal{C} as the following subset of $\text{Runs}_f(\mathcal{A}, q)$:

$$\pi_{\mathcal{C}}(q_0, e_1, \dots, e_n) = \{\rho = q_0 \xrightarrow{t_1, e_1} q_1 \cdots \xrightarrow{t_n, e_n} q_n \mid (t_1, \dots, t_n) \in \mathcal{C}\}.$$

Now if π is a (constrained) symbolic path of \mathcal{A} , we define the *cylinder* generated by π , denoted by $\text{Cyl}(\pi)$, as the set of infinite runs ρ such that a prefix ρ' of ρ is in π . In other words, if $\pi = \pi_{\mathcal{C}}(q_0, e_1, \dots, e_n)$ where $q_0 \in Q$, $e_1, \dots, e_n \in E$ and $\mathcal{C} \subseteq \mathbb{R}_+^n$ then

$$\begin{aligned} \text{Cyl}(\pi) = \{ & \rho \in \text{Runs}(\mathcal{A}, q_0) \mid \rho = q_0 \xrightarrow{t_1, e_1} q_1 \cdots \xrightarrow{t_n, e_n} q_n \rightarrow \cdots \\ & \text{and } (t_1, \dots, t_n) \in \mathcal{C}\}. \end{aligned}$$

For each state $q_0 \in Q$, we can then define $(\text{Runs}(\mathcal{A}, q_0), \Omega_{q_0}^{\mathcal{A}}, \text{Prob}_{q_0}^{\mathcal{A}})$ a probability space where

- $\Omega_{q_0}^{\mathcal{A}}$ is the σ -algebra generated by the cylinders starting in q_0 ,
- $\text{Prob}_{q_0}^{\mathcal{A}}$ is defined inductively as follows:

$$\begin{aligned} \text{Prob}_{q_0}^{\mathcal{A}}(\text{Cyl}(\pi(q_0, e_1, \dots, e_n))) \\ = \int_{t \in I(q_0, e_1)} p_{q_0+t}(e_1) \text{Prob}_{q_t}^{\mathcal{A}}(\text{Cyl}(\pi(q_t, e_2, \dots, e_n))) \, d\mu_{q_0}(t) \quad (3.1) \end{aligned}$$

where e_1, \dots, e_n are in E , q_t is such that $q_0 \xrightarrow{t, e_1} q_t$, and we initialize with $\text{Prob}_q^{\mathcal{A}}(\text{Cyl}(\pi(q))) = 1$ for each state q ; and thanks to Caratheodory's extension theorem, it extends to all measurable sets.

Remark 3.1.6. As said in [BBJM12], one can see that “the formula for $\text{Prob}_{q_0}^A$ relies on the fact that the probability of taking transition e_1 at time t coincides with the probability of waiting t time units and then choosing e_1 among the enabled transitions, i.e. $p_{q_0+t}(e_1) d\mu_{q_0}(t)$.”

Now given $q \in Q$ and $e_1, \dots, e_n \in E$, the value of $\text{Prob}_q^A(\text{Cyl}(\pi))$, where $\pi = \pi(q_0, e_1, \dots, e_n)$, can be expressed as follows:

$$\begin{aligned} & \text{Prob}_{q_0}^A(\text{Cyl}(\pi)) \\ &= \int_{t_1 \in I(q, e_1)} p_{q+t_1}(e_1) \int_{t_2 \in I(q_{t_1}, e_2)} p_{q_{t_1}+t_2}(e_2) \dots \\ & \quad \int_{t_n \in I(q_{t_1 \dots t_{n-1}}, e_n)} p_{q_{t_1 \dots t_{n-1}}+t_n}(e_n) d\mu_{q_{t_1 \dots t_{n-1}}}(t_n) \dots d\mu_{q_{t_1}}(t_2) d\mu_q(t_1) \end{aligned}$$

where for every $i \geq 2$, the state $q_{t_1 \dots t_i}$ is such that $q_{t_1 \dots t_{i-1}} \xrightarrow{t_i, e_i} q_{t_1 \dots t_i}$ and the state q_{t_1} is such that $q \xrightarrow{t_1, e_1} q_{t_1}$. Hence one can express the probability of cylinders generated by constrained symbolic paths as follows: given a state q and a constraint $\mathcal{C} \in \mathcal{B}(\mathbb{R}_+^n)$, we get that

$$\begin{aligned} & \text{Prob}_q^A(\text{Cyl}(\pi)) \\ &= \int_{t_1 \in I(q, e_1)} p_{q+t_1}(e_1) \int_{t_2 \in I(q_{t_1}, e_2)} p_{q_{t_1}+t_2}(e_2) \dots \\ & \quad \int_{t_n \in I(q_{t_1 \dots t_{n-1}}, e_n)} p_{q_{t_1 \dots t_{n-1}}+t_n}(e_n) \mathbb{1}_{\mathcal{C}}(t_1, \dots, t_n) d\mu_{q_{t_1 \dots t_{n-1}}}(t_n) \dots d\mu_q(t_1) \end{aligned} \tag{3.2}$$

where $\pi = \pi(q, e_1, \dots, e_n)$ and $\mathbb{1}_{\mathcal{C}}(t_1, \dots, t_n)$ is defined as follows:

$$\mathbb{1}_{\mathcal{C}}(t_1, \dots, t_n) = \begin{cases} 1 & \text{if } (t_1, \dots, t_n) \in \mathcal{C}, \\ 0 & \text{otherwise.} \end{cases}$$

Constrained symbolic paths are required to measure rather complex sets like the set of zeno runs (Section 2.1.2). In fact, we can show that $\text{Prob}_q^A(\text{Cyl}(\pi))$ can be defined by induction as follows:

$$\begin{aligned} & \text{Prob}_q^A(\text{Cyl}(\pi_{\mathcal{C}}(q, e_1, \dots, e_n))) \\ &= \int_{t_1 \in I(q, e_1)} p_{q+t_1}(e_1) \text{Prob}_{q_{t_1}}^A(\text{Cyl}(\pi_{\mathcal{C}_{t_1}}(q, e_2, \dots, e_n))) d\mu_q(t_1) \end{aligned} \tag{3.3}$$

where for every $t_1 \geq 0$, $\mathcal{C}_{t_1} = \{(t_2, \dots, t_n) \in \mathbb{R}_+^{n-1} \mid (t_1, t_2, \dots, t_n) \in \mathcal{C}\}$. This is due to the fact that for every $t_1 \geq 0$, $(t_1, t_2, \dots, t_n) \in \mathcal{C}$ if and only if $(t_2, \dots, t_n) \in \mathcal{C}_{t_1}$.

Remark 3.1.7. Like for DMC in Section 2.2, Remark 2.2.4, it is possible to extend this probability distribution to the set of all infinite runs of \mathcal{A} , $\text{Runs}(\mathcal{A})$. Here, instead of fixing an initial state q_0 , we consider some initial distribution μ on the set of states Q . Then, one can define similarly as in Remark 2.2.4, a probability distribution $\text{Prob}_\mu^{\mathcal{A}}$ over $\text{Runs}(\mathcal{A})$. The difference lies in the fact that here, the set of states Q is non-denumerable and henceforth, the value $\mu(q_0)$ could be null for any state q_0 leading to the fact that for each symbolic path π , $\text{Prob}_\mu^{\mathcal{A}}(\text{Cyl}(\pi)) = 0$. Hence in this case, a symbolic path should be defined with a measurable set of initial states. We do not define it here formally as it will be of no use to us in this chapter. We will come back to this in Chapter 7.

It requires some work to prove that $(\text{Runs}(\mathcal{A}, q), \Omega_{q_0}^{\mathcal{A}}, \text{Prob}_q^{\mathcal{A}})$ is a probability space for each state $q \in Q$. It has been proven in [BBB⁺14] (Proposition 3.2) and we will not give the details here.

Proposition 3.1.8 ([BBB⁺14]). *For every state $q \in Q$, $\text{Prob}_q^{\mathcal{A}}$ is a probability measure over $(\text{Runs}(\mathcal{A}, q), \Omega_{q_0}^{\mathcal{A}})$.*

Example 3.1.9. Consider again the STA \mathcal{A} of Example 3.1.4 (depicted in Figure 2.1), where we make the assumption that in each state, μ_q is the uniform distribution over $I(q)$ (which from the invariants, and the uniform distribution has a sense), or a Dirac distribution if $I(q)$ is a single point. Recall that if A is a bounded set of \mathbb{R}_+ such that $\Lambda(A) > 0$ (where Λ is the Lebesgue measure), $\mathcal{U}(A)$ is defined by the density function $f(t) = \frac{1}{\Lambda(A)} \mathbf{1}_A(t)$. If μ_q is a uniform distribution, we will write f_q for its density function. Assume also that p_q is the uniform discrete distribution over the edges enabled in q . Then we can make the following computations. We consider $q_0 = (l_0, (0, 0))$ as the initial state (*i.e.* we consider the set of runs $\text{Runs}(\mathcal{A}, q_0)$). We have observed that $I(q_0, e_1) =]0, 1]$ and $I(q_0, e_2) = \{2\}$ and thus $I(q_0) =]0, 1] \cup \{2\}$. Then, μ_{q_0} is given as $\mathcal{U}([0, 1])$ and observe that for each $0 < t \leq 1$, $p_{q_0+t}(e_1) = 1$ and $p_{q_0+2}(e_2) = 1$. This gives us:

- $\text{Prob}_{q_0}^{\mathcal{A}}(\text{Cyl}(\pi(q_0, e_1))) = \int_{t \in I(q_0, e_1)} p_{q_0+t}(e_1) \cdot f_{q_0}(t) dt = 1$, and
- $\text{Prob}_{q_0}^{\mathcal{A}}(\text{Cyl}(\pi(q_0, e_2))) = 0$.

Now for each $t \in I(q_0, e_1)$, we write q_t for the state such that $q_0 \xrightarrow{t, e_1} q_t$, *i.e.* $q_t = (l_0, (0, t))$. Then it holds that $I(q_t, e_1) = I(q_0, e_1)$ and $I(q_t, e_2) = [2 - t, 2]$. Observe also that $2 - t \geq 1$ since $t \leq 1$. We thus get that

$$I(q_t) = \begin{cases}]0, 2] & \text{if } t = 1 \\]0, 1] \cup [2 - t, 2] & \text{otherwise} \end{cases}$$


where in the second case, $I(q_t)$ is thus a disjoint union of two intervals and $\Lambda([2-t, 2]) > 0$ for each $0 < t < 1$. Hence μ_{q_t} is a uniform distribution over $I(q_t)$ for each $0 < t \leq 1$. In any case, $\Lambda(]0, 1] \cup]2-t, 2]) = \frac{1}{1+t}$. Finally observe that for each $t \in]0, 1]$, if $t' \in]0, 1]$ then e_1 is the only edge enabled in $q_t + t'$ and if $t' \in [2-t, 2]$, then e_2 is the only edge enabled in $q_t + t'$, except if $t = 1$ and $t' = 1$: then e_1 and e_2 are enabled. However, it holds that for each $t \in]0, 1]$, $p_{q_t+t'}(e_1) = 1$ almost-surely for each $t' \in]0, 1]$. We thus finally compute:

$$\begin{aligned}
\text{Prob}_{q_0}^A(\text{Cyl}(\pi(q_0, e_1, e_1))) &= \int_{t_1 \in I(q_0, e_1)} p_{q_0+t_1}(e_1) \text{Prob}_{q_{t_1}}^A(\text{Cyl}(\pi(q_{t_1}, e_1))) \, d\mu_{q_0}(t_1) \\
&= \int_{t_1=0}^1 \int_{t_2 \in I(q_{t_1}, e_1)} p_{q_{t_1}+t_2}(e_1) \, d\mu_{q_{t_1}}(t_2) \, dt_1 \\
&= \int_{t_1=0}^1 \frac{1}{1+t_1} \left(\int_{t_2=0}^1 dt_2 + \int_{t_2=2-t_1}^2 0 \, dt_2 \right) dt_1 \\
&= \int_{t_1=0}^1 \frac{1}{1+t_1} dt_1 \\
&= \ln(2) - \ln(1) = \ln(2).
\end{aligned}$$

We give now a classical example of probabilistic system as a STA, allowing thus to add timed aspects.

Example 3.1.10. We model the IPv4 Zeroconf protocol [BvdSHV03] using STA as done in [BBB⁺14] (see Figure 3.1). The task of this protocol is to configure IP addresses in a local network of appliances in the following way. When a new appliance is connected to the network, a unique IP address has to be configured with this appliance. The protocol selects randomly an IP address and then sends N messages to the network in order to verify if the IP address is already used. If one of the messages receives an answer in a bounded time, then the IP address is already used and a new one is selected. If none of the messages get an answer, then either the IP address is not used and the appliance is well plugged, or the IP address is used and there was an error when sending the messages. In [BK08], a simple model for the IPv4 Zeroconf protocol is given as a DMC, which abstracts away timing constraints.

In Figure 3.1, we model the protocol for $N = 3$ as a STA with a single clock x . We can describe it as follows. In location IP, an IP address has been randomly selected. Then a first message is broadcast and the system moves to location Wait₁ after resetting clock x to 0. From there, either an answer is received before T time units and the system goes back to location IP where a new IP address is randomly chosen, or no answer is given before T time units and then a new message is sent, clock x is reset to 0 and location Wait₂ is reached where we

Figure 3.1: The IPv4 Zeroconf STA for $N = 3$.

observe a similar behaviour. In location Wait_3 the third and last message has been broadcast. Then again, either an answer is received before T time units and location IP is again visited, or no answer is sent. If there was an error in the broadcast of the messages and the IP address is already used, the system moves to location Error , otherwise location OK is reached.

We finally have to equip each state with probability distributions over the edges and the delays. From state (IP, ν) with $\nu \in \mathbb{R}_+$, $I((\text{IP}, \nu)) = \mathbb{R}_+$ so that we could equip the state with an exponential distribution of parameter μ , written $\text{Exp}(\mu)$, over the delays. And since there is only one edge enabled, the probability over the edges is trivial. For states of the form (Wait_i, ν) , the set of delays is $\mathbb{R}_+ \setminus \{T - \nu\}$ so that we can define the probability over the delays in (Wait_i, ν) as $\text{Exp}(\lambda)$. Now, if $\nu < T$, there is only one enabled edge and so the probability over the edges is trivial. If $\nu = T$, then no edge is enabled. Finally, if $\nu > T$ then there is a unique enabled edge in (Wait_1, ν) and (Wait_2, ν) so that the probability is trivial. In state (Wait_3, ν) , we can either move to OK or to Error , and we assume that the probability to go to Error is given by $0 < \beta < 1$.

We now illustrate another classical probabilistic example as a STA, a queuing system.

Example 3.1.11. We consider a queuing system in which jobs arrive and wait until they are executed, and we assume the capacity of the queue is k . We assume the interarrival time is given by distribution μ_{it} , and that the serving time of each job is given by distribution μ_{st} . We will use f 's for density functions and F 's for cumulative functions, with respective indices. This is what is known as a $G/G/1/k$ -queue (the two first G 's stand for ‘‘Generalized’’ distributions for arrival times and process times).

We propose an STA model for this queuing system. To do so we need to use some of the techniques that will appear later in the thesis when tackling the definition of parallel composition operator in Chapter 9. However we believe

this can be understood without going into the details. Firstly, we assume the STA to have two clocks x and y . Clock x will be used for arrivals time and clock y for process time. The set of locations is given by $\{0, \dots, k\}$, and edges between locations explicitly describe arrival of jobs or processed jobs. Guards on the edges are chosen respectfully with the supports of μ_{it} and μ_{st} . Distributions over the edges and the delays are defined via the following construction. In the sequel ν_x and ν_y will represent values of respectively clock x and clock y .

- For every $\nu_x \geq 0$, we write μ_{it, ν_x} for the distribution obtained from μ_{it} by conditioning over the fact that the job has not arrived within the first ν_x time units. It is obtained using techniques of Chapter 9 and is expressed by the following density function:

$$f_{it, \nu_x}(t) = \frac{f_{it}(\nu_x + t)}{1 - F_{it}(\nu_x)},$$

where f_{it} is thus the density function of μ_{it} . Note that this is only defined for ν_x smaller than the upper bound of the support of μ_{it} .

- Similarly, for every $\nu_y \geq 0$, we write μ_{st, ν_y} for the distribution obtained from μ_{st} by conditioning over the fact that the job has not arrived within the first ν_y time units. The expression is identical to the first item.
- For every $\nu_x, \nu_y \geq 0$, we write $\mu_{\min, (\nu_x, \nu_y)}$ or simply $\mu_{\min, \nu}$ for the distribution $\min(\mu_{it, \nu_x}, \mu_{st, \nu_y})$, representing a race between the two distributions representing arrival of job, and processing of job. It can again be computed using the technical developments of Chapter 9 and can be expressed by the following density function:

$$f_{\min, \nu}(t) = f_{it, \nu_x}(t) \cdot (1 - F_{st, \nu_y}(t)) + f_{st, \nu_y}(t) \cdot (1 - F_{it, \nu_x}(t)).$$

Distribution $\mu_{\min, \nu}$ corresponds thus to the distribution over the delays in each state (i, ν) of the STA, with $i \in \{1, \dots, k-1\}$. The cases of locations 0 and k are special since jobs can respectively only arrive and only be served in those locations.

- For every $\nu_x, \nu_y \geq 0$, we define probabilities $p_{(\nu_x, \nu_y)}^x(t)$ and $p_{(\nu_x, \nu_y)}^y(t)$ (or simply $p_{\nu}^x(t)$ and $p_{\nu}^y(t)$) for the probabilities of having a job arrival (resp. processing) at that time, under the assumption that the delay is t since

one arrived in the location with valuation $\nu = (\nu_x, \nu_y)$. Following rules described in Chapter 9, it can be expressed as follows:

$$p_\nu^x(t) = \frac{f_{it, \nu_x}(t) \cdot (1 - F_{st, \nu_y}(t))}{f_{\min, \nu}(t)} \quad \text{and} \quad p_\nu^y(t) = \frac{f_{st, \nu_y}(t) \cdot (1 - F_{it, \nu_x}(t))}{f_{\min, \nu}(t)}.$$

Note that it is the case that $p_\nu^*(t) = p_{\nu'}^*(t')$ whenever $\nu + t = \nu' + t'$. Value $p_\nu^x(t)$ (resp. $p_\nu^y(t)$) corresponds thus to the probability that a job arrives (resp. is processed) from state $q + t$ with $q = (i, \nu)$ for some $i \in \{1, \dots, k-1\}$. In location 0, jobs can only arrive, and in location k , jobs can only be done.

The STA for the G/G/1/ k -queue is now depicted in Figure 3.2. It should be noted that in the case where μ_{it} is an exponential distribution of parameter λ and μ_{st} is an exponential distribution of parameter ν , then we get exactly the CTMC of Example 2.3.2 (due to the memoryless property of exponential distributions, clocks x and y can be collapsed into a single clock).




Figure 3.2: A STA modelling a G/G/1/ k -queue.

Measuring zeno runs. As seen in Section 2.1.2, a run $\rho = q_0 \xrightarrow{t_0, e_0} q_1 \xrightarrow{t_1, e_1} q_2 \xrightarrow{t_2, e_2} \dots$ is zeno if $\sum_{i \geq 0} t_i < \infty$. We have observed that such runs are unrealistic: they describe behaviours that perform infinitely many actions in a finite amount of time, which made zeno runs as undesirable. We would thus like to consider STA in which the set of zeno runs is negligible, *i.e.* it has a null probability.

The σ -algebra that we consider on the set of infinite runs, allows us to express the set of zeno runs. Indeed given $M \in \mathbb{N}$ and $n \in \mathbb{N}_0$, we write $\mathcal{C}_{M,n} = \{(t_1, \dots, t_n) \in \mathbb{R}_+^n \mid t_1 + \dots + t_n \leq M\}$. Given an initial state q_0 and writing $\text{Zeno}(q_0)$ for the set of zeno runs starting in q_0 , it holds that

$$\text{Zeno}(q_0) = \bigcup_{M \in \mathbb{N}} \bigcap_{n \in \mathbb{N}_0} \bigcup_{(e_1, \dots, e_n) \in E^n} \text{Cyl}(\pi_{\mathcal{C}_{M,n}}(q_0, e_1, \dots, e_n)).$$

It thus makes sense to compute the probability of zeno runs. Again as argued before, one would like to avoid such runs as they are problematic. The model

allows us hence to consider *almost-surely non-zeno* STA: we call a STA \mathcal{A} *almost-surely non-zeno* if for each state q , $\text{Prob}_q^{\mathcal{A}}(\text{Zeno}(q)) = 0$.

An hypothesis on STA. Before going further into this chapter, we will make the following assumption: we require that all STA \mathcal{A} considered, satisfies the following condition. Writing Λ for the Lebesgue measure, for each state q , if $\Lambda(I(q)) > 0$ then μ_q is equivalent¹ to the restriction² of Λ on $I(q)$, written $\Lambda_{I(q)}$. We write this condition (\ddagger). It should be noted that in [BBB⁺14], the STA model is already defined with condition (\ddagger). However, we relaxed the hypothesis in our definition for convenience.

3.2. Thick region graph and Markov chain

In Section 2.1.1, we have defined a region graph of a given timed automaton. In this section, we will present a slightly different version of this region graph and show how, given a STA, we can construct a finite Markov chain for the region graph. But before that, we will interpret the new region graph as a timed automaton: the *timed region automaton*. All those notions are taken from [BBB⁺14] and [BBJM12].

We refer to Section 2.1.1 for the notion of region in timed automata. Recall that given a timed automaton \mathcal{A} , $R_{\mathcal{A}}$ denotes its set of regions. Given a region $r \in R_{\mathcal{A}}$, we write $\text{cell}(r)$ for the tightest guard containing r , and we write $\text{cell}(R_{\mathcal{A}}) = \{\text{cell}(r) \mid r \in R_{\mathcal{A}}\}$.

Definition 3.2.1. Given a timed automaton $\mathcal{A} = (L, X, \text{Act}, E, \text{Inv}, \text{AP}, \mathcal{L})$, the *timed region automaton* is the timed automaton $\mathbf{R}(\mathcal{A}) = (L', X, E', \text{Inv}', \text{AP}, \mathcal{L}')$ where:

- $L' = L \times R_{\mathcal{A}}$,
- $E' \subseteq L' \times \text{Act} \times \text{cell}(R_{\mathcal{A}}) \times 2^X \times L'$ is such that $((l, r), a, \text{cell}(r''), Y, (l', r')) \in E'$ if and only if there exist $\nu \in r$, $\nu' \in r'$, $t \geq 0$ and $e = (l, a, g, Y, l') \in E$ such that $\nu + t \in r''$ and $(l, \nu) \xrightarrow{t, e} (l', \nu')$,
- for each $(l, r) \in L'$, $\text{Inv}'((l, r)) = \text{Inv}(l)$ and $\mathcal{L}'((l, r)) = \mathcal{L}(l)$.

¹Two measures μ and μ' are equivalent whenever for each measurable set A , $\mu(A) = 0$ iff $\mu'(A) = 0$.

²If B is a Borel set, then the restriction of Λ on B is defined as follows: for each Borel set A , $\Lambda_B(A) = \Lambda(A \cap B)$.

We do not give here a detailed example of the region timed automaton as it has no interest for the purpose of the next chapters. It can be found in [BBB⁺14]. However, we give the first steps of the region timed automaton of \mathcal{A} from Example 2.1.2.

Example 3.2.2. Consider the timed automaton \mathcal{A} of Example 2.1.2 depicted in Figure 2.1. We describe the first steps of the timed region automaton $R(\mathcal{A})$ on Figure 3.3, starting from the location $l'_0 = (l_0, \{(0, 0)\})$.




Figure 3.3: The first steps of $R(\mathcal{A})$

Observe that edges f_1 and f_2 correspond to edge e_1 of \mathcal{A} where in the first case we cross the region $(0 < x_1 < 1)$, and in the second case we cross region $(x_1 = 1)$; while edge f_3 corresponds to edge e_2 of \mathcal{A} where we cross region $(x_1 = 2 \wedge x_2 = 2)$.

Remark 3.2.3. In the sequel, an edge f of the timed region automaton will be abusively denoted by its corresponding edge e in the initial timed automaton. The set of states of the timed region automaton is written $Q' \subseteq L' \times \mathbb{R}_+^X$.

Observe also that there is an edge between (l, r) and (l', r') in $R(\mathcal{A})$ if and only if there is an edge between (l, r) and (l', r') in the region graph $\mathcal{R}_{\mathcal{A}}$ (see Definition 2.1.23).

The previous definition can be extended to STA leading to the stochastic timed region automaton. We first introduce some notations. Given a state $q = (l, \nu) \in Q$, we write $\iota(q)$ for its unique image in Q' , *i.e.* $\iota(q) = ((l, [\nu]_{\mathcal{A}}), \nu)$. Then given a symbolic path $\pi((l, \nu), e_1, \dots, e_n)$ in \mathcal{A} , we can associate a finite set of symbolic paths $\pi((l, [\nu]_{\mathcal{A}}), \nu), f_1, \dots, f_n)$ in $R(\mathcal{A})$, each one corresponding to the regions we chose to visit before taking an edge. Given a run ρ in \mathcal{A} , we write $\iota(\rho)$ for its unique image in $R(\mathcal{A})$ and $\iota(\pi((l, \nu), e_1, \dots, e_n))$ is given by the finite union of its corresponding symbolic paths in $R(\mathcal{A})$. Finally, given a set of infinite runs $S \in \text{Runs}(\mathcal{A})$ we write $\iota(S) \in \text{Runs}(R(\mathcal{A}))$ for its image in $R(\mathcal{A})$.

We now extend Definition 3.2.1 to STA.

Definition 3.2.4. Given a STA $\mathcal{A} = (L, X, Act, E, Inv, AP, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X})$, the *stochastic timed region automaton* is the STA $R(\mathcal{A}) = (L', X, Act, E', Inv', AP, \mathcal{L}', (\mu'_q, p'_q)_{q \in L' \times \mathbb{R}_+^X})$ where

- $(L', X, Act, E', Inv', AP, \mathcal{L}')$ is the timed region automaton of \mathcal{A} ,
- for each $q \in Q$, $\mu'_{\iota(q)} = \mu_q$ and for each $t \geq 0$, $p'_{\iota(q)+t} = p_{q+t}$.

This leads to this important result of [BBB⁺14] (Lemme 3.4) linking \mathcal{A} and $R(\mathcal{A})$.

Proposition 3.2.5 ([BBB⁺14]). *Let \mathcal{A} be a STA and $R(\mathcal{A})$ be its stochastic timed region automaton. Then for every state $q \in Q$ it holds that $S \in \Omega_q^{\mathcal{A}}$ if and only if $\iota(S) \in \Omega_{\iota(q)}^{R(\mathcal{A})}$ and in this case, $\text{Prob}_q^{\mathcal{A}}(S) = \text{Prob}_{\iota(q)}^{R(\mathcal{A})}(\iota(S))$.*

The main work of [BBB⁺14] consists in considering the almost-sure model-checking problem of STA for LTL formulas: similarly as in Definition 2.2.6, given a STA \mathcal{A} , an initial state q_0 and an LTL formula φ , it aims at checking whether $\text{Prob}_{q_0}^{\mathcal{A}}(\varphi) = 1$ or not. The answer is not easily computed in a context where the set of states is infinite and even non-denumerable. Proposition 3.2.5 shows that it can be reduced to the almost-sure model-checking problem on the stochastic timed region automaton, but this automaton still has a non-denumerable set of states. In [BBB⁺14], the authors reduce this problem to the almost-sure model-checking of a corresponding LTL formula in a finite Markov chain built on the basis of the previous region graph. But before constructing this finite Markov chain, we need to remove some problematic edges: *singular edges*.

Definition 3.2.6. An edge e of $R(\mathcal{A})$ is said *singular* whenever, writing $(l, r) = \text{source}(e)$, there exists $\nu \in r$ such that $I((l, \nu), e)$ is a single point but there is an edge e' of \mathcal{A} such that $I((l, \nu), e')$ is not a single point (nor the empty set).

Remark 3.2.7. Observe that as stated in Remark 3.1.1, $I((l, \nu), e')$ is thus an interval of \mathbb{R}_+ with $\Lambda(I((l, \nu), e')) > 0$. Considering a STA \mathcal{A} , if e is a singular edge with $\text{source}(e) = (l, r)$, then for each $\nu \in r$, $\mu_{(l, \nu)}(I((l, \nu), e)) = 0$ (it is here that we importantly need \mathcal{A} to satisfy hypothesis (\ddagger)). This is why we need to remove singular edges: they are taken with probability 0.

We now define the *thick* region graph of a given STA \mathcal{A} , which is a slightly different version of the region graph of Definition 2.1.23. We will then be able to construct a finite Markov chain.

Definition 3.2.8. Let \mathcal{A} be a (stochastic) timed automaton. The thick region graph $\mathcal{R}_{\mathcal{A}}^{\text{tk}} = (V, F')$ is an oriented graph where:

- $V = L \times R_{\mathcal{A}}$ is the set of vertices, and
- $F' \subseteq V \times V$ is the set of edges such that $(l, r) \rightarrow (l', r')$ if and only if there is a non-singular edge f in $R(\mathcal{A})$ such that $\text{source}(f) = (l, r)$ and $\text{target}(f) = (l', r')$.

Given a STA \mathcal{A} , we then construct a finite Markov chain based on the thick region graph $\mathcal{R}_{\mathcal{A}}^{\text{tk}} = (V, F')$ as follows: $\text{MC}(\mathcal{A}) = (S, P)$ where $S = V$, $P((l, r), (l', r')) > 0$ if and only $(l, r) \rightarrow (l', r')$ in $\mathcal{R}_{\mathcal{A}}^{\text{tk}}$ and for each (l, r) , (l', r') and (l'', r'') such that $P((l, r), (l', r')) > 0$ and $P((l, r), (l'', r'')) > 0$, it holds that $P((l, r), (l', r')) = P((l, r), (l'', r''))$.

The question is then, given a STA \mathcal{A} , an initial state $q_0 \in Q$, an LTL formula φ on \mathcal{A} and its lifting $\tilde{\varphi}$ to $\text{MC}(\mathcal{A})$, does it hold true that $\text{Prob}_{q_0}^{\mathcal{A}}(\varphi) = 1$ if and only if $\text{Prob}_{[q_0]_{\mathcal{A}}}^{\text{MC}(\mathcal{A})}(\tilde{\varphi}) = 1$? The answer is not true in general and we exhibit a counter-example thanks to the timed automaton \mathcal{A}_{cvg} of Example 2.1.28.

Remark 3.2.9. We do not formally define the grammar of LTL formula here, nor the notion of lifting of an LTL formula from a STA to its corresponding Markov chain as it is not the purpose of this chapter. However it is quite natural and details can be found in [BBB⁺14]. We just point out that it is done through some product between STA and a deterministic Muller automaton that characterises the formula, and it uses the labelling function on the locations. This will be defined formally in a more general setting in Chapter 4. We here only intend to get an overview of some results of this paper.

Example 3.2.10 (Counter-example). Consider the timed automaton \mathcal{A}_{cvg} of Example 2.1.28 depicted in Figure 2.6. We extend it into a STA that satisfies (\ddagger) . Note that in each state q of the automaton, $I(q)$ is either a bounded interval or a single point. In the first case, we assume that μ_q is a uniform distribution, and in the second case that it is a Dirac distribution. We consider $q_0 = (\ell_0, (0, 0))$ as the initial state. We consider the LTL formula $\mathbf{F} p_2$. It can be shown that in the finite Markov chain $\text{MC}(\mathcal{A}_{\text{cvg}})$, it holds that p_2 is reachable from each state in the underlying graph $\mathcal{R}_{\mathcal{A}}^{\text{tk}}$. It follows that $\text{Prob}_{[q_0]_{\mathcal{A}}}^{\text{MC}(\mathcal{A}_{\text{cvg}})}(\mathbf{F} p_2) = 1$.

However $\text{Prob}_{q_0}^{\mathcal{A}_{\text{cvg}}}(\mathbf{F} p_2) < 1$. We do not prove it here, all details are given in [BBB⁺14]. But we give the intuition. This is due to the time-converging aspect discussed in Example 2.1.28. It should be noted that each time we come back to location ℓ_0 , the value of clock x is null while the value of clock y increases but is always less than 1. Hence at the n -th visit to ℓ_0 , we are in state $q_n = (\ell_0, (0, t_n))$ with $t_{n-1} < t_n < 1$. Now observe that $I(q_n) = [0, 1 - t_n[\cup]1 - t_n, 2 - t_n[$ where $I(q_n, e_1) = [0, 1 - t_n[$ and $I(q_n, e_4) =]1 - t_n, 2 - t_n[$. Then with the distributions on the delays chosen, it can be shown that almost-surely, the sequence $(t_n)_{n \geq 0}$ converges to 1. It follows that the probability to take edge e_1 converges towards 0 (since $I(q_n, e_1)$ becomes negligible compared to $I(q_n, e_4)$), and thus the probability to always take edge e_4 is positive. This gives $\text{Prob}_{q_0}^{\mathcal{A}_{\text{cvg}}}(\mathbf{F} p_2) < 1$.

3.3. Fairness and classes of STA

In this section, we present the notion of fairness as defined in [BBB⁺14]. It is motivated by Counter-example 3.2.10. It deletes undesired behaviours and allows one to reduce the almost-sure model-checking problem of LTL formulas from STA to finite Markov chains. Fix a STA \mathcal{A} for the section and recall all previous notations.

Definition 3.3.1. An infinite path $(l_0, r_0) \xrightarrow{f_1} (l_1, r_1) \xrightarrow{f_2} (l_2, r_2) \dots$ in $\text{R}(\mathcal{A})$ is said *fair* if for each non-singular edge e , if there are infinitely many i 's such that $\text{source}(e) = (l_i, r_i)$ then there are infinitely many i 's such that $f_i = e$.

We skip the details, but fairness extends in a natural way to symbolic paths and infinite runs of $\text{R}(\mathcal{A})$ as explained in [BBB⁺14]. A infinite run ρ of \mathcal{A} is then said fair whenever its unique image $\iota(\rho)$ in $\text{R}(\mathcal{A})$ is fair. Given an initial state $q_0 \in \mathcal{Q}$, we write $\text{fair}(q_0) \subseteq \text{Runs}(\mathcal{A}, q_0)$ for the set of infinite runs of \mathcal{A} starting in q_0 that are fair. It should be noted that for each state q , $\text{fair}(q)$ is a

measurable set ([Var85]). We then get the following strong result: for almost-surely fair STA, the almost-sure model-checking problem of LTL formulas can be done through the finite Markov chain constructed in Section 3.2. We recall that is written $\text{MC}(\mathcal{A})$. This was done in [BBB⁺14] (Theorem 6.6).

Theorem 3.3.2 ([BBB⁺14]). *Let \mathcal{A} be a STA. For every state $q \in Q$, if $\text{Prob}_q^{\mathcal{A}}(\text{fair}(q)) = 1$ then for each LTL formula φ ,*

$$\text{Prob}_q^{\mathcal{A}}(\varphi) = 1 \text{ iff } \text{Prob}_{[q]_{\mathcal{A}}}^{\text{MC}(\mathcal{A})}(\tilde{\varphi}) = 1,$$

where $\tilde{\varphi}$ is the lifting of φ to $\text{MC}(\mathcal{A})$.

We can show that the STA \mathcal{A}_{cvg} of Example 3.2.10 is not almost-surely fair.

Example 3.3.3 (Counter-example). Consider the STA of Example 3.2.10 depicted in Figure 2.6. We showed that from state $q_0 = (\ell_0, (0, 0))$, $\text{Prob}_{q_0}^{\mathcal{A}_{\text{cvg}}}(\mathbf{F} p_2) < 1$ and thus $\text{Prob}_{q_0}^{\mathcal{A}_{\text{cvg}}}(\mathbf{G} p_1) > 0$. Now observe that every run starting from q_0 and satisfying $\mathbf{G} p_1$ is unfair as e_1 is enabled infinitely many times but never taken. Hence, $\text{Prob}_{q_0}^{\mathcal{A}_{\text{cvg}}}(\text{fair}(q_0)) < 1$.

This reduces the almost-sure model-checking problem of LTL formulas in STA to the almost-sure model-checking problem of LTL formulas in finite Markov chains, for almost-surely fair STA. In [BBB⁺14], the authors identified several classes of STA that are almost-surely fair and thus for which, Theorem 3.3.2 can be applied. We present two of them.

Reactive STA. A STA \mathcal{A} is said *reactive* whenever the following conditions hold true:

(H1) for every state q , $I(q) = \mathbb{R}_+$ and for every $l \in L$, there exists a distribution μ_l equivalent to the Lebesgue measure on \mathbb{R}_+ , such that for every $\nu \in \mathbb{R}_+^X$, $\mu_{(l, \nu)} = \mu_l$;

(H2) for every edge e there exists $w_e \in \mathbb{N}_0$ such that for every state q ,

$$p_q(e) = \begin{cases} \frac{w_e}{\sum_{e' \text{ enabled in } q} w_{e'}} & \text{if } e \text{ is enabled in } q, \\ 0 & \text{otherwise.} \end{cases}$$

It has been shown in [BBB⁺14] (Proposition 7.11) that reactive STA are almost-surely fair from each state.

Single-clock STA. We consider the class of single-clock STA satisfying condition (H2) and the following hypotheses:

- (H3) for all $l \in L$, for all $[a, b] \subseteq \mathbb{R}_+$, the function $\nu \rightarrow \mu_{(l, \nu)}([a, b])$ is continuous;
- (H4) if $q' = q + t$ for some $t \geq 0$, and if $0 \notin I(q + t', e)$ for each $0 \leq t' \leq t$, then $\mu_q(I(q, e)) \leq \mu_{q'}(I(q', e))$;
- (H5) there is $0 < \lambda_0 < 1$ such that for every state q with $I(q)$ unbounded, $\mu_q([0, \frac{1}{2}]) \leq \lambda_0$.

It has been shown in [BBB⁺14] (Theorem 7.2) that if \mathcal{A} is a single-clock STA that satisfies hypotheses (H2), (H3), (H4) and (H5), then \mathcal{A} is almost-surely fair from each state $q \in Q$.

Note that in both classes, the proof to show the almost-sure fairness is very technical and *ad hoc*. In Part I, we will again be interested in the qualitative model-checking problem of LTL formulas, but also in the quantitative model-checking problem for more general probabilistic systems. We will present a unifying way to show similar results, *i.e.* the reduction of the almost-sure model-checking problem to a finite or denumerable abstraction for those general probabilistic systems. In Chapter 7, we will show how STA can be incorporated into this general setting and thus how the same results can be implied with this new unifying way to prove it. Finally, we will also get new results for the quantitative model-checking problem.

Part I

Qualitative and Quantitative Analysis of Stochastic Transition Systems and Application to Stochastic Timed Automata

Stochastic Transition Systems

In this chapter, we define the notion of stochastic transition systems (STSs for short) as defined in [BBBC17] and introduce several notions that will be needed in the sequel.

In Section 4.1, we define and illustrate the model of STSs. STSs are general stochastic systems with a possible continuous set of states and defined with a Markov kernels. They correspond to the labelled Markov process model of [Pan01] with only one label. We define the logic that we will consider and the qualitative and quantitative model-checking problems for this logic in STSs.

In Section 4.2, we introduce several notions that will be useful when tackling the qualitative and quantitative model-checking problems in STSs. In particular, we introduce and illustrate several notions of *decisiveness* inspired from [ABM07]. We recall that decisive DMCs comes with nice decidability results and approximation schemes for reachability and repeated reachability properties [ABM07] (see Section 2.2.2 for some details). The objective is thus to extend those results to decisive STSs. We also define the notions of attractors and fairness and end in Section 4.2.4 with the links between those notions. This will be the object of a brief discussion in Section 4.3.

We already said it earlier but we repeat it: we assume the reader to be familiar with basic notions of probability theory. One may for instance refer to [Fel66] and [Fel69]

4.1. Definition and illustration of the model

In this section, we define the general model of STSs, which are somehow Markov chains with a continuous state-space. This model corresponds to labelled Markov processes of [Pan01] with a single action (hence removing non-determinism). We then define several probability measures, on infinite runs (like for DMCs in Section 2.2 and STA in Section 3.1), but also on the state-space, which gives different point-of-views over the behaviours of the systems. We continue by defining the logic that we will consider leading to measurable events, and by defining deterministic Muller automata and technical material for handling properties specified by these automata (which include LTL formulas). We end the section with the definition of the qualitative and quantitative model-checking problems in STSs.

Given (S, Σ) a measurable space (Σ is a σ -algebra over S), we write $\text{Dist}(S, \Sigma)$ for the set of probability distributions over (S, Σ) . In the sequel, when the context is clear, we will omit the σ -algebra and simply write this set as $\text{Dist}(S)$.

Definition 4.1.1. A *stochastic transition system* (STS) is a tuple $\mathcal{T} = (S, \Sigma, \kappa)$ consisting of a measurable analytic space (S, Σ) , and a function $\kappa : S \times \Sigma \rightarrow [0, 1]$ satisfying that for every fixed $s \in S$, $\kappa(s, \cdot)$ is a probability measure (*i.e.* $\kappa(s, \cdot) \in \text{Dist}(S)$) and for each fixed $A \in \Sigma$, $\kappa(\cdot, A)$ is a measurable function. Function κ is the *Markov kernel* of \mathcal{T} .

Note that it is sufficient to define $\kappa(s, \cdot)$ (for every $s \in S$) over a subset which generates the σ -algebra Σ . The assumption that (S, Σ) should be analytic is for the STS to have smooth properties [Pan09, Section 7.5].

Observe that if S is a denumerable set and $\Sigma = 2^S$, then \mathcal{T} is a DMC: it corresponds to $\mathcal{M} = (S, P)$ where for each $s, s' \in S$, $P(s, s') = \kappa(s, \{s'\})$. We now give two examples of STS. The first example will be called \mathcal{T}_2 and the second \mathcal{T}_1 . This seems odd, but these notations will suit better for the notations of Chapter 5 and we find it more logical to present our two example in the presented order.

The first one is the DMC representing the random-walk of Example 2.2.3 but expressed in this new formalism.

Example 4.1.2 (Denumerable Markov chain). The first example is the DMC depicted in Figure 2.10 of Example 2.2.3. We consider thus here $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2)$ where

- $S_2 = \mathbb{N}$ and $\Sigma_2 = 2^{S_2}$,
- for each $i \geq 1$, $\kappa_2(i, \{i + 1\}) = p$ and $\kappa_2(i, \{i - 1\}) = 1 - p$ with $p \in]0, 1[$, and

- $\kappa_2(0, \{1\}) = 1$.

This represents a random walk over the natural numbers.

In the sequel, given a DMC $\mathcal{T} = (S, \Sigma, \kappa)$ and two states $s, s' \in S$, we will write $\kappa(s, s')$ instead of $\kappa(s, \{s'\})$.

Example 4.1.3 (Continuous-space Markov chain). We now give a continuous variant of the previous random walk which models a simple queueing system. Precisely, we consider a queueing system with a single queue, a parameter λ for arrivals and ν for serving times. Each state $i \in \mathbb{N}$ is equipped with a non-negative real number that corresponds to the time that has elapsed since the beginning. Formally, we consider $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1)$ with $S_1 = \mathbb{N} \times \mathbb{R}_+$. We equip S_1 with the σ -algebra generated by $2^{\mathbb{N}} \times \mathcal{B}(\mathbb{R}_+)$ where $\mathcal{B}(\mathbb{R}_+)$ is the Borel σ -algebra on \mathbb{R}_+ . Then intuitively, κ_1 describes how the length of the queue evolves with time. Formally, for each $s, t \in \mathbb{R}_+$, for each $i \geq 1$,

$$\begin{aligned} \kappa_1((0, t), (1, [0, s + t])) &= \kappa_1((0, t), (1, [t, s + t])) = \int_0^s \lambda e^{-\lambda x} dx \\ \kappa_1((i, t), (i + 1, [0, s + t])) &= \kappa_1((i, t), (i + 1, [t, s + t])) = \int_0^s \lambda e^{-(\lambda + \nu)x} dx \text{ and} \\ \kappa_1((i, t), (i - 1, [0, s + t])) &= \kappa_1((i, t), (i - 1, [t, s + t])) = \int_0^s \nu e^{-(\lambda + \nu)x} dx. \end{aligned}$$

Remark 4.1.4. Observe that Example 4.1.3 does not represent a CTMC (see Section 2.3) since CTMCs can only have a finite set of locations (while here the corresponding set of locations would be \mathbb{N}). However it is easily seen that the CTMC of Example 2.3.2 can be encoded as a STS in a similar way as in Example 4.1.3. In fact, all CTMCs can be interpreted by a STS.

Again, like in DMCs (Section 2.2) or STA (Section 3.1), the objective of such a model is to define a probability measure on the set of infinite runs. In the sequel, we fix an STS $\mathcal{T} = (S, \Sigma, \kappa)$.

We follow the approach of [DP03]. A *finite (resp. infinite) run* of \mathcal{T} is a finite (resp. infinite) sequence of states $\rho = s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_n$ with $n \in \mathbb{N}$ (resp. $\rho = s_0 \rightarrow s_1 \rightarrow s_2 \rightarrow \dots$). We will sometimes denote runs by $\rho = (s_n)_{n \in \{0, \dots, n\}}$ or $\rho = (s_n)_{n \in \mathbb{N}}$. We write $\text{Runs}(\mathcal{T})$ for the set of infinite runs of \mathcal{T} . In order to get a probability measure over $\text{Runs}(\mathcal{T})$, we need to equip this set with a σ -algebra. We therefore define for each finite sequence of measurable sets $(A_i)_{0 \leq i \leq n} \in \Sigma^{n+1}$ the following set of infinite runs:

$$\text{Cyl}(A_0, A_1, \dots, A_n) = \{\rho = s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_n \rightarrow \dots \mid \forall 0 \leq i \leq n, s_i \in A_i\} .$$

This set is called a *cylinder*. Now fix an initial distribution μ over Σ , i.e. $\mu \in \text{Dist}(S)$. We then define the probability space $(\text{Runs}(\mathcal{T}), \mathcal{F}_{\mathcal{T}}, \text{Prob}_{\mu}^{\mathcal{T}})$ where:

- $\mathcal{F}_{\mathcal{T}}$ is the σ -algebra generated by cylinders,
- $\text{Prob}_{\mu}^{\mathcal{T}}$ is defined as follows: for every finite sequence of measurable subsets $(A_i)_{0 \leq i \leq n} \in \Sigma^{n+1}$, we set:

$$\text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_0, A_1, \dots, A_n)) = \int_{s_0 \in A_0} \text{Prob}_{\kappa(s_0, \cdot)}^{\mathcal{T}}(\text{Cyl}(A_1, \dots, A_n)) d\mu(s_0),$$

and we initialize with $\text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_0)) = \mu(A_0)$.

From now on, we will use the classical notation $\mu(ds_0) = d\mu(s_0)$. It should be noted that the value $\text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_0, A_1, \dots, A_n))$ is the result of n successive integrals and can be expressed as follows:

$$\begin{aligned} \text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_0, A_1, \dots, A_n)) &= \int_{s_0 \in A_0} \int_{s_1 \in A_1} \dots \\ &\dots \int_{s_{n-1} \in A_{n-1}} \kappa(s_0, ds_1) \cdot \kappa(s_1, ds_2) \cdot \dots \cdot \kappa(s_{n-2}, ds_{n-1}) \cdot \kappa(s_{n-1}, A_n) \cdot \mu(ds_0). \end{aligned}$$

Using the classical Caratheodory's extension theorem, $\text{Prob}_{\mu}^{\mathcal{T}}$ can be extended in a unique way to the σ -algebra $\mathcal{F}_{\mathcal{T}}$.

Lemma 4.1.5. For each initial distribution $\mu \in \text{Dist}(S)$, $\text{Prob}_{\mu}^{\mathcal{T}}$ defines a probability measure over $(\text{Runs}(\mathcal{T}), \mathcal{F}_{\mathcal{T}})$.

The proof of this lemma is classical and we omit it here. The interested reader may *e.g.* refer to the proof of [BBB⁺14, Proposition 3.2] (Proposition 3.1.8 here), which can easily be adapted to this context of STS.

Remark 4.1.6. Observe that if the initial distribution is a Dirac distribution δ_s over a single state $s \in S$, then we have that

$$\text{Prob}_{\delta_s}^{\mathcal{T}}(\text{Cyl}(A_0, \dots, A_n)) = \begin{cases} 0 & \text{if } s \notin A_0, \\ \text{Prob}_{\kappa(s, \cdot)}^{\mathcal{T}}(\text{Cyl}(A_1, \dots, A_n)) & \text{otherwise.} \end{cases}$$

We write this distribution $\text{Prob}_s^{\mathcal{T}}$ recovering the notation of Sections 2.2 and 3.1. It follows that for every $\mu \in \text{Dist}(S)$, we can write

$$\text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_0, \dots, A_n)) = \int_{s_0 \in A_0} \text{Prob}_{s_0}^{\mathcal{T}}(\text{Cyl}(A_0, \dots, A_n)) \mu(ds_0)$$

and thus for every $\varpi \in \mathcal{F}_{\mathcal{T}}$,

$$\text{Prob}_{\mu}^{\mathcal{T}}(\varpi) = \int_{s_0 \in S} \text{Prob}_{s_0}^{\mathcal{T}}(\varpi) \mu(ds_0).$$

If we write $\text{Runs}(\mathcal{T}, s_0)$ for the set of infinite runs starting in s_0 , it holds that $\text{Prob}_{s_0}^{\mathcal{T}}$ is a probability measure over $\text{Runs}(\mathcal{T}, s_0)$.

Finally, one can observe that if \mathcal{T} is a DMC, then $\text{Cyl}(A_0, A_1, \dots, A_n)$ corresponds to a disjoint union of cylinders as defined in Section 2.2: given a finite run $\rho = (s_k)_{k \in \{0, \dots, n\}}$ of the DMC, $\text{Cyl}(\rho)$ as defined in Section 2.2 corresponds to $\text{Cyl}(\{s_0\}, \{s_1\}, \dots, \{s_n\})$ in this new formalism and here, A_i is an at most denumerable set of states for each i ; and $\text{Prob}_{s_0}^{\mathcal{T}}$ corresponds exactly to the definition of $\text{Prob}_{s_0}^{\mathcal{M}}$ of Section 2.2 (and hence, $\text{Prob}_{\mu}^{\mathcal{T}}$ corresponds to the definition of $\text{Prob}_{\mu}^{\mathcal{M}}$ of Remark 2.2.4).

Recall that given two probability distributions μ and ν over some probability space (S, Σ) , μ and ν are *equivalent* if for each $A \in \Sigma$, $\mu(A) = 0 \iff \nu(A) = 0$. We can then prove the following lemma.

Lemma 4.1.7. Let μ and ν be two probability measures over (S, Σ) . If μ and ν are equivalent, then $\text{Prob}_{\mu}^{\mathcal{T}}$ and $\text{Prob}_{\nu}^{\mathcal{T}}$ are also equivalent.

Proof. We have to show that for each $\varpi \in \mathcal{F}_T$, $\text{Prob}_{\mu}^{\mathcal{T}}(\varpi) = 0 \iff \text{Prob}_{\nu}^{\mathcal{T}}(\varpi) = 0$. Since the complementary of each cylinder is a finite union of cylinders and since each denumerable union of cylinders can be written as a denumerable disjoint union of cylinders, it suffices to show this for each cylinder $\text{Cyl}(A_0, \dots, A_n)$ with $A_0, \dots, A_n \in \Sigma$. We have to show that for each $A_0, \dots, A_n \in \Sigma$,

$$\text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_0, \dots, A_n)) = 0 \iff \text{Prob}_{\nu}^{\mathcal{T}}(\text{Cyl}(A_0, \dots, A_n)) = 0.$$

We prove it by induction over n . It should be observed that, by symmetry, it suffices to show one of the implications. First, assume $n = 0$ and fix $A_0 \in \Sigma$. Then from the definition of $\text{Prob}_{\mu}^{\mathcal{T}}$ and $\text{Prob}_{\nu}^{\mathcal{T}}$ and from the hypothesis, we get that:

$$\text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_0)) = 0 \iff \mu(A_0) = 0 \iff \nu(A_0) = 0 \iff \text{Prob}_{\nu}^{\mathcal{T}}(\text{Cyl}(A_0)) = 0.$$

Now consider $n = 1$ and fix $A_0, A_1 \in \Sigma$. Suppose that $\text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_0, A_1)) = 0$, *i.e.* from the definition:

$$\int_{s_0 \in A_0} \kappa(s_0, A_1) \mu(ds_0) = 0. \quad (4.1)$$

Write $B = \{s_0 \in A_0 \mid \kappa(s_0, A_1) > 0\}$. We can write $B = \kappa(\cdot, A_1)^{-1}([0, 1]) \cap A_0$ which is in Σ from the hypotheses over κ (see Definition 4.1.1). From (4.1), we can easily check that $\mu(B) = 0$, which implies that $\nu(B) = 0$ and thus

$$\int_{s_0 \in A_0} \kappa(s_0, A_1) \nu(ds_0) = 0.$$

Using again the definition, it follows that $\text{Prob}_\nu^\mathcal{T}(\text{Cyl}(A_0, A_1)) = 0$. Now, assume that $n \geq 2$, fix $A_0, \dots, A_n \in \Sigma$ and assume that $\text{Prob}_\mu^\mathcal{T}(\text{Cyl}(A_0, \dots, A_n)) = 0$. Recall that

$$\begin{aligned} \text{Prob}_\mu^\mathcal{T}(\text{Cyl}(A_0, \dots, A_n)) &= \int_{s_0 \in A_0} \left(\int_{s_1 \in A_1} \dots \right. \\ &\quad \left. \dots \left(\int_{s_{n-1} \in A_{n-1}} \kappa(s_{n-1}, A_n) \kappa(s_{n-2}, ds_{n-1}) \right) \dots \kappa(s_0, ds_1) \right) \mu(ds_0). \end{aligned}$$

We inductively define:

$$\begin{cases} B_{n-1} = \kappa(\cdot, A_n)^{-1}([0, 1]) \cap A_{n-1} \\ B_i = \kappa(\cdot, B_{i+1})^{-1}([0, 1]) \cap A_i \quad \forall 0 \leq i \leq n-2. \end{cases}$$

From the hypotheses over κ , it is easily seen that for each $0 \leq i \leq n-1$, $B_i \in \Sigma$. Let us consider the value $\int_{s_{n-1} \in A_{n-1}} \kappa(s_{n-1}, A_n) \kappa(s_{n-2}, ds_{n-1})$. From the definition of B_{n-1} , it holds that

$$\begin{aligned} \int_{s_{n-1} \in A_{n-1}} \kappa(s_{n-1}, A_n) \kappa(s_{n-2}, ds_{n-1}) &= \int_{s_{n-1} \in B_{n-1}} \kappa(s_{n-1}, A_n) \kappa(s_{n-2}, ds_{n-1}) \\ &= \text{Prob}_{\kappa(s_{n-2}, \cdot)}^\mathcal{T}(\text{Cyl}(B_{n-1}, A_n)). \end{aligned}$$

We thus get that

$$\begin{aligned} \text{Prob}_\mu^\mathcal{T}(\text{Cyl}(A_0, \dots, A_n)) &= \\ &= \int_{s_0 \in A_0} \dots \left(\int_{s_{n-2} \in A_{n-2}} \text{Prob}_{\kappa(s_{n-2}, \cdot)}^\mathcal{T}(\text{Cyl}(B_{n-1}, A_n)) \kappa(s_{n-3}, ds_{n-2}) \right) \dots \mu(ds_0). \end{aligned}$$

We prove the two following statements: for each $0 \leq i \leq n-2$,

- (a) $\{s_i \in S \mid \text{Prob}_{\kappa(s_i, \cdot)}^\mathcal{T}(\text{Cyl}(B_{i+1}, \dots, B_{n-1}, A_n)) > 0\} \cap A_i = B_i$ and
- (b)

$$\begin{aligned} \int_{s_i \in A_i} \text{Prob}_{\kappa(s_i, \cdot)}^\mathcal{T}(\text{Cyl}(B_{i+1}, \dots, B_{n-1}, A_n)) \kappa(s_{i-1}, ds_i) \\ = \text{Prob}_{\kappa(s_{i-1}, \cdot)}^\mathcal{T}(\text{Cyl}(B_i, \dots, B_{n-1}, A_n)), \end{aligned}$$

where if $i = 0$, $\kappa(s_{i-1}, \cdot)$ will stand for the initial distribution μ . Point (a) is here in order to establish that the sets $\{s_i \in S \mid \text{Prob}_{\kappa(s_i, \cdot)}^\mathcal{T}(\text{Cyl}(B_{i+1}, \dots, B_{n-1}, A_n)) > 0\} \cap A_i$ are measurable, and point (b) aims at reducing our integrals to sets whose images have positive values. It should be observed that the second point is an

immediate consequence of the first point. We thus only need to prove point (a). We do this by induction over i . First, if $i = n - 2$, we show that

$$\begin{aligned} \{s_{n-2} \in S \mid \text{Prob}_{\kappa(s_{n-2}, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{n-1}, A_n)) > 0\} \\ = \{s_{n-2} \in S \mid \kappa(s_{n-2}, B_{n-1}) > 0\} \end{aligned}$$

which will ensure that (a) is satisfied. First assume that $s_{n-2} \in S$ is such that

$$\text{Prob}_{\kappa(s_{n-2}, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{n-1}, A_n)) > 0 .$$

Towards a contradiction, assume that $\kappa(s_{n-2}, B_{n-1}) = 0$. Then it holds that

$$\begin{aligned} 0 = \kappa(s_{n-2}, B_{n-1}) &= \text{Prob}_{\kappa(s_{n-2}, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{n-1})) \\ &\geq \text{Prob}_{\kappa(s_{n-2}, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{n-1}, A_n)) > 0 \end{aligned}$$

which is the needed contradiction. Now assume that $\kappa(s_{n-2}, B_{n-1}) > 0$. Then from the definitions of B_{n-1} and of $\text{Prob}_{\kappa(s_{n-2}, \cdot)}^{\mathcal{T}}$, and from classical properties on integrals, it is straightforward to check that the second inclusion holds. Now suppose that point (a) holds for each $i + 1 \leq j \leq n - 2$ for some $i \geq 0$, and let us show that it is still true for i . As before, it suffices to establish that

$$\{s_i \in S \mid \text{Prob}_{\kappa(s_i, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{i+1}, \dots, B_{n-1}, A_n)) > 0\} = \{s_i \in S \mid \kappa(s_i, B_{i+1}) > 0\}.$$

The first inclusion can be verified just like in the first case. Now assume that $\kappa(s_i, B_{i+1}) > 0$. We know that

$$\begin{aligned} \text{Prob}_{\kappa(s_i, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{i+1}, \dots, B_{n-1}, A_n)) &= \\ \int_{s_{i+1} \in B_{i+1}} \text{Prob}_{\kappa(s_{i+1}, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{i+2}, \dots, B_{n-1}, A_n)) \kappa(s_i, ds_{i+1}). \end{aligned}$$

Using the induction hypothesis over $i + 1$, we get that for each $s_{i+1} \in B_{i+1}$,

$$\text{Prob}_{\kappa(s_{i+1}, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{i+2}, \dots, B_{n-1}, A_n)) > 0 .$$

And since $\kappa(s_i, B_{i+1}) > 0$, this induces that

$$\text{Prob}_{\kappa(s_i, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{i+1}, \dots, B_{n-1}, A_n)) > 0$$

which concludes that point (a) is satisfied. Hence from points (a) and (b), we get that

$$\begin{aligned} \text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_0, \dots, A_n)) &= \text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(B_0, \dots, B_{n-1}, A_n)) \\ &= \int_{s_0 \in B_0} \text{Prob}_{\kappa(s_0, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_1, \dots, B_{n-1}, A_n)) \mu(ds_0). \end{aligned}$$

From the fact that $B_0 = \{s_0 \in A_0 \mid \text{Prob}_{\kappa(s_0, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_1, \dots, B_{n-1}, A_n)) > 0\}$ and that $\text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_0, \dots, A_n)) = 0$, it follows that $\mu(B_0) = 0$. From the hypothesis, we thus get that $\nu(B_0) = 0$. Now observing that we can prove similarly that $\text{Prob}_{\nu}^{\mathcal{T}}(\text{Cyl}(A_0, \dots, A_n)) = \text{Prob}_{\nu}^{\mathcal{T}}(\text{Cyl}(B_0, \dots, B_{n-1}, A_n))$, we can establish that $\text{Prob}_{\nu}^{\mathcal{T}}(\text{Cyl}(A_0, \dots, A_n)) = 0$ which concludes the proof. \square

One can also interpret the dynamics of \mathcal{T} as a transformer of probability measures over (S, Σ) . Compared to the previous point-of-view, here one is interested not in states the system can be in, but rather in how the probability mass evolves along steps. It has been considered for both discrete-time Markov chains [AAGT12] and continuous-time models, *e.g.* those induced by stochastic time Petri nets [HPRV12]. This second point-of-view the other side of the same coin of the previous one. It will be of some use in the sequel, although we will ultimately only consider the semantics through the measure of infinite runs, which suits better for the properties we want to check.

For μ a probability measure over Σ , its transformation through \mathcal{T} can be defined as the measure $\Omega_{\mathcal{T}}(\mu)$ defined for every $A \in \Sigma$ by:

$$\Omega_{\mathcal{T}}(\mu)(A) = \int_{s_0 \in S} \kappa(s_0, A) \cdot \mu(ds_0).$$

It can be shown that $\Omega_{\mathcal{T}}(\mu)$ is also a probability measure over (S, Σ) .

This interpretation offers a dual view on the STS \mathcal{T} . Indeed, roughly speaking, $\Omega_{\mathcal{T}}(\mu)(A)$ is the probability of being in A after one step, when μ is the initial distribution on \mathcal{T} . Given a distribution $\mu \in \text{Dist}(S)$ and given $A \in \Sigma$ such that $\mu(A) > 0$, we write μ_A for the conditional probability of μ given A , that is for each $B \in \Sigma$ $\mu_A(B) = \frac{\mu(A \cap B)}{\mu(A)}$. It should be observed that $\mu_A \in \text{Dist}(S)$. There is a strong relation between the transformer $\Omega_{\mathcal{T}}(\mu)$ and the probability measure $\text{Prob}_{\mu}^{\mathcal{T}}$ over runs defined previously, which we formalize in Lemma 4.1.8.

Lemma 4.1.8. Let $\mu \in \text{Dist}(S)$ be an initial distribution and let $(A_i)_{0 \leq i \leq n}$ be a sequence of measurable sets. Write $\nu_0 = \mu_{A_0}$, the conditional probability of μ given A_0 , and for every $1 \leq j \leq n-1$, write $\nu_j = (\Omega_{\mathcal{T}}(\nu_{j-1}))_{A_j}$. Then, for every $0 \leq j \leq n$:

$$\begin{aligned} \text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_0, A_1, \dots, A_n)) = \\ \mu(A_0) \cdot \prod_{i=1}^j (\Omega_{\mathcal{T}}(\nu_{i-1}))(A_i) \cdot \text{Prob}_{\Omega_{\mathcal{T}}(\nu_j)}^{\mathcal{T}}(\text{Cyl}(A_{j+1}, \dots, A_n)). \end{aligned}$$

Proof. The proof is by induction on j . Assume that $j = 0$, we have to show: $\text{Prob}_\mu^\mathcal{T}(\text{Cyl}(A_0, A_1, \dots, A_n)) = \mu(A_0) \cdot \text{Prob}_{\Omega_\mathcal{T}(\nu_0)}^\mathcal{T}(\text{Cyl}(A_1, \dots, A_n))$. First,

$$\begin{aligned} & \text{Prob}_\mu^\mathcal{T}(\text{Cyl}(A_0, \dots, A_n)) \\ &= \text{Prob}_\mu^\mathcal{T}(\text{Cyl}(A_0) \cap \text{Cyl}(S, A_1, \dots, A_n)) \\ &= \text{Prob}_\mu^\mathcal{T}(\text{Cyl}(A_0)) \cdot \text{Prob}_\mu^\mathcal{T}(\text{Cyl}(S, A_1, \dots, A_n) \mid \text{Cyl}(A_0)) \\ &= \mu(A_0) \cdot \text{Prob}_{\mu_{A_0}}^\mathcal{T}(\text{Cyl}(A_0, \dots, A_n)). \end{aligned}$$

Now let us unfold $\text{Prob}_{\Omega_\mathcal{T}(\nu_0)}^\mathcal{T}(\text{Cyl}(A_1, \dots, A_n))$:

$$\begin{aligned} & \text{Prob}_{\Omega_\mathcal{T}(\nu_0)}^\mathcal{T}(\text{Cyl}(A_1, \dots, A_n)) \\ &= \int_{s_1 \in A_1} \text{Prob}_{\kappa(s_1, \cdot)}^\mathcal{T}(\text{Cyl}(A_2, \dots, A_n))(\Omega_\mathcal{T}(\nu_0))(ds_1) \\ &= \int_{s_1 \in A_1} \text{Prob}_{\kappa(s_1, \cdot)}^\mathcal{T}(\text{Cyl}(A_2, \dots, A_n)) \int_{s_0 \in S} \kappa(s_0, ds_1) \nu_0(ds_0) \\ &= \int_{s_0 \in A_0} \left(\int_{s_1 \in A_1} \text{Prob}_{\kappa(s_1, \cdot)}^\mathcal{T}(\text{Cyl}(A_2, \dots, A_n)) \kappa(s_0, ds_1) \right) \mu_{A_0}(ds_0) \\ &= \int_{s_0 \in A_0} \text{Prob}_{\kappa(s_0, \cdot)}^\mathcal{T}(\text{Cyl}(A_1, \dots, A_n)) \mu_{A_0}(ds_0) \\ &= \text{Prob}_{\mu_{A_0}}^\mathcal{T}(\text{Cyl}(A_0, \dots, A_n)) . \end{aligned}$$

Now fix $0 < j \leq n$ and assume that for each for each $0 \leq i < j$ the equality above holds. We will prove that it is still the case for j . First, observe that if $j = n$ then the induction hypothesis states that

$$\begin{aligned} & \text{Prob}_\mu^\mathcal{T}(\text{Cyl}(A_0, A_1, \dots, A_n)) \\ &= \mu(A_0) \cdot \prod_{i=1}^{n-1} (\Omega_\mathcal{T}(\nu_{i-1}))(A_i) \cdot \text{Prob}_{\Omega_\mathcal{T}(\nu_{n-1})}^\mathcal{T}(\text{Cyl}(A_n)) \\ &= \mu(A_0) \cdot \prod_{i=1}^{n-1} (\Omega_\mathcal{T}(\nu_{i-1}))(A_i) \cdot \Omega_\mathcal{T}(\nu_{n-1})(A_n) \\ &= \mu(A_0) \cdot \prod_{i=1}^n (\Omega_\mathcal{T}(\nu_{i-1}))(A_i) \end{aligned}$$

which is what we wanted. Otherwise, if $j < n$, then the hypothesis induction

states that

$$\begin{aligned} \text{Prob}_\mu^\mathcal{T}(\text{Cyl}(A_0, A_1, \dots, A_n)) = \\ \mu(A_0) \cdot \prod_{i=1}^{j-1} (\Omega_{\mathcal{T}}(\nu_{i-1}))(A_i) \cdot \text{Prob}_{\Omega_{\mathcal{T}}(\nu_{j-1})}^\mathcal{T}(\text{Cyl}(A_j, \dots, A_n)). \end{aligned}$$

Then using a similar argument as in the first case, we get that

$$\text{Prob}_{\Omega_{\mathcal{T}}(\nu_{j-1})}^\mathcal{T}(\text{Cyl}(A_j, \dots, A_n)) = \Omega_{\mathcal{T}}(\nu_{j-1})(A_j) \cdot \text{Prob}_{\Omega_{\mathcal{T}}(\nu_j)}^\mathcal{T}(\text{Cyl}(A_{j+1}, \dots, A_n))$$

since $\Omega_{\mathcal{T}}(\nu_j) = (\Omega_{\mathcal{T}}(\nu_{j-1}))_{A_j}$. This concludes the proof. \square

Remark 4.1.9. From this result, we can express the probability to reach A in n steps from the initial distribution μ :

$$(\Omega_{\mathcal{T}}^{(n)}(\mu))(A) = \text{Prob}_\mu^\mathcal{T}(\text{Cyl}(\overbrace{S, \dots, S}^{n \text{ times}}, A)).$$

4.1.1 Formulas for STSs

Like for DMCs in Section 2.2, the σ -algebra $\mathcal{F}_{\mathcal{T}}$ allows to express a rich variety of properties. To define properties on the STS \mathcal{T} , we use again LTL-like notations, that will be interpreted as measurable subsets of $\text{Runs}(\mathcal{T})$. Let $\mathcal{L}_{S,\Sigma}$ be the set of formulas defined by the following grammar:

$$\varphi ::= B \mid \varphi_1 \mathbf{U}_{\bowtie k} \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \wedge \varphi_2 \mid \neg\varphi,$$

where $B \in \Sigma$, $\bowtie \in \{\geq, \leq, =\}$ is a comparison operator and $k \in \mathbb{N}$ is a natural number. Observe that this logic can express all LTL formulas (see Section 2.2). Given a run $\rho = (s_n)_{n \geq 0} \in \text{Runs}(\mathcal{T})$, we write $\rho_{\geq i} = (s_n)_{n \geq i} \in \text{Runs}(\mathcal{T})$ for each $i \geq 0$. Then the satisfaction relation of formulas is given as follows:

$$\begin{aligned} \rho \models B & \iff s_0 \in B \\ \rho \models \varphi_1 \mathbf{U}_{\bowtie k} \varphi_2 & \iff \exists i \geq 0, i \bowtie k, \text{ s.t. } \rho_{\geq i} \models \varphi_2 \text{ and } \forall 0 \leq j < i, \rho_{\geq j} \models \varphi_1 \\ \rho \models \varphi_1 \vee \varphi_2 & \iff \rho \models \varphi_1 \text{ or } \rho \models \varphi_2 \\ \rho \models \varphi_1 \wedge \varphi_2 & \iff \rho \models \varphi_1 \text{ and } \rho \models \varphi_2 \\ \rho \models \neg\varphi & \iff \rho \not\models \varphi. \end{aligned}$$

We write $\text{Ev}_{\mathcal{T}}(\varphi)$ for the set of infinite runs ρ in \mathcal{T} such that $\rho \models \varphi$. It is standard to show that the event $\text{Ev}_{\mathcal{T}}(\varphi)$ is a measurable subset of $(\text{Runs}(\mathcal{T}), \mathcal{F}_{\mathcal{T}})$ (see e.g. [Var85]). In particular, for each initial probability measure μ , $\text{Prob}_\mu^\mathcal{T}(\text{Ev}_{\mathcal{T}}(\varphi))$

is well-defined. In the sequel, for simplicity, we often write $\text{Prob}_\mu^{\mathcal{T}}(\varphi)$ instead of $\text{Prob}_\mu^{\mathcal{T}}(\text{Ev}_{\mathcal{T}}(\varphi))$.

We will also use classical notations like $\top = S$; $\perp = \emptyset$; $\varphi_1 \mathbf{U} \varphi_2 = \varphi_1 \mathbf{U}_{\geq 0} \varphi_2$; $\mathbf{F} \varphi = \top \mathbf{U} \varphi$; $\mathbf{F}_{\bowtie k} \varphi = \top \mathbf{U}_{\bowtie k} \varphi$; $\mathbf{G} \varphi = \neg \mathbf{F}(\neg \varphi)$.

As we will often use them, we give the semantics of some formulas in $\mathcal{L}_{S,\Sigma}$ in terms of events: we get inductively that

$$\begin{aligned} \text{Ev}_{\mathcal{T}}(B \mathbf{U}_{\bowtie k} B') &= \bigcup_{i \bowtie k} (\text{Ev}_{\mathcal{T}}(\mathbf{F}_{=i} B') \cap \bigcap_{0 \leq j < i} \text{Ev}_{\mathcal{T}}(\mathbf{F}_{=j} B)) \\ \text{Ev}_{\mathcal{T}}(\mathbf{G} \mathbf{F} B) &= \bigcap_{i \geq 0} \bigcup_{j \geq i} \text{Ev}_{\mathcal{T}}(\mathbf{F}_{=j} B) \\ \text{Ev}_{\mathcal{T}}(\varphi_1 \vee \varphi_2) &= \text{Ev}_{\mathcal{T}}(\varphi_1) \cup \text{Ev}_{\mathcal{T}}(\varphi_2) \\ \text{Ev}_{\mathcal{T}}(\varphi_1 \wedge \varphi_2) &= \text{Ev}_{\mathcal{T}}(\varphi_1) \cap \text{Ev}_{\mathcal{T}}(\varphi_2) \\ \text{Ev}_{\mathcal{T}}(\neg \varphi) &= \text{Runs}(\mathcal{T}) \setminus \text{Ev}_{\mathcal{T}}(\varphi) \end{aligned}$$

where for each $n \in \mathbb{N}$ and each $B \in \Sigma$,

$$\text{Ev}_{\mathcal{T}}(\mathbf{F}_{=n} B) = \text{Cyl}(\overbrace{S, \dots, S}^{n \text{ times}}, B).$$

4.1.2 Labelled STSs and their properties

To ease the expression of rich properties over STSs, we extend the model with a labelling with atomic propositions.

Definition 4.1.10. A *labelled stochastic transition system* (LSTS for short) is a tuple $\mathcal{T} = (S, \Sigma, \kappa, \text{AP}, \mathcal{L})$, where (S, Σ, κ) is an STS, AP is a finite set of atomic propositions, and $\mathcal{L} : S \rightarrow 2^{\text{AP}}$ is a measurable labelling function.

Measures and other notions are extended in a straightforward way from STSs to LSTSs. We fix a finite set AP of atomic propositions and an LSTS $\mathcal{T} = (S, \Sigma, \kappa, \text{AP}, \mathcal{L})$.

A property over AP is a subset P of $(2^{\text{AP}})^\omega$ (the set of infinite sequences of elements in 2^{AP}). An infinite run $\rho = s_0 s_1 \dots$ of \mathcal{T} satisfies the property P whenever $\mathcal{L}(s_0)\mathcal{L}(s_1)\mathcal{L}(s_2)\dots \in P$, written $\rho \models P$. Whenever the set of runs satisfying property P is measurable (*i.e.* $\{\rho \in \text{Runs}(\mathcal{T}) \mid \rho \models P\} \in \mathcal{F}_{\mathcal{T}}$), for each $\mu \in \text{Dist}(S)$, we write $\text{Prob}_\mu^{\mathcal{T}}(P)$ for the probability of this set of runs.

The notion of ω -regularity is standard in computer science to characterise simple sets of infinite behaviours, and typical ω -regular properties are Büchi and Muller properties. In order to express such properties, we introduce a new notation for the set of atomic propositions that are true infinitely often along a

sequence of labels: for $\varpi = u_0u_1u_2\dots \in (2^{\text{AP}})^\omega$, we define $\text{Inf}(\varpi) = \{a \in \text{AP} \mid |\{j \in \mathbb{N} \mid a \in u_j\}| = \infty\}$. We extend this notation to runs in a natural way: if $\rho = s_0s_1s_2\dots \in S^\omega$, writing $\varpi = \mathcal{L}(s_0)\mathcal{L}(s_1)\mathcal{L}(s_2)\dots$, we define (with a slight abuse of notation) $\text{Inf}(\rho) = \text{Inf}(\varpi)$.

A *Büchi property* P over AP can be specified by a subset of atomic propositions $F \subseteq \text{AP}$ as $P = \{\varpi \in (2^{\text{AP}})^\omega \mid \text{Inf}(\varpi) \cap F \neq \emptyset\}$. A *Muller property* over AP is a property P such that there exists $\mathcal{F} \subseteq 2^{\text{AP}}$ with $P = \{\varpi \in (2^{\text{AP}})^\omega \mid \text{Inf}(\varpi) \in \mathcal{F}\}$.

Given a Muller property P defined by $\mathcal{F} \subseteq 2^{\text{AP}}$ and given $\mu \in \text{Dist}(S)$, we will sometimes write $\text{Prob}_\mu^T(\text{Inf} \in \mathcal{F})$ instead of $\text{Prob}_\mu^T(P)$.

Remark 4.1.11. It should be noted that the set of infinite runs satisfying Büchi or Muller properties can be expressed using events as in Section 4.1.1. Indeed, for $F \subseteq \text{AP}$ we write $2_F^{\text{AP}} = \{u \in 2^{\text{AP}} \mid u \cap F \neq \emptyset\}$ and given $a \in \text{AP}$, $2_a^{\text{AP}} = \{u \in 2^{\text{AP}} \mid a \in u\}$. Then,

- the set of runs satisfying the Büchi property with acceptance condition F is

$$\text{Ev}_{\mathcal{T}}\left(\mathbf{GF}\left(\bigvee_{u \in 2_F^{\text{AP}}} \mathcal{L}^{-1}(u)\right)\right);$$

- the set of runs satisfying the Muller property with acceptance condition \mathcal{F} is

$$\text{Ev}_{\mathcal{T}}\left(\bigvee_{F \in \mathcal{F}} \left(\bigwedge_{a \in F} \left(\mathbf{GF} \bigvee_{u \in 2_a^{\text{AP}}} \mathcal{L}^{-1}(u)\right) \wedge \bigwedge_{a \notin F} \bigwedge_{u \in 2_a^{\text{AP}}} \mathbf{FG}(\mathcal{L}^{-1}(u))^c\right)\right).$$

It is well known that automata equipped with Büchi or Muller acceptance conditions capture all ω -regular properties, and this also holds for deterministic Muller automata.

Definition 4.1.12. A *deterministic Muller automaton* (DMA) over AP is a tuple $M = (Q, q_0, E, \mathcal{F})$ where:

- Q is a finite set of locations, and $q_0 \in Q$ is the initial location;
- $E \subseteq Q \times 2^{\text{AP}} \times Q$ is a finite set of edges;
- \mathcal{F} is a Muller condition over Q ;

and such that

- M is deterministic: for all pair of edges (q, u, q_1) and (q, u, q_2) in E , $q_1 = q_2$;

- M is complete: for every $q \in Q$, for every $u \in 2^{\text{AP}}$, there exists $(q, u, q') \in E$.

A DMA M naturally gives rise to a property P_M defined by the language (over 2^{AP}) accepted by M . A run in a DMA M is any infinite sequence of the form $q_0 u_0 q_1 u_1 \dots$ with $q_i \in Q$ and $u_i \in 2^{\text{AP}}$ for each $i \geq 0$ and such that for each $i \geq 0$, $(q_i, u_i, q_{i+1}) \in E$. A run $q_0 u_0 q_1 u_1 \dots$ is accepted by M if $\text{Inf}((q_n)_{n \geq 0}) \in \mathcal{F}$. An infinite sequence $(u_n)_{n \geq 0} \in (2^{\text{AP}})^\omega$ is accepted by M if the unique corresponding run in M (the existence and the uniqueness of this run are due to the fact that M is deterministic and complete) is accepted by M . The property P_M is thus defined as follows: $P_M = \{\varpi \in (2^{\text{AP}})^\omega \mid \varpi \text{ is accepted by } M\}$. The semantics of a DMA M over infinite runs of \mathcal{T} is derived from that property P_M : if $\rho \in \text{Runs}(\mathcal{T})$, we write $\rho \models M$ whenever $\rho \models P_M$. Expanding Remark 4.1.11, one derives the standard fact that the set $\mathcal{T}[M] \stackrel{\text{def}}{=} \{\rho \in \text{Runs}(\mathcal{T}) \mid \rho \models M\}$ is measurable, and we write $\text{Prob}_\mu^\mathcal{T}(M)$ for $\text{Prob}_\mu^\mathcal{T}(\mathcal{T}[M])$.

Remark 4.1.13. It is well known that for any LTL formula φ over AP (the syntax given in Section 4.1.1, where we replace sets B by inverse images by \mathcal{L} of atomic propositions from AP), there is a DMA M_φ that characterises φ , that is: for every run ρ , $\rho \models \varphi$ iff $\rho \models M_\varphi$. See [VW94] and [GTW02, Chapter 3].

Remark 4.1.14. Observe that property P_M is not a Muller property as defined above and thus one cannot express P_M like in the second item of Remark 4.1.11. This is due to the fact that in Definition 4.1.12, the winning condition of a DMA is given by $\mathcal{F} \subseteq 2^Q$ (and not $\mathcal{F} \subseteq 2^{\text{AP}}$). However, each Muller property can be expressed as a DMA.

Now, in order to measure the probability of properties specified by a DMA $M = (Q, q_0, E, \mathcal{F})$, it is convenient to build a new STS consisting of the product of \mathcal{T} with M . To this aim, we consider the discrete σ -algebra 2^Q on the finite set of locations Q of M . The product $S \times Q$ can then be equipped with the product σ -algebra $\Sigma \times 2^Q$ defined as the smallest σ -algebra containing all rectangles, that is, all sets of the form $A_1 \times A_2$ with $A_1 \in \Sigma$ and $A_2 \in 2^Q$. Then, the product σ -algebra $\Sigma \times 2^Q$ coincides with $\Sigma' = \{\bigcup_{q \in Q} C_q \times \{q\} \mid \forall q \in Q, C_q \in \Sigma\}$. Note that in the sequel, we will sometimes write (C_q, q) instead of $C_q \times \{q\}$.

We now define the product of \mathcal{T} with M as follows.

Definition 4.1.15. Given $\mathcal{T} = (S, \Sigma, \kappa, \text{AP}, \mathcal{L})$ an LSTS and $M = (Q, q_0, E, \mathcal{F})$ a DMA over AP, we define the product of \mathcal{T} with M as the LSTS $\mathcal{T} \times M = (S', \Sigma', \kappa', \text{AP}', \mathcal{L}')$ such that:

- $S' = S \times Q$;
- Σ' is the product σ -algebra $\Sigma \times 2^Q$;

- $\kappa'((s, q), (A, q')) = \begin{cases} \kappa(s, A) & \text{if } (q, \mathcal{L}(s), q') \in E, \text{ and} \\ 0 & \text{otherwise;} \end{cases}$
- $\text{AP}' = Q$;
- $\mathcal{L}'(s, q) = q$.

Note that the above definition of κ' extends naturally to all elements of the σ -algebra Σ' : for each pair (q, u) with $q \in Q$ and $u \in 2^{\text{AP}}$, there is a unique $q' \in Q$ such that $(q, u, q') \in E$. Fix $(s, q) \in S \times Q$, write q' for the unique location such that $(q, \mathcal{L}(s), q') \in E$. Then for each $A = \bigcup_{q \in Q} C_q \times \{q\}$, $\kappa'((s, q), A) = \kappa'((s, q), (C_{q'}, q')) = \kappa(s, C_{q'})$.

Example 4.1.16. We consider the random walk over \mathbb{N} of Example 4.1.2. We assume that it is equipped with the simple set of atomic propositions $\text{AP} = \{a\}$ and we assume that each state of the STS is labelled with a . Let M be the DMA depicted on the left-hand side of Figure 4.1. The winning condition is given by $\mathcal{F} = \{\{q_1, q_2\}\}$. The product $\mathcal{T}_2 \times M$ is then depicted on the right-hand side of Figure 4.1. It should be noted that we assume here that the system starts at $(0, q_0)$ however, there should be similar chains starting in (i, q_0) for each $i \geq 1$. Note also that we did not specify the labels on the states: according to Definition 4.1.15, each state is labelled with its current position in M .




Figure 4.1: A Muller automaton M and the product $\mathcal{T}_2 \times M$.

The product $\mathcal{T} \times M$ gives rise to a new LSTS which has a labelling function \mathcal{L}' over Q . Hence property $P_M \in (2^{\text{AP}})^\omega$ cannot directly be derived on $\mathcal{T} \times M$. However, DMA M gives also rise to a property P'_M over Q : P'_M corresponds to the set of infinite runs $q_0 q_1 q_2 \dots \in Q^\omega$ accepted by M , which is a Muller property over Q . Then, one can thus use the previous semantics in order to state that a run ρ satisfies the Muller property P'_M whenever $\mathcal{L}'(\rho) \in P_M$ and one can use Remark 4.1.11 to show that P'_M is measurable in $\mathcal{T} \times M$ (*i.e.* $P'_M \in \mathcal{F}_{\mathcal{T} \times M}$).

We now explain how initial distributions for \mathcal{T} are lifted to the product $\mathcal{T} \times M$. The idea is simple: the \mathcal{T} -component is inherited from \mathcal{T} , and the M -component

is a Dirac measure over q_0 , the initial state of M . In other words, when an initial distribution $\mu \in \text{Dist}(S)$ is fixed for \mathcal{T} , the initial distribution of $\mathcal{T} \times M$ will be $\mu \times \delta_{q_0}$.

Here given $\mu \in \text{Dist}(S)$, we can thus express $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times M}(\{\rho \in \text{Runs}(\mathcal{T} \times M) \mid \rho \models \mathcal{F}\}) = \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times M}(P'_M)$ and as before, we will write it $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times M}(\text{Inf} \in \mathcal{F})$.

We show that this allows to properly compute the probability of a property described by a DMA, with the following correspondence.

Proposition 4.1.17. *Let $\mu \in \text{Dist}(S)$ be an initial distribution for \mathcal{T} , and $M = (Q, q_0, E, \mathcal{F})$ be a DMA. Then:*

$$\text{Prob}_{\mu}^{\mathcal{T}}(M) = \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times M}(\text{Inf} \in \mathcal{F}) .$$

In order to prove this result, we first introduce some notations. Given $A_0, A_1, \dots, A_n \in \Sigma'$ we write for each i , $A_i = \bigcup_{q \in Q} A_{i,q} \times \{q\}$. Also given $u_1, \dots, u_n \in 2^{\text{AP}}$ and $q \in Q$ we inductively define

$$\begin{cases} q_{u_1} = q' \in Q & \text{such that } (q, u_1, q') \in E \\ q_{u_1 \dots u_i} = q' \in Q & \text{such that } (q_{u_1 \dots u_{i-1}}, u_i, q') \in E, \forall 2 \leq i \leq n. \end{cases}$$

Observe that since M is deterministic and complete, those states are uniquely defined. We then have the following result.

Lemma 4.1.18. *For each initial distribution $\mu \in \text{Dist}(S)$ for \mathcal{T} , for each state $q \in Q$ of M , for each $n \in \mathbb{N}$ and for each $A_0, \dots, A_n \in \Sigma'$, it holds that*

$$\begin{aligned} \text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times M}(\text{Cyl}(A_0, A_1, \dots, A_n)) = \\ \sum_{u_1, \dots, u_n \in 2^{\text{AP}}} \text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_{0,q} \cap \mathcal{L}^{-1}(u_1), A_{1,q_{u_1}} \cap \mathcal{L}^{-1}(u_2), \dots, \\ A_{n-1,q_{u_1 \dots u_{n-1}}} \cap \mathcal{L}^{-1}(u_n), A_{n,q_{u_1 \dots u_n}})). \end{aligned}$$

Proof. We prove it by induction over n . First if $n = 0$, we have to show that for every $\mu \in \text{Dist}(S)$, every $q \in Q$ and every $A_0 \in \Sigma'$,

$$\text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times M}(\text{Cyl}(A_0)) = \text{Prob}_{\mu}^{\mathcal{T}}(A_{0,q})$$

which is trivial from the definition of $\mu \times \delta_q$. Now fix $n \geq 0$. Assume that for each $0 \leq i \leq n$, the above property holds true and show that it is still the case

for $i = n + 1$. Let $\mu \in \text{Dist}(S)$, $q \in Q$ and $A_0, \dots, A_{n+1} \in \Sigma'$. We have that

$$\begin{aligned}
& \text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times \mathcal{M}}(\text{Cyl}(A_0, \dots, A_{n+1})) \\
&= \int_{(s_0, q') \in A_0} \text{Prob}_{\kappa'((s_0, q'), \cdot)}^{\mathcal{T} \times \mathcal{M}}(\text{Cyl}(A_1, \dots, A_{n+1})) d(\mu \times \delta_q)((s_0, q')) \\
&= \int_{s_0 \in A_{0,q}} \text{Prob}_{\kappa'((s_0, q), \cdot)}^{\mathcal{T} \times \mathcal{M}}(\text{Cyl}(A_1, \dots, A_{n+1})) d\mu(s_0) \\
&= \sum_{u_1 \in 2^{\text{AP}}} \int_{s_0 \in A_{0,q} \cap \mathcal{L}^{-1}(u_1)} \text{Prob}_{\kappa'((s_0, q), \cdot)}^{\mathcal{T} \times \mathcal{M}}(\text{Cyl}(A_1, \dots, A_{n+1})) d\mu(s_0) \\
&= \sum_{u_1 \in 2^{\text{AP}}} \int_{s_0 \in A_{0,q} \cap \mathcal{L}^{-1}(u_1)} \text{Prob}_{\kappa(s_0, \cdot) \times \delta_{qu_1}}^{\mathcal{T} \times \mathcal{M}}(\text{Cyl}(A_1, \dots, A_{n+1})) d\mu(s_0) \quad (4.2)
\end{aligned}$$

where the last equality comes from the unicity of q_{u_1} . Using the induction hypothesis, we get that

$$\begin{aligned}
& \text{Prob}_{\kappa(s_0, \cdot) \times \delta_{qu_1}}^{\mathcal{T} \times \mathcal{M}}(\text{Cyl}(A_1, \dots, A_{n+1})) = \\
& \sum_{u_2, \dots, u_{n+1} \in 2^{\text{AP}}} \text{Prob}_{\kappa(s_0, \cdot)}^{\mathcal{T}}(\text{Cyl}(A_{1, q_{u_1}} \cap \mathcal{L}^{-1}(u_2), \dots \\
& \qquad \qquad \qquad \dots, A_{n, q_{u_1 \dots u_n}} \cap \mathcal{L}^{-1}(u_{n+1}), A_{n+1, q_{u_1 \dots u_{n+1}}}).
\end{aligned}$$

Combining with (4.2), we thus obtain that

$$\begin{aligned}
& \text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times \mathcal{M}}(\text{Cyl}(A_0, \dots, A_{n+1})) = \\
& \sum_{u_1, \dots, u_{n+1} \in 2^{\text{AP}}} \text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_{0,q} \cap \mathcal{L}^{-1}(u_1), \dots \\
& \qquad \qquad \qquad \dots, A_{n, q_{u_1 \dots u_n}} \cap \mathcal{L}^{-1}(u_{n+1}), A_{n+1, q_{u_1 \dots u_{n+1}}}).
\end{aligned}$$

which concludes the proof. \square

Proposition 4.1.17 is then an immediate consequence of Lemma 4.1.18.

4.1.3 Qualitative and quantitative model-checking problem

The objective of this part of the thesis, is to define a nice framework on STSs in order to get strong results on the qualitative and quantitative model-checking problems in STSs. In this section, we adapt the Definitions 2.2.6 and 2.2.7 of the qualitative and quantitative problems in DMCs to this richest context.

Definition 4.1.19. Given an STS \mathcal{T} , an initial distribution μ and a property φ that can be expressed in the σ -algebra $\mathcal{F}_{\mathcal{T}}$, the *qualitative model-checking problem* aims at verifying whether $\text{Prob}_{\mu}^{\mathcal{T}}(\varphi) = 1$ or not.

Definition 4.1.20. Given an STS \mathcal{T} , an initial distribution μ and a property φ that can be expressed in the σ -algebra $\mathcal{F}_{\mathcal{T}}$, the *quantitative model-checking problem* aims at computing an approximation of $\text{Prob}_{\mu}^{\mathcal{T}}(\varphi)$.

In Section 2.2, we have presented the work of [ABM07] that shows that decisiveness plays a key role in proving the different results over the qualitative and quantitative model-checking problems in DMCs, while in Section 3.3 we have presented the work of [BBB⁺14] that shows that fairness plays this key role for the qualitative model-checking problem results for STA.

In Chapter 6, we will present a unifying way to tackle those qualitative and quantitative model-checking problems for STSs, recovering then the results of [ABM07] on DMCs and the results of [BBB⁺14] on STA (which can be seen as STSs, as we will see in Chapter 7). In this last case, it will moreover be shown that the *ad hoc* proofs of [BBB⁺14] can be simplified in our new formalism and that we get new results for the quantitative model-checking problem.

In Section 4.2, we thus introduce different notions on STSs: decisiveness and fairness as they were important for DMCs and STA, but also the notion of attractor as it will be of a main interest to us.

4.2. Properties of STSs

In this section, we introduce several notions on STSs that will be needed in the sequel in Chapters 5 and 6 in order to tackle the qualitative and quantitative model-checking problems in STSs. We then show the relationships between those notions. We will also make parallels with some notions of Section 2.2 since DMCs and STSs are related in the sense that a STS can be seen as a DMC with a continuous set of states (and that all DMCs are STSs).

We fix an STS $\mathcal{T} = (S, \Sigma, \kappa)$ for all this section.

4.2.1 Several decisiveness notions

We have presented in Section 2.2 the elegant concept of decisive Markov chain which has been introduced in [ABM07]. As said before, it has been defined as a desirable property of DMCs, since it implies that they behave essentially like finite Markov chains. In this section, we show how to extend and refine this notion of decisiveness to general STSs.

For $B \in \Sigma$ a measurable set of states, we define its *avoid-set* $\widetilde{B} = \{s \in S \mid \text{Prob}_s^{\mathcal{T}}(\mathbf{F} B) = 0\}$. It corresponds to the set of states from which the system will always avoid the set B with probability 1. It should be noted that in the case where \mathcal{T} is a DMC, \widetilde{B} corresponds to the definition of Section 2.2.2.

Remark 4.2.1. Recalling the notations of Section 2.2.2, if $\mathcal{T} = \mathcal{M}$ is a DMC, then $\{s \in S \mid s \not\models \exists \mathbf{F} B\} = \{s \in S \mid \text{Prob}_s^{\mathcal{T}}(\mathbf{F} B) = 0\}$. Observe that if $s \not\models \exists \mathbf{F} B$, then $\text{Runs}(\mathcal{T}, s) \cap \text{Ev}_{\mathcal{T}}(\mathbf{F} B) = \emptyset$. The left-to-right inclusion is thus immediate. Now assume that $\text{Prob}_s^{\mathcal{T}}(\mathbf{F} B) = 0$ and towards a contradiction, suppose that $s \models \exists \mathbf{F} B$ i.e. there is $\rho = (s_n)_{n \geq 0} \in \text{Runs}(\mathcal{M}, s)$ such that there is $n \geq 0$, $s_n \in B$. It follows that $\text{Cyl}(s, s_1, \dots, s_n) \subseteq \text{Ev}_{\mathcal{T}}(\mathbf{F} B)$ and thus $\text{Prob}_s^{\mathcal{T}}(\mathbf{F} B) \geq \text{Prob}_s^{\mathcal{T}}(\text{Cyl}(s, s_1, \dots, s_n)) > 0$ leading to the desired contradiction.

The set \widetilde{B} enjoys the following properties, that obviously hold also in the context of DMCs, but require proofs in our general context of STSs.

Lemma 4.2.2. Given $B \in \Sigma$, it holds that:

- \widetilde{B} belongs to the σ -algebra Σ ;
- for every $\mu \in \text{Dist}(\widetilde{B})$, $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B) = 0$;
- for every $\mu \in \text{Dist}(S)$, if $\mu((\widetilde{B})^c) > 0$, then $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B) > 0$;
- for every $\mu \in \text{Dist}(S)$, $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} \widetilde{B}) = \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} \mathbf{G} \widetilde{B}) = \text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} \widetilde{B})$;
- for every $\mu \in \text{Dist}(S)$, $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B \vee \mathbf{F} \widetilde{B}) = \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B \vee (\neg B \mathbf{U} \widetilde{B}))$.

Let us comment on the third and fourth properties stated in this lemma. The third item indicates that if we start from outside \widetilde{B} , then we will always have a positive probability to hit B . The fourth property says that \widetilde{B} is some kind of a sink set of states: once we hit \widetilde{B} , we cannot escape it. This comes from the definition of \widetilde{B} . We get that for almost-surely each state s reachable from \widetilde{B} , it is also the case that $\text{Prob}_s^{\mathcal{T}}(\mathbf{F} B) = 0$. The other properties are rather straightforward to understand (even though proving the first property requires some technical developments).

Proof. We begin with the first point. Recall that given $B \in \Sigma$, $\widetilde{B} = \{s \in S \mid \text{Prob}_s^{\mathcal{T}}(\mathbf{F} B) = 0\}$. Observe that we can write:

$$\widetilde{B} = \bigcap_{n \geq 0} \{s \in S \mid \text{Prob}_s^{\mathcal{T}}(\overbrace{\text{Cyl}(S, \dots, S, B)}^{n \text{ times}}) = 0\}.$$

It thus suffices to show that for each $n \geq 0$,

$$\{s \in S \mid \text{Prob}_s^{\mathcal{T}}(\overbrace{\text{Cyl}(S, \dots, S, B)}^{n \text{ times}}) = 0\} \in \Sigma.$$

We will use similar arguments as in the proof of Lemma 4.1.7. Recall that if $n \geq 1$, it holds that $\text{Prob}_s^{\mathcal{T}}(\text{Cyl}(\overbrace{S, \dots, S}^{n \text{ times}}, B)) = \text{Prob}_{\kappa(s, \cdot)}^{\mathcal{T}}(\text{Cyl}(\overbrace{S, \dots, S}^{n-1 \text{ times}}, B))$ from Remark 4.1.6.

First, if $n = 0$ then this set corresponds to the set $\{s \in S \mid \delta_s(B) = 0\} = B^c$ which is in Σ . Now if $n = 1$ then

$$\{s \in S \mid \text{Prob}_{\kappa(s, \cdot)}(\text{Cyl}(B)) = 0\} = (\kappa(\cdot, B))^{-1}(\{0\})$$

which is in Σ from the hypotheses over κ (see Definition 4.1.1). Now assume that $n \geq 2$, it holds that

$$\begin{aligned} \text{Prob}_{\kappa(s, \cdot)}^{\mathcal{T}}(\text{Cyl}(\overbrace{S, \dots, S}^{n-1 \text{ times}}, B)) = \\ \int_{s_1 \in S} \cdots \int_{s_{n-1} \in S} \kappa(s_{n-1}, B) \kappa(s_{n-2}, ds_{n-1}) \cdots \kappa(s_1, ds_2) \kappa(s, ds_1). \end{aligned}$$

We inductively define:

$$\begin{cases} B_{n-1} = \kappa(\cdot, B)^{-1}([0, 1]) \\ B_i = \kappa(\cdot, B_{i+1})^{-1}([0, 1]) \quad \forall 0 \leq i \leq n-2. \end{cases}$$

From the hypotheses over κ , it holds that $B_i \in \Sigma$ for each $0 \leq i < n$. In the sequel, s_0 denotes s . As in the proof of Lemma 4.1.7, we can show that firstly, $\int_{s_{n-1} \in S} \kappa(s_{n-1}, B) \kappa(s_{n-2}, ds_{n-1}) = \text{Prob}_{\kappa(s_{n-2}, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{n-1}, B))$ and that for each $1 \leq i \leq n-2$,

(a) $\{s_i \in S \mid \text{Prob}_{\kappa(s_i, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{i+1}, \dots, B_{n-1}, B)) > 0\} = B_i$ and

(b)

$$\begin{aligned} \int_{s_i \in S} \text{Prob}_{\kappa(s_i, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{i+1}, \dots, B_{n-1}, A_n)) \kappa(s_{i-1}, ds_i) = \\ \text{Prob}_{\kappa(s_{i-1}, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_i, \dots, B_{n-1}, A_n)). \end{aligned}$$

It follows that

$$\begin{aligned} \text{Prob}_{\kappa(s, \cdot)}^{\mathcal{T}}(\text{Cyl}(\overbrace{S, \dots, S}^{n-1 \text{ times}}, B)) &= \text{Prob}_{\kappa(s, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_1, \dots, B_{n-1}, B)) \\ &= \int_{s_1 \in B_1} \text{Prob}_{\kappa(s_1, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_2, \dots, B_{n-1}, B)) \kappa(s, ds_1) \end{aligned}$$

Now since for each $s_1 \in B_1$, $\text{Prob}_{\kappa(s_1, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_2, \dots, B_{n-1}, B)) > 0$, it holds that

$$\text{Prob}_{\kappa(s, \cdot)}^{\mathcal{T}}(\text{Cyl}(\overbrace{S, \dots, S}^{n-1 \text{ times}}, B)) = 0$$

if and only if $\kappa(s, B_1) = 0$, *i.e.* if and only if $s \notin B_0$. And since $B_0 \in \Sigma$, it follows that $B_0^c \in \Sigma$ and thus

$$B_0^c = \{s \in S \mid \text{Prob}_s^{\mathcal{T}}(\text{Cyl}(\overbrace{S, \dots, S}^{n \text{ times}}, B)) = 0\} \in \Sigma.$$

The second property is a direct consequence of the definition of \tilde{B} .

We now focus on the third property. Towards a contradiction, assume that there is $\mu \in \text{Dist}(S)$ such that $\mu((\tilde{B})^c) > 0$ but $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} B) = 0$. It follows that there is $s \in (\tilde{B})^c$ such that $\text{Prob}_s^{\mathcal{T}}(\mathbf{F} B) = 0$ and thus $s \in \tilde{B}$ which is the wanted contradiction.

Let us show the fourth item. It should be observed that given $\mu \in \text{Dist}(S)$, $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \mathbf{G} \tilde{B}) \leq \text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} \tilde{B}) \leq \text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \tilde{B})$. It thus suffices to show that $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \mathbf{G} \tilde{B}) = \text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \tilde{B})$. Since $\text{Ev}_{\mathcal{T}}(\mathbf{F} \mathbf{G} \tilde{B}) \subseteq \text{Ev}_{\mathcal{T}}(\mathbf{F} \tilde{B})$, towards a contradiction, we assume that $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \tilde{B} \wedge \mathbf{G} \mathbf{F} (\tilde{B})^c) > 0$. As

$$\begin{aligned} \text{Ev}_{\mathcal{T}}(\mathbf{F} \tilde{B} \wedge \mathbf{G} \mathbf{F} (\tilde{B})^c) &\subseteq \text{Ev}_{\mathcal{T}}\left(\bigvee_{n \geq 0} (\mathbf{F}_{=n} \tilde{B} \wedge \mathbf{F}_{>n} (\tilde{B})^c)\right) \\ &= \bigcup_{n \geq 0} \bigcup_{m \geq 0} \text{Cyl}(\overbrace{S, \dots, S}^{n \text{ times}}, \tilde{B}, \overbrace{S, \dots, S}^{m \text{ times}}, (\tilde{B})^c) \end{aligned}$$

it follows that there is $n, m \in \mathbb{N}$ such that

$$\text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(\overbrace{S, \dots, S}^{n \text{ times}}, \tilde{B}, \overbrace{S, \dots, S}^{m \text{ times}}, (\tilde{B})^c)) > 0.$$

From Lemma 4.1.8, writing $\nu = \Omega_{\mathcal{T}}^{(n)}(\mu)$, we get that

$$\text{Prob}_{\nu}^{\mathcal{T}}(\text{Cyl}(\tilde{B}, \overbrace{S, \dots, S}^{m \text{ times}}, (\tilde{B})^c)) > 0.$$

And from the third property proven previously, we deduce that

$$\text{Prob}_{\nu_{\tilde{B}}}^{\mathcal{T}}(\mathbf{F} B) > 0$$

with $\nu_{\tilde{B}} \in \text{Dist}(\tilde{B})$ which contradicts the second property of this lemma.

Finally, we prove the last property. It is straightforward by observing that the two events measured in this equality are exactly the same:

$$\text{Ev}_{\mathcal{T}}(\mathbf{F} B \vee \mathbf{F} \widetilde{B}) = \text{Ev}_{\mathcal{T}}(\mathbf{F} B \vee (\neg B \mathbf{U} \widetilde{B})).$$

□

We are now ready to define different decisiveness concepts. Two of them comes from [ABM07] (though no initial distribution was fixed) while the third one comes from [BBBC16] where we identified the notion to be a useful alternative.

Definition 4.2.3. Let μ be an initial probability distribution ($\mu \in \text{Dist}(S)$). Then:

- \mathcal{T} is *decisive w.r.t. B from μ* whenever $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$; we then write that \mathcal{T} is $\text{Dec}(\mu, B)$.
- \mathcal{T} is *strongly decisive w.r.t. B from μ* whenever $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$; we then write that \mathcal{T} is $\text{StrDec}(\mu, B)$.
- \mathcal{T} is *persistently decisive w.r.t. B from μ* whenever for every $k \geq 0$, $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F}_{\geq k} B \vee \mathbf{F}_{\geq k} \widetilde{B}) = 1$; we then write that \mathcal{T} is $\text{PersDec}(\mu, B)$.

Furthermore: \mathcal{T} is (strongly, persistently) decisive w.r.t. B whenever it is (strongly, persistently) decisive w.r.t. B from every initial distribution μ ; We then write that \mathcal{T} is $\text{Dec}(B)$ (resp. $\text{StrDec}(B)$, $\text{PersDec}(B)$). Also, given $\mathcal{B} \subseteq \Sigma$, \mathcal{T} is (strongly, persistently) decisive w.r.t. \mathcal{B} from μ if it is $\text{Dec}(\mu, B)$ (resp. $\text{StrDec}(\mu, B)$, $\text{PersDec}(\mu, B)$) for each $B \in \mathcal{B}$. We write \mathcal{T} is $\text{Dec}(\mu, \mathcal{B})$ (resp. $\text{StrDec}(\mu, \mathcal{B})$, $\text{PersDec}(\mu, \mathcal{B})$). Similarly \mathcal{T} is (strongly, persistently) decisive w.r.t. \mathcal{B} if it is $\text{Dec}(B)$ (resp. $\text{StrDec}(B)$, $\text{PersDec}(B)$) for each $B \in \mathcal{B}$. We write \mathcal{T} is $\text{Dec}(\mathcal{B})$ (resp. $\text{StrDec}(\mathcal{B})$, $\text{PersDec}(\mathcal{B})$).

Remark 4.2.4. Observe that in the case where \mathcal{T} is a DMC, it is straightforward to get that \mathcal{T} is (strongly) decisive in the sense of Definition 2.2.12 if and only if \mathcal{T} is $\text{Dec}(2^S)$ (resp. $\text{StrDec}(2^S)$).

Intuitively, the (simple) decisiveness property says that, almost-surely, either B will eventually be visited, or states from which B can no more be reached will eventually be visited. It denotes a dichotomy between the behaviours of the STS \mathcal{T} : there are those behaviours that visit B , and those that do not visit B , but then visit \widetilde{B} ; other behaviours have probability 0 to occur. Strong decisiveness imposes a similar dichotomy, but between behaviours that visit B infinitely often and behaviours that visit \widetilde{B} . Persistent decisiveness refines simple decisiveness, except that here we look at an arbitrary horizon. It can also be seen as being decisive from $\Omega_{\mathcal{T}}^{(n)}(\mu)$ for each $n \geq 0$.

Example 4.2.5. Let us consider again the STS \mathcal{T}_2 of Example 4.1.2, representing the random walk over \mathbb{N} as a DMC. We have already shown in Example 2.2.14 that $\widetilde{B} = \emptyset$ for any set of states B and that if $p > 1/2$, \mathcal{T}_2 is not $\text{Dec}(2^{S_2})$ nor $\text{StrDec}(2^{S_2})$. However here, we can prove decisiveness from some initial distribution μ . We consider $\mu = \delta_0$, the Dirac distribution over state 0. Then it can be shown that for each set of states B , $\text{Prob}_\mu^{\mathcal{T}_2}(\mathbf{F} B) = 1$ and thus, \mathcal{T}_2 is $\text{Dec}(\mu, B)$. Now if $\mu' = \delta_1$ and $B' = \{0\}$ then $\text{Prob}_{\mu'}^{\mathcal{T}_2}(\mathbf{F} \{0\}) < 1$ as already observed in Example 2.2.11; but since $\widetilde{B}' = \emptyset$, we derive that \mathcal{T}_2 is not $\text{Dec}(\mu', B')$. We now give an example where strong decisiveness is contradicted from an initial distribution. For each $i \geq 0$, we consider $B_i = \{i\}$ and we still consider $\mu = \delta_0$. Since $p > 1/2$, classical results on random walks imply that for each i , $\text{Prob}_\mu^{\mathcal{T}_2}(\mathbf{G} \mathbf{F} B_i) = 0$. And since $\widetilde{B}_i = \emptyset$, we obtain that \mathcal{T}_2 is not $\text{StrDec}(\mu, B_i)$. Finally, one can compute $\Omega_{\mathcal{T}_2}(\mu) = \mu'$. Since \mathcal{T}_2 is not $\text{Dec}(\mu', B')$, we get that \mathcal{T}_2 is not $\text{PersDec}(\mu, B')$.

Consider now the STS \mathcal{T}_1 of Example 4.1.3. Assume that $\lambda > \nu$ and that $\mu = \delta_{(0,0)}$ and fix some $T > 0$. We consider $B_1 = \{1\} \times [0, T]$. Then one can compute $\widetilde{B} = \mathbb{N} \times]T, \infty[$ (since the time component can only increase, while all natural numbers are reachable from each other by discrete jumps). Note that here, as time almost-surely always progresses, $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \widetilde{B}) = 1$ and even for each $k \geq 0$, $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F}_{k \geq 0} \widetilde{B}) = 1$. It thus follows that \mathcal{T}_1 is $\text{Dec}(\mu, B)$, $\text{StrDec}(\mu, B)$ and $\text{PersDec}(\mu, B)$.

4.2.2 Attractors

The notion of finite attractor has been used in several contexts like probabilistic lossy channel systems (see *e.g.* [ABRS05]) but in [ABM07] in the context of DMCs (see Section 2.2) where finite attractors were proved to imply decisiveness (see Proposition 2.2.13). A finite attractor is a finite set of states which is reached almost-surely from every state of the system. We lift this definition to our context, obviously relaxing the finiteness assumption, since it is very unlikely that systems with a continuous state-space will have finite attractors. Since the whole set of states is a trivial attractor, this general definition will appear to be useful once we are able to define attractors with some finiteness property, which will be done through *abstractions* in Chapter 5.

Definition 4.2.6. Let $\mu \in \text{Dist}(S)$ be an initial distribution. We say that $B \in \Sigma$ is a μ -attractor for \mathcal{T} if $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B) = 1$. Furthermore, B is an *attractor for \mathcal{T}* if it is a μ -attractor for every $\mu \in \text{Dist}(S)$.

Observe that if \mathcal{T} is a DMC, Definition 2.2.8 is obviously equivalent to Defi-

inition 4.2.6 for the case without a fixed initial distribution.

Example 4.2.7. We illustrate the notion on the random-walk of Example 4.1.2. Recall that we have already illustrated the notion of attractor on the random-walk in Example 2.2.11. However, we have here also a notion of μ -attractor. Assuming that $p > 1/2$, we have stated that there is no finite attractor. However, for instance for $B = \{5\}$, it can be shown as stated in Example 4.2.5, that B is a μ -attractor for $\mu = \delta_0$. However, for any distribution $\mu' \in \text{Dist}(\mathbb{N}_{\geq 6})$ over natural numbers greater than 6, $\text{Prob}_{\mu'}^{\mathcal{T}_2}(\mathbf{F} B) < 1$ and thus B is not a μ' -attractor.

On the other hand as stated in Example 2.2.11, if we assume $p \leq 1/2$, it is a well-known property of random walks that $\{0\}$ is reached almost-surely from every state, hence we can infer that any bounded subset A of \mathbb{N} is an attractor (for every initial distribution).

The notion of attractor is very strong in STSs in the sense that it can be reached with probability 1 from any state. Even in our general context, the following strong property is then satisfied.

Lemma 4.2.8. If B is an attractor for \mathcal{T} then for every initial distribution $\mu \in \text{Dist}(S)$,

$$\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B) = 1.$$

Proof. Let B be an attractor for \mathcal{T} , *i.e.* for each initial distribution $\mu \in \text{Dist}(S)$, $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} B) = 1$. Towards a contradiction, assume that there is $\mu \in \text{Dist}(S)$ such that $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B) < 1$. Then, $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \mathbf{G} B^c) > 0$. Now recall that from the definitions, we have that

$$\text{Ev}_{\mathcal{T}}(\mathbf{F} \mathbf{G} B^c) = \bigcup_{n \geq 0} \bigcap_{m \geq 0} \text{Cyl}(\overbrace{S, \dots, S}^{n \text{ times}}, \overbrace{B^c, \dots, B^c}^{m \text{ times}}).$$

It follows that there is $n \in \mathbb{N}$ such that

$$\lim_{m \rightarrow \infty} \text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(\overbrace{S, \dots, S}^{n \text{ times}}, \overbrace{B^c, \dots, B^c}^{m \text{ times}})) > 0.$$

From Lemma 4.1.8, if we write $\nu_0 = \mu$ and $\nu_j = \Omega_{\mathcal{T}}(\nu_{j-1})$ for each $1 \leq j \leq n-1$, we get that for each $m \geq 1$,

$$\text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(\overbrace{S, \dots, S}^{n \text{ times}}, \overbrace{B^c, \dots, B^c}^{m \text{ times}})) = \text{Prob}_{\Omega_{\mathcal{T}}(\nu_{n-1})}^{\mathcal{T}}(\text{Cyl}(\overbrace{B^c, \dots, B^c}^{m \text{ times}}))$$

since $\mu(S) = 1$ and for each $0 \leq j \leq n - 2$, $(\Omega_{\mathcal{T}}(\nu_j))(S) = 1$. It can be seen that in this case, for each $0 \leq j \leq n - 1$, $\nu_j = \Omega_{\mathcal{T}}^{(j+1)}(\mu)$. We write $\nu = \Omega_{\mathcal{T}}(\nu_{n-1}) = \Omega_{\mathcal{T}}^{(n)}(\mu) \in \text{Dist}(S)$. We thus get that

$$\lim_{m \rightarrow \infty} \text{Prob}_{\nu}^{\mathcal{T}}(\text{Cyl}(\overbrace{B^c, \dots, B^c}^{m \text{ times}})) = \text{Prob}_{\nu}^{\mathcal{T}}(\mathbf{G} B^c) > 0,$$

which contradicts the fact that B is an attractor, hence a ν -attractor, for \mathcal{T} . \square

4.2.3 Fairness

Fairness is a standard notion in probabilistic systems [Pnu83], saying that something which is allowed infinitely often should happen infinitely often almost-surely. This can for instance be instantiated in DMCs as follows: if a state s is visited infinitely often, and the probability to move from s to s' is positive then, almost-surely, the state s' is visited infinitely often. It is well-known that not all Markov chains are fair, but finitely-branching Markov chains are fair. Fairness cannot be lifted directly to continuous state-space STSs (since for two states s and s' , the probability to move from s to s' is likely to be 0). A more careful definition of this notion must be provided for general STSs. It has also been studied in STA [BBB⁺14] (see Section 3.3) but the following definition is a generalization as we will see in Chapter 7.

For $B \in \Sigma$, we define

$$\text{PreProb}^{\mathcal{T}}(B) = \{B' \in \Sigma \mid \forall \mu' \in \text{Dist}(B'), \text{Prob}_{\mu'}^{\mathcal{T}}(\text{Cyl}(B', B)) > 0\}$$

as the set of measurable sets of states B' from which B can be reached with positive probability. Note that, ideally we would like to define the maximal set that allows one to reach B , but the union of all such sets may not be measurable in our general context.

Definition 4.2.9. Let $\mu \in \text{Dist}(S)$ be some initial distribution, and $B \in \Sigma$. We say that \mathcal{T} is *fair w.r.t. B from μ* , written \mathcal{T} is $\text{fair}(\mu, B)$, if for every $B' \in \text{PreProb}^{\mathcal{T}}(B)$, $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B') > 0$ implies

$$\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B \mid \mathbf{G} \mathbf{F} B') = 1.$$

Furthermore: if $\mathcal{B} \subseteq \Sigma$, \mathcal{T} is *fair w.r.t. \mathcal{B} from μ* whenever it is $\text{fair}(\mu, B)$ for each $B \in \mathcal{B}$. We write \mathcal{T} is $\text{fair}(\mu, \mathcal{B})$. We say that \mathcal{T} is *fair w.r.t. \mathcal{B}* if it is $\text{fair}(\mu, \mathcal{B})$ for each initial distribution μ . We write \mathcal{T} is $\text{fair}(\mathcal{B})$. Finally, we say that \mathcal{T} is *strongly fair* whenever it is fair wr.t. B from μ for every $B \in \Sigma$ and every $\mu \in \text{Dist}(S)$.

Example 4.2.10. Consider again the random walk of Example 4.1.2. We can show that \mathcal{T}_2 is strongly fair by observing that there is a positive lower bound on the non-zero probabilities to reach any set of states. Formally there exists $\varepsilon > 0$ such that for each $B \subseteq S_2$, for each $B' \in \text{PreProb}^{\mathcal{T}_2}(B)$ and for each $s \in B'$, $\kappa_2(s, B) \geq \varepsilon$. It suffices to choose $\varepsilon = \min(p, 1 - p) > 0$.

Example 4.2.11 (Counter-example). Consider now the DMC \mathcal{T}_3 depicted in Figure 4.2. The DMC can be described as follows. The set of states is denumerable and composed of: state b that is reachable from each state of the DMC and the states a_1, a_2, a_3, \dots that are connected sequentially: $a_1 \rightarrow a_2, a_2 \rightarrow a_3$ and so on. And thus each a_i can reach state b in one step. In b , the system moves to a_1 with probability 1. Once in the sequence of a_i 's, the probability to stay within the sequence gets bigger and bigger and converges towards 1, or said otherwise the probability to reach state b decreases and converges towards 0. More precisely, if we enter state a_n , the probability to move to b is given by $\frac{1}{3^n}$ while the probability to move to a_{n+1} is given by $1 - \frac{1}{3^n}$.

Consider $B = \{b\}$, $\mu = \delta_b$ and $B' = \{a_n \mid n \in \mathbb{N}\}$, $B' \in \text{PreProb}^{\mathcal{T}_3}(B)$. It holds that $\text{Prob}_\mu^{\mathcal{T}_3}(\mathbf{G F} B') > 0$, however, $\text{Prob}_\mu^{\mathcal{T}_3}(\mathbf{G F} B \mid \mathbf{G F} B') < 1$ and thus \mathcal{T}_3 is not fair(μ, B).




Figure 4.2: A DMC \mathcal{T}_3 that is not strongly fair.

4.2.4 Relationships between the various properties

In this section, we compare all the notions, and give the precise links between all these notions. We first analyse the general case, and reinforce the results in the case of DMCs obtaining the results of [ABM07].

We can establish the following links between the notions of decisiveness and fairness. The first result is straightforward.

Lemma 4.2.12. For each $\mathcal{B} \subseteq \Sigma$ and for each $\mu \in \text{Dist}(S)$, it holds that $\text{Dec}(\mathcal{B})$ (resp. $\text{StrDec}(\mathcal{B})$, $\text{PersDec}(\mathcal{B})$) implies $\text{Dec}(\mu, \mathcal{B})$ (resp. $\text{StrDec}(\mu, \mathcal{B})$, $\text{PersDec}(\mu, \mathcal{B})$), and $\text{fair}(\mathcal{B})$ implies $\text{fair}(\mu, \mathcal{B})$.

We also get straightforwardly from the definitions, the following implication.

Lemma 4.2.13. For each $\mathcal{B} \subseteq \Sigma$ and for each $\mu \in \text{Dist}(S)$, it holds that $\text{StrDec}(\mu, \mathcal{B})$ implies $\text{Dec}(\mu, \mathcal{B})$, and $\text{PersDec}(\mu, \mathcal{B})$ implies $\text{Dec}(\mu, \mathcal{B})$.

It then turns out that strong decisiveness and persistent decisiveness are two equivalent notions.

Lemma 4.2.14. For each $\mathcal{B} \subseteq \Sigma$ and for each $\mu \in \text{Dist}(S)$, it holds that $\text{StrDec}(\mu, \mathcal{B})$ is equivalent to $\text{PersDec}(\mu, \mathcal{B})$.

Proof. Fix $\mathcal{B} \subseteq \Sigma$, $\mu \in \text{Dist}(S)$ and $B \in \mathcal{B}$. Assume that \mathcal{T} is $\text{PersDec}(\mu, B)$, *i.e.* for each $k \geq 0$, $\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq k} B \vee \mathbf{F}_{\geq k} \widetilde{B}) = 1$. We want to show that \mathcal{T} is $\text{StrDec}(\mu, B)$, *i.e.* that $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$, or equivalently that $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \mathbf{G} B^c \wedge \mathbf{G} (\widetilde{B})^c) = 0$. We have that:

$$\begin{aligned} \text{Prob}_\mu^\mathcal{T}(\mathbf{F} \mathbf{G} B^c \wedge \mathbf{G} (\widetilde{B})^c) &\leq \sum_{k \geq 0} \text{Prob}_\mu^\mathcal{T}(\mathbf{G}_{\geq k} (B^c \cap (\widetilde{B})^c)) \\ &= \sum_{k \geq 0} (1 - \text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq k} B \vee \mathbf{F}_{\geq k} \widetilde{B})) \\ &= 0 \text{ from the hypothesis.} \end{aligned}$$

Hence we get that $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$ and thus \mathcal{T} is $\text{StrDec}(\mu, B)$ and $\text{StrDec}(\mu, \mathcal{B})$ as it holds true for each $B \in \mathcal{B}$.

Now fix again $B \in \mathcal{B}$ and assume that \mathcal{T} is $\text{StrDec}(\mu, B)$, *i.e.* $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$. From Lemma 4.2.2 (fourth item), we get that $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \vee \mathbf{G} \mathbf{F} \widetilde{B}) = 1$ and it is then straightforward to establish that for each $k \geq 0$, $\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq k} B \vee \mathbf{F}_{\geq k} \widetilde{B}) = 1$. We hence deduce that \mathcal{T} is $\text{PersDec}(\mu, B)$ and thus $\text{PersDec}(\mu, \mathcal{B})$ as it holds true for each $B \in \mathcal{B}$. This concludes the proof. \square

Now, we have the following equivalences between the decisiveness notions.

Lemma 4.2.15. For each $\mathcal{B} \subseteq \Sigma$, it holds that all three notions $\text{PersDec}(\mathcal{B})$, $\text{StrDec}(\mathcal{B})$ and $\text{Dec}(\mathcal{B})$ are equivalent.

Proof. Fix $\mathcal{B} \subseteq \Sigma$. From Lemmas 4.2.13 and 4.2.14, it remains to prove that $\text{Dec}(\mathcal{B}) \implies \text{StrDec}(\mathcal{B})$ or $\text{Dec}(\mathcal{B}) \implies \text{PersDec}(\mathcal{B})$. We prove the last one. We

pick $B \in \mathcal{B}$ and assume that \mathcal{T} is $\text{Dec}(B)$, *i.e.* for each $\mu \in \text{Dist}(S)$, $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$. Pick $\mu \in \text{Dist}(S)$ and $k \geq 0$. We get that

$$\begin{aligned} \text{Prob}_\mu^\mathcal{T}(\mathbf{G}_{\geq k} B^c \wedge \mathbf{G}_{\geq k} (\widetilde{B})^c) &\leq \text{Prob}_{\mu_k}^\mathcal{T}(\mathbf{G}(B^c \cap (\widetilde{B})^c)) \\ &\leq 0 \quad \text{since } \mathcal{T} \text{ is } \text{Dec}(B), \end{aligned}$$

where the first inequality stands from Lemma 4.1.8 with $\mu_k = \Omega_\mathcal{T}^{(k)}(\mu)$ and from a similar argument as in the proof of Lemma 4.2.8. Hence for each $k \geq 0$, $\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq k} B \vee \mathbf{F}_{\geq k} \widetilde{B}) = 1$ and since it holds true for each $\mu \in \text{Dist}(S)$ and each $B \in \mathcal{B}$, we get that \mathcal{T} is $\text{PersDec}(\mathcal{B})$. \square

Finally, we show the following links between fairness and decisiveness.

Lemma 4.2.16. For each $\mathcal{B} \subseteq \Sigma$ and for each $\mu \in \text{Dist}(S)$, it holds that $\text{StrDec}(\mu, \mathcal{B})$ implies $\text{fair}(\mu, \mathcal{B})$, and $\text{StrDec}(\mathcal{B})$ implies $\text{fair}(\mathcal{B})$.

Proof. Fix $\mathcal{B} \subseteq \Sigma$ and $\mu \in \text{Dist}(S)$. Assume that \mathcal{T} is strongly decisive w.r.t. \mathcal{B} from μ , that is for each $B \in \mathcal{B}$, $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$. We want to prove that for each $B \in \mathcal{B}$, for each $B' \in \text{PreProb}(B)$ with $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B') > 0$, we have that $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \mid \mathbf{G} \mathbf{F} B') = 1$.

Fix $B \in \mathcal{B}$ and $B' \in \text{PreProb}(B)$ such that $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B') > 0$. We can notice that

$$\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B' \wedge \mathbf{F} \widetilde{B}) = 0. \quad (4.3)$$

Indeed, towards a contradiction, assume that $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B' \wedge \mathbf{F} \widetilde{B}) > 0$. Observe that

$$\text{Ev}_\mathcal{T}(\mathbf{G} \mathbf{F} B' \wedge \mathbf{F} \widetilde{B}) = \bigcup_{n \geq 0} \bigcap_{m \geq 0} \bigcup_{l \geq m} \text{Cyl}(\overbrace{S, \dots, S}^{n \text{ times}}, \widetilde{B}, \overbrace{S, \dots, S}^{l \text{ times}}, B').$$

Then, there are $n, m \in \mathbb{N}$ such that

$$\text{Prob}_\mu^\mathcal{T}(\text{Cyl}(\overbrace{S, \dots, S}^{n \text{ times}}, \widetilde{B}, \overbrace{S, \dots, S}^{m \text{ times}}, B')) > 0.$$

It follows, from Lemma 4.1.8 like seen previously, that there is $\nu \in \text{Dist}(S)$ ($\nu = \Omega_\mathcal{T}^{(n)}(\mu)$), such that

$$\text{Prob}_\nu^\mathcal{T}(\text{Cyl}(\widetilde{B}, \overbrace{S, \dots, S}^{m \text{ times}}, B')) > 0.$$

And since $B' \in \text{PreProb}(B)$, we get that

$$\text{Prob}_\nu^{\mathcal{T}}(\text{Cyl}(\widetilde{B}, \overbrace{S, \dots, S}^{m \text{ times}}, B', B)) > 0.$$

Hence, $\nu(\widetilde{B}) > 0$ and we can apply Lemma 4.2.2 (second item) to obtain a contradiction. Hence, equation (4.3) holds. We then write:

$$\begin{aligned} 1 &= \text{Prob}_\mu^{\mathcal{T}}(\mathbf{G F B} \vee \mathbf{F} \widetilde{B} \mid \mathbf{G F B}') \quad \text{from strong decisiveness} \\ &= \frac{\text{Prob}_\mu^{\mathcal{T}}((\mathbf{G F B} \vee \mathbf{F} \widetilde{B}) \wedge \mathbf{G F B}')}{\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G F B}')} \\ &= \frac{\text{Prob}_\mu^{\mathcal{T}}((\mathbf{G F B} \wedge \mathbf{G F B}') \vee (\mathbf{F} \widetilde{B} \wedge \mathbf{G F B}'))}{\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G F B}')} \\ &= \frac{\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G F B} \wedge \mathbf{G F B}')}{\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G F B}')} \quad \text{from (4.3)} \\ &= \text{Prob}_\mu^{\mathcal{T}}(\mathbf{G F B} \mid \mathbf{G F B}') \end{aligned}$$

which proves that $\text{StrDec}(\mu, \mathcal{B}) \implies \text{fair}(\mu, \mathcal{B})$. The implication $\text{StrDec}(\mathcal{B}) \implies \text{fair}(\mathcal{B})$ is then immediate since the previous implication holds for any initial distribution $\mu \in \text{Dist}(S)$. \square

We can summarize the previous implications as follows:

Proposition 4.2.17. *For each $\mathcal{B} \subseteq \Sigma$ and for each $\mu \in \text{Dist}(S)$, it holds that*

$$\mathcal{T} \text{ is } \text{Dec}(\mu, \mathcal{B}) \iff \mathcal{T} \text{ is } \text{StrDec}(\mu, \mathcal{B}) \iff \mathcal{T} \text{ is } \text{PersDec}(\mu, \mathcal{B}) \implies \mathcal{T} \text{ is } \text{fair}(\mu, \mathcal{B})$$

$$\mathcal{T} \text{ is } \text{Dec}(\mathcal{B}) \iff \mathcal{T} \text{ is } \text{StrDec}(\mathcal{B}) \iff \mathcal{T} \text{ is } \text{PersDec}(\mathcal{B}) \implies \mathcal{T} \text{ is } \text{fair}(\mathcal{B})$$

The three missing implications in the above proposition do actually not hold, as witnessed by the following example. We also illustrate the fact that $\text{Dec}(\mu, \mathcal{B})$ and $\text{fair}(\mu, \mathcal{B})$ are incomparable.

Example 4.2.18 (Counter-example). Consider the random walk \mathcal{T}_2 of Example 4.1.2. We have shown in Example 4.2.10 that \mathcal{T}_2 is strongly fair. Now let us assume that $p > 1/2$ and let us consider the initial distribution $\mu = \delta_0$, the Dirac distribution over 0. Then from Example 4.2.5, \mathcal{T}_2 is decisive from μ w.r.t. any set of states. Again in this example, we have observed that it is not strongly decisive w.r.t. any set of the form $B = \{i\}$ with $i \geq 0$ from μ . This shows that

we do not have $\text{Dec}(\mu, \mathcal{B}) \implies \text{StrDec}(\mu, \mathcal{B})$, nor $\text{fair}(\mu, \mathcal{B}) \implies \text{StrDec}(\mu, \mathcal{B})$ and $\text{fair}(\mathcal{B}) \implies \text{StrDec}(\mathcal{B})$. And since \mathcal{T}_2 is not decisive from δ_1 w.r.t. $\{0\}$, this also proves that $\text{fair}(\mu, \mathcal{B})$ does not imply $\text{Dec}(\mu, \mathcal{B})$.

In order to illustrate that $\text{Dec}(\mu, \mathcal{B})$ does not imply $\text{fair}(\mu, \mathcal{B})$ in general, we consider the DMC \mathcal{T}_3 of Example 4.2.11. We consider $B = \{b\}$ and $\mu = \delta_b$. It is easily observed that \mathcal{T}_3 is $\text{Dec}(\mu, B)$ as we start in b with probability 1, but we have shown that \mathcal{T}_3 is not $\text{fair}(\mu, B)$.

If \mathcal{T} is a DMC, *i.e.* if S is at most denumerable and $\Sigma = 2^S$, we can complete the picture using the result of [ABM07] presented in Section 2.2.2 as Proposition 2.2.13 bonding the notion of finite attractor and decisiveness in DMCs. We can then sum up the implications in DMCs as follows:

$$\left. \begin{array}{l} \mathcal{T} \text{ DMC with a} \\ \text{finite attractor} \end{array} \right\} \implies \mathcal{T} \text{ is } \text{Dec}(2^S) \iff \mathcal{T} \text{ is } \text{StrDec}(2^S) \iff \mathcal{T} \text{ is } \text{PersDec}(2^S)$$

$$\Downarrow$$

$$\mathcal{T} \text{ is strongly fair}$$

Note that for the equivalence between decisiveness and strong decisiveness, we recover Lemma 3.2 of [ABM07].

4.3. Concluding remarks

The objective is now to study the qualitative and quantitative model-checking problems of STSs (see Definitions 4.1.19 and 4.1.20). Inspired from the work of [ABM07] on DMC, we would like to get similar results as the ones presented in Section 2.2.2. First of all, it should be noted we will not be able to obtain as nice results as Propositions 2.2.15, 2.2.16, 2.2.17 and 2.2.18 where basically, the almost-sure model-checking problem on DMCs (Definition 2.2.6) of (repeated) reachability properties was reduced to structural properties of the underlying graph: in our more general context of STSs, we cannot consider an underlying graph.

Also we have observed that the notion of decisiveness played a key part in proving all those results on DMCs. Hence the work here consists of, in a first time adapting the results to our general context, and then getting new results for all LTL formulas. Hence we will need to identify classes of decisive STSs on which we get results for the qualitative and quantitative model-checking problems on STSs of LTL formulas. This will be done in Chapter 6.

However proving that a general STS is decisive can be very technical. Hence the purpose of Chapter 5 will be to deal with this difficulty by defining a notion of abstraction.

Abstraction Between STSs

As said before, in order to tackle the qualitative and quantitative model-checking problems for STSs (see Definitions 4.1.19 and 4.1.20), inspired from the work of [ABM07] (see Section 2.2.2), we will need to identify classes of decisive STSs.

While decisiveness is well-defined for general STSs, proving that a given STS \mathcal{T} is decisive might be technical in general. A standard approach in model-checking to avoid such difficulties is to abstract the system into a simpler one, that can be analysed and that provides guarantees on the concrete system. In this chapter, we thus propose a notion of abstraction, which will help proving properties of general STSs. Also, through abstractions, we will be able to characterise meaningful attractors.

In Section 5.1 we define the notion of α -abstraction for STSs. Roughly speaking, STS \mathcal{T}_2 is an α -abstraction of STS \mathcal{T}_1 if it preserves the qualitative one-step behaviour of \mathcal{T}_1 . We prove basic results of abstractions and we introduce the notions of *sound* and *complete* α -abstractions. Those notions will allow to transfer richest properties from the abstraction to the STS and vice-versa.

In Section 5.2, have a look at decisiveness properties and show under which conditions those properties are transferred through abstractions. We will have a particular interest for denumerable abstractions (*i.e.* abstractions that are DMCs) as those are easier to analyse. We will also be invested in attractors and fairness.

Finally, in Section 5.3, we identify conditions for the soundness and the completeness of abstractions.

5.1. Abstraction

In this section, we introduce the notion of α -abstraction. As said before, an STS \mathcal{T}_2 is an α -abstraction of STS \mathcal{T}_1 if it preserves the positive probabilities of one-step moves. In Section 9.2, we will show basic properties of α -abstraction. Those define the properties that are preserved through abstractions. As they are not rich enough for what we intend to analyse, we define the notions of decisiveness and completeness in Section 5.1.2.

Let $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1)$ and $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2)$ be two STSs. Let $\alpha : (S_1, \Sigma_1) \rightarrow (S_2, \Sigma_2)$ be a measurable function. A set $B \in \Sigma_1$ is said *α -closed* whenever $B = \alpha^{-1}(\alpha(B))$: for every $s, s' \in S_1$, if $s \in B$ and $\alpha(s) = \alpha(s')$, then $s' \in B$. Following [GBK16], we define the pushforward of α as $\alpha_{\#} : \text{Dist}(S_1) \rightarrow \text{Dist}(S_2)$ by $\alpha_{\#}(\mu)(M_2) = \mu(\alpha^{-1}(M_2))$ for every $\mu \in \text{Dist}(S_1)$ and for every $M_2 \in \Sigma_2$. The role of the pushforward $\alpha_{\#}$ is to transfer the measures from (S_1, Σ_1) to (S_2, Σ_2) .

Definition 5.1.1. We say that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 if

$$\forall \mu \in \text{Dist}(S_1), \alpha_{\#}(\Omega_{\mathcal{T}_1}(\mu)) \text{ is equivalent to } \Omega_{\mathcal{T}_2}(\alpha_{\#}(\mu)).$$

From the definitions of $\Omega_{\mathcal{T}}$ (see Section 4.1), $\alpha_{\#}$ and equivalent measures, the notion of α -abstraction equivalently requires that for every $\mu \in \text{Dist}(S_1)$ and every $A \in \Sigma_2$,

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(S_1, \alpha^{-1}(A))) > 0 \iff \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(S_2, A)) > 0 .$$

Intuitively, the two STSs have the same one-step qualitative behaviour.

The notion of α -abstraction naturally extends to LSTSs as follows.

Definition 5.1.2. LSTS $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2, \text{AP}_2, \mathcal{L}_2)$ is an α -abstraction of LSTS $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1, \text{AP}_1, \mathcal{L}_1)$ whenever:

- $(S_2, \Sigma_2, \kappa_2)$ is an α -abstraction of $(S_1, \Sigma_1, \kappa_1)$;
- $\text{AP}_1 = \text{AP}_2$;
- for every $s_1, s'_1 \in S_1$, $\alpha(s_1) = \alpha(s'_1) \implies \mathcal{L}_1(s_1) = \mathcal{L}_1(s'_1)$;
- for every $s \in S_1$, $\mathcal{L}_1(s) = a \implies \mathcal{L}_2(\alpha(s)) = a$.

The two last conditions imply that for each $a \in 2^{\text{AP}}$, $\mathcal{L}_1^{-1}(\{a\})$ is α -closed. Moreover, for each $a \in 2^{\text{AP}}$, $\alpha^{-1}(\mathcal{L}_2^{-1}(\{a\})) = \mathcal{L}_1^{-1}(\{a\})$.

We illustrate the notion of α -abstraction on STSs \mathcal{T}_1 and \mathcal{T}_2 of Examples 4.1.2 and 4.1.3.

Example 5.1.3. Consider again the STS \mathcal{T}_2 with parameter $p \in]0, 1[$ and \mathcal{T}_1 with parameters λ and $\nu > 0$ of Examples 4.1.2 and 4.1.3. Let $\alpha : S_1 \rightarrow S_2$ be the mapping defined as follows: for every $i \in \mathbb{N}$ and every $t \in \mathbb{R}_+$, $\alpha((i, t)) = i$. It can be shown that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . We will illustrate it by considering any distribution $\mu \in \text{Dist}(\{0\} \times \mathbb{R}_+)$. Then, observe that the pushforward of μ corresponds to $\alpha_{\#}(\mu) = \delta_0 \in \text{Dist}(S_2)$ (note that this will be formally proven in Lemma 5.1.4). It follows that for any $A \subseteq S_2$, $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(S_2, A)) > 0$ if and only if $1 \in A$ if and only if $\{1\} \times \mathbb{R}_+ \subseteq \alpha^{-1}(A)$. From the definition of κ_1 (see Example 4.1.3), it is then obvious that this last condition is equivalent to $\text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(S_1, \alpha^{-1}(A))) > 0$.

It should be noted that a similar reasoning can be applied to any initial distribution $\mu_i \in \text{Dist}(\{i\} \times \mathbb{R}_+)$ with $i \in \mathbb{N}$ and thus to any initial distribution $\mu' \in \text{Dist}(S_1)$.

We fix two STSs $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1)$ and $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2)$ for the rest of this chapter.

5.1.1 Properties of abstractions

We now establish several technical results, which explicit how STSs are related through an α -abstraction. The relationship is only qualitative, in the sense that it only relates positive reachability probabilities, but does not relate almost-sure or lower-bounded probabilities.

Lemma 5.1.4. Let $\alpha : (S_1, \Sigma_1) \rightarrow (S_2, \Sigma_2)$ be a measurable function. Then for every $s \in S_2$ and every $\mu \in \text{Dist}(\alpha^{-1}(\{s\}))$, $\alpha_{\#}(\mu) = \delta_s$.

Proof. Fix $s \in S_2$ and $\mu \in \text{Dist}(\alpha^{-1}(\{s\}))$. For each $A \in \Sigma_2$, we have that $(\alpha_{\#}(\mu))(A) = \mu(\alpha^{-1}(A))$. If $s \in A$, then trivially $\alpha^{-1}(\{s\}) \subseteq \alpha^{-1}(A)$ and thus $\mu(\alpha^{-1}(A)) = 1$. Otherwise, if $s \notin A$, then $\alpha^{-1}(\{s\}) \cap \alpha^{-1}(A) = \emptyset$ and thus $\mu(\alpha^{-1}(A)) = 0$. This directly implies that $\alpha_{\#}(\mu) = \delta_s$. \square

Lemma 5.1.5. Assume that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Then, for every $i \in \mathbb{N}$, for every $\mu \in \text{Dist}(s_1)$, $\alpha_{\#}(\Omega_{\mathcal{T}_1}^{(i)}(\mu))$ is equivalent to $\Omega_{\mathcal{T}_2}^{(i)}(\alpha_{\#}(\mu))$.

Proof. We show this by induction on i . Case $i = 1$ is by definition of α -abstraction (Definition 5.1.1). Fix some $i \geq 1$ and assume that the statement holds true for each $1 \leq j \leq i$. By induction hypothesis, we have that $\alpha_{\#}(\Omega_{\mathcal{T}_1}^{(i)}(\mu))$ is equivalent to $\Omega_{\mathcal{T}_2}^{(i)}(\alpha_{\#}(\mu))$. We want to show that $\alpha_{\#}(\Omega_{\mathcal{T}_1}^{(i+1)}(\mu))$ is equivalent to $\Omega_{\mathcal{T}_2}^{(i+1)}(\alpha_{\#}(\mu))$.

We first notice that $\Omega_{\mathcal{T}_2}(\alpha_{\#}(\Omega_{\mathcal{T}_1}^{(i)}(\mu)))$ is equivalent to $\Omega_{\mathcal{T}_2}^{(i+1)}(\alpha_{\#}(\mu))$. Indeed write $\nu = \alpha_{\#}(\Omega_{\mathcal{T}_1}^{(i)}(\mu))$ and $\nu' = \Omega_{\mathcal{T}_2}^{(i)}(\alpha_{\#}(\mu))$. From the induction hypothesis, we know that ν and ν' are equivalent. Following a similar argument as in the proof of Lemma 4.1.7 and from the definition of $\Omega_{\mathcal{T}_2}$ (see Section 4.1, we can deduce that $\Omega_{\mathcal{T}_2}(\nu)$ is equivalent to $\Omega_{\mathcal{T}_2}(\nu')$. So it remains to show that $\Omega_{\mathcal{T}_2}(\alpha_{\#}(\mu'))$ is equivalent to $\alpha_{\#}(\Omega_{\mathcal{T}_1}(\mu'))$, where $\mu' = \Omega_{\mathcal{T}_1}^{(i)}(\mu)$. This is by definition of an α -abstraction. \square

In other words, Lemma 5.1.5 states that for each $A \in \Sigma_2$ and for each $i \in \mathbb{N}$,

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F}_{=i} \alpha^{-1}(A)) > 0 \iff \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F}_{=i} A) > 0,$$

i.e. α -abstractions preserve the positive qualitative behaviour in any step. This can even be generalised to cylinders.

Lemma 5.1.6. Assume that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Then for every $\mu \in \text{Dist}(S_1)$, for every $(A_i)_{0 \leq i \leq n} \in \Sigma_2^{n+1}$,

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_0), \dots, \alpha^{-1}(A_n))) > 0 \iff \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(A_0, \dots, A_n)) > 0.$$

Proof. We do the proof by induction on n . The case $n = 0$ is obvious from the definition of $\alpha_{\#}$. Now fix $n \geq 1$ and assume that for each $0 \leq k \leq n - 1$, for each $\mu \in \text{Dist}(S_1)$ and for each $(A_i)_{0 \leq i \leq k} \in \Sigma_2^{k+1}$,

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_0), \dots, \alpha^{-1}(A_k))) > 0 \iff \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(A_0, \dots, A_k)) > 0.$$

We show that it is still the case for n . Fix $\mu \in \text{Dist}(S_1)$ and $(A_i)_{0 \leq i \leq n+1} \in \Sigma_2^{n+2}$. We let $\nu_0 = \mu_{\alpha^{-1}(A_0)}$ and $\nu'_0 = (\alpha_{\#}(\mu))_{A_0}$. Note that we hence assume that $\mu(\alpha^{-1}(A_0)) > 0$. We first realise that $\nu'_0 = \alpha_{\#}(\nu_0)$. Indeed for each $A \in \Sigma_2$,

$$(\alpha_{\#}(\nu_0))(A) = \nu_0(\alpha^{-1}(A)) = \frac{\mu(\alpha^{-1}(A \cap A_0))}{\mu(\alpha^{-1}(A_0))} = \frac{(\alpha_{\#}(\mu))(A \cap A_0)}{(\alpha_{\#}(\mu))(A_0)} = \nu'_0(A).$$

Then, applying Lemma 4.1.8, we get:

$$\begin{aligned} & \text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_0), \alpha^{-1}(A_1), \dots, \alpha^{-1}(A_n))) \\ &= \mu(\alpha^{-1}(A_0)) \cdot \text{Prob}_{\Omega_{\mathcal{T}_1}(\nu_0)}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_1), \dots, \alpha^{-1}(A_n))) \end{aligned}$$

and

$$\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(A_0, A_1, \dots, A_n)) = (\alpha_{\#}(\mu))(A_0) \cdot \text{Prob}_{\Omega_{\mathcal{T}_2}(\nu'_0)}^{\mathcal{T}_1}(\text{Cyl}(A_1, \dots, A_n)).$$

By definition of an α -abstraction (Definition 5.1.1), the measures $\Omega_{\mathcal{T}_2}(\nu'_0)$ and $\alpha_{\#}(\Omega_{\mathcal{T}_1}(\nu_0))$ are equivalent. Hence from Lemma 4.1.7,

$$\text{Prob}_{\Omega_{\mathcal{T}_2}(\nu'_0)}^{\mathcal{T}_2}(\text{Cyl}(A_1, \dots, A_n)) > 0 \iff \text{Prob}_{\alpha_{\#}(\Omega_{\mathcal{T}_1}(\nu_0))}^{\mathcal{T}_2}(\text{Cyl}(A_1, \dots, A_n)) > 0.$$

From the hypothesis of induction, we get that

$$\begin{aligned} \text{Prob}_{\alpha_{\#}(\Omega_{\mathcal{T}_1}(\nu_0))}^{\mathcal{T}_2}(\text{Cyl}(A_1, \dots, A_n)) > 0 \\ \iff \text{Prob}_{\Omega_{\mathcal{T}_1}(\nu_0)}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_1), \dots, \alpha^{-1}(A_n))) > 0. \end{aligned}$$

Since $(\alpha_{\#}(\mu))(A_0) = \mu(\alpha^{-1}(A_0))$, we conclude:

$$\begin{aligned} \text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_0), \alpha^{-1}(A_1), \dots, \alpha^{-1}(A_n))) > 0 \\ \iff \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(A_0, A_1, \dots, A_n)) > 0. \end{aligned}$$

We still have to consider the case where $\mu(\alpha^{-1}(A_0)) = 0$. In that case, it holds that $(\alpha_{\#}(\mu))(A_0) = 0$ and thus

$$\begin{aligned} \text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_0), \alpha^{-1}(A_1), \dots, \alpha^{-1}(A_n))) = 0 \\ = \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(A_0, A_1, \dots, A_n)) \end{aligned}$$

which terminates the proof. \square

Lemma 5.1.6 states thus that α -abstractions preserve the positive qualitative behaviour. As an immediate consequence, the positivity of probabilities of “Until” formulas (and thus of reachability properties) are preserved through α -abstractions.

Corollary 5.1.7. *Assume that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Then for every $\mu \in \text{Dist}(S_1)$, for every $A, B \in \Sigma_2$:*

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Ev}_{\mathcal{T}_1}(\alpha^{-1}(A) \mathbf{U} \alpha^{-1}(B))) > 0 \iff \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Ev}_{\mathcal{T}_2}(A \mathbf{U} B)) > 0.$$

Note that this however does not apply to liveness properties, like $\text{Ev}_{\mathcal{T}_2}(\mathbf{G F} A)$ with $A \in \Sigma_2$. To ensure that these more involved properties are preserved via abstractions, we will strengthen the assumptions on the abstraction and on the STS.

5.1.2 Soundness and completeness of abstractions

In this section, we introduce the notions of *sound* and *complete* α -abstraction. An α -abstraction is sound if it transfers the almost-sure behaviour of reachability properties to the STS; it is complete if it satisfies the other implication. As the final objective is to reduce the analyse of STSs to simpler systems through abstractions, we will have a peculiar interest for sound abstractions.

We assume that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 .

Definition 5.1.8. Let $\mu \in \text{Dist}(S_1)$. The α -abstraction \mathcal{T}_2 is μ -*sound* whenever for every $B \in \Sigma_2$:

$$\text{Prob}_{\alpha\#(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) = 1 \implies \text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1 .$$

We say that \mathcal{T}_2 is a *sound* α -abstraction of \mathcal{T}_1 if it is μ -sound for every $\mu \in \text{Dist}(S_1)$.

Definition 5.1.9. Let $\mu \in \text{Dist}(S_1)$. The α -abstraction \mathcal{T}_2 is μ -*complete* whenever for every $B \in \Sigma_2$,

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1 \implies \text{Prob}_{\alpha\#(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) = 1$$

We say that \mathcal{T}_2 is a *complete* α -abstraction of \mathcal{T}_1 if it is μ -complete for every $\mu \in \text{Dist}(S_1)$.

Sound and complete abstractions will guarantee that, up to α , the same properties are satisfied almost-surely in \mathcal{T}_1 and \mathcal{T}_2 (provided some properties are satisfied by \mathcal{T}_1 and \mathcal{T}_2).

Example 5.1.10. We go back to Example 5.1.3 where we have shown that \mathcal{T}_2 (see Example 4.1.2) is an α -abstraction of \mathcal{T}_1 (see Example 4.1.3). It can be shown moreover that \mathcal{T}_2 is sound and complete whenever $p > 1/2 \iff \lambda > \nu$. Indeed it suffices to observe that for any $t \geq 0$,

- $\kappa_1((0, t), \{1\} \times \mathbb{R}_+) = 1$, and
- for any $i \geq 1$, $\kappa_1((i, t), \{i+1\} \times \mathbb{R}_+) = \frac{\lambda}{\lambda+\nu}$ and $\kappa_1((i, t), \{i-1\} \times \mathbb{R}_+) = \frac{\nu}{\lambda+\nu}$.

Hence if we consider only α -closed sets in \mathcal{T}_1 , one can see that its behaviour is very similar to \mathcal{T}_2 and it is more than just what can be deduced from an α -abstraction: we have a constant probability to go up in the queue, and a constant probability to go down just like in \mathcal{T}_2 . Now observe that $\frac{\lambda}{\lambda+\nu} > 1/2$ if and only if $\lambda > \nu$. The value $\frac{\lambda}{\lambda+\nu}$ corresponds to the probability to go up in the queue, which is the equivalent part of probability p in \mathcal{T}_2 . From classical results on random-walks, it is then trivial to get that \mathcal{T}_2 is sound and complete whenever $p > 1/2 \iff \lambda > \nu$.

When \mathcal{T}_2 is a DMC, then α -abstraction, soundness and completeness have a simpler characterisation, which will be useful in the proofs.

Lemma 5.1.11. Assume that \mathcal{T}_2 is a DMC. Then:

- \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 iff for every $s, s' \in S_2$,

$$\kappa_2(s, s') > 0 \iff \forall \mu \in \text{Dist}(\alpha^{-1}(\{s\})), \text{Prob}_\mu^{\mathcal{T}_1}(\text{Cyl}(S_1, \alpha^{-1}(\{s'\}))) > 0;$$

- \mathcal{T}_2 is sound iff for every $s \in S_2$ and every $B \in \Sigma_2$,

$$\text{Prob}_s^{\mathcal{T}_2}(\mathbf{F} B) = 1 \implies \forall \mu \in \text{Dist}(\alpha^{-1}(\{s\})), \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1;$$

- \mathcal{T}_2 is complete iff for every $s \in S_2$ and every $B \in \Sigma_2$,

$$\forall \mu \in \text{Dist}(\alpha^{-1}(\{s\})), \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1 \implies \text{Prob}_s^{\mathcal{T}_2}(\mathbf{F} B) = 1.$$

Proof. We handle the case of soundness, other cases are similar. Observe that the implication from left to right is obvious. Now assume that for each $s \in S_2$ and for each $B \in \Sigma_2$, the condition presented in the statement (second item) holds true. Then fix $\mu \in \text{Dist}(S_1)$, $B \in \Sigma_2$ and assume that $\text{Prob}_{\alpha_\#(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) = 1$. We have to show that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1$. Towards a contradiction, assume that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) < 1$. Then, since \mathcal{T}_2 is a DMC (and thus S_2 is denumerable), there is $s \in S_2$ such that $(\alpha_\#(\mu))(s) > 0$ and

$$\text{Prob}_{\mu_{\alpha^{-1}(s)}}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) < 1.$$

From the hypothesis, it follows that $\text{Prob}_s^{\mathcal{T}_2}(\mathbf{F} B) < 1$. Observe that since $\mu(\alpha^{-1}(s)) > 0$, we have that $(\alpha_\#(\mu))(s) > 0$. Hence we get a contradiction by noticing:

$$\begin{aligned} \text{Prob}_{\alpha_\#(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) &= \sum_{s' \in S_2} (\alpha_\#(\mu))(s') \cdot \text{Prob}_{s'}^{\mathcal{T}_2}(\mathbf{F} B) \\ &< \sum_{s' \in S_2} (\alpha_\#(\mu))(s') = (\alpha_\#(\mu))(S_2) = 1 \end{aligned}$$

where the first equality comes from the fact that S_2 is at most denumerable and the strict inequality holds from the fact that $\text{Prob}_s^{\mathcal{T}_2}(\mathbf{F} B) < 1$, $(\alpha_\#(\mu))(s) > 0$ and $\text{Prob}_{s'}^{\mathcal{T}_2}(\mathbf{F} B) \leq 1$ for each $s' \in S_2$. \square

5.2. Transfer of properties through abstractions

In this section, we explain how and under which conditions one can transfer interesting decisiveness, attractor and fairness properties of STSs through abstractions.

5.2.1 The case of sound abstractions

We first consider sound abstractions. We will see that the soundness assumption allows one to transfer decisiveness and attractors properties from the abstraction to the concrete model.

The first result establishes that decisiveness is preserved through sound abstractions.

Proposition 5.2.1. *If \mathcal{T}_2 is a μ -sound α -abstraction of \mathcal{T}_1 , then for every $B \in \Sigma_2$:*

$$\mathcal{T}_2 \text{ is Dec}(\alpha_{\#}(\mu), B) \implies \mathcal{T}_1 \text{ is Dec}(\mu, \alpha^{-1}(B)).$$

This result will play a key role in the sequel. Indeed it states that if you have to prove that a general STS \mathcal{T} is decisive, it suffices to have a sound α -abstractions that is decisive. The idea will then be to find an abstraction on which it is easy to prove decisiveness (for instance, recall that all finite Markov chains, or more generally DMCs with a finite attractor, are decisive w.r.t. any set of states!) in order to imply the decisiveness of \mathcal{T} . We will also see why the notion of decisiveness in STSs is important to us (see Chapter 6).

In order to prove Proposition 5.2.1, we first show the following technical lemma, which relates avoid-sets in \mathcal{T}_1 and in \mathcal{T}_2 .

Lemma 5.2.2. *Let \mathcal{T}_2 be an α -abstraction of \mathcal{T}_1 . Then, for every $B \in \Sigma_2$: $\widetilde{\alpha^{-1}(B)} = \alpha^{-1}(\widetilde{B})$.*

Proof. Fix $B \in \Sigma_2$. We have the following series of equivalences:

$$\begin{aligned} s \in \widetilde{\alpha^{-1}(B)} &\iff \text{Prob}_s^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 0 \\ &\iff \text{Prob}_{\alpha_{\#}(\delta_s)}^{\mathcal{T}_2}(\mathbf{F} B) = 0 \quad (\text{Corollary 5.1.7}). \end{aligned}$$

Now from Lemma 5.1.4, one can show that $\alpha_{\#}(\delta_s) = \delta_{\alpha(s)}$ by noticing that $\delta_s \in \text{Dist}(\alpha^{-1}(\alpha(s)))$. Hence $s \in \widetilde{\alpha^{-1}(B)}$ iff $\alpha(s) \in \widetilde{B}$ (i.e. $s \in \alpha^{-1}(\widetilde{B})$), which concludes the proof. \square

We are now ready to prove Proposition 5.2.1.

Proof of Proposition 5.2.1. Fix $B \in \Sigma_2$ and assume that \mathcal{T}_2 is $\text{Dec}(\alpha_{\#}(\mu), B)$, i.e.

$$\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F} B \vee \mathbf{F} \tilde{B}) = 1. \quad (5.1)$$

In order to show that \mathcal{T}_1 is $\text{Dec}(\mu, \alpha^{-1}(B))$, Lemma 5.2.2 states that it suffices to prove that

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\tilde{B})) = 1.$$

The latter is immediate by (5.1) since \mathcal{T}_2 is μ -sound. \square

Thanks to Proposition 4.2.17, this result obviously extends to stronger decisiveness notions.

Corollary 5.2.3. *If \mathcal{T}_2 is a sound α -abstraction of \mathcal{T}_1 , then for every $B \in \Sigma_2$:*

$$\begin{aligned} \mathcal{T}_2 \text{ is } \text{Dec}(B) \text{ (or equiv. } \text{StrDec}(B), \text{ PersDec}(B)) &\implies \\ \mathcal{T}_1 \text{ is } \text{Dec}(\alpha^{-1}(B)) \text{ (or equiv. } \text{StrDec}(\alpha^{-1}(B)), \text{ PersDec}(\alpha^{-1}(B))) & \end{aligned}$$

The definitions of attractor and of sound α -abstraction yield a similar result, which is straightforward in this case.

Proposition 5.2.4. *If \mathcal{T}_2 is a sound α -abstraction of \mathcal{T}_1 and if $A \in \Sigma_2$ is an attractor for \mathcal{T}_2 , then $\alpha^{-1}(A)$ is an attractor for \mathcal{T}_1 .*

As a direct consequence of Proposition 2.2.13 and Corollary 5.2.3, we get the following result for denumerable abstractions, which will be crucial for designing approximation algorithms taking advantage of abstractions.

Proposition 5.2.5. *Let \mathcal{T}_2 be a DMC with a finite attractor. If \mathcal{T}_2 is a sound α -abstraction of \mathcal{T}_1 , then \mathcal{T}_1 is decisive w.r.t. every α -closed set.*

We summarize the interesting results on denumerable abstractions, making a parallel with Proposition 4.2.17. Assume \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 , and write $\mathcal{B} = \{\alpha^{-1}(B) \mid B \in \Sigma_2\}$, the set of α -closed sets of Σ_1 . The following implications hold true:

$$\left. \begin{array}{l} \mathcal{T}_2 \text{ sound and DMC} \\ \text{with finite attractor} \end{array} \right\} \implies \mathcal{T}_1 \text{ is } \text{Dec}(\mathcal{B}) \iff \mathcal{T}_1 \text{ is } \text{StrDec}(\mathcal{B}) \iff \mathcal{T}_1 \text{ is } \text{PersDec}(\mathcal{B})$$

$$\downarrow$$

$$\mathcal{T}_1 \text{ is } \text{fair}(\mathcal{B})$$

5.2.2 Trickier transfers of properties

We established that decisiveness properties could be transferred through sound abstractions. However in the next section, we will also see that soundness of an abstraction can be proved via decisiveness properties. It is therefore relevant to explore alternatives to prove decisiveness properties. In this section, we give two frameworks where this can be done.

First, we assume a denumerable abstraction and lower bounds on probabilities of reachability properties.

Proposition 5.2.6. *Let \mathcal{T}_2 be a DMC such that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Assume that there is a finite set $A_2 = \{s_1, \dots, s_n\} \subseteq S_2$ such that A_2 is an attractor for \mathcal{T}_2 and $A_1 = \bigcup_{i=1}^n \alpha^{-1}(s_i) = \alpha^{-1}(A_2)$ is an attractor for \mathcal{T}_1 . Assume moreover that for every $1 \leq i \leq n$, for every α -closed set B in Σ_1 , there exist $p > 0$ and $k \in \mathbb{N}$ such that:*

- for every $\mu \in \text{Dist}(\alpha^{-1}(s_i))$, $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F}_{\leq k} B) \geq p$, or
- for every $\mu \in \text{Dist}(\alpha^{-1}(s_i))$, $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} B) = 0$.

Then \mathcal{T}_1 is decisive w.r.t. every α -closed set.

We write (\dagger) for the hypotheses over \mathcal{T}_1 in this proposition. The idea behind this result is that, with probability 1, the attractor of \mathcal{T}_1 will be visited infinitely often and, if at each visit of the attractor, there is a positive probability to reach some (α -closed) set B , since that probability is by assumption bounded from below, then B will indeed be visited infinitely often with probability 1. This will allow to show the dichotomy between B and \widetilde{B} that is required for proving the decisiveness property: either \widetilde{B} will be visited, or \widetilde{B} will never be visited in which case, using the fact that the attractor is visited infinitely often with probability 1 and the fact that from \widetilde{B}^c , B is visited with a positive probability, and finally from hypothesis (\dagger) , we will be able to conclude that B is visited with probability 1. We give here the full proof. Note that this kind of proof appears quite often in the literature (see *e.g.* [ABM07, Lemma 3.4], but we have to do it carefully here, since the framework is rather general).

Proof. Fix $B \subseteq S_2$ and $\mu \in \text{Dist}(S_1)$. We want to show that \mathcal{T}_1 is μ -decisive w.r.t. $\alpha^{-1}(B)$. We therefore have to show that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B})) = 1$. Towards a contradiction we assume that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G}(\neg \alpha^{-1}(B)) \wedge \mathbf{G}(\neg \alpha^{-1}(\widetilde{B}))) > 0$, i.e. $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \alpha^{-1}(\widetilde{B}^c)) > 0$.

Since $A_1 = \alpha^{-1}(A_2)$ is an attractor of \mathcal{T}_1 , we deduce from Lemma 4.2.8 that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G F} \alpha^{-1}(A_2)) = 1$, hence:

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \alpha^{-1}((\widetilde{B})^c) \wedge \mathbf{G F} \alpha^{-1}(A_2)) > 0. \quad (5.2)$$

We let $A'_2 \subseteq A_2$ be the subset of states $s \in A_2$ such that:

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \alpha^{-1}((\widetilde{B})^c) \wedge \mathbf{G F} \alpha^{-1}(\{s\})) > 0.$$

Due to inequality (5.2) and finiteness of A_2 , A'_2 is non-empty and furthermore, every $s \in A'_2$ belongs to B^c and $(\widetilde{B})^c$. We now set $A'_1 = \alpha^{-1}(A'_2)$.

In particular, $A'_1 \subseteq \alpha^{-1}((\widetilde{B})^c)$, hence from Lemma 4.2.2 (third item) we get that for every $\nu \in \text{Dist}(A'_1)$, $\text{Prob}_\nu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) > 0$. According to hypothesis (\dagger), for every $s \in A'_2$, we can find $p_s > 0$ and $k_s \in \mathbb{N}$ such that for every $\nu_s \in \text{Dist}(\alpha^{-1}(s))$,

$$\text{Prob}_{\nu_s}^{\mathcal{T}_1}(\mathbf{F}_{\leq k_s} \alpha^{-1}(B)) \geq p_s.$$

Then taking $p = \min\{p_s \mid s \in A'_2\} > 0$ and $k = \max\{k_s \mid s \in A'_2\} \in \mathbb{N}$ (since A'_2 is finite), it holds that for every $\nu \in \text{Dist}(A'_1)$,

$$\text{Prob}_\nu^{\mathcal{T}_1}(\mathbf{F}_{\leq k} \alpha^{-1}(B)) \geq p \quad \text{hence} \quad \text{Prob}_\nu^{\mathcal{T}_1}(\mathbf{G}_{\leq k} \alpha^{-1}(B^c)) \leq 1 - p, \quad (5.3)$$

where

$$\text{Ev}_{\mathcal{T}}(\mathbf{G}_{\leq k} \alpha^{-1}(B^c)) = \text{Cyl}(\overbrace{\alpha^{-1}(B^c), \dots, \alpha^{-1}(B^c)}^{k \text{ times}}).$$

From (5.2), we can deduce that:

$$\begin{aligned} 0 < \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \alpha^{-1}((\widetilde{B})^c) \wedge \mathbf{G F} A'_1) \\ \leq \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G F} A'_1) \leq \lim_{n \rightarrow \infty} (1 - p)^n = 0. \end{aligned}$$

It remains to show the last inequality. It is this part of the proof that is quite classical in the literature (see *e.g.* [ABM07, Lemma 3.4]), but we have to prove it carefully in our general setting. We will prove it by induction as follows.

First we introduce some useful notations. Observe that from the definition of A'_1 , it holds that $A'_1 \subseteq \alpha^{-1}(B^c)$. Then for each $j \in \mathbb{N}$, we will write B_j^c for the finite sequence $\alpha^{-1}(B^c), \dots, \alpha^{-1}(B^c)$ where $\alpha^{-1}(B^c)$ occurs exactly j times, and similarly we will write $(B^c \setminus A'_1)_j$ for the finite sequence $\alpha^{-1}(B^c) \setminus A'_1, \dots, \alpha^{-1}(B^c) \setminus A'_1$ where $\alpha^{-1}(B^c) \setminus A'_1$ occurs exactly j times. Then observe that

$$\begin{aligned} \text{Ev}_{\mathcal{T}_1}(\mathbf{G F} A'_1 \wedge \mathbf{G} \alpha^{-1}(B^c)) = \\ \bigcap_{n \in \mathbb{N}} \bigcup_{j_0 \in \mathbb{N}} \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_{j_0}^c, A'_1, B_{j_1}^c, A'_1, B_{j_2}^c, \dots, B_{j_{n-1}}^c, A'_1, B_{j_n}^c), \quad (5.4) \end{aligned}$$

where $\mathbb{N}_{\geq k}$ denotes the set of natural numbers greater than k . We scheme the behaviour of this set and what we can infer on the probabilities on Figure 5.1. As all behaviours are always in $\alpha^{-1}(B^c)$, the big rectangle represents this set, while the small one represents $A'_1 \subseteq \alpha^{-1}(B^c)$ which we know is reached infinitely often with probability 1. The behaviours are thus decomposed accordingly to each visit in A'_1 followed by k moves (while staying in A'_1). The dashed arrows represent these k steps. Note that within those k steps, A'_1 could be reached but it has no importance. What matters here is the fact that from A'_1 , the probability of the next k steps within $\alpha^{-1}(B^c)$ is upper bounded by $1 - p$. The curled arrows hold for the next visit to A'_1 which we hence know that it will happen with probability 1.




Figure 5.1: Scheme for the proof of Proposition 5.2.6.

We will prove by induction over n that for each $n \geq 0$ and for each $\nu \in \text{Dist}(S_1)$,

$$\text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j_0 \in \mathbb{N}} \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_{j_0}^c, A'_1, B_{j_1}^c, A'_1, B_{j_2}^c, \dots, B_{j_{n-1}}^c, A'_1, B_{j_n}^c) \right) \leq (1-p)^n.$$

Observe that for each $n \geq 0$, it holds that

$$\bigcup_{j_0 \in \mathbb{N}} \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_{j_0}^c, A'_1, B_{j_1}^c, A'_1, B_{j_2}^c, \dots, B_{j_{n-1}}^c, A'_1, B_{j_n}^c) \subseteq \bigcup_{j_0 \in \mathbb{N}} \bigcup_{(j_1, \dots, j_{n-1}) \in \mathbb{N}_{\geq k}^{n-1}} \text{Cyl}(B_{j_0}^c, A'_1, B_{j_1}^c, A'_1, B_{j_2}^c, \dots, B_{j_{n-1}}^c, A'_1, B_k^c).$$

Somehow, Figure 5.1 represents one of these cylinders. Hence it is enough to demonstrate that for each $n \geq 0$ and for each $\nu \in \text{Dist}(S_1)$,

$$\text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j_0 \in \mathbb{N}} \bigcup_{(j_1, \dots, j_{n-1}) \in \mathbb{N}_{\geq k}^{n-1}} \text{Cyl}(B_{j_0}^c, A'_1, B_{j_1}^c, A'_1, B_{j_2}^c, \dots, B_{j_{n-1}}^c, A'_1, B_k^c) \right) \leq (1-p)^n. \quad (5.5)$$

First fix $n = 0$ and $\nu \in \text{Dist}(S_1)$. It corresponds to the two first arrows on Figure 5.1. We will show that for each $m \geq 0$,

$$\text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \text{Cyl}(B_j^c, A'_1, B_k^c) \right) \leq 1-p,$$

that is we decompose Figure 5.1 according to the length of the first curled arrow. We first prove cases $m = 0$ and $m = 1$ in order to illustrate what is happening, and then we will make the general case. If $m = 0$, it then holds that

$$\begin{aligned} \text{Prob}_\nu^{\mathcal{T}_1}(\text{Cyl}(A'_1, B_k^c)) &= \nu(A'_1) \cdot \text{Prob}_{\nu_{A'_1}^{\mathcal{T}_1}}(\text{Cyl}(A'_1, B_k^c)) \\ &\leq \text{Prob}_{\nu_{A'_1}^{\mathcal{T}_1}}(\text{Cyl}(B_{k+1}^c)) \leq 1-p \end{aligned} \quad (5.6)$$

where the first inequality holds from the fact that $A'_1 \subseteq \alpha^{-1}(B^c)$, and the second one from (5.3). Note that we assumed here that $\nu(A'_1) > 0$, but it has no importance since if $\nu(A'_1) = 0$, then the inequality trivially holds. Now if $m = 1$, first observe that

$$\text{Cyl}(A'_1, B_k^c) \cup \text{Cyl}(B^c, A'_1, B_k^c) = \text{Cyl}(A'_1, B_k^c) \cup \text{Cyl}(B^c \setminus A'_1, A'_1, B_k^c)$$

where in the second member of the equality, the union is disjoint. It follows that, writting $\nu'_0 = \nu_{B^c \setminus A'_1}$ and $\nu_1 = (\Omega_{\mathcal{T}_1}(\nu'_0))_{A'_1}$:

$$\begin{aligned} &\text{Prob}_\nu^{\mathcal{T}_1}(\text{Cyl}(A'_1, B_k^c) \cup \text{Cyl}(B^c, A'_1, B_k^c)) \\ &= \text{Prob}_\nu^{\mathcal{T}_1}(\text{Cyl}(A'_1, B_k^c)) + \text{Prob}_\nu^{\mathcal{T}_1}(\text{Cyl}(B^c \setminus A'_1, A'_1, B_k^c)) \\ &\leq \nu(A'_1) \cdot (1-p) + \nu(B^c \setminus A'_1) \cdot (\Omega_{\mathcal{T}_1}(\nu'_0))(A'_1) \cdot \text{Prob}_{\nu_1^{\mathcal{T}_1}}(\text{Cyl}(A'_1, B_k^c)) \\ &\quad \text{from Lemma 4.1.8} \\ &\leq \nu(A'_1) \cdot (1-p) + \nu(B^c \setminus A'_1) \cdot (\Omega_{\mathcal{T}_1}(\nu'_0))(A'_1) \cdot (1-p) \leq (1-p). \end{aligned}$$

Note that we again assumed here that $\nu(B^c \setminus A'_1) > 0$ and $(\Omega_{\mathcal{T}_1}(\nu'_0))(A'_1) > 0$, which has again no importance since otherwise, the probability of one of the

cylinders would be equal to 0 and which would thus not interfere on the above inequality. We now prove the general case for $m \geq 2$. Again, we can decompose the union of the cylinders into a disjoint one as follows:

$$\bigcup_{j=0}^m \text{Cyl}(B_j^c, A'_1, B_k^c) = \bigcup_{j=0}^m \text{Cyl}((B^c \setminus A'_1)_j, A'_1, B_k^c).$$

We use the following notations: $\nu'_0 = \nu_{B^c \setminus A'_1}$, $\nu_0 = \nu_{A'_1}$, and

- for each $1 \leq i \leq m-1$, $\nu'_i = (\Omega_{\mathcal{T}_1}(\nu'_{i-1}))_{B^c \setminus A'_1}$ and
- for each $1 \leq i \leq m$, $\nu_i = (\Omega_{\mathcal{T}_1}(\nu'_{i-1}))_{A'_1}$.

Note that we assume again that the conditional probability are well-defined, but like in cases $m=0$ and $m=1$, we can make this supposition w.l.o.g. Then using Lemma 4.1.8 and the observation (5.3), we get that:

$$\begin{aligned} \text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \text{Cyl}(B_j^c, A'_1, B_k^c) \right) &= \sum_{j=0}^m \text{Prob}_\nu^{\mathcal{T}_1} (\text{Cyl}(B_j^c, A'_1, B_k^c)) \\ &= \nu(A'_1) \cdot \text{Prob}_{\nu'_0}^{\mathcal{T}_1} (\text{Cyl}(A'_1, B_k^c)) \\ &\quad + \sum_{j=1}^m \left(\nu(B^c \setminus A'_1) \cdot \prod_{i=1}^{j-1} (\Omega_{\mathcal{T}_1}(\nu'_i))(B^c \setminus A'_1) \cdot (\Omega_{\mathcal{T}_1}(\nu'_{j-1}))(A'_1) \right. \\ &\quad \left. \cdot \underbrace{\text{Prob}_{\nu'_j}^{\mathcal{T}_1} (\text{Cyl}(A'_1, B_k^c))}_{\leq 1-p} \right) \\ &\leq (1-p) \cdot \left(\nu(A'_1) + \sum_{j=1}^m \left(\nu(B^c \setminus A'_1) \cdot \prod_{i=1}^{j-1} (\Omega_{\mathcal{T}_1}(\nu'_i))(B^c \setminus A'_1) \cdot (\Omega_{\mathcal{T}_1}(\nu'_{j-1}))(A'_1) \right) \right) \\ &= (1-p) \cdot \text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \text{Cyl}(B_j^c, A'_1) \right) \leq 1-p \end{aligned}$$

where the last equality comes again from Lemma 4.1.8, but in the other sense this time. Finally through the limit over m , we obtain that (5.5) is true when $n=0$.

Now fix $n \geq 0$ and assume that for $0 \leq l \leq n$ and for each $\nu \in \text{Dist}(S_1)$, the inequality (5.5) holds true. We get in particular that for each $\nu \in \text{Dist}(S_1)$,

$$\begin{aligned} \text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j_0 \in \mathbb{N}} \bigcup_{(j_1, \dots, j_{n-1}) \in \mathbb{N}_{\geq k}^{n-1}} \text{Cyl}(B_{j_0}^c, A'_1, B_{j_1}^c, A'_1, B_{j_2}^c, \dots, B_{j_{n-1}}^c, A'_1, B_k^c) \right) \\ \leq (1-p)^n. \end{aligned}$$

We want to show that (5.5) is still satisfied for $n + 1$. Like in case $n = 0$, we will show that for each $m \geq 0$,

$$\text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_j^c, A'_1, B_{j_1}^c, A'_1, B_{j_2}^c, \dots, B_{j_n}^c, A'_1, B_k^c) \right) \leq (1 - p)^{n+1}.$$

We thus again decompose the scheme of Figure 5.1 accordingly to the length of the first arrow. In fact the proof is very similar to the case $n = 0$ as once you hit for the second time $\alpha^{-1}(B^c)$ in the scheme (*i.e.* after the first dashed arrow), the induction hypothesis can be applied. What happens before is the exact same behaviour as in the case for $n = 0$. For each $m \geq 0$ this finite union of cylinders can be decomposed into a finite union of disjoint sets as follows:

$$\begin{aligned} \bigcup_{j=0}^m \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_j^c, A'_1, B_{j_1}^c, A'_1, B_{j_2}^c, \dots, B_{j_n}^c, A'_1, B_k^c) = \\ \bigcup_{j=0}^m \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}((B^c \setminus A'_1)_j, A'_1, B_{j_1}^c, A'_1, B_{j_2}^c, \dots, B_{j_n}^c, A'_1, B_k^c). \end{aligned}$$

Then using Lemma 4.1.8 and this decomposition into a disjoint union, it holds that

$$\begin{aligned} \text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_j^c, A'_1, B_{j_1}^c, A'_1, B_{j_2}^c, \dots, B_{j_n}^c, A'_1, B_k^c) \right) = \\ \sum_{j=0}^m \alpha_j \cdot \text{Prob}_{\mu_j}^{\mathcal{T}_1} \left(\bigcup_{j' \in \mathbb{N}} \bigcup_{(j_2, \dots, j_n) \in \mathbb{N}_{\geq k}^{n-1}} \text{Cyl}(B_{j'}^c, A'_1, B_{j_2}^c, \dots, B_{j_n}^c, A'_1, B_k^c) \right), \end{aligned}$$

where for each $0 \leq j \leq m$, $0 < \alpha_j < 1$ and $\mu_j \in \text{Dist}(S_1)$ are given by Lemma 4.1.8, where α_j corresponds to:

$$\alpha_j = \text{Prob}_\nu^{\mathcal{T}_1}(\text{Cyl}((B^c \setminus A'_1)_j, A'_1, B_k^c)).$$

Note that this is possible due to the fact that we look at the union of all $j_1 \geq k$. Using the induction hypothesis and this last equality, we get that

$$\begin{aligned} \text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_j^c, A'_1, B_{j_1}^c, A'_1, B_{j_2}^c, \dots, B_{j_n}^c, A'_1, B_k^c) \right) \\ \leq (1 - p)^n \cdot \text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \text{Cyl}(B_j^c, A'_1, B_k^c) \right) \\ \leq (1 - p)^{n+1} \end{aligned}$$

where the last inequality stands from what we have done in case $n = 0$. Through the limit over m , we can thus deduce that (5.5) is still true for $n + 1$.

Finally coming back to (5.4), through the limit over n this time, we conclude that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G F} A'_1 \wedge \mathbf{G} \alpha^{-1}(B^c)) \leq \lim_{n \rightarrow \infty} (1 - p)^n = 0$. This yields a contradiction and concludes the proof. \square

This result gives us thus an alternative to Proposition 5.2.5: here we assume again an α -abstraction that is a DMC with a finite attractor, but this time we assume hypothesis (\dagger) on the attractor instead of the soundness of the α -abstraction. And it is sufficient to get decisiveness. We will see however in Section 5.3 how this implies soundness (see Proposition 5.3.4).

Now in the second proposition, we strengthen the hypothesis on the α -abstraction and assume that it is finite. The condition which applies in this case is the weakest property that we have seen (see Proposition 4.2.17), namely fairness!

Proposition 5.2.7. *Let \mathcal{T}_2 be a finite Markov chain such that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Fix $\mu \in \text{Dist}(S_1)$, and assume that \mathcal{T}_1 is μ -fair w.r.t. every α -closed set. Then \mathcal{T}_1 is μ -decisive w.r.t. every α -closed set.*

A key element of the proof relies on the fact that, since \mathcal{T}_2 is a finite Markov Chain, it can be viewed as a graph and we can talk of the *bottom strongly connected components (BSCC)* of \mathcal{T}_2 , see Section 2.2.1 for some details. Recovering the notations of Section 2.2.1, we write $\text{BSCC}(\mathcal{T}_2)$ for the set of BSCCs of \mathcal{T}_2 .

The first step of the proof aims at showing that, roughly speaking, the union of all BSCCs of \mathcal{T}_2 is a μ -attractor for \mathcal{T}_1 . More precisely, if $\mathcal{C} = \{s \in S_2 \mid \exists C \in \text{BSCC}(\mathcal{T}_2), s \in C\}$, we prove that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(\mathcal{C})) = 1$. This is shown thanks to the following arguments:

- for each $s \in S_2$, $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G F} \alpha^{-1}(s)) > 0$ implies that $s \in \mathcal{C}$ – this uses the μ -fairness assumption of \mathcal{T}_1 w.r.t. α -closed sets, and the core property of BSCCs (we cannot escape from them);
- using Bayes formula, one can decompose the set of paths according to the states which are visited infinitely often (which corresponds to a decomposition according to the BSCC the path ultimately visit).

Once we have shown that $\alpha^{-1}(\mathcal{C})$ is a μ -attractor for \mathcal{T}_1 , it suffices to observe that for each $B \subseteq S_2$ and each BSCC C of \mathcal{T}_2 , either $B \cap C \neq \emptyset$, or $C \subseteq \widetilde{B}$. Transferring those observations to \mathcal{T}_1 and using Bayes formula to decompose $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B}))$ according to which BSCC is reached, it is easy to check that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B})) = 1$.

Proof. We define $\mathcal{C} = \{s \in S_2 \mid \exists C \in \text{BSCC}(\mathcal{T}_2), s \in C\}$. We first prove that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(\mathcal{C})) = 1$. In order to establish this, we show that for each $s \in S_2$, $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(s)) > 0$ implies that $s \in \mathcal{C}$. Indeed, pick $s \in S_2$ such that:

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(\{s\})) > 0.$$

We can state that for each $k \geq 1$ and for each $s_0, s_1, \dots, s_k \in S_2$ with $s_0 = s$ and such that for each $0 \leq i < k$, $\kappa_2(s_i, s_{i+1}) > 0$, it holds that

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(s_k) \mid \mathbf{G} \mathbf{F} \alpha^{-1}(s)) = 1.$$

We prove this by induction over k . First fix $k = 1$ and let $s_1 \in S_2$ such that $\kappa_2(s, s_1) > 0$. Then, from Definition 5.1.1, for every $\nu \in \text{Dist}(\alpha^{-1}(s))$, $\text{Prob}_\nu^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(\{s\}), \alpha^{-1}(\{s_1\}))) > 0$. Hence $\alpha^{-1}(s) \in \text{PreProb}^{\mathcal{T}}(\{\alpha^{-1}(s_1)\})$. And since \mathcal{T}_1 is fair w.r.t. α -closed sets, we get that

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(\{s_1\}) \mid \mathbf{G} \mathbf{F} \alpha^{-1}(\{s\})) = 1.$$

Now fix $k > 1$ and assume that for each $1 \leq j < k$ and for each $s_0, \dots, s_j \in S_2$ with $s_0 = s$ and such that for each $0 \leq i < j$, $\kappa_2(s_i, s_{i+1}) > 0$, it holds that

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(s_j) \mid \mathbf{G} \mathbf{F} \alpha^{-1}(s)) = 1.$$

We want to show that it is still the case for k . Fix $s_0, s_1, \dots, s_k \in S_2$ satisfying all the desired hypotheses. Using the induction hypothesis, we know that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(s_{k-1}) \mid \mathbf{G} \mathbf{F} \alpha^{-1}(s_0)) = 1$ and that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(s_k) \mid \mathbf{G} \mathbf{F} \alpha^{-1}(s_{k-1})) = 1$. We can then compute:

$$\begin{aligned} & \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(s_k) \mid \mathbf{G} \mathbf{F} \alpha^{-1}(s_0)) \\ &= \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(s_k) \wedge \mathbf{G} \mathbf{F} \alpha^{-1}(s_{k-1}) \mid \mathbf{G} \mathbf{F} \alpha^{-1}(s_0)) \\ &= \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(s_k) \mid \mathbf{G} \mathbf{F} \alpha^{-1}(s_{k-1}) \wedge \mathbf{G} \mathbf{F} \alpha^{-1}(s_0)) \\ &\quad \cdot \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(s_{k-1}) \mid \mathbf{G} \mathbf{F} \alpha^{-1}(s_0)) \\ &= 1 \end{aligned}$$

from the induction hypothesis. This shows that for every state s' which is reachable from s in \mathcal{T}_2 ,

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(\{s'\}) \mid \mathbf{G} \mathbf{F} \alpha^{-1}(\{s\})) = 1.$$

Then fix s' reachable from s in \mathcal{T}_2 . We can show that s is also reachable from s' . Towards a contradiction, assume that it is not the case. It follows that

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(\{s'\}) \wedge \mathbf{G} \mathbf{F} \alpha^{-1}(\{s\})) = 0$$

which is a contradiction with $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G F} \alpha^{-1}(\{s'\}) \mid \mathbf{G F} \alpha^{-1}(\{s\})) = 1$ and $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G F} \alpha^{-1}(\{s\})) > 0$. We deduce thus that s belongs to a BSCC of \mathcal{T}_2 : if $C = \{s' \in S_2 \mid s' \text{ is reachable from } s\}$, then

- firstly, you cannot leave C - otherwise you would reach a state s' out of C which would then be reachable from s and thus should be in C ;
- secondly, for each pair of states $s', s'' \in C$, we have $s' \rightarrow^* s''$ and $s'' \rightarrow^* s'$, where \rightarrow^* denotes the existence of a finite path between two states - indeed, it suffices to notice that $s \rightarrow^* s' \rightarrow^* s \rightarrow^* s'' \rightarrow^* s \rightarrow^* s'$.

We can now prove that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(\mathcal{C})) = 1$. Indeed observe first that from the finiteness of \mathcal{T}_2 , it holds that for every paths $\rho = t_0 t_1 t_2 \dots \in \text{Runs}(\mathcal{T}_1)$, there is $s \in S_2$ such that $\{i \in \mathbb{N} \mid t_i \in \alpha^{-1}(s)\}$ is infinite. Keeping this in mind, we write $S_2 = \{s_1, \dots, s_k, s_{k+1}, \dots, s_n\}$ where $k \geq 1$ and $\{s_1, \dots, s_k\} = \mathcal{C}$. Then we can write

$$\begin{aligned} \text{Runs}(\mathcal{T}_1) = & \text{Ev}_{\mathcal{T}_1}(\mathbf{G F} \alpha^{-1}(s_1)) \cup \text{Ev}_{\mathcal{T}_1}(\mathbf{G F} \alpha^{-1}(s_2) \wedge \mathbf{F G} \neg \alpha^{-1}(s_1)) \\ & \cup \dots \cup \text{Ev}_{\mathcal{T}_1}(\mathbf{G F} \alpha^{-1}(s_n) \wedge \bigwedge_{i=1}^{n-1} \mathbf{F G} \neg \alpha^{-1}(s_i)). \end{aligned}$$

From what we have shown previously, we now get that for each $j \geq k + 1$,

$$0 = \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G F} \alpha^{-1}(s_j)) \geq \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G F} \alpha^{-1}(s_j) \wedge \bigwedge_{i=1}^{j-1} \mathbf{F G} \neg \alpha^{-1}(s_i)).$$

And we conclude that

$$\begin{aligned} 1 &= \text{Prob}_\mu^{\mathcal{T}_1}(\text{Runs}(\mathcal{T}_1)) \\ &= \sum_{j=1}^k \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G F} \alpha^{-1}(s_j) \wedge \bigwedge_{i=1}^{j-1} \mathbf{F G} \neg \alpha^{-1}(s_i)) \\ &\leq \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(\mathcal{C})). \end{aligned}$$

We are now able to prove that \mathcal{T}_1 is $\text{Dec}(\mu, \mathcal{B})$. Fix $B \subseteq S_2$, we want to show that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\tilde{B})) = 1$. We have that

$$\begin{aligned} \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\tilde{B})) = \\ \sum_{\substack{C \in \text{BSCC}(\mathcal{T}_2) \text{ s.t.} \\ \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(C)) > 0}} \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(C)) \cdot \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\tilde{B}) \mid \mathbf{F} \alpha^{-1}(C)). \end{aligned}$$

Now we fix some $C \in \text{BSCC}(\mathcal{T}_2)$ such that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(C)) > 0$. There are two cases:

- first if there is $s \in C$ such that $s \in B$, then $\alpha^{-1}(s) \subseteq \alpha^{-1}(B)$ and thus $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B}) \mid \mathbf{F} \alpha^{-1}(C)) = 1$;
- or for each $s \in C$, $s \in \widetilde{B}$ which implies that $\alpha^{-1}(C) \subseteq \alpha^{-1}(\widetilde{B})$ and in that case again, $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B}) \mid \mathbf{F} \alpha^{-1}(C)) = 1$.

We finally conclude that

$$\begin{aligned} \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B})) &= \sum_{\substack{C \in \text{BSCC}(\mathcal{T}_2) \text{ s.t.} \\ \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(C)) > 0}} \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(C)) \\ &= \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(C)) = 1. \end{aligned}$$

□

This result gives us thus again an alternative to Proposition 5.2.5, but when we restricted the hypothesis to an α -abstraction that is a finite Markov chain: here, instead of the soundness assumption, we only need fairness, which was proven to be the weakest notion defined in Chapter 4 (see Proposition 4.2.17), in order to imply decisiveness!

5.3. Conditions for completeness and soundness

In this section, we give conditions that ensure completeness and soundness of α -abstractions. The first result shows that a simple condition on the abstraction implies completeness.

Lemma 5.3.1. If \mathcal{T}_2 is a finite Markov chain and an α -abstraction of \mathcal{T}_1 , then \mathcal{T}_2 is complete.

Proof. Fix $s_0 \in S_2$ and $\mu \in \text{Dist}(\alpha^{-1}(\{s_0\}))$. Recall that from Lemma 5.1.4, it holds true that $\alpha_\#(\mu) = \delta_{s_0}$. Towards a contradiction, we assume that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1$ but $\text{Prob}_{\alpha_\#(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) < 1$.

Since \mathcal{T}_2 is a finite Markov chain, there are $s_1, \dots, s_n \in S_2$ such that

$$\text{Prob}_{s_0}^{\mathcal{T}_2}(\text{Cyl}(s_0, s_1, \dots, s_n)) > 0$$

and for each $\rho = (s_i)_{i \geq 0} \in \text{Cyl}(s_0, \dots, s_n)$ and for each $i \geq 0$, $s_i \notin B$.

For each $0 \leq i \leq n$, we write $A_i = \alpha^{-1}(\{s_i\})$. Then, by Lemma 5.1.6, we get that $\text{Prob}_\mu^{\mathcal{T}_1}(\text{Cyl}(A_0, A_1, \dots, A_n)) > 0$. However, $\text{Cyl}(A_0, A_1, \dots, A_n) \cap \text{Ev}_{\mathcal{T}}(\mathbf{F} \alpha^{-1}(B)) = \emptyset$, yielding a contradiction. □

Note that the above lemma does not hold for denumerable abstractions. To illustrate this, any two random walks over \mathbb{N} (see Example 4.1.2) are abstractions of each other (similar to the observations made in Example 5.1.3), and it is well-known that almost-sure reachability depends on the probability values (as stated in Example 4.2.7).

In general, completeness can be guaranteed by some decisiveness condition. Note that, since finite Markov chains are always decisive, the next lemma actually subsumes the latter one, that we however found interesting to have as such.

Lemma 5.3.2. Let $\mu \in \text{Dist}(S_1)$. Assume that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 and that \mathcal{T}_2 is $\text{Dec}(\alpha_{\#}(\mu))$. Then, \mathcal{T}_2 is a μ -complete α -abstraction.

Proof. Fix $B \in \Sigma_2$ and towards a contradiction assume that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1$ but $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) < 1$. Therefore, since \mathcal{T}_2 is $\text{Dec}(\alpha_{\#}(\mu))$, we infer from Lemma 4.2.2 (fifth item) that $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}((\neg B) \mathbf{U} \widetilde{B}) > 0$, and applying Corollary 5.1.7, we get that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\alpha^{-1}(\neg B) \mathbf{U} \alpha^{-1}(\widetilde{B})) > 0$. This contradicts the hypothesis that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1$. \square

Remark 5.3.3. Observe that thanks to Proposition 2.2.13 and Lemma 5.3.2, it holds that if \mathcal{T}_2 is a DMC with a finite attractor and if \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 , then it is a complete α -abstraction.

Given the results of Section 5.2, we realise that soundness is often more critical than completeness, and showing it may require some effort. Below, we give a condition under which soundness holds.

Proposition 5.3.4. Let \mathcal{T}_2 be an α -abstraction of \mathcal{T}_1 . Assume that \mathcal{T}_1 is decisive w.r.t. every α -closed set. Then \mathcal{T}_2 is a sound α -abstraction of \mathcal{T}_1 .

Proof. Fix $B \in \Sigma_2$. Towards a contradiction assume that $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) = 1$ but $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) < 1$. Hence, since \mathcal{T}_1 is decisive w.r.t. $\alpha^{-1}(B)$ from μ , it holds from Lemma 4.2.2 (fifth item) that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\neg \alpha^{-1}(B) \mathbf{U} \alpha^{-1}(\widetilde{B})) > 0$. Applying Corollary 5.1.7, we get that $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}((\neg B) \mathbf{U} \widetilde{B}) > 0$, which contradicts the assumption that $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) = 1$. \square

Remark 5.3.5. For \mathcal{T}_2 an α -abstraction of \mathcal{T}_1 , notice that completeness is ensured by a decisiveness assumption on \mathcal{T}_2 , whereas soundness requires \mathcal{T}_1 being decisive w.r.t. every α -closed set. While these conditions look very similar, the condition for soundness on \mathcal{T}_1 is actually harder to check since the abstract STS \mathcal{T}_2 will be simpler than the original concrete STS \mathcal{T}_1 .

However observe that in the case where \mathcal{T}_2 is a DMC, combining Propositions 5.2.6 and 5.3.4, we get another assumption ensuring the soundness: the existence of a finite attractor in \mathcal{T}_2 satisfying hypothesis (†) (see page 114).

CHAPTER 6

Qualitative and Quantitative Analysis

In this chapter we are interested in the qualitative and quantitative model-checking problem of STSs (Definitions 4.1.19 and 4.1.20). As said before, the objective is to extend the results of [ABM07] to more general stochastic systems and to richer properties, *i.e.* to STSs and to properties recognised by DMA, also called ω -regular properties. The Chapter is divided as follows.

In Section 6.1, we consider the qualitative model-checking problem. We first aim at extending the results of [ABM07] to our context and we hence only consider reachability and repeated reachability properties in Section 6.1.1. Using decisiveness assumptions, we get very similar results as in the case of DMCs, however the qualitative model-checking problem cannot be reduced to the satisfaction of structural properties on the underlying graph (to the contrary of decisive DMCs) as in this continuous context, there is no underlying graph. But those lead to simpler problems. We show moreover that the qualitative analysis can equivalently be done on an α -abstraction that is sound, under decisiveness assumptions. This will thus be particularly interesting for abstractions that are DMCs.

Then, we extend the results to ω -regular properties. We first consider the simpler case of DMCs in Section 6.1.2, that is we consider the same context of [ABM07] but we enrich the class of properties we want to check. The almost-sure satisfaction of properties given as a DMA can be characterised by finitely many reachability properties [ABRS05] and [Ber06] for DMCs with a finite attractor. We prove the characterisation in our new formalism. It will allow to use, in part, the results of [ABM07] but it will also require to prove some new

structural properties on some graph. In Section 6.1.3, we extend those results to STSs that have a sound α -abstraction that is a DMC and that has a finite attractor.

In Section 6.2, we consider the quantitative model-checking problem. Sections 6.2.1 and 6.2.2 present approximation schemes for reachability and repeated reachability properties for STSs, that generalize the algorithms of [ABM07]. We show also that under decisiveness assumptions, the procedures are correct and terminates. In Sections 6.2.3 and 6.2.4, we come back to the case of ω -regular properties and show how the characterisations of Sections 6.1.2 and 6.1.3 can be used in order to get approximation schemes of properties given by a DMA for DMCs with a finite attractor and for STSs that have a sound α -abstraction that is a DMC and that has a finite attractor. Those procedures amount to compute reachability probabilities and for which the approximation scheme of Section 6.2.1 can thus be applied!

We end the chapter with a summary of all the qualitative and quantitative results in Section 6.3.

6.1. Qualitative analysis

In this section, we are interested in the almost-sure model-checking problem of STSs (see Definition 4.1.19). We rely on the notions previously introduced and studied to design generic procedures for the qualitative analysis of properties of STSs, under some assumptions that will be made precise.

In Section 6.1.1, we are concerned with the extension of the qualitative results of [ABM07] to STSs and we thus only consider reachability and repeated reachability properties. We get similar results, however we cannot reduce the problem to structural properties of the underlying graph as this does not exist in the continuous setting; but those lead to simpler problems. We show moreover that under some assumptions, the analysis can be done on an α -abstraction.

In Sections 6.1.2 and 6.1.3 we enrich the class of properties we want to check. We consider properties specified by a DMA, *i.e.* ω -regular properties. We first consider DMCs in Section 6.1.2. We are interested in the approach of [ABRS05] and [Ber06] which consider DMCs with a finite attractor. The authors provide a characterisation of the almost-sure satisfaction of such properties through reachability probabilities. It is done through the construction of some graph based on the attractor. We provide the proofs in our new formalism. In Section 6.1.3, we extend those results to STSs that have a sound α -abstraction that is a DMC and that has a finite attractor.

6.1.1 Reachability and repeated reachability properties

Inspired from [ABM07] (see Section 2.2.2), our objective here is to describe generic procedures that capture the qualitative (almost-sure, positive) satisfaction of reachability and repeated reachability properties.

We fix a STS $\mathcal{T} = (S, \Sigma, \kappa)$. Given $B \in \Sigma$ a measurable set, recall that $\widetilde{B} = \{s \in S \mid \text{Prob}_s^{\mathcal{T}}(\mathbf{F} B) = 0\}$ denotes its avoid-set. We start with two technical lemmas that will be useful to show various results thereafter.

Lemma 6.1.1. For every $\mu \in \text{Dist}(S)$

$$(i) \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B \wedge (\neg B \mathbf{U} \widetilde{B})) = 0;$$

$$(ii) \text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B \wedge \mathbf{F} \widetilde{B}) = 0.$$

Proof. We first prove point (i). Since B cannot be reached while we are in $\neg B$, it holds that

$$\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B \wedge (\neg B \mathbf{U} \widetilde{B})) = \text{Prob}_\mu^{\mathcal{T}}(\neg B \mathbf{U} (\widetilde{B} \wedge \mathbf{F} B)).$$

Relaxing the constraint on the “Until” formula, we get $\text{Prob}_\mu^{\mathcal{T}}(\neg B \mathbf{U} (\widetilde{B} \wedge \mathbf{F} B)) \leq \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} (\widetilde{B} \wedge \mathbf{F} B))$, and the latter is null by definition of \widetilde{B} . This proves the first item.

Point (ii) is straightforward from the definition of \widetilde{B} and by observing that $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B \wedge \mathbf{F} \widetilde{B}) \leq \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} (\widetilde{B} \wedge \mathbf{F} B)) = 0$. \square

Lemma 6.1.2. For every $\mu \in \text{Dist}(S)$, if \mathcal{T} is PersDec(μ, B), then $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} \widetilde{B} \wedge \widetilde{\widetilde{B}}) = 0$.

Proof. Assume that \mathcal{T} is PersDec(μ, B), i.e. for each $p \geq 0$, $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F}_{\geq p} B \vee \mathbf{F}_{\geq p} \widetilde{B}) = 1$. Towards a contradiction, we suppose that $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} \widetilde{B} \wedge \widetilde{\widetilde{B}}) > 0$. Since

$$\text{Ev}_{\mathcal{T}}(\mathbf{F} \widetilde{B} \wedge \widetilde{\widetilde{B}}) = \bigcup_{n \geq 0} \bigcup_{m \geq 0} \text{Ev}_{\mathcal{T}}(\mathbf{F}_{=n} \widetilde{B}) \cap \text{Ev}_{\mathcal{T}}(\mathbf{F}_{=m} \widetilde{\widetilde{B}}),$$

we deduce that there are $n, m \geq 0$ such that $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F}_{=n} \widetilde{B} \wedge \mathbf{F}_{=m} \widetilde{\widetilde{B}}) > 0$. We write e for the event $e = \text{Ev}_{\mathcal{T}}(\mathbf{F}_{=n} \widetilde{B} \wedge \mathbf{F}_{=m} \widetilde{\widetilde{B}})$. We can show that $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F}_{\geq n} B \mid$

$e) = 0$ and $\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq m} \widetilde{B} \mid e) = 0$. Indeed we get that:

$$\begin{aligned} \text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq n} B \mid e) &= \frac{\text{Prob}_\mu^\mathcal{T}((\mathbf{F}_{\geq n} B) \wedge e)}{\text{Prob}_\mu^\mathcal{T}(e)} \\ &\leq \frac{\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq n} B \wedge \mathbf{F}_{=n} \widetilde{B})}{\text{Prob}_\mu^\mathcal{T}(e)} \\ &= 0 \end{aligned}$$

from the definition of \widetilde{B} . The equality $\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq m} \widetilde{B} \mid e) = 0$ is proved similarly. Writing $q = \max(m, n)$, it follows that

$$\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq q} B \vee \mathbf{F}_{\geq q} \widetilde{B} \mid e) = 0.$$

And since $\text{Prob}_\mu^\mathcal{T}(e) > 0$, this contradicts the fact that \mathcal{T} is $\text{PersDec}(\mu, B)$, which concludes the proof. \square

Extending the approach of [ABM07] (see Section 2.2.2), we establish characterisations of the qualitative satisfaction of (repeated) reachability properties in terms of the positive satisfaction of reachability-like properties. We advocate that these are simpler to check on STSs: positive reachability amounts to guess a cylinder leading to the target, and to show that this path has a positive measure.

Proposition 6.1.3 (Almost-sure reachability). *For every $\mu \in \text{Dist}(S)$,*

- if $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} B) = 1$, then $\text{Prob}_\mu^\mathcal{T}(\neg B \mathbf{U} \widetilde{B}) = 0$;
- if \mathcal{T} is $\text{Dec}(\mu, B)$ and $\text{Prob}_\mu^\mathcal{T}(\neg B \mathbf{U} \widetilde{B}) = 0$, then $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} B) = 1$.

Proof. We start with the first item. Since the event $\text{Ev}_\mathcal{T}(\mathbf{F} B)$ is almost-sure, we have

$$\text{Prob}_\mu^\mathcal{T}(\neg B \mathbf{U} \widetilde{B}) = \text{Prob}_\mu^\mathcal{T}((\neg B \mathbf{U} \widetilde{B}) \wedge \mathbf{F} B)$$

and then it is straightforward from point (i) of Lemma 6.1.1.

In order to prove the other implication, we need the assumption that \mathcal{T} is $\text{Dec}(\mu, B)$. We have that:

$$\begin{aligned} 1 &= \text{Prob}_\mu^\mathcal{T}(\mathbf{F} B \vee \mathbf{F} \widetilde{B}) = \text{Prob}_\mu^\mathcal{T}(\mathbf{F} B \vee (\neg B \mathbf{U} \widetilde{B})) \\ &= \text{Prob}_\mu^\mathcal{T}(\mathbf{F} B) + \text{Prob}_\mu^\mathcal{T}(\neg B \mathbf{U} \widetilde{B}), \end{aligned}$$

where the first equality comes from Lemma 4.2.2 (fifth item) and the second equality from Lemma 6.1.1 (point (i)). Then from $\text{Prob}_\mu^\mathcal{T}(\neg B \mathbf{U} \widetilde{B}) = 0$, we derive that $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} B) = 1$. \square

This reduces the almost-sure model-checking of a reachability property to the 0-model-checking of an “Until” formula, a slight generalization of reachability properties.

Remark 6.1.4. Note that in the case where \mathcal{T} is a DMC, Proposition 6.1.3 recovers a result of [ABM07] (corresponding to Proposition 2.2.15 here). Indeed, given a state s , we have that $\text{Prob}_s^{\mathcal{T}}(\neg B \mathbf{U} \tilde{B}) > 0$ if and only if there exists $n \geq 0$, $s_0, s_1, \dots, s_n \in S$ such that $s_0 = s$, $s_0, s_1, \dots, s_{n-1} \in B^c$ and $s_n \in \tilde{B}$ (since in this case, $\text{Ev}_{\mathcal{T}}(\neg B \mathbf{U} \tilde{B})$ is a denumerable union of cylinders of the form $\text{Cyl}(s_0, s_1, \dots, s_n)$ with s_0, s_1, \dots, s_n satisfying the previous hypothesis), if and only if $s \models \exists(\neg B \mathbf{U} \tilde{B})$.

Proposition 6.1.5 (Almost-sure repeated reachability). *For every $\mu \in \text{Dist}(S)$,*

- *if $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B) = 1$, then $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \tilde{B}) = 0$;*
- *if \mathcal{T} is $\text{StrDec}(\mu, B)$ and $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \tilde{B}) = 0$, then $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B) = 1$.*

Proof. We first show the first item. Since the event $\text{Ev}_{\mathcal{T}}(\mathbf{G} \mathbf{F} B)$ is almost-sure, we have

$$\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \tilde{B}) = \text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \tilde{B} \wedge \mathbf{G} \mathbf{F} B)$$

and then it is straightforward from point (ii) of Lemma 6.1.1.

In order to prove the second item, we assume that \mathcal{T} is $\text{StrDec}(\mu, B)$, i.e. $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B \vee \mathbf{F} \tilde{B}) = 1$. By assumption, the event $\text{Ev}_{\mathcal{T}}(\mathbf{F} \tilde{B})$ has probability 0, and thus $\text{Ev}_{\mathcal{T}}(\mathbf{G} \mathbf{F} B)$ is almost-sure. \square

This reduces the almost-sure model-checking of a repeated reachability property to the 0-model-checking of a reachability property.

Remark 6.1.6. Observe again that if \mathcal{T} is a DMC, Proposition 6.1.5 is equivalent to a result of [ABM07], which is given here as Proposition 2.2.17. Indeed in that case, as similarly noticed in Remark 6.1.4, it holds that for each state s , $\text{Prob}_s^{\mathcal{T}}(\mathbf{F} \tilde{B}) > 0$ if and only if $s \models \exists \mathbf{F} \tilde{B}$ and this is equivalent to $s \not\models \forall \mathbf{G} \exists \mathbf{F} B$.

Proposition 6.1.7 (Positive repeated reachability). *For every $\mu \in \text{Dist}(S)$,*

- *if \mathcal{T} is $\text{Dec}(\mu, \tilde{B})$ and if $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B) > 0$, then $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \tilde{\tilde{B}}) > 0$;*
- *if \mathcal{T} is $\text{PersDec}(\mu, B)$ and if $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \tilde{\tilde{B}}) > 0$, then $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B) > 0$.*

Proof. We begin with the first item. As \mathcal{T} is $\text{Dec}(\mu, \tilde{B})$, it holds that $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \tilde{B} \vee \mathbf{F} \tilde{\tilde{B}}) = 1$. Since the event $\text{Ev}_{\mathcal{T}}(\mathbf{F} \tilde{B} \vee \mathbf{F} \tilde{\tilde{B}})$ is almost-sure, we derive the equality:

$$\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B) = \text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B \wedge (\mathbf{F} \tilde{B} \vee \mathbf{F} \tilde{\tilde{B}})) .$$

Now from point (ii) of Lemma 6.1.1, we get that $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B \wedge (\mathbf{F} \widetilde{B} \vee \mathbf{F} \widetilde{\widetilde{B}})) = \text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B \wedge \mathbf{F} \widetilde{\widetilde{B}})$. Therefore $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B \wedge \mathbf{F} \widetilde{\widetilde{B}}) = \text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B) > 0$, and thus $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} \widetilde{\widetilde{B}}) > 0$.

Assume now that \mathcal{T} is $\text{PersDec}(\mu, B)$ and that $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} \widetilde{\widetilde{B}}) > 0$. Lemma 6.1.2 implies that $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} \widetilde{B}) < 1$. Since $\text{PersDec}(\mu, B)$ implies $\text{StrDec}(\mu, B)$, it follows that $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$ and thus, $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B) > 0$. \square

This reduces the positive model-checking of a repeated reachability property to the positive model-checking of a reachability property.

Remark 6.1.8. As previously observed in Remarks 6.1.4 and 6.1.6, Proposition 6.1.7 recovers partially a result of [ABM07] when \mathcal{T} is a DMC, namely here Theorem 2.2.18. Indeed, it can be established that given a state s , $\text{Prob}_s^{\mathcal{T}}(\mathbf{F} \widetilde{\widetilde{B}}) > 0$ if and only if $s \models \exists \mathbf{F} \widetilde{\widetilde{B}}$. Observe that the assumptions here are slightly different from the ones in Theorem 2.2.18. For the first item, the authors of [ABM07] needed moreover for the DMC to be decisive w.r.t. B , but they proved a stronger result that we do not need here. For the second item, the authors required for the DMC to be decisive w.r.t. B . However, remember that in Definition 2.2.12 a DMC being decisive w.r.t. B is equivalent in our context of Definition 4.2.3 to $\text{Dec}(\delta_{s'}, B)$ for each state s' . Note that for this result, it is in fact only needed to be decisive from each state s' reachable from s which is somehow what $\text{PersDec}(\mu, B)$ means as stated on page 95.

Hence, in all cases, under some specific assumptions, the properties one wants to check are reduced to checking whether a specific reachability (or Until) property has positive probability. These are the simplest properties one can hope to be decidable in a class of models. Effectiveness hence relies here on the computation of avoid-sets, avoid-sets of avoid-sets, and on the decidability of the positive reachability (or Until) problem.

Through abstractions, one can get nicer results. Indeed, via abstractions, one can reduce the qualitative analysis of basic properties (reachability and repeated reachability) from the concrete model to the abstract model. Indeed, one can use the previous results, in order to show the following proposition.

Proposition 6.1.9. *Assume \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 , and fix $B \in \Sigma_2$.*

- *Let $\mu \in \text{Dist}(S_1)$ be an initial distribution for \mathcal{T}_1 . Assume that \mathcal{T}_2 is μ -sound and μ -complete. Then:*

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1 \text{ iff } \text{Prob}_{\alpha\#(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) = 1 .$$

- Assume that \mathcal{T}_2 is sound and that \mathcal{T}_2 is $\text{Dec}(B)$. Then for every $\mu \in \text{Dist}(S_1)$:

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1 \text{ iff } \text{Prob}_{\alpha_\#(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) = 1 .$$

- Assume that \mathcal{T}_2 is sound and that \mathcal{T}_2 is $\text{StrDec}(B)$. Then for every $\mu \in \text{Dist}(S_1)$:

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(B)) = 1 \text{ iff } \text{Prob}_{\alpha_\#(\mu)}^{\mathcal{T}_2}(\mathbf{G} \mathbf{F} B) = 1 .$$

- Assume that \mathcal{T}_2 is sound and that \mathcal{T}_2 is $\text{PersDec}(B)$ and $\text{Dec}(\tilde{B})$. Then for every $\mu \in \text{Dist}(S_1)$:

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(B)) > 0 \text{ iff } \text{Prob}_{\alpha_\#(\mu)}^{\mathcal{T}_2}(\mathbf{G} \mathbf{F} B) > 0 .$$

Proof. The first item is just a consequence of the definitions of μ -sound and μ -complete abstractions (Definitions 5.1.8 and 5.1.9).

In order to get the following items, observe first that since \mathcal{T}_2 is sound, from Corollary 5.2.3 we can induce the decisiveness hypotheses of \mathcal{T}_2 on \mathcal{T}_1 . We also refer the reader to Corollary 5.1.7 and Lemma 5.2.2 for some properties of α -abstractions that will be useful here.

The first item is then immediate by using Proposition 6.1.3 for \mathcal{T}_1 and \mathcal{T}_2 , since they are respectively $\text{Dec}(\alpha^{-1}(B))$ and $\text{Dec}(B)$. Indeed, we obtain that

$$\begin{aligned} \text{Prob}_{\alpha_\#(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) = 1 &\iff \text{Prob}_{\alpha_\#(\mu)}^{\mathcal{T}_2}(\neg B \mathbf{U} \tilde{B}) = 0 \\ &\iff \text{Prob}_\mu^{\mathcal{T}_1}(\neg \alpha^{-1}(B) \mathbf{U} \alpha^{-1}(\tilde{B})) = 0 \iff \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1. \end{aligned}$$

We get similarly the third item by using Proposition 6.1.5 since \mathcal{T}_1 and \mathcal{T}_2 are respectively $\text{StrDec}(\alpha^{-1}(B))$ and $\text{StrDec}(B)$; and the fourth item by using Proposition 6.1.7 since \mathcal{T}_1 and \mathcal{T}_2 are respectively $\text{PersDec}(\alpha^{-1}(B))$, $\text{Dec}(\alpha^{-1}(\tilde{B}))$ and $\text{PersDec}(B)$, $\text{Dec}(\tilde{B})$. \square

Note that in particular, if \mathcal{T}_2 is DMC with a finite attractor and a sound α -abstraction of \mathcal{T}_1 , Proposition 6.1.9 can be applied thanks to Proposition 2.2.13. Then results of [ABM07] on the qualitative model-checking problem (see Section 2.2.2), induce that the qualitative model-checking problem of reachability and repeated reachability properties of α -closed sets for \mathcal{T}_1 is reduced to check that some structural properties are satisfied on the underlying graph of the Markov chain \mathcal{T}_2 . We will see that our applications enter in this framework.

6.1.2 Properties given by a DMA in DMCs

The objective of this section is to extend the results of [ABM07] to a richer class of properties. We consider properties given by a DMA, also called ω -regular properties. Remember that all LTL formulas can be encoded as a DMA. In this section, we consider thus the framework of DMCs.

In the case of DMCs with a finite attractor, the almost-sure satisfaction of properties given as a DMA can be characterised by finitely many reachability properties [ABRS05] and [Ber06]. In many cases, this characterisation yields an effective algorithm to decide the almost-sure satisfaction of ω -regular properties. Note that, since DMA are closed under complement, the positive probability of properties specified by DMA reduces to the (non-)almost-sure model-checking of the property specified by the complement automaton. We therefore concentrate on almost-sure model-checking in the sequel.

We aim here at adapting the approach of [ABRS05] and [Ber06] to our new formalism.

We fix a finite set of atomic propositions AP and we fix a DMA $M = (Q, q_0, E, \mathcal{F})$. Remember that given a LSTS $\mathcal{T} = (S, \Sigma, \kappa, AP, \mathcal{L})$, the product $\mathcal{T} \times M$ has been defined in Section 4.1.2. We first show a general result for LSTS, we will then consider labelled DMCs.

Since M has finitely many states, attractors transfer from LSTS \mathcal{T} to the product $\mathcal{T} \times M$, as formally stated below.

Lemma 6.1.10. Fix a LSTS $\mathcal{T} = (S, \Sigma, \kappa, AP, \mathcal{L})$ and assume that A is an attractor for \mathcal{T} . Then $A \times Q$ is an attractor for $\mathcal{T} \times M$.

In order to show the previous lemma, we first prove the following one:

Lemma 6.1.11. Let $\mathcal{T} = (S, \Sigma, \kappa, AP, \mathcal{L})$ be a LSTS. Fix $\mu \in \text{Dist}(S)$ and assume that $A \in \Sigma$ is a μ -attractor for \mathcal{T} . Then for each $q \in Q$, $A \times Q$ is a $(\mu \times \delta_q)$ -attractor for $\mathcal{T} \times M$.

Proof. Fix $\mu \in \text{Dist}(S)$ and $A \in \Sigma$ such that $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} A) = 1$. Fix $q \in Q$. We know that

$$\text{Ev}_{\mathcal{T} \times M}(\mathbf{F} A \times Q) = \text{Ev}_{\mathcal{T} \times M} \left(\bigcup_{n \in \mathbb{N}} \overbrace{\text{Cyl}(S', \dots, S', A \times Q)}^{n \text{ times}} \right).$$

Then from Lemma 4.1.18, we know that for each $n \in \mathbb{N}$

$$\begin{aligned} \text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times \mathbb{M}}(\text{Cyl}(\overbrace{S', \dots, S'}^{n \text{ times}}, A \times Q)) = \\ \sum_{u_1, \dots, u_n \in 2^{\mathcal{A}^P}} \text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(\mathcal{L}^{-1}(u_1), \dots, \mathcal{L}^{-1}(u_n), A)) = \text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(\overbrace{S, \dots, S}^{n \text{ times}}, A \times Q)). \end{aligned}$$

As this holds true for each $n \geq 0$, we thus get that $\text{Prob}_{\mu \times \delta_q}(\mathbf{F} A \times Q) = \text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} A) = 1$ from the hypothesis. This concludes the proof. \square

Proof of Lemma 6.1.10. Fix $A \in \Sigma$ such that for each $\mu \in \text{Dist}(S)$, $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} A) = 1$. We want to prove that for each $\nu \in \text{Dist}(S \times Q)$, $\text{Prob}_{\nu}^{\mathcal{T} \times \mathbb{M}}(\mathbf{F} A \times Q) = 1$. Fix $\nu \in \text{Dist}(S \times Q)$ and compute:

$$\text{Prob}_{\nu}^{\mathcal{T} \times \mathbb{M}}(\mathbf{F} A \times Q) = \sum_{q \in Q} \nu(S \times \{q\}) \cdot \text{Prob}_{\nu_{S \times \{q\}}}^{\mathcal{T} \times \mathbb{M}}(\mathbf{F} A \times Q).$$

Note that $\nu_{S \times \{q\}}$ induces a distribution $\nu_q \in \text{Dist}(S)$ as follows: for each $B \in \Sigma$, $\nu_q(B) = \nu_{S \times \{q\}}(B \times \{q\})$. Writing $\mu = \nu_q$ it then holds that $\nu_{S \times \{q\}} = \mu \times \delta_q$. We then get, from the hypothesis and Lemma 6.1.11, that $\text{Prob}_{\nu_{S \times \{q\}}}^{\mathcal{T} \times \mathbb{M}}(\mathbf{F} A \times Q) = 1$ for each $q \in Q$. Hence, $\text{Prob}_{\nu}^{\mathcal{T} \times \mathbb{M}}(\mathbf{F} A \times Q) = \sum_{q \in Q} \nu(S \times \{q\}) = 1$ which concludes the proof. \square

In the rest of the section, we assume \mathcal{T} to be a labelled DMC with a finite attractor. Observe that if A denotes this finite attractor, then Lemma 6.1.10 implies that $A \times Q$ is an attractor for $\mathcal{T} \times \mathbb{M}$. From finiteness of Q , it is furthermore obvious that $\mathcal{T} \times \mathbb{M}$ admits a finite attractor. We write it B . We write $\text{Graph}_{\mathcal{T} \times \mathbb{M}}(B)$ (or simply $\text{Graph}(B)$ when \mathcal{T} and \mathbb{M} are clear from the context) for the finite graph whose vertices are states of B , and in which there is an edge from (s_1, q_1) to (s_2, q_2) if there exists a finite run from (s_1, q_1) to (s_2, q_2) in $\mathcal{T} \times \mathbb{M}$ (written $(s_1, q_1) \rightarrow^* (s_2, q_2)$).

The BSCCs (see Section 2.2.1 for a brief reminder on the subject) of the graph $\text{Graph}_{\mathcal{T} \times \mathbb{M}}(B)$ play a central role in the model-checking problem of ω -regular properties of \mathcal{T} . Let us first discuss the relationships between the BSCCs and attractors for $\mathcal{T} \times \mathbb{M}$.

Lemma 6.1.12. The following properties are satisfied.

- The set $\{(s, q) \in C \mid C \text{ BSCC of } \text{Graph}_{\mathcal{T} \times \mathbb{M}}(B)\}$ is an attractor of $\mathcal{T} \times \mathbb{M}$.

- If C and C' are two distinct BSCCs of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$, then for every $(s, q) \in S \times Q$, $\text{Prob}_{(s,q)}^{\mathcal{T} \times \mathcal{M}}(\mathbf{F} C \wedge \mathbf{F} C') = 0$.
- If C is a BSCC of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$, then for each $(s, q) \in C$, it holds that $\text{Prob}_{(s,q)}^{\mathcal{T} \times \mathcal{M}}(\mathbf{G} \mathbf{F} C) = 1$.

Proof. The first property is obvious thanks to Proposition 2.2.10. The second property is a consequence of the fact that there is no run between two states of two different BSCCs: otherwise, from the definition of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$, there would be a vertex between those two states. This second property implies that for each BSCC $C' \neq C$ of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$ and for each $(s, q) \in C$, $\text{Prob}_{(s,q)}^{\mathcal{T} \times \mathcal{M}}(\mathbf{F} C') = 0$. From the first item and Lemma 4.2.8, we know that for each $(s, q) \in S \times Q$,

$$\text{Prob}_{(s,q)}^{\mathcal{T} \times \mathcal{M}}(\mathbf{G} \mathbf{F} \bigvee_{C' \in \text{BSCC}(\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B))} C') = 1.$$

This holds true in particular for each $(s, q) \in C$ and thus, from the previous observation for such initial distributions, we get that $\text{Prob}_{(s,q)}^{\mathcal{T} \times \mathcal{M}}(\mathbf{G} \mathbf{F} C) = 1$ for each $(s, q) \in C$. \square

From Lemma 6.1.12, the BSCCs of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$ form an attractor, and once the system enters a BSCC C , only that BSCC will be visited again, and this will happen infinitely often. In particular, the satisfaction of the Muller condition in $\mathcal{T} \times \mathcal{M}$, inherited from \mathcal{F} , can be characterised by the BSCCs satisfying the Muller condition \mathcal{F} (in a sense that we will make precise).

Definition 6.1.13 (Good BSCC). A BSCC C of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$ is *good for \mathcal{F}* , written $C \in \text{Good}_{\mathcal{T} \times \mathcal{M}}^B(\mathcal{F})$, if there exists $F \in \mathcal{F}$ such that

- for every state $(s, q) \in S \times Q$, if there exists $(r, p) \in C$ with $(r, p) \rightarrow^* (s, q)$, then $q \in F$; and
- for every $q \in F$ there exists $s \in S$, there exists a state $(r, p) \in C$ such that $(r, p) \rightarrow^* (s, q)$.

Let C be an arbitrary BSCC of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$. We define the set $F_C = \{q \in Q \mid \exists s \in S \exists (r, p) \in C \text{ s.t. } (r, p) \rightarrow^* (s, q)\}$ as the set of states of the DMA that can be reached from C . One can get the following obvious characterisation for a BSCC to be good.

Lemma 6.1.14. A BSCC C of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$ is good for \mathcal{F} if and only if $F_C \in \mathcal{F}$.

Within a BSCC, all reachable states will actually be visited infinitely often almost-surely. It is stated more precisely, in Lemma 6.1.15. We first recall some notations of Section 4.1.2. We write P'_M for the property over Q induced by a DMA M , which is a Muller property. Given an initial state $s \in S$ (resp. an initial distribution $\mu \in \text{Dist}(S)$), we then write $\text{Prob}_{(s,q_0)}^{\mathcal{T} \times M}(\text{Inf} \in \mathcal{F})$ (resp. $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times M}(\text{Inf} \in \mathcal{F})$) for the probability of the set of runs in $\mathcal{T} \times M$ satisfying the property P'_M , i.e. $\{\rho \in \text{Runs}(\mathcal{T} \times M) \mid \text{Inf}(\mathcal{L}'(\rho)) \in \mathcal{F}\}$ (\mathcal{L}' is the labelling function over $\mathcal{T} \times M$ such that for each state (s, q) , $\mathcal{L}'((s, q)) = q$; and if $\rho = (s_0, q_0)(s_1, q_1) \dots$, $\text{Inf}(\mathcal{L}'(\rho)) = \{q \in Q \mid |\{j \in \mathbb{N} \mid q = q_j\}| = \infty\}$). Similarly, given $F \subseteq Q$ and $s \in S$ (resp. $\mu \in \text{Dist}(S)$), we will write $\text{Prob}_{(s,q_0)}^{\mathcal{T} \times M}(\text{Inf} \sim F)$ (resp. $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times M}(\text{Inf} \sim F)$) with $\sim \in \{\subseteq, \supseteq, =\}$, for the set of runs $\rho \in \text{Runs}(\mathcal{T} \times M)$ such that $\text{Inf}(\mathcal{L}'(\rho)) \sim F$.

Lemma 6.1.15. Fix a BSCC C of $\text{Graph}_{\mathcal{T} \times M}(B)$. For every $(s, q) \in C$, it holds that $\text{Prob}_{(s,q)}^{\mathcal{T} \times M}(\text{Inf} = F_C) = 1$.

Proof. Fix $(s, q) \in C$ and a run $\rho = (s, q)(s_1, q_1)(s_2, q_2) \dots \in \text{Runs}(\mathcal{T} \times M)$ starting at (s, q) . By definition of F_C , it holds that for each $i \geq 1$, $q_i \in F_C$. Hence $\text{Prob}_{(s,q)}^{\mathcal{T} \times M}(\text{Inf} \subseteq F_C) = 1$.

We now argue why all elements of F_C are actually almost-surely visited infinitely often. Fix $p \in F_C$ and $(r, p) \in S \times Q$ that is reachable from C . Since all two states of C are reachable one from each other, we get that from every state of C , (r, p) is reachable through a finite path. Hence, as C is finite, there is some $\iota > 0$ and $k \in \mathbb{N}$ such that for every state $(s', q') \in C$,

$$\text{Prob}_{(s',q')}^{\mathcal{T} \times M}(\mathbf{F}_{\leq k}(r, p)) \geq \iota .$$

Applying a reasoning similar to the proof of Proposition 5.2.6, we get that $\text{Prob}_{(s,q)}^{\mathcal{T} \times M}(\mathbf{G} \mathbf{F}(r, p) \mid \mathbf{G} \mathbf{F} C) = 1$. Indeed, $\text{Prob}_{(s,q)}^{\mathcal{T} \times M}(\mathbf{F} \mathbf{G} \neg(r, p) \wedge \mathbf{G} \mathbf{F} C) \leq \lim_{n \rightarrow \infty} (1 - \iota)^n = 0$. Thanks to the third item of Lemma 6.1.12, we obtain that

$$\text{Prob}_{(s,q)}^{\mathcal{T} \times M}(\mathbf{G} \mathbf{F}(r, p)) = 1 .$$

As this last equality holds for each (r, p) reachable from (s, q) , we conclude that $\text{Prob}_{(s,q)}^{\mathcal{T} \times M}(\text{Inf} \supseteq F_C) = 1$, which completes the proof. \square

As a consequence, we get the following result.

Corollary 6.1.16. For every BSCC C of $\text{Graph}(B)$ and for every initial distribution $\mu \in \text{Dist}(S)$ for \mathcal{T} , $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times M}(\mathbf{F} C) > 0$ implies $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times M}(\text{Inf} = F_C \mid \mathbf{F} C) = 1$.

We can now completely characterise the probability of satisfying an ω -regular property.

Theorem 6.1.17. *Let \mathcal{T} be a labelled DMC with a finite attractor A , and $\mathbf{M} = (Q, q_0, E, \mathcal{F})$ be a DMA. Let $B = A \times Q$ be the corresponding finite attractor of $\mathcal{T} \times \mathbf{M}$. Then, for every initial distribution $\mu \in \text{Dist}(S)$ for \mathcal{T} :*

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathbf{M}}(\text{Inf} \in \mathcal{F}) = \sum_{C \in \text{Good}_{\mathcal{T} \times \mathbf{M}}^B(\mathcal{F})} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathbf{M}}(\mathbf{F} C).$$

Proof. Fix $\mu \in \text{Dist}(S)$. As stated in Lemma 6.1.12, the BSCCs of $\text{Graph}(B)$ form an attractor for $\mathcal{T} \times \mathbf{M}$, and two BSCCs are probabilistically disjoint. Using Bayes formula with a disjunction over the BSCCs, we can write:

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathbf{M}}(\text{Inf} \in \mathcal{F}) = \sum_{\substack{C \in \text{BSCC}(\text{Graph}(B)) \\ C \text{ } \mu \times \delta_{q_0}\text{-reachable}}} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathbf{M}}(\mathbf{F} C) \cdot \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathbf{M}}(\text{Inf} \in \mathcal{F} \mid \mathbf{F} C)$$

where we say that C is $\mu \times \delta_{q_0}$ -reachable whenever $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathbf{M}}(\mathbf{F} C) > 0$. Hence we deduce that:

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathbf{M}}(\text{Inf} \in \mathcal{F}) = \sum_{C \in \text{BSCC}(\text{Graph}(B))} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathbf{M}}(\mathbf{F} C) \cdot \mathbf{1}_{\mathcal{F}}(F_C)$$

thanks to Corollary 6.1.16, where $\mathbf{1}_{\mathcal{F}}$ is the characteristic function of \mathcal{F} (that is, $\mathbf{1}_{\mathcal{F}}(F) = 1$ if $F \in \mathcal{F}$, and $\mathbf{1}_{\mathcal{F}}(F) = 0$ otherwise). Lemma 6.1.14 allows then us to conclude the proof of the theorem. \square

As an immediate corollary of Theorem 6.1.17, we obtain a characterisation for the almost-sure model-checking problem of properties given as a DMA.

Corollary 6.1.18 (Almost-sure ω -regular property). *Let \mathcal{T} be a labelled DMC with a finite attractor, and $\mathbf{M} = (Q, q_0, E, \mathcal{F})$ be a DMA. Let B be a finite attractor for $\mathcal{T} \times \mathbf{M}$. For every initial distribution $\mu \in \text{Dist}(S)$ for \mathcal{T} :*

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathbf{M}}(\text{Inf} \in \mathcal{F}) = 1 \text{ if and only if for every BSCC } C \text{ of } \text{Graph}_{\mathcal{T} \times \mathbf{M}}(B) \\ \text{such that } \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathbf{M}}(\mathbf{F} C) > 0, C \text{ is good for } \mathcal{F}.$$

Proof. This is immediate from Theorem 6.1.17 and Lemma 6.1.12. \square

Observe that since $\mathcal{T} \times \mathbf{M}$ is a DMC, the positive reachability model-checking problem is reduced to the satisfaction of a reachability property on the underlying

graph [ABM07] (see Proposition 2.2.16 here). Checking whether BSCC C is good for \mathcal{F} is also a structural property on the underlying graph. Hence, the qualitative model-checking problem of ω -regular properties in DMCs with a finite attractor is, like in [ABM07], reduced to the satisfaction of some structural properties on the underlying graph.

In order to turn this characterisation into a decision procedure, we need to be able to compute the attractor B for $\mathcal{T} \times \mathbb{M}$, and to build the graph $\text{Graph}_{\mathcal{T} \times \mathbb{M}}(B)$; also one needs to be able to compute for every BSCC C the set F_C . Observe also that we did not need decisiveness in order to get this procedure, however it was hidden in the fact that \mathcal{T} has a finite attractor: this implies decisiveness [ABM07] (see Proposition 2.2.13).

6.1.3 Properties given by a DMA in general STSs via denumerable abstractions

While the approach of Section 6.1.2 is adapted to DMCs, it does not apply directly to general STSs: indeed, it is unlikely that general STSs have *finite* attractors, and finiteness of the attractor is fundamental for the correctness of the approach. The idea will then be to rely on an abstraction, and to transfer properties through that abstraction.

Let $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1, \text{AP}, \mathcal{L}_1)$ and $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2, \text{AP}, \mathcal{L}_2)$ be two labelled STSs such that \mathcal{T}_2 is a DMC, which is an α -abstraction of \mathcal{T}_1 . Under certain conditions, we show how to perform the qualitative model-checking of ω -regular properties on \mathcal{T}_1 by transferring the same analysis on \mathcal{T}_2 . We assume the ω -regular property is given by a DMA $\mathbb{M} = (Q, q_0, E, \mathcal{F})$. We consider both the product $\mathcal{T}_1 \times \mathbb{M}$ and the product $\mathcal{T}_2 \times \mathbb{M}$ (we refer to Section 4.1.2 for the definitions and notations).

First we justify why, within a slight abuse of terminology, $\mathcal{T}_2 \times \mathbb{M}$ can be viewed as an α -abstraction of $\mathcal{T}_1 \times \mathbb{M}$. We also exhibit a sufficient condition under which it is sound.

Lemma 6.1.19. Let $\alpha_{\mathbb{M}} : S_1 \times Q \rightarrow S_2 \times Q$ be the unique lifting of α such that $\alpha_{\mathbb{M}}(s, q) = (\alpha(s), q)$. If \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 , then $\mathcal{T}_2 \times \mathbb{M}$ is an $\alpha_{\mathbb{M}}$ -abstraction of $\mathcal{T}_1 \times \mathbb{M}$. Furthermore, if $\mathcal{T}_1 \times \mathbb{M}$ is $\text{Dec}(\mathcal{B})$ where $\mathcal{B} = \{\alpha_{\mathbb{M}}^{-1}(B) \mid B \in \Sigma'_2\}$, then $\mathcal{T}_2 \times \mathbb{M}$ is a sound $\alpha_{\mathbb{M}}$ -abstraction of $\mathcal{T}_1 \times \mathbb{M}$.

While the proof of the first part of the lemma is technical, the second part of the lemma is a direct consequence of Proposition 5.3.4.

Proof. We first show that $\mathcal{T}_2 \times \mathbb{M}$ is an $\alpha_{\mathbb{M}}$ -abstraction of $\mathcal{T}_1 \times \mathbb{M}$. It suffices to

show that for each $\mu \in \text{Dist}(S_1)$, for each $q, q' \in Q$ and for each $B_{q'} \in \Sigma_2$,

$$\begin{aligned} \text{Prob}_{\mu \times \delta_q}^{\mathcal{T}_1 \times \mathbb{M}}(\text{Cyl}(S_1 \times Q, \alpha_M^{-1}(B_{q'} \times \{q'\}))) &> 0 \\ \iff \text{Prob}_{(\alpha_M)_\#(\mu \times \delta_q)}^{\mathcal{T}_2 \times \mathbb{M}}(\text{Cyl}(S_2 \times Q, B_{q'} \times \{q'\})) &> 0. \end{aligned} \quad (6.1)$$

Fix $\mu \in \text{Dist}(S_1)$, $q, q' \in Q$ and $B_{q'} \in \Sigma_2$. Write $u \in 2^{\text{AP}}$ for the unique label such that $(q, u, q') \in E$. In order to prove (6.1), we will use the fact that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . And in order to make the link with the wanted equivalence, we will use Lemma 4.1.18. We can establish that $(\alpha_M)_\#(\mu \times \delta_q) = \alpha_\#(\mu) \times \delta_q$. Indeed given $p \in Q$ and $C_p \in \Sigma_2$, it holds that

$$\begin{aligned} (\alpha_M)_\#(\mu \times \delta_q)(C_p \times \{p\}) &= (\mu \times \delta_q)(\alpha^{-1}(C_p) \times \{p\}) \\ &= \mu(\alpha^{-1}(C_p)) \cdot \delta_q(p) \\ &= \alpha_\#(\mu)(\alpha^{-1}(C_p)) \cdot \delta_q(p) = (\alpha_\#(\mu) \times \delta_q)(C_p \times \{p\}). \end{aligned}$$

Hence we get that

$$\begin{aligned} \text{Prob}_{(\alpha_M)_\#(\mu \times \delta_q)}^{\mathcal{T}_2 \times \mathbb{M}}(\text{Cyl}(S_2 \times Q, B_{q'} \times \{q'\})) &> 0 \\ \iff \text{Prob}_{\alpha_\#(\mu)}^{\mathcal{T}_2}(\text{Cyl}(\mathcal{L}_2^{-1}(u), B_{q'})) &> 0 \\ \iff \text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(\mathcal{L}_1^{-1}(u), \alpha^{-1}(B_{q'}))) &> 0 \\ \iff \text{Prob}_{\mu \times \delta_q}^{\mathcal{T}_1 \times \mathbb{M}}(\text{Cyl}(S_1 \times Q, \alpha_M^{-1}(B_{q'} \times \{q'\}))) &> 0 \end{aligned}$$

where the first and third equivalences hold from Lemma 4.1.18, and the second equivalence holds from the fact that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 (Definition 5.1.2).

Finally, since $\mathcal{T}_1 \times \mathbb{M}$ is decisive w.r.t $\alpha_M^{-1}(B)$ for each $B \in \Sigma'_2$ and since $\mathcal{T}_2 \times \mathbb{M}$ is an α_M -abstraction of $\mathcal{T}_1 \times \mathbb{M}$, Proposition 5.3.4 allows us to conclude that $\mathcal{T}_2 \times \mathbb{M}$ is a sound α_M -abstraction of $\mathcal{T}_1 \times \mathbb{M}$. \square

Remark 6.1.20. In the sequel, our applications will be smooth enough to meet the hypothesis: $\mathcal{T}_1 \times \mathbb{M}$ is decisive w.r.t. α_M -closed sets. However we still have several open questions. The first one is the following: does soundness between \mathcal{T}_2 and \mathcal{T}_1 imply soundness between $\mathcal{T}_2 \times \mathbb{M}$ and $\mathcal{T}_1 \times \mathbb{M}$? While this seems quite natural, it is surprisingly tricky. Although we did not manage to find a counter-example for this general question, we found one for a fixed initial distribution. It is described in Example 6.1.21 below and highlights some difficulties we encounter when aiming at transferring analysis from the abstraction to the concrete model.

This justifies the fact that we assumed decisiveness. As we already know, if \mathcal{T}_2 is a sound α -abstraction of \mathcal{T}_1 and \mathcal{T}_2 is decisive w.r.t. any set of states, then \mathcal{T}_1

is decisive w.r.t. any α -closed sets. We do not know if the product preserves the soundness, however a second natural question is the following: does decisiveness w.r.t. α -closed sets for \mathcal{T}_1 imply decisiveness w.r.t. α_M -closed sets for $\mathcal{T}_1 \times M$? Again, we do not have a general counter-example, but we have one for a fixed initial distribution. This is also described in Example 6.1.21 below.

Recall though that a condition for decisiveness is given in Proposition 5.2.6 for the case where the abstraction is a DMC with a finite attractor. While we know from Lemma 6.1.10 that if \mathcal{T}_2 has finite attractor, then so has $\mathcal{T}_2 \times M$, another still open question is whether the satisfaction of hypothesis (\dagger) is preserved through the product with a DMA, *i.e.* if \mathcal{T}_1 satisfies (\dagger) , does $\mathcal{T}_1 \times M$ satisfy it too? The tricky part comes from the lower bounded reachability probabilities as it will also be somehow enlightened in Example 6.1.21.

Proposition 5.2.7 gives another condition for decisiveness in the case of a finite Markov chain for the abstraction, which is fairness of $\mathcal{T}_1 \times M$. We will discuss this particular case later and show that fairness is preserved through the product with a DMA in Lemma 6.1.27.

Example 6.1.21. We illustrate Remark 6.1.20 by exhibiting an example where soundness (w.r.t. a fixed distribution) as well as decisiveness properties do not transfer to the product with a deterministic Muller automaton.

Consider the DMC \mathcal{T}_1 depicted on the left of Figure 6.1 which corresponds to the random walk over \mathbb{N} from Example 4.1.2, when $p = 2/3$. Consider also the finite MC \mathcal{T}_2 on the right of the same figure. Clearly enough, \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 for the mapping $\alpha : \mathbb{N} \rightarrow \{s_0, s_1, s_2\}$ defined as follows: $\alpha(0) = s_0$, $\alpha(1) = s_1$ and $\alpha(i) = s_2$ for any $i \geq 2$.

Define $\mu = \delta_0$ as the initial distribution in \mathcal{T}_1 . For any $B \subseteq \mathbb{N}$, $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} B) = 1$ and it follows that \mathcal{T}_2 is a μ -sound α -abstraction of \mathcal{T}_1 . It should be noted that it is however not sound when considering $\mu' = \delta_1$ as initial distribution. Indeed, $\text{Prob}_{\mu'}^{\mathcal{T}_1}(\mathbf{F} \{0\}) < 1$ although $\text{Prob}_{s_1}^{\mathcal{T}_2}(\mathbf{F} \{s_0\}) = 1$ (and $\delta_{s_1} = \alpha_\#(\mu')$).




Figure 6.1: Left, \mathcal{T}_1 a random walk over \mathbb{N} and right, its sound finite abstraction \mathcal{T}_2 .

Consider now the Muller automaton of Section 4.1.2 on the left of Figure 4.1. As stated in Lemma 6.1.19, it holds that $\mathcal{T}_2 \times M$ is an α_M -abstraction of $\mathcal{T}_1 \times M$

where for each $n \in \mathbb{N}$ and each $q \in Q$, $\alpha_M((n, q)) = (\alpha(n), q)$. Consider $\mu \times \delta_{q_0} = \delta_{(0, q_0)}$ and $B = \{(s_0, q_2)\}$. It then holds that $(\alpha_M)_\#(\mu \times \delta_{q_0}) = \delta_{(s_0, q_0)}$ and that $\alpha_M^{-1}(B) = \{(0, q_2)\}$. It is easily observed that starting in state $(0, q_0)$ (resp. (s_0, q_0)) in $\mathcal{T}_1 \times M$ (resp. $\mathcal{T}_2 \times M$), if we visit in the future a state $(0, q)$ (resp. (s_0, q)) we will necessarily get that $q = q_2$. Keeping this in mind, one can see that $\text{Prob}_{(s_0, q_0)}^{\mathcal{T}_2 \times M}(\mathbf{F} B) = 1$ while

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times M}(\mathbf{F} \alpha_M^{-1}(B)) = \text{Prob}_{(1, q_1)}^{\mathcal{T}_1 \times M}(\mathbf{F} \alpha_M^{-1}(B)) = \text{Prob}_{\mu'}^{\mathcal{T}_1}(\mathbf{F} \{0\}) < 1$$

where the first equality holds from Lemma 4.1.8 and the second equality holds from Lemma 4.1.18. This proves that $\mathcal{T}_2 \times M$ is not $(\mu \times \delta_{q_0})$ -sound for $\mathcal{T}_1 \times M$. This difference in the reachability probabilities somehow shows also the difficulties encountered when trying to show that hypothesis (†) of Proposition 5.2.6 is preserved through the product.

Now, observe that \mathcal{T}_1 is decisive w.r.t. any set of states $B \subseteq \mathbb{N}$ from μ as we have seen that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} B) = 1$ for any set of states B . It should be noted that \mathcal{T}_1 is not decisive by considering μ' as the initial distribution and $B = \{0\}$. In this case, $\widetilde{\{0\}} = \emptyset$ and thus $\text{Prob}_{\mu'}^{\mathcal{T}_1}(\mathbf{F} \{0\} \vee \mathbf{F} \widetilde{\{0\}}) = \text{Prob}_{\mu'}^{\mathcal{T}_1}(\mathbf{F} \{0\}) < 1$. Consider now $\mathcal{T}_1 \times M$, we have already shown that $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times M}(\mathbf{F} \{(0, q_2)\}) < 1$. It can be established that $\widetilde{\{(0, q_2)\}} = (2\mathbb{N} + 1) \times \{q_0, q_2\} \cup 2\mathbb{N} \times \{q_1\}$ which are states not reachable from $(0, q_0)$. We deduce that $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times M}(\mathbf{F} \{(0, q_2)\} \vee \mathbf{F} \widetilde{\{(0, q_2)\}}) = \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times M}(\mathbf{F} \{(0, q_2)\}) < 1$. This shows that $\mathcal{T}_1 \times M$ is not decisive w.r.t. $\{(0, q_2)\}$ from $\mu \times \delta_{q_0}$.

From now on, whenever $\mathcal{T}_1 \times M$ is decisive w.r.t. α_M -closed sets (or whenever $\mathcal{T}_2 \times M$ is a sound α_M -abstraction of $\mathcal{T}_1 \times M$) and thus Lemma 6.1.19 is applicable, we will abusively write α for α_M .

We focus now on the case where \mathcal{T}_2 has a finite attractor.¹ Thanks to Lemma 6.1.10, $\mathcal{T}_2 \times M$ has also a finite attractor, which we denote B_2 . We reuse notations of the previous subsection, in particular the graph of the attractor $\text{Graph}_{\mathcal{T}_2 \times M}(B_2)$, and the set F_C of recurring states when C is a BSCC of that graph.

The following lemma is a counterpart to Lemma 6.1.12 for \mathcal{T}_1 . Under the hypothesis that $\mathcal{T}_1 \times M$ is decisive w.r.t. α -closed sets or, since it is what is truly important for the next result but it is implied by this decisiveness assumption (Lemma 6.1.19), under the hypothesis that $\mathcal{T}_2 \times M$ is a sound α -abstraction of $\mathcal{T}_1 \times$

¹As \mathcal{T}_2 has a finite attractor it is decisive and thus \mathcal{T}_2 is a complete α -abstraction of \mathcal{T}_1 by Lemma 5.3.2.

M , even though $\mathcal{T}_1 \times M$ does not have a finite attractor, it has an attractor with an interesting structure inherited from $\mathcal{T}_2 \times M$. In the sequel, we write $\mathcal{B} = \{\alpha^{-1}(B) \mid B \in \Sigma'_2\}$ (Σ'_2 is the σ -algebra of $S'_2 = S_2 \times Q$ defined in Definition 4.1.15).

Lemma 6.1.22. Assume that \mathcal{T}_2 has a finite attractor, and assume that $\mathcal{T}_2 \times M$ is a sound α -abstraction of $\mathcal{T}_1 \times M$. The following properties are satisfied.

- The set $\alpha^{-1}(\{(s, q) \in C \mid C \text{ BSCC of } \text{Graph}_{\mathcal{T}_2 \times M}(B_2)\})$ is an attractor of $\mathcal{T}_1 \times M$.
- If C and C' are two distinct BSCCs of $\text{Graph}_{\mathcal{T}_2 \times M}(B_2)$, for every $\mu \in \text{Dist}(S_1 \times Q)$, $\text{Prob}_\mu^{\mathcal{T}_1 \times M}(\mathbf{F} \alpha^{-1}(C) \wedge \mathbf{F} \alpha^{-1}(C')) = 0$.
- If C is a BSCC of $\text{Graph}_{\mathcal{T}_2 \times M}(B_2)$, for every $\mu \in \text{Dist}(\alpha^{-1}(C))$, it holds that $\text{Prob}_\mu^{\mathcal{T}_1 \times M}(\mathbf{G} \mathbf{F} \alpha^{-1}(C)) = 1$.

Proof. Since $\mathcal{T}_2 \times M$ is a sound α -abstraction of $\mathcal{T}_1 \times M$, the first property derives from Proposition 5.2.4 and Lemma 6.1.12. The second property is a consequence of Lemma 6.1.12, and of the fact that $\mathcal{T}_2 \times M$ is an α -abstraction of $\mathcal{T}_1 \times M$. Finally, the third property is, as in the proof of Lemma 6.1.12, a consequence of the second point and of Lemma 4.2.8. \square

We then prove a counterpart to Lemma 6.1.15 for \mathcal{T}_1 , which shows that a BSCC C is characterised by the set F_C of states that are visited infinitely often from C .

Lemma 6.1.23. Assume that \mathcal{T}_2 has a finite attractor, and assume that $\mathcal{T}_2 \times M$ is a sound α -abstraction of $\mathcal{T}_1 \times M$. Let C be a BSCC of $\text{Graph}_{\mathcal{T}_2 \times M}(B_2)$, and $\mu \in \text{Dist}(\alpha^{-1}(C))$. Then:

$$\text{Prob}_\mu^{\mathcal{T}_1 \times M}(\text{Inf} = F_C) = 1.$$

Proof. As already argued in the proof of Lemma 6.1.15, for every $p \in F_C$, for every state $\mathbf{s}_2 \in C$, $\text{Prob}_{\mathbf{s}_2}^{\mathcal{T}_2 \times M}(\mathbf{F} p) = 1$ and thus in particular, $\text{Prob}_{\mathbf{s}_2}^{\mathcal{T}_2 \times M}(\mathbf{F} p) = 1$ (where we abusively write p for the measurable set $S_2 \times \{p\}$). Since $\mathcal{T}_2 \times M$ is a sound α -abstraction of $\mathcal{T}_1 \times M$, we derive for every $\nu \in \text{Dist}(\alpha^{-1}(C))$ that $\text{Prob}_\nu^{\mathcal{T}_1 \times M}(\mathbf{F} p) = 1$ (as before we abusively write p for $S_1 \times \{p\} = \alpha^{-1}(S_2 \times \{p\})$). We can then show that for each $\nu \in \text{Dist}(\alpha^{-1}(C))$ and for each $p \in F_C$,

$$\text{Prob}_\nu^{\mathcal{T}_1 \times M}(\mathbf{G} \mathbf{F} p) = 1.$$

Indeed, towards a contradiction, we assume that there is a distribution $\nu \in \text{Dist}(\alpha^{-1}(C))$ such that $\text{Prob}_\nu^{\mathcal{T}_1 \times M}(\mathbf{G} \mathbf{F} p) < 1$, i.e. $\text{Prob}_\nu^{\mathcal{T}_1 \times M}(\mathbf{F} \mathbf{G} \neg p) > 0$. The

third point of Lemma 6.1.22 implies that $\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{G} \mathbf{F} \alpha^{-1}(C) \wedge \mathbf{F} \mathbf{G} \neg p) > 0$. Now, observe that

$$\text{Ev}_{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{G} \mathbf{F} \alpha^{-1}(C) \wedge \mathbf{F} \mathbf{G} \neg p) \subseteq \text{Ev}_{\mathcal{T}_1 \times \mathcal{M}}\left(\bigcup_{n \in \mathbb{N}} (\mathbf{F}_{=n} \alpha^{-1}(C) \wedge \mathbf{G}_{\geq n} \neg p)\right).$$

It follows that there is $n \in \mathbb{N}$ such that $\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F}_{=n} \alpha^{-1}(C) \wedge \mathbf{G}_{\geq n} \neg p) > 0$. From Lemma 4.1.8, we get that there is $\nu' \in \text{Dist}(S'_1)$ (with $S'_1 = S_1 \times Q$, see Definition 4.1.15) such that

$$\begin{aligned} & \text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F}_{=n} \alpha^{-1}(C) \wedge \mathbf{G}_{\geq n} \neg p) \\ &= \lim_{m \rightarrow \infty} \text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\overbrace{\text{Cyl}(S'_1, \dots, S'_1)}^{n \text{ times}}, \alpha^{-1}(C) \wedge \neg p, \overbrace{\neg p, \dots, \neg p}^{m \text{ times}}) \\ &\leq \lim_{m \rightarrow \infty} \text{Prob}_{\nu'}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Cyl}(\alpha^{-1}(C) \wedge \neg p, \overbrace{\neg p, \dots, \neg p}^{m \text{ times}})) \quad \text{from Lemma 4.1.8} \\ &= \lim_{m \rightarrow \infty} \nu'(\alpha^{-1}(C)) \cdot \text{Prob}_{\nu'}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Cyl}(\neg p, \overbrace{\neg p, \dots, \neg p}^{m \text{ times}})) \\ &= \nu'(\alpha^{-1}(C)) \cdot \text{Prob}_{\nu'}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{G} \neg p). \end{aligned}$$

From the assumption, we thus get that $\text{Prob}_{\nu'}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{G} \neg p) > 0$ where $\nu'_{\alpha^{-1}(C)} \in \text{Dist}(\alpha^{-1}(C))$ which is the required contradiction. We conclude that for each $\nu \in \text{Dist}(\alpha^{-1}(C))$ and for each $p \in F_C$, $\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{G} \mathbf{F} p) = 1$.

It now suffices to show that, from any $\nu \in \text{Dist}(\alpha^{-1}(C))$, no other state is visited almost-surely infinitely often. Fix $p' \notin F_C$. Then, by definition of F_C , we have that $\text{Prob}_{\alpha_{\#}(\nu)}^{\mathcal{T}_2 \times \mathcal{M}}(\mathbf{F} p') = 0$. Since $\mathcal{T}_2 \times \mathcal{M}$ is an α -abstraction of $\mathcal{T}_1 \times \mathcal{M}$, we deduce from Corollary 5.1.7 that $\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} p') = 0$.

We conclude that $\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Inf} = F_C) = 1$, which is the claim of the lemma. \square

We get as a consequence, a similar result as Corollary 6.1.16.

Corollary 6.1.24. *For each BSCC C of $\text{Graph}_{\mathcal{T}_2 \times \mathcal{M}}(B_2)$ and for each initial distribution $\mu \in \text{Dist}(S_1)$ for \mathcal{T}_1 , $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \alpha^{-1}(C)) > 0$ implies $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Inf} = F_C \mid \mathbf{F} \alpha^{-1}(C))$.*

We are now in a position to decompose the probability to satisfy the Muller condition \mathcal{F} in $\mathcal{T}_1 \times \mathcal{M}$ into the reachability probability of good BSCCs.

Theorem 6.1.25. *Let \mathcal{T}_1 and \mathcal{T}_2 be two LSTs such that \mathcal{T}_2 is a DMC with a finite attractor B_2 , and \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Let $\mathcal{M} = (Q, q_0, E, \mathcal{F})$ be a*

DMA. Assume moreover that $\mathcal{T}_2 \times \mathbf{M}$ is a sound α -abstraction of $\mathcal{T}_1 \times \mathbf{M}$. Then, for every initial distribution $\mu \in \text{Dist}(S_1)$ for \mathcal{T}_1 :

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathbf{M}}(\text{Inf} \in \mathcal{F}) = \sum_{C \in \text{Good}_{\mathcal{T}_2 \times \mathbf{M}}^{B_2}(\mathcal{F})} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathbf{M}}(\mathbf{F} \alpha^{-1}(C)) .$$

Proof. Fix an initial distribution $\mu \in \text{Dist}(S_1)$. We get the result similarly as in the proof of Theorem 6.1.17. By the two first properties of Lemma 6.1.22, we can write the following Bayes formula, with a disjunction over the BSCCs of $\text{Graph}_{\mathcal{T}_2 \times \mathbf{M}}(B_2)$:

$$\begin{aligned} & \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathbf{M}}(\text{Inf} \in \mathcal{F}) \\ &= \sum_{\substack{C \in \text{BSCC}(\text{Graph}_{\mathcal{T}_2 \times \mathbf{M}}(B_2)) \\ C \text{ } \mu \times \delta_{q_0}\text{-reachable}}} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathbf{M}}(\mathbf{F} \alpha^{-1}(C)) \cdot \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathbf{M}}(\text{Inf} \in \mathcal{F} \mid \mathbf{F} \alpha^{-1}(C)) \\ &= \sum_{C \in \text{BSCC}(\text{Graph}_{\mathcal{T}_2 \times \mathbf{M}}(B_2))} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathbf{M}}(\mathbf{F} \alpha^{-1}(C)) \cdot \mathbb{1}_{\mathcal{F}}(F_C) \\ &= \sum_{C \in \text{Good}_{\mathcal{T}_2 \times \mathbf{M}}^{B_2}(\mathcal{F})} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathbf{M}}(\mathbf{F} \alpha^{-1}(C)), \end{aligned}$$

where C is said $\mu \times \delta_{q_0}$ -reachable whenever $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathbf{M}}(\mathbf{F} \alpha^{-1}(C)) > 0$, and where the second equality stands from Corollary 6.1.24 and the third equality stands from Lemma 6.1.14. This concludes the proof of the theorem. \square

In particular, regarding the qualitative model-checking problem of ω -regular properties, we conclude with the following characterisation of the almost-sure satisfaction of properties specified by a DMA.

Corollary 6.1.26. *Let \mathcal{T}_1 and \mathcal{T}_2 be two LSTSs such that \mathcal{T}_2 is a DMC with a finite attractor, and \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Let $\mathbf{M} = (Q, q_0, E, \mathcal{F})$ be a DMA. Assume moreover that $\mathcal{T}_2 \times \mathbf{M}$ is a sound α -abstraction of $\mathcal{T}_1 \times \mathbf{M}$. Then, for every initial distribution μ for \mathcal{T}_1 and for every $q \in Q$:*

$$\text{Prob}_{\mu \times \delta_q}^{\mathcal{T}_1 \times \mathbf{M}}(\text{Inf} \in \mathcal{F}) = 1 \quad \text{if and only if} \quad \text{Prob}_{\alpha\#(\mu \times \delta_q)}^{\mathcal{T}_2 \times \mathbf{M}}(\text{Inf} \in \mathcal{F}) = 1 .$$

Proof. We get first a similar characterisation as the one of Corollary 6.1.18: $\text{Prob}_{\mu \times \delta_q}^{\mathcal{T}_1 \times \mathbf{M}}(\text{Inf} \in \mathcal{F}) = 1$ if and only if for each BSCC C of $\text{Graph}_{\mathcal{T}_2 \times \mathbf{M}}(B_2)$ such that $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathbf{M}}(\mathbf{F} \alpha^{-1}(C)) > 0$, $C \in \text{Good}_{\mathcal{T}_2 \times \mathbf{M}}^{B_2}(\mathcal{F})$ (it comes from Lemma 6.1.22 and Theorem 6.1.25). We conclude from Corollary 5.1.7 and Corollary 6.1.18. \square

Hence, this reduces the almost-sure model-checking of a property given by M in \mathcal{T}_1 to the almost-sure model-checking of a reachability property (applying Corollary 6.1.18) in a DMC and this is reduced to structural properties on the underlying graph. For the approach to be effective, it is sufficient that the analysis at the level of $\mathcal{T}_2 \times M$ is effective.

As already quickly mentioned, under the hypotheses of Corollary 6.1.26, the abstraction $\mathcal{T}_2 \times M$ is complete (since it has a finite attractor). Though it is not explicitly used, we could not have such an equivalence without some completeness of the abstraction.

The particular of a finite Markov chain abstraction

In this section, we consider the particular case where \mathcal{T}_2 is a finite Markov chain. Observe that in particular, \mathcal{T}_2 is a DMC with a finite attractor, and thus results of Section 6.1.2 can be applied.

The results of Section 6.1.3 rely very much on the assumption that $\mathcal{T}_2 \times M$ is a sound α -abstraction of $\mathcal{T}_1 \times M$ which can be implied, in this case, by the assumption that $\mathcal{T}_1 \times M$ is fair w.r.t. any α_M -closed sets thanks to Proposition 5.2.7.

In Remark 6.1.20, we have discussed several open questions for the transfer of properties from a LSTS to its product with some DMA. Those forced us to add the hypothesis that $\mathcal{T}_2 \times M$ is sound (or more simply $\mathcal{T}_1 \times M$ is decisive w.r.t. α -closed sets, see Lemma 6.1.19). However, we show in this section that fairness is preserved through the product with a DMA.

Hence, like in Proposition 5.2.7, the only assumption needed in the case where the α -abstraction \mathcal{T}_2 is a finite Markov chain in order to get that $\mathcal{T}_2 \times M$ is sound, is for \mathcal{T}_1 to be fair w.r.t. α -closed sets! This simple assumption allows thus to apply the results of Section 6.1.3 in this framework.

We immediately show the mentioned result.

Lemma 6.1.27. Let $\mathcal{T} = (S, \Sigma, \kappa, AP, \mathcal{L})$ be a LSTS and let $M = (Q, q_0, E, \mathcal{F})$ be a DMA. Fix $\mathcal{B} \subseteq \Sigma$ and write $\mathcal{B}_M = \{\bigcup_{q \in Q} B_q \times \{q\} \mid B_q \in \mathcal{B} \forall q \in Q\}$. It holds that if \mathcal{T} is fair(\mathcal{B}) then $\mathcal{T} \times M$ is fair(\mathcal{B}_M).

Like in the proof of Lemma 6.1.10, this is a consequence of the following lemma.

Lemma 6.1.28. Let $\mathcal{T} = (S, \Sigma, \kappa, AP, \mathcal{L})$ be a LSTS and let $M = (Q, q_0, E, \mathcal{F})$ be a DMA. Fix $\mu \in \text{Dist}(S)$, $\mathcal{B} \subseteq \Sigma$ and write $\mathcal{B}_M = \{\bigcup_{q \in Q} B_q \times \{q\} \mid B_q \in \mathcal{B} \forall q \in Q\}$. It holds that if \mathcal{T} is fair(μ, \mathcal{B}) then for any $q \in Q$, $\mathcal{T} \times M$ is fair($\mu \times \delta_q, \mathcal{B}_M$).

Proof. We have to show that for any $B \in \mathcal{B}_M$ and for any $B' \in \text{PreProb}^{\mathcal{T} \times M}(B)$ such that $\text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times M}(\mathbf{G F} B') > 0$,

$$\text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times M}(\mathbf{G F} B \mid \mathbf{G F} B') = 1.$$

We first prove this for $B = B_{q'} \times \{q'\}$ for some $q' \in Q$ and $B_{q'} \in \mathcal{B}$ and we will then show how this extends to any $B \in \mathcal{B}_M$. We write $B = (B_{q'}, q')$. Fix $B' = \cup_{q'' \in Q} (B'_{q''}, q'') \in \text{PreProb}^{\mathcal{T} \times M}(B)$ and w.l.o.g., assume that $B'_{q''} \neq \emptyset$ for each $q'' \in Q$. Then,

$$\begin{aligned} & \forall \mu' \in \text{Dist}(B'), \text{Prob}_{\mu'}^{\mathcal{T} \times M}(\text{Cyl}(B', (B_{q'}, q'))) > 0 \\ \iff & \forall q'' \in Q, \forall \mu' \in \text{Dist}(B'_{q''}), \text{Prob}_{\mu' \times \delta_{q''}}^{\mathcal{T} \times M}(\text{Cyl}((B'_{q''}, q''), (B_{q'}, q'))) > 0 \\ \iff & \forall q'' \in Q, \forall s, s' \in B'_{q''}, \mathcal{L}(s) = \mathcal{L}(s'), q'' \xrightarrow{\mathcal{L}(s)} q' \text{ in } M, \text{ and} \quad (6.2) \\ & \forall q'' \in Q, \forall \mu' \in \text{Dist}(B'_{q''}), \text{Prob}_{\mu'}^{\mathcal{T}}(\text{Cyl}(B'_{q''}, B_{q'})) > 0 \end{aligned}$$

where the last equivalence comes from the the definition of $\mathcal{T} \times M$ (Definition 4.1.15) and from Lemma 4.1.18. Note that the very last line states that $B'_{q''} \in \text{PreProb}^{\mathcal{T}}(B_{q'})$.

Assume now that $\text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times M}(\mathbf{G F} B') > 0$. Write $Q' \subseteq Q$ for the set of states q'' such that $\text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times M}(\mathbf{G F} (B'_{q''}, q'')) > 0$ (which is non-empty from finiteness of Q). Observe that from Proposition 4.1.17, we get that for any $q'' \in Q$, $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G F} B'_{q''}) > 0$. Still applying Proposition 4.1.17 and from (6.2), it holds that

$$\text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times M}(\mathbf{G F} (B_{q'}, q') \mid \mathbf{G F} (B'_{q''}, q'')) = \text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G F} B_{q'} \mid \mathbf{G F} B'_{q''}) = 1$$

where the last equality comes from the fairness assumption on \mathcal{T} . We conclude that

$$\text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times M}(\mathbf{G F} B \mid \mathbf{G F} B') = 1$$

by observing that

$$\text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times M}(\mathbf{G F} B \mid \mathbf{G F} B') = \text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times M} \left(\mathbf{G F} B \mid \mathbf{G F} \left(\bigcup_{q'' \in Q'} (B'_{q''}, q'') \right) \right)$$

and using Bayes formula.

We finally fix $B = \cup_{q' \in Q} (B_{q'}, q') \in \mathcal{B}_M$ and $B' \in \text{PreProb}^{\mathcal{T} \times M}(B)$ such that $\text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times M}(\mathbf{G F} B') > 0$. We can write B' as the union $B' = \cup_{q' \in Q} B^{(q')}$ such that for any $q' \in Q$, $B^{(q')} \neq \emptyset$ if and only if $B^{(q')} \in \text{PreProb}^{\mathcal{T} \times M}((B_{q'}, q'))$. Write

$Q' \subseteq Q$ for the set of states q' such that $\text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times \mathbb{M}}(\mathbf{GF} B^{(q')}) > 0$ (which is non-empty from finiteness of Q). We can then conclude by observing that for any $q' \in Q'$,

$$\text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times \mathbb{M}}(\mathbf{GF}(B_{q'}, q') \mid \mathbf{GF} B^{(q')}) = 1$$

from the the first part of the proof, and by applying the fact that: given a probability space $(\Omega, \mathcal{E}, \text{Prob})$ if for any $1 \leq k \leq n$, $\text{Prob}(A_k \mid B_k) = 1$, then $\text{Prob}(\bigcup_{k=1}^n A_k \mid \bigcup_{k=1}^n B_k) = 1$. \square

This allows thus to define a framework in which the qualitative analysis of ω -regular properties is reduced to the qualitative analysis in a finite Markov chain. We only need for the hypotheses of Proposition 5.2.7 to be met: \mathcal{T}_2 is a finite Markov chain and an α -abstraction of \mathcal{T}_1 , and \mathcal{T}_1 is fair w.r.t. α -closed sets. We thus get that $\mathcal{T}_2 \times \mathbb{M}$ is also a finite Markov chain and an $\alpha_{\mathbb{M}}$ abstraction of $\mathcal{T}_1 \times \mathbb{M}$ from Lemma 6.1.19, and $\mathcal{T}_1 \times \mathbb{M}$ is fair w.r.t. $\alpha_{\mathbb{M}}$ -closed sets from Lemma 6.1.27. It follows from Proposition 5.2.7 that $\mathcal{T}_1 \times \mathbb{M}$ is decisive w.r.t. $\alpha_{\mathbb{M}}$ -closed sets and therefore, $\mathcal{T}_2 \times \mathbb{M}$ is a sound α -abstraction of $\mathcal{T}_1 \times \mathbb{M}$ from Proposition 5.3.4. Conditions of Theorem 6.1.25 are thus met since \mathcal{T}_2 has trivially a finite attractor, allowing a nice qualitative analysis for properties given by a DMA.

Remark 6.1.29 (Discussion on the approach of [BBB⁺14]). While the notion of abstraction was not precisely defined in [BBB⁺14] for STA, it was implicitly already there. Also, decidability of the almost-sure satisfaction was ensured thanks to a fairness condition. Using the terminology of the current paper, the framework was the following: \mathcal{T}_1 and \mathcal{T}_2 are two STSs such that \mathcal{T}_2 is a *finite* Markov chain which is an α -abstraction of \mathcal{T}_1 . Then the condition for the abstraction to yield interesting results was that \mathcal{T}_1 should be fair w.r.t. every α -closed sets. Applying the reasoning mentioned above in this section, we immediately get that the conditions of Theorem 6.1.25 are satisfied, and the approach of [BBB⁺14] was then a particular case of that theorem, when applied to specific subclasses of STA (more details are provided in Chapter 7).

6.2. Quantitative analysis

In this section, we are interested in the quantitative model-checking problem for STSs (see Definition 4.1.20). Beyond the qualitative analysis performed in the previous section, we will see that, under reasonable assumptions, one may derive approximation schemes to compute, within arbitrary precision, the probability of a given property.

In Sections 6.2.1 and 6.2.2, we consider first reachability, then repeated reachability properties. We use the algorithms of [ABM07] in order to get approximation schemes in our more general context. We get similar results for the correctness and the termination of the procedures.

In Sections 6.2.3 and 6.2.4, we study properties given by a DMA first in DMCs with a finite attractor, and then in general STSs who have a DMC with a finite attractor as α -abstraction. We show how the procedures of Sections 6.1.2 and 6.1.3 allow us to get approximation schemes for ω -regular properties, that amount to approach finitely many reachability probabilities, applying thus the scheme of Section 6.2.1.

For the next two subsections, fix an STS $\mathcal{T} = (S, \Sigma, \kappa)$, and an initial distribution $\mu \in \text{Dist}(S)$.

6.2.1 Quantitative reachability analysis

In order to approximate the reachability probability of a set $B \in \Sigma$ in \mathcal{T} , we define the two following sequences, similar to the ones given for decisive Markov chains [ABM07] (see Section 2.2.2 for some details). For every $n \in \mathbb{N}$:

$$\begin{cases} p_n^{\text{Yes}} &= \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F}_{\leq n} B); \\ p_n^{\text{No}} &= \text{Prob}_\mu^{\mathcal{T}}(\neg B \mathbf{U}_{\leq n} \tilde{B}). \end{cases}$$

Since the sequences of events $\text{Ev}_{\mathcal{T}}(\mathbf{F}_{\leq n} B)_{n \in \mathbb{N}}$ and $\text{Ev}_{\mathcal{T}}(\neg B \mathbf{U}_{\leq n} \tilde{B})_{n \in \mathbb{N}}$ are non-decreasing and converge respectively towards $\text{Ev}_{\mathcal{T}}(\mathbf{F} B)$ and $\text{Ev}_{\mathcal{T}}(\neg B \mathbf{U} \tilde{B})$, it is easy to determine the limit of the sequences $(p_n^{\text{Yes}})_n$ and $(p_n^{\text{No}})_n$, with no assumption on the model.

Lemma 6.2.1. The sequences $(p_n^{\text{Yes}})_n$ and $(p_n^{\text{No}})_n$ are non-decreasing and converge respectively to $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B)$ and $\text{Prob}_\mu^{\mathcal{T}}(\neg B \mathbf{U} \tilde{B})$.

Assuming now that \mathcal{T} is decisive w.r.t. B , the two limits are related.

Corollary 6.2.2. If \mathcal{T} is $\text{Dec}(\mu, B)$, then $\lim_{n \rightarrow +\infty} p_n^{\text{Yes}} + p_n^{\text{No}} = 1$.

Proof. From Lemma 6.2.1, it holds that

$$\begin{aligned} \lim_{n \rightarrow +\infty} p_n^{\text{Yes}} + p_n^{\text{No}} &= \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B) + \text{Prob}_\mu^{\mathcal{T}}(\neg B \mathbf{U} \tilde{B}) \\ &= \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B \vee (\neg B \mathbf{U} \tilde{B})) \quad \text{from point (i) of Lemma 6.1.1} \\ &= \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B \vee \mathbf{F} \tilde{B}) \quad \text{from Lemma 4.2.2 (fifth item)} \\ &= 1 . \end{aligned}$$

The last equality comes from the decisiveness assumption. □

Observe that if \mathcal{T} is a DMC, we recover the algorithm of [ABM07] described here in Section 2.2.2 and Proposition 2.2.20.

Corollary 6.2.2 can be used in order to derive an approximation scheme for $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} B)$. To obtain an ε -approximation, it suffices to evaluate p_n^{Yes} and p_n^{No} for larger and larger values of n , until $1 - p_n^{\text{Yes}} - p_n^{\text{No}} < \varepsilon$, and to return p_n^{Yes} . This scheme is effective as soon as one can compute \widetilde{B} , and the probability (from μ) of cylinders of the forms $\text{Cyl}(\underbrace{S, \dots, S}_{n \text{ times}}, B)$ and $\text{Cyl}(\underbrace{\neg B, \dots, \neg B}_{n \text{ times}}, \widetilde{B})$. In case p_n^{Yes} and p_n^{No} cannot be computed exactly, but can only be approximated up to any desired error bound, this scheme can be refined to obtain a 2ε -approximation for $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} B)$.

6.2.2 Quantitative repeated reachability analysis

We now define two sequences that will yield an approximation scheme for a repeated reachability probability, under stronger assumptions on the model. For every $n \in \mathbb{N}$:

$$\begin{cases} q_n^{\text{Yes}} &= \text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\leq n} \widetilde{\widetilde{B}}); \\ q_n^{\text{No}} &= \text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\leq n} \widetilde{\widetilde{B}}). \end{cases}$$

Here again, with no assumption on \mathcal{T} :

Lemma 6.2.3. The sequences $(q_n^{\text{Yes}})_n$ and $(q_n^{\text{No}})_n$ are non-decreasing and converge respectively to $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{\widetilde{B}})$ and $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{\widetilde{B}})$.

Assuming now that \mathcal{T} is persistently decisive w.r.t. B and decisive w.r.t. \widetilde{B} , the two sequences are closely related.

Corollary 6.2.4. If \mathcal{T} is $\text{PersDec}(\mu, B)$ and $\text{Dec}(\mu, \widetilde{B})$, then the two sequences $(q_n^{\text{Yes}})_n$ and $(1 - q_n^{\text{No}})_n$ converge towards $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B)$ and $\lim_{n \rightarrow +\infty} q_n^{\text{Yes}} + q_n^{\text{No}} = 1$.

Proof. Since \mathcal{T} is $\text{Dec}(\mu, \widetilde{B})$, it holds that $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{B} \vee \mathbf{F} \widetilde{\widetilde{B}}) = 1$. Since \mathcal{T} is $\text{PersDec}(\mu, B)$, one derives from Lemma 6.1.2 that

$$\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{\widetilde{B}}) = 1 - \text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{B}).$$

We can now show that

$$1 - \text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{B}) = \text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B).$$

It comes from the fact that $\text{PersDec}(\mu, B)$ is equivalent to $\text{StrDec}(\mu, B)$ and from point (ii) of Lemma 6.1.1. This proves the first part of the corollary, thanks to Lemma 6.2.3.

Finally, we can directly establish from Lemmas 6.2.3 and 6.1.2 and from the hypothesis $\text{Dec}(\mu, \widetilde{B})$, that $\lim_{n \rightarrow +\infty} q_n^{\text{Yes}} + q_n^{\text{No}} = 1$. \square

Observe again that if \mathcal{T} is a DMC, this procedure corresponds to the algorithm of [ABM07] described here in Section 2.2.2 and Proposition 2.2.21.

Here also, under the assumptions of Corollary 6.2.4, we obtain an approximation scheme for the value $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B)$. Effectiveness of the scheme relies on the computability of the avoid sets \widetilde{B} and $\widetilde{\widetilde{B}}$, and on the effective computation of the probability of cylinders of the forms $\text{Cyl}(\underbrace{S, \dots, S}_{n \text{ times}}, \widetilde{\widetilde{B}})$ and $\text{Cyl}(\underbrace{S, \dots, S}_{n \text{ times}}, \widetilde{B})$.

Similarly as before, in case q_n^{Yes} and q_n^{No} cannot be computed exactly, but can only be approximated up to any desired error bound, this scheme can be refined to obtain a 2ε -approximation for $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B)$.

6.2.3 Properties given by a DMA in DMCs

To go beyond reachability and repeated reachability, we now consider an ω -regular property given by a DMA $\mathbf{M} = (Q, q_0, E, \mathcal{F})$. We assume that $\mathcal{T} = (S, \Sigma, \kappa, \text{AP}, \mathcal{L})$ is a labelled DMC.

In order to approximate the probability that the model satisfies this external specification, we assume that \mathcal{T} has a finite attractor. Following Section 6.1.2, we consider the finite attractor B of $\mathcal{T} \times \mathbf{M}$, and we apply Theorem 6.1.17: for each $\mu \in \text{Dist}(S)$,

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathbf{M}}(\text{Inf} \in \mathcal{F}) = \sum_{C \in \text{Good}_{\mathcal{T} \times \mathbf{M}}^B(\mathcal{F})} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathbf{M}}(\mathbf{F} C) .$$

Thus, the computation of the probability that a given model satisfies a given external specification is reduced to the computation of finitely many reachability probabilities. Now, given that \mathcal{T} and hence $\mathcal{T} \times \mathbf{M}$ has a finite attractor, $\mathcal{T} \times \mathbf{M}$ is $\text{Dec}(\mu \times \delta_{q_0}, B)$ for any measurable set B (see Proposition 2.2.13), so that we can apply the approximation scheme of Section 6.2.1 to obtain an approximation of the desired value (see Corollary 6.2.2).

The effectiveness of the approach relies on the effectiveness of the scheme for reachability, but also on the computability of an attractor for \mathcal{T} , and of the set of good BSCCs of the graph of the attractor.

6.2.4 Properties given by a DMA in general STSs via denumerable abstractions

We assume to be in the same framework as in Section 6.1.3, that is $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1, \text{AP}, \mathcal{L}_1)$ and $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2, \text{AP}, \mathcal{L}_2)$ are two LSTSs such that:

- \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1
- \mathcal{T}_2 is a DMC with a finite attractor B_2 .

We consider again a DMA $\mathbf{M} = (Q, q_0, E, \mathcal{F})$, as well as the products $\mathcal{T}_1 \times \mathbf{M}$ and $\mathcal{T}_2 \times \mathbf{M}$. Writing $\mathcal{B} = \{\alpha_{\mathbf{M}}^{-1}(B) \mid B \in \Sigma'_2\}$, we assume that $\mathcal{T}_1 \times \mathbf{M}$ is $\text{Dec}(\mathcal{B})$. Remember that this implies, from Lemma 6.1.19, that $\mathcal{T}_2 \times \mathbf{M}$ is a sound $\alpha_{\mathbf{M}}$ -abstraction of $\mathcal{T}_1 \times \mathbf{M}$.

Fix an initial distribution μ for \mathcal{T}_1 . Thanks to Theorem 6.1.25, we have that:

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathbf{M}}(\text{Inf} \in \mathcal{F}) = \sum_{C \in \text{Good}_{\mathcal{T}_2 \times \mathbf{M}}^{B_2}(\mathcal{F})} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathbf{M}}(\mathbf{F} \alpha^{-1}(C)) .$$

Thus, as previously, the computation of the probability that a given model satisfies a given external specification is reduced to the computation of finitely many reachability probabilities. Since we assumed $\mathcal{T}_1 \times \mathbf{M}$ to be $\text{Dec}(\mathcal{B})$, we can use the approximation scheme from Section 6.2.1 to approximate the searched value (see Corollary 6.2.2).

Recall that in Section 6.1.3, we discussed the particular case when \mathcal{T}_2 is a finite Markov chain. In that case, from Lemma 6.1.27 and from Proposition 5.2.7, it suffices for \mathcal{T}_1 be fair w.r.t. α -closed sets in order to get correctness and termination of the approximation scheme (see Corollary 6.2.2).

Effectiveness of the approach requires effective numerical computations for the distributions, as well as the ability of constructing various sets, like the BSCCs of the graph of the attractor, and avoid-sets of these, etc.

6.3. Summary of the results on STSs

In this section, we provide a summary of the different results presented in the previous sections and chapters. We do not enter in the details! We give the main idea of the principle results and point towards the section or the precise lemma, proposition or theorem.

Qualitative analysis. We begin with the qualitative analysis (Section 6.1). We first fix some arbitrary (L)STS \mathcal{T} . Given a property φ , the qualitative problem amounts to decide whether $\text{Prob}_\mu^{\mathcal{T}}(\varphi) = 0$ or not or $\text{Prob}_\mu^{\mathcal{T}}(\varphi) = 1$ or not. We have the following results.

- If \mathcal{T} satisfies decisiveness properties and if φ is a (repeated) reachability property, then the qualitative problem is reduced to the simple 0-model-checking problem of some “Until” or reachability property (Section 6.1.1).
- If \mathcal{T} is a DMC with a finite attractor and φ is given by a DMA M , then the almost-sure model-checking problem is reduced to a simple 0-model-checking problem of finitely many reachability properties in $\mathcal{T} \times M$ and to check some structural properties on the underlying graph of the attractor (see developments in Section 6.1.2 and more precisely Corollary 6.1.18).

We now fix a (L)STS \mathcal{T}_1 and \mathcal{T}_2 , an α -abstraction of \mathcal{T}_1 . The question is the same as before, except that now we consider \mathcal{T}_1 instead of \mathcal{T} . We have the following results.

- If \mathcal{T}_2 is sound and satisfies decisiveness properties and if φ is a (repeated) reachability property, the qualitative problem for \mathcal{T}_1 is reduced to the same qualitative problem in \mathcal{T}_2 (Section 6.1.1, Proposition 6.1.9).
- If \mathcal{T}_2 is a DMC with a finite attractor, if φ is given by a DMA M and if $\mathcal{T}_2 \times M$ is a sound α_M -abstraction of $\mathcal{T}_1 \times M$, then the almost-sure model-checking problem in \mathcal{T}_1 is reduced to the same problem in \mathcal{T}_2 (Section 6.1.3 and more precisely Corollary 6.1.26).

Quantitative analysis. We are now concerned with the quantitative analysis (Section 6.2). We fix a (L)STS \mathcal{T} . Given a property φ , we are interested in a scheme that approaches $\text{Prob}_\mu^{\mathcal{T}}(\varphi)$.

- If \mathcal{T} satisfies decisiveness properties and if φ is a (repeated) reachability property, then there is an approximation scheme that is correct and that terminates (Sections 6.2.1 and 6.2.2).
- If \mathcal{T} is a DMC with a finite attractor and if φ is given by a DMA M , then there is an approximation scheme for $\mathcal{T} \times M$ that requires to approach finitely many reachability probabilities (using scheme of Section 6.2.1) and to check some structural properties on the underlying graph of the attractor of $\mathcal{T} \times M$ (Section 6.2.3).

We now fix a (L)STS \mathcal{T}_1 and \mathcal{T}_2 an α -abstraction of \mathcal{T}_1 . We are again interested in an approximation scheme, but this time for \mathcal{T}_1 .

- If \mathcal{T}_2 is a DMC with a finite attractor, if φ is given by a DMA M and if $\mathcal{T}_2 \times M$ is a sound α_M -abstraction of $\mathcal{T}_1 \times M$, then there is an approximation scheme for $\mathcal{T}_1 \times M$ that requires to approach finitely many reachability probabilities (using scheme of Section 6.2.1) and to check some structural properties on the underlying graph of the attractor of $\mathcal{T}_2 \times M$ (Section 6.2.4).

Properties of STSs. We are here interested in the basic properties of STSs described in Section 4.2. We fix a STS $\mathcal{T} = (S, \Sigma, \kappa)$ and $\mathcal{B} \subseteq \Sigma$.

- It holds that \mathcal{T} is $\text{Dec}(\mathcal{B}) \iff \mathcal{T}$ is $\text{StrDec}(\mathcal{B}) \iff \mathcal{T}$ is $\text{PersDec}(\mathcal{B}) \implies \mathcal{T}$ is $\text{fair}(\mathcal{B})$ (Proposition 4.2.17).
- If \mathcal{T} is a DMC with a finite attractor, then \mathcal{T} is $\text{Dec}(2^S)$ ([ABM07], see Proposition 2.2.13 here).

Transfer of properties through the product with a DMA. We now look at the different properties that are preserved through the product with a DMA. We fix a LSTS $\mathcal{T} = (S, \Sigma, \kappa)$, a DMA M and $\mathcal{B} \subseteq \Sigma$

- If \mathcal{T} has a finite attractor, then $\mathcal{T} \times M$ has also a finite attractor (see Lemma 6.1.10). Moreover if \mathcal{T} is a DMC, then so is $\mathcal{T} \times M$.
- If \mathcal{T} is $\text{fair}(\mathcal{B})$, then $\mathcal{T} \times M$ is $\text{fair}(\mathcal{B}_M)$ (Lemma 6.1.27).

Still **open questions** are whether decisiveness properties or soundness are preserved through the product with a DMA (Remark 6.1.20). More precisely for soundness, we do not know whether \mathcal{T}_2 sound α -abstraction of \mathcal{T}_1 implies $\mathcal{T}_2 \times M$ sound α_M -abstraction of $\mathcal{T}_1 \times M$.

Transfer of properties through abstractions. We summarize now which properties are transferred through α -abstractions (Section 5.2). We first consider two STSs \mathcal{T}_1 and \mathcal{T}_2 such that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . We have the following results (Section 5.2.2).

- If \mathcal{T}_2 is a DMC with a finite attractor and if \mathcal{T}_1 satisfies hypotheses (†) of page 114, then \mathcal{T}_1 is decisive w.r.t. α -closed sets (Proposition 5.2.6).
- If \mathcal{T}_2 is a finite Markov chain and if \mathcal{T}_1 is fair w.r.t. α -closed sets, then \mathcal{T}_1 is decisive w.r.t. α -closed sets (Proposition 5.2.7).

If we assume moreover that the α -abstraction \mathcal{T}_2 is sound, we get the results of Section 5.2.1.

- If \mathcal{T}_2 is decisive w.r.t. B , then \mathcal{T}_1 is decisive w.r.t. $\alpha^{-1}(B)$ (Corollary 5.2.3).
- If A_2 is an attractor of \mathcal{T}_2 , then $\alpha^{-1}(A_2)$ is an attractor of \mathcal{T}_1 (Proposition 5.2.4).

Conditions for soundness and completeness. We end this summary with the two main results that give conditions for soundness and completeness of α -abstractions (Section 5.3). We fix two STSs \mathcal{T}_1 and \mathcal{T}_2 such that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . We have the following results.

- If \mathcal{T}_2 is decisive, then \mathcal{T}_2 is complete (Lemma 5.3.2).
- If \mathcal{T}_1 is decisive w.r.t. α -closed sets, then \mathcal{T}_2 is sound (Proposition 5.3.4).

Application to STA

In this chapter, we show how the results of Chapter 6 can be applied to the STA model (see Chapter 3). This was again done in [BBBC17], a first version can be found [BBBC16] however in this paper, we did not have the results on ω -regular properties (see Sections 6.1.2, 6.1.3, 6.2.3 and 6.2.4).

As already stated in Chapter 3, several decidability results have been proven for subclasses of STA, requiring the development of ad-hoc methods [BBB⁺07, BBB⁺08, BBBM08, BBJM12]. In [BBB⁺14], the authors proposed the first unifying method capturing all known decidability results for the qualitative model-checking problem: the thick region graph (see Definition 3.2.8) is a finite graph based on the standard region automaton construction for timed automata [AD94], which allows one, through the construction of a finite Markov chain, to infer good transfer properties from this finite graph to the original STA when some fairness property is satisfied (see Sections 3.2 and 3.3). The current work allows us both to unify all decidability and approximation results that were known, and to get new approximation schemes for the quantitative model-checking problem (of ω -regular properties).

In Section 7.1, we show how the semantics of a STA can be given as a STS. We define formally the Markov kernel and the probability distribution on the set of runs that results from this kernel. We then define the α -abstraction that will be used for STA, which is the thick region graph defined in Section 3.2.

We then consider two subclasses of STA that were proven to have nice qualitative results in [BBB⁺14] (see Section 3.3). In Section 7.2, we consider the class of reactive STA while in Section 7.3, we consider the class of one-clock STA. In both classes, we present a new way of proving qualitative results of [BBB⁺14] through the results of Section 6.1, which is more unifying than the proofs of [BBB⁺14];

and we show that the approximation schemes of Section 6.2 can be applied, leading to new quantitative results for STA!

It should be noted that in [BBBC17], we also applied the previous results to a subclass of the GSMP model, based on technical lemmas of [BKKR11]. We do not describe it here as it is not the subject of the thesis.

7.1. From STA to STS

In this section, we show how the semantics of a STA can be given as a STS. We fix for the section a STA $\mathcal{A} = (L, X, E, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X})^1$ and we assume that it satisfies the hypothesis (\ddagger) (see page 65): for each state q , μ_q is equivalent to the Lebesgue measure restricted on $I(q)$, *i.e.* $\mu_q(A) > 0$ if and only if $\Lambda(A \cap I(q)) > 0$ for each Borel set A , where Λ denotes the Lebesgue measure.

The STA \mathcal{A} can be interpreted as a LSTS $\mathcal{T}_{\mathcal{A}} = (S_{\mathcal{A}}, \Sigma_{\mathcal{A}}, \kappa_{\mathcal{A}}, \text{AP}, \mathcal{L}_{\mathcal{A}})$ where $S_{\mathcal{A}} = Q = L \times \mathbb{R}_+^X$ is the set of states of \mathcal{A} , $\Sigma_{\mathcal{A}}$ is the σ -algebra product $2^L \times \mathcal{B}(\mathbb{R}_+^{|X|})^2$ where $\mathcal{B}(\mathbb{R}_+^{|X|})$ is the Borel σ -algebra on $\mathbb{R}_+^{|X|}$, the labelling function is given as $\mathcal{L}_{\mathcal{A}}((l, \nu)) = \mathcal{L}((l, \nu)) = \mathcal{L}(l)$ for any state (l, ν) , and the kernel $\kappa_{\mathcal{A}}$ is defined as follows: for each $q = (l, \nu) \in S_{\mathcal{A}}$ and each $B \in \Sigma_{\mathcal{A}}$,

$$\kappa_{\mathcal{A}}(q, B) = \sum_{e=(l,g,Y,l') \in E} \int_{t \in \mathbb{R}_+} p_{q+t}(e) \cdot \mathbb{1}_B(l', [Y \leftarrow 0](\nu + t)) \, d\mu_q(t)$$

where $\mathbb{1}_B$ is the characteristic function of B . It gives the probability to hit set $B \subseteq S_{\mathcal{A}}$ from state q in one step (composed of a delay transition followed by a discrete transition).

Remark 7.1.1. The probability measure on runs derived from $\mathcal{T}_{\mathcal{A}}$ in Section 4.1 coincides with the original definition of [BBB⁺14] (see Section 3.1), under a slight hypothesis on the STA: we need that there are no edges with the same source and the same target but with non-disjoint guards, *i.e.* we need in some sense, determinism. Observe that the assumption is not too restrictive: if there are two edges $e_1 = (l, g_1, Y_1, l')$ and $e_2 = (l, g_2, Y_2, l')$ with $g_1 \wedge g_2 \neq \mathbf{false}$, then one can separate those edges into several new edges with disjoint guards, notably edge $(l, g_1 \wedge g_2, Y_1 \cup Y_2, l')$.

We now show the correspondence between $\text{Prob}_q^{\mathcal{A}}$ and $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}$. First it holds

¹We omit here the invariant to simplify notations, however it has no cost to add it; the labelling on edges is not needed.

²Recall that $2^L \times \mathcal{B}(\mathbb{R}_+^{|X|}) = \{\bigcup_{l \in L} \{l\} \times C_l \mid \forall l \in L, C_l \in \mathcal{B}(\mathbb{R}_+^{|X|})\}$.

that for each state $q \in S_{\mathcal{A}}$ and each $B = \cup_{l \in L} \{l\} \times B_l$,

$$\kappa_{\mathcal{A}}(q, B) = \text{Prob}_q^{\mathcal{A}} \left(\bigcup_{e \in E} \text{Cyl}(\pi_{\mathcal{C}_{\nu, e}}(q, e)) \right)$$

where for each $e = (l, g, Y, l') \in E$, $\mathcal{C}_{\nu, e} = \{t \in \mathbb{R}_+ \mid [Y \leftarrow 0](\nu + t) \in B_{l'}\}$ is the adequate constraint ensuring that we end up in B . This due to the fact that

$$\text{Cyl}(\{q\}, B) = \bigcup_{e \in E} \text{Cyl}(\pi_{\mathcal{C}_{\nu, e}}(q, e))$$

and from the definitions of $\kappa_{\mathcal{A}}$ and $\text{Prob}_q^{\mathcal{A}}$. We can decompose similarly each cylinder of $\mathcal{F}_{\mathcal{T}_{\mathcal{A}}}$ into a disjoint union of constrained cylinders of $\Omega_q^{\mathcal{A}}$. This holds also without the new assumption described above.

Secondly, we can also establish a similar link from a cylinder of $\Omega_q^{\mathcal{A}}$ to a union of cylinders in $\mathcal{F}_{\mathcal{T}_{\mathcal{A}}}$. But this time the above condition is needed. If this condition did not hold, it could be the case to have two edges with the same source and same target, but with non-disjoint guards. Then in $\mathcal{T}_{\mathcal{A}}$, we would not be able to distinguish the two edges when at the intersection of the two guards (because cylinders in $\mathcal{F}_{\mathcal{T}_{\mathcal{A}}}$ are defined through the set of states that are visited, not through the edges that are taken like in $\Omega_q^{\mathcal{A}}$). With the assumption that those edges are prohibited, we get the same probability spaces on the runs in both semantics.

We fix for the rest of the chapter a STA $\mathcal{A} = (L, X, E, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X})$ and its corresponding LSTS $\mathcal{T}_{\mathcal{A}} = (S_{\mathcal{A}}, \Sigma_{\mathcal{A}}, \kappa_{\mathcal{A}}, \text{AP}, \mathcal{L}_{\mathcal{A}})$. The objective of this chapter, is to apply the results of Chapter 6 to STA. As explained in Chapters 5 and 6, abstractions can help to apply those results, if they satisfy some nice properties. We show next that STA have an α -abstraction that is a finite Markov chain.

7.1.1 The thick region graph abstraction

We refer to Section 2.1.1 for a reminder on regions in timed automata. We recall that $R_{\mathcal{A}}$ denotes the set of regions of STA \mathcal{A} , and that given $\nu \in \mathbb{R}_+^X$, we write $[\nu]_{\mathcal{A}}$ for the region containing ν . We also send the reader to Section 3.2 for the thick region graph $\mathcal{R}_{\mathcal{A}}^{\text{tk}}$ (see Definition 3.2.8) and the corresponding finite Markov chain $\text{MC}(\mathcal{A}) = (S, P)$; we recall that it is defined as follows:

- the set of states $S = L \times R_{\mathcal{A}}$,
- $P((l, r), (l', r')) > 0$ if and only if there is an edge $(l, r) \rightarrow (l', r')$ in $\mathcal{R}_{\mathcal{A}}^{\text{tk}}$, and

- for each state $(l, r) \in L \times R_{\mathcal{A}}$, $P((l, r), \cdot)$ is the discrete uniform distribution over $\{(l', r') \in L \times R_{\mathcal{A}} \mid P((l, r), (l', r')) > 0\}$.

Observe that from Remark 3.2.7 and Definition 3.2.8 of the thick region graph, the second item is equivalent to: $P((l, r), (l', r')) > 0$ if and only if there exists some $\nu \in r$ such that $\kappa_{\mathcal{A}}((l, \nu), \{l'\} \times r') > 0$. Recall also that since $\text{MC}(\mathcal{A})$ is a finite Markov chain, *e.g.* it is a STS. To keep some consistency in the notations, we will write in the sequel $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ instead of $\text{MC}(\mathcal{A})$.

We finally define the function $\alpha : L \times \mathbb{R}_+^X \rightarrow L \times R_{\mathcal{A}}$ by $\alpha((l, \nu)) = (l, [\nu]_{\mathcal{A}})$ for each $(l, \nu) \in L \times \mathbb{R}_+^X$. It holds that $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ is an α -abstraction of $\mathcal{T}_{\mathcal{A}}$.

Lemma 7.1.2. It holds true that $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ is a finite α -abstraction of $\mathcal{T}_{\mathcal{A}}$.

The proof is immediate by construction of $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$. Let us notice that finiteness of the abstraction implies completeness (Lemma 5.3.1).

As witnessed in [BBB⁺14, Appendix D.2], see here Chapter 3, this abstraction may not give much information in general about the probability of linear-time properties in the original STA: in particular the abstraction is not always sound, this has been illustrated in Example 3.2.10 with STA \mathcal{A}_{cvg} of Figure 2.6. We come briefly back to this STA in Example 7.1.3. However it has been shown in [BBB⁺14] (see Section 3.3 here) that in several cases, it helps to obtain decidability results (see Theorem 3.3.2): fairness was identified as the key notion to get those results.

Example 7.1.3. We come back to STA \mathcal{A}_{cvg} of Examples 2.1.28 and 3.2.10 depicted on Figure 2.6. Lemma 7.1.2 ensures that the corresponding finite Markov chain $\mathcal{T}_{\mathcal{A}_{\text{cvg}}}^{\text{tg}}$ is an α -abstraction of $\mathcal{T}_{\mathcal{A}_{\text{cvg}}}$. However, Example 3.2.10 witnesses that $\mathcal{T}_{\mathcal{A}_{\text{cvg}}}^{\text{tg}}$ is not sound and Example 3.3.3 shows that $\mathcal{T}_{\mathcal{A}_{\text{cvg}}}$ is not almost-surely fair.

Remark 7.1.4. As said before, a condition for $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ to be a useful abstraction was identified as *fairness* in [BBB⁺14]. It should be observed that fairness in the sense of Definition 3.3.1 in STA corresponds to fairness w.r.t. α -closed sets of Definition 4.2.9 for $\mathcal{T}_{\mathcal{A}}$, the corresponding STS. Recall that fairness is the weakest notion defined here (see Proposition 4.2.17). Later, in [BBBC16], we identified the condition as (strong) decisiveness, but this was not sufficient for decidability and approximability results of ω -regular properties as seen in Chapter 6, Sections 6.1.3 and 6.2.3. In [BBBC17], we realised that we have a finite-attractor property (through an abstraction). Observe that in the case of STA, since the abstraction is finite, we always get a finite attractor (through the abstraction). From there,

- either we can get a finite attractor (through an abstraction) which satisfies conditions (†) of Proposition 5.2.6,
- or we have a finite abstraction (which is the case here) and we can show that the STA is fair w.r.t. α -closed sets (which is thus obviously again the case here for almost-surely fair STA), meeting thus the conditions of Proposition 5.2.7,

allowing us to infer, in both cases, the whole class of decidability and approximability results through Proposition 5.3.4. As already mentioned in Remark 6.1.29, observe thus that in the case of almost-surely fair STA, we immediately get all those results through Propositions 5.2.7 and 5.3.4, but also through Lemmas 6.1.19 and 6.1.27 which state that through the product with a DMA, finite abstractions and fairness are preserved.

However fairness is in general (*i.e.* when the abstraction is not finite) not sufficient to get all the qualitative and quantitative results. We therefore prove stronger properties for subclasses of STA: in the sequel, we will consider both classes of STA introduced in Section 3.3 and which were proved to be almost-surely fair in [BBB⁺14], and we will show that for both classes, the conditions (†) of Proposition 5.2.6 are always fulfilled. We believe that it moreover simplifies the proofs of [BBB⁺14].

7.2. Reactive STA

In this section, we recall the notion of reactive STA (see Section 3.3) and show that the results of Chapter 6 can be applied.

Following [BBJM12], the STA \mathcal{A} is said *reactive* whenever it satisfies conditions (H1) and (H2) of Section 3.3, we recall them here:

- (H1) for every state q , $I(q) = \mathbb{R}_+$ and for every $l \in L$, there exists a distribution μ_l equivalent to the Lebesgue measure on \mathbb{R}_+ , such that for every $\nu \in \mathbb{R}_+^X$, $\mu_{(l,\nu)} = \mu_l$;
- (H2) for every edge e there exists $w_e \in \mathbb{N}_0$ such that for every state q ,

$$p_q(e) = \begin{cases} \frac{w_e}{\sum_{e' \text{ enabled in } q} w_{e'}} & \text{if } e \text{ is enabled in } q, \\ 0 & \text{otherwise.} \end{cases}$$

We take the notations used in Sections 2.1.1, 3.2 but also 7.1.1 for the thick region graph. We recall that $M_{\mathcal{A}}$ denotes the maximal constant appearing in a

guard of \mathcal{A} and that a region $r \in R_{\mathcal{A}}$ is said *memoryless* whenever for each clock $x \in X$, either $\nu(x) = 0$ for each $\nu \in r$, or $\nu(x) > M_{\mathcal{A}}$ for each $\nu \in r$. We write $\mathcal{R}_{\mathcal{A}}^{\text{mem}}$ for the set of memoryless regions.

From [BBB⁺14, Lemma 13], which states that memoryless regions are visited infinitely often almost-surely from every state $q \in S_{\mathcal{A}}$, we thus get a finite attractor through the abstraction.

Proposition 7.2.1. *The set $\alpha^{-1}(L \times \mathcal{R}_{\mathcal{A}}^{\text{mem}})$ is an attractor for $\mathcal{T}_{\mathcal{A}}$.*

We give an idea of the main argument. It can be shown that, in one step, one can ensure reaching a memoryless region by delaying at least $M_{\mathcal{A}} + 1$ time units; since there is one single distribution which is applied at every state of a given location, the probability to do so is uniformly bounded from below from every state (the bound is given by $\min_{l \in L}([M_{\mathcal{A}} + 1, \infty]) > 0$ from hypothesis (H1)). And you can conclude by similar arguments as in the proof of Proposition 5.2.6.

Using Propositions 5.2.6 and 5.3.4, we also get the following result.

Proposition 7.2.2. *It holds that $\mathcal{T}_{\mathcal{A}}$ is decisive w.r.t. α -closed sets and that $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ is a sound α -abstraction of $\mathcal{T}_{\mathcal{A}}$.*

Proof. It suffices to prove that the hypotheses (†) of Proposition 5.2.6 are met. It can easily be shown that $L \times \mathcal{R}_{\mathcal{A}}^{\text{mem}}$ is a finite attractor of $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$. Thanks to Proposition 7.2.1, $\alpha^{-1}(L \times \mathcal{R}_{\mathcal{A}}^{\text{mem}})$ is an attractor for $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$. It remains to show the last condition of the hypotheses (†) of Proposition 5.2.6. We therefore need to prove that for each $(l_{\text{m}}, r_{\text{m}}) \in L \times \mathcal{R}_{\mathcal{A}}^{\text{mem}}$, there are $p > 0$ and $k \in \mathbb{N}$ such that for each region $(l, r) \in L \times R_{\mathcal{A}}$:

- for each $\mu \in \text{Dist}(\alpha^{-1}(l_{\text{m}}, r_{\text{m}}))$, $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F}_{\leq k} \alpha^{-1}(l, r)) \geq p$, or
- for each $\mu \in \text{Dist}(\alpha^{-1}(l_{\text{m}}, r_{\text{m}}))$, $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F} \alpha^{-1}(l, r)) = 0$.

This is a consequence of [BBB⁺14, Lemma F.4] which says that from a memoryless region, the future (and its probability) is independent of the precise current state. This in particular implies that for two states $q, q' \in \alpha^{-1}(l_{\text{m}}, r_{\text{m}})$, for every α -closed set B , for every integer k , $\text{Prob}_q^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F}_{=k} B) = \text{Prob}_{q'}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F}_{=k} B)$. By extension, for every $\mu \in \text{Dist}(\alpha^{-1}(l_{\text{m}}, r_{\text{m}}))$, $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F}_{=k} B) = \text{Prob}_q^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F}_{=k} B)$. This implies the expected bounds, by taking $B = \alpha^{-1}(l, r)$.

Thus Proposition 5.2.6 implies that $\mathcal{T}_{\mathcal{A}}$ is decisive w.r.t. α -closed sets and therefore, Proposition 5.3.4 implies that $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ is a sound α -abstraction of $\mathcal{T}_{\mathcal{A}}$. \square

In order to deal with the qualitative and quantitative model-checking problems for reactive STA of ω -regular properties (see Definitions 4.1.19 and 4.1.20),

we will consider the product of the reactive STA \mathcal{A} with a DMA M . We refer to Lemma 6.1.19 for notation α_M and the fact that $\mathcal{T}_{\mathcal{A}}^{\text{tg}} \times M$ is an α_M -abstraction of $\mathcal{T}_{\mathcal{A}} \times M$. As consequences of Chapter 6, we get the following decidability and approximability results for reactive STA.

Corollary 7.2.3. *Let \mathcal{A} be a reactive STA, and M a DMA. Then:*

1. *we can decide whether \mathcal{A} satisfies almost-surely M ;*
2. *for every initial distribution μ which is numerically amenable w.r.t. \mathcal{A} ³, we can compute arbitrarily close approximations of $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(M)$.*

Proof. This is an application of Theorem 6.1.25, Corollary 6.1.26 and of Sections 6.2.1 and 6.2.4. It should be noted that all the hypotheses are met:

- $\mathcal{T}_{\mathcal{A}}^{\text{tg}} \times M$ has a finite attractor: since $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ is a finite MC then so is $\mathcal{T}_{\mathcal{A}}^{\text{tg}} \times M$ and we get a trivial finite attractor;
- $\mathcal{T}_{\mathcal{A}} \times M$ is decisive w.r.t. any α_M -closed sets.

This second point is a little more tricky. First, one should realise that since $\mathcal{T}_{\mathcal{A}}$ is reactive, then so is $\mathcal{T}_{\mathcal{A}} \times M$ since the condition to be reactive concerns only the distributions over the delays on each location of the STA and those distributions are not modified from the product with M . It should be noted that $\mathcal{T}_{\mathcal{A}}^{\text{tg}} \times M$ corresponds to the thick region graph abstraction of $\mathcal{T}_{\mathcal{A}} \times M$ since M does not influence the behaviour of $\mathcal{T}_{\mathcal{A}}$. Then from Proposition 7.2.2, we know that $\mathcal{T}_{\mathcal{A}} \times M$ is decisive w.r.t. α_M -closed sets (and thus that $\mathcal{T}_{\mathcal{A}}^{\text{tg}} \times M$ is a sound α -abstraction of $\mathcal{T}_{\mathcal{A}} \times M$). \square

Remark 7.2.4. As already mentioned in Remark 7.1.4, we believe that the proposed approach through abstractions and finite attractors simplifies drastically the proof of decidability of the almost-sure model-checking problem of reactive STA, and in particular avoids the ad-hoc but long and technical proof of [BBB⁺14, Lemma 7.14]. Furthermore, we obtain interesting approximability results, some of them being consequences of [BBBC16], but the general case of ω -regular properties being particular to [BBBC17].

Remark 7.2.5. Corollary 7.2.3 can be extended to properties expressed as deterministic and complete Muller *timed* automata (DCMTA), which are standard

³We say that a distribution μ is numerically amenable w.r.t. \mathcal{A} if, given $k \in \mathbb{N}$, given $\varepsilon > 0$ and given a sequence of locations and regions $(l_0, r_0), (l_1, r_1), \dots, (l_k, r_k)$, one can approximate $\text{Prob}_{\mu}^{\mathcal{A}}(\text{Cyl}((l_0, r_0), (l_1, r_1), \dots, (l_k, r_k)))$ up to ε .

deterministic and complete⁴ timed automata [AD94] with a Muller accepting condition. Indeed, the product of a reactive STA with such a DCMTA is reactive. Hence, the whole theory that we have developed applies: the STS of the product has a finite sound abstraction. This allows to express rich properties with timing constraints and evaluate their likelihood in the STA.

7.3. Single-clock STA

In this section, we recall the subclass of STA defined in Section 3.3 and we apply a similar reasoning to this subclass.

We therefore assume that \mathcal{A} is now a single-clock STA. As in [BBB⁺14, Section 7.1], we assume that \mathcal{A} satisfies conditions (H2), (H3), (H4) and (H5) of Section 3.3. We recall them here:

(H2) for every edge e there exists $w_e \in \mathbb{N}_0$ such that for every state q ,

$$p_q(e) = \begin{cases} \frac{w_e}{\sum_{e' \text{ enabled in } q} w_{e'}} & \text{if } e \text{ is enabled in } q, \\ 0 & \text{otherwise.} \end{cases}$$

(H3) for all $l \in L$, for all $[a, b] \subseteq \mathbb{R}_+$, the function $\nu \rightarrow \mu_{(l, \nu)}([a, b])$ is continuous;

(H4) if $q' = q + t$ for some $t \geq 0$, and if $0 \notin I(q + t', e)$ for each $0 \leq t' \leq t$, then $\mu_q(I(q, e)) \leq \mu_{q'}(I(q', e))$;

(H5) there is $0 < \lambda_0 < 1$ such that for every state q with $I(q)$ unbounded, $\mu_q([0, \frac{1}{2}]) \leq \lambda_0$.

These requirements are technical, but they are rather natural and easily satisfiable. For instance, a timed automaton equipped with uniform (resp. exponential) distributions on bounded (resp. unbounded) intervals satisfy these conditions. If we assume exponential distributions on unbounded intervals, the very last requirement corresponds to the bounded transition rate condition in [DP03], required to have reasonable and realistic behaviours, *i.e.* the authors prove that the set of Zeno runs has a measure null.

In [BBB⁺14, Section 7.1], there is no clear attractor property. From the details of the proofs we can nevertheless define $A_{\mathcal{A}}^{\max} = \{(l, r_0) \mid l \in L\} \cup \{(l, r) \in L \times R_{\mathcal{A}} \mid \forall (l', r') \in L \times R_{\mathcal{A}}, (l, r) \rightarrow^* (l', r') \text{ in } \mathcal{T}_{\mathcal{A}}^{\text{tg}} \text{ implies } r' = r\}$ where r_0 is the

⁴In this context, complete means that from every configuration, for every subset of AP, and every $t \in \mathbb{R}_+$, there is an edge labelled by that subset which is enabled after t time units. So this is complete w.r.t time and actions.

region composed of the single null valuation and where $(l, r) \rightarrow^* (l', r')$ in $\mathcal{T}_A^{\text{tg}}$ stands for “there is a finite run in the thick region graph $\mathcal{R}_A^{\text{tk}}$ from (l, r) to (l', r') ” (see Section 3.2).

Proposition 7.3.1. *The set $\alpha^{-1}(A_{\mathcal{A}}^{\text{max}})$ is an attractor for \mathcal{T}_A .*

Proof. Let $C = \{0\} \cup \{c \mid c \text{ constant appearing in a guard of } \mathcal{A}\} \stackrel{\text{def}}{=} \{c_0 < c_1 < \dots < c_h\}$. The set of regions for \mathcal{A} can be chosen as $\{\{c_i\} \mid 0 \leq i \leq h\} \cup \{[c_{i-1}; c_i[\mid 1 \leq i \leq h\} \cup \{[c_h, +\infty[$ (see [LMS04]).

Following the proof of [BBB⁺14, Theorem 7.2], the set of infinite paths in \mathcal{A} can be divided into (a) the set of paths that take resetting edges infinitely often, and (b) the set of paths that take resetting edges only finitely often. We will condition the set of runs with both events, and we will show that in both cases, we almost-surely reach $\alpha^{-1}(A_{\mathcal{A}}^{\text{max}})$. We will then conclude by using Bayes formula.

The first case is easy. If we assume that the probability that (a) happens is positive, and we reason now on the σ -algebra which is conditioned by (a), then under condition (a), $\alpha^{-1}(\{(l, r_0) \mid l \in L\})$ is obviously reached almost-surely. Hence for each $\mu \in \text{Dist}(S_{\mathcal{A}})$ with $\text{Prob}_{\mu}^{\mathcal{T}_A}((a)) > 0$,

$$\text{Prob}_{\mu}^{\mathcal{T}_A}(\mathbf{F} \alpha^{-1}(A_{\mathcal{A}}^{\text{max}}) \mid (a)) = 1.$$

We now assume that the probability that (b) happens is positive, and we reason in the σ -algebra which is conditioned by (b). Under condition (b), almost-surely the value of the clock is non-decreasing along the path, and almost-surely a final region r is reached (since once time c_h is reached, the only region is $[c_h, +\infty[$), that is, ultimately the value of the clock along the path belongs to r forever. We now divide (b) according to the “final region” that is reached. We fix such a region r , and we condition again with regard to that final region r . We write E_r for the event (b) intersected with “the run ends up in r ”. Observe that the sets E_r with $r \in R_{\mathcal{A}}$ partition event (b). We assume that event E_r happens with non-null probability.

From the assumption on E_r and finiteness of L , there exists $l \in L$ such that (l, r) is visited infinitely often with a positive probability. Again if we write $E_{(l,r)}$ for the event E_r intersected with “ (l, r) is visited infinitely often”, for each $r \in R_{\mathcal{A}}$, the sets $E_{(l,r)}$ with $l \in L$ partition E_r . We fix $l \in L$ and we assume that the probability of event $E_{(l,r)}$ is positive.

Let (l', r') be a successor of (l, r) in $\mathcal{T}_A^{\text{tg}}$. We want to prove that necessary, $r' = r$, and thus $(l, r) \in A_{\mathcal{A}}^{\text{max}}$. We distinguish two cases: the special case where $r = [c_h, +\infty[$ and the other cases

We first assume that $r =]c_h, +\infty[$. In this case given ν and $\nu' \in r$, edge e is enabled in (l, ν) if and only if e is enabled in (l, ν') and the probability to fire edge e is the same from both states (hypothesis (H2)). Here, we get that r' is either the region $\{0\}$ or still the region r . We are thus only interested in edges e such that $(l, r) \xrightarrow{e} (l', \{0\})$. We get that e is infinitely often enabled along $E_{(l,r)}$, and it is bounded from below to fire e from any state (l, ν) with $\nu \in r$. From classical results and *e.g.* a similar argument as in the proof of Proposition 5.2.6, it follows that for each $\mu \in \text{Dist}(S_{\mathcal{A}})$ with $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(E_{(l,r)}) > 0$,

$$\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{GF}(l', \{0\}) \mid E_{(l,r)}) = 1$$

which contradicts condition (b). This implies that $(l, r) \in A_{\mathcal{A}}^{\max}$. Recall that E_r is partitioned by the sets $E_{(l,r)}$ with $l \in L$. Therefore using Bayes formula, we get that for each $\mu \in \text{Dist}(S_{\mathcal{A}})$ with $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(E_r) > 0$,

$$\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F} \alpha^{-1}(A_{\mathcal{A}}^{\max}) \mid E_r) = 1.$$

It remains to show this same equality for any other region r .

We finally consider some region r that is not $]c_h, +\infty[$. Then there exists a non-singular edge e (see Definition 3.2.6) such that $(l, r) \xrightarrow{e} (l', r')$. Towards a contradiction, assume that r' is a strict successor of r . Fix some run $\rho = q_1 \rightarrow q_2 \rightarrow \dots$ in $E_{(l,r)}$ and assume that from q_n with $n \geq 1$, we are always in r . Then if $n \leq j_1 \leq j_2$ are such that we are in location l is q_{j_1} and q_{j_2} , hypotheses (H2) and (H4) and the definition of E_r (and more precisely (b)) ensure us that the probability to fire edge e in q_{j_2} is greater than the probability to fire edge e in q_{j_1} . It follows that there is $\beta > 0$ such that for each $\nu \in r$, $\text{Prob}_{(l,\nu)}^{\mathcal{T}_{\mathcal{A}}}(\text{Cyl}((l, r), (l', r')))) > \beta$. Hence, using standard techniques like in the proof of Proposition 5.2.6, we show that for each $\mu \in \text{Dist}(S_{\mathcal{A}})$ with $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(E_{(l,r)}) > 0$,

$$\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{GF}(l', r') \mid E_{(l,r)}) = 1.$$

Since $E_{(l,r)} \subseteq E_r$ which is the set of runs that ultimately ends in region r it follows that necessarily, $r' = r$. This is the required contradiction. Observe that it could be the case that e resets the unique clock and thus that $r' = \{0\}$. This is dealt in the same way as the situation in which $r =]c_h, +\infty[$. This implies that $(l, r) \in A_{\mathcal{A}}^{\max}$. Recall that E_r is partitioned by the sets $E_{(l,r)}$ with $l \in L$. Therefore using Bayes formula, we get that for each $\mu \in \text{Dist}(S_{\mathcal{A}})$ with $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(E_r) > 0$,

$$\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F} \alpha^{-1}(A_{\mathcal{A}}^{\max}) \mid E_r) = 1$$

for each region r that is an open interval. Using Bayes formula on events E_r with $r \in R_{\mathcal{A}}$ that partition (b) , we deduce that for each $\mu \in \text{Dist}(S_{\mathcal{A}})$ with $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(b) > 0$,

$$\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F} \alpha^{-1}(A_{\mathcal{A}}^{\max}) \mid (b)) = 1$$

and using again Bayes formula on (a) and (b), we get that for each $\mu \in \text{Dist}(S_{\mathcal{A}})$,

$$\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F} \alpha^{-1}(A_{\mathcal{A}}^{\max})) = 1$$

which concludes the proof. \square

As before in Section 7.2, we get the following result from Propositions 5.2.6 and 5.3.4.

Proposition 7.3.2. *It holds that $\mathcal{T}_{\mathcal{A}}$ is decisive w.r.t. α -closed sets and that $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ is a sound α -abstraction of $\mathcal{T}_{\mathcal{A}}$.*

Proof. It suffices to prove that the hypotheses (\dagger) of Proposition 5.2.6 are met. We easily get that $A_{\mathcal{A}}^{\max}$ is a finite attractor for $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$, whereas $\alpha^{-1}(A_{\mathcal{A}}^{\max})$ is an attractor for $\mathcal{T}_{\mathcal{A}}$ (Proposition 7.3.1). As for reactive STA, it remains to show the last property appearing in the hypotheses (\dagger) of Proposition 5.2.6. The required bounds obviously exist for the region r_0 (since only a single valuation belongs to r_0). Furthermore, once we reach $\alpha^{-1}(\{(l, r) \in L \times R_{\mathcal{A}} \mid \forall (l', r') \in L \times R_{\mathcal{A}}, (l, r) \rightarrow^* (l', r') \text{ in } \mathcal{T}_{\mathcal{A}}^{\text{tg}} \text{ implies } r' = r\})$, ultimately, the runs almost surely end up in the same region r . Hence ultimately, from hypothesis (H2), the STA behaves like a finite Markov chain. The required bounds can be inferred.

This allows us to conclude from Proposition 5.2.6 that $\mathcal{T}_{\mathcal{A}}$ is decisive w.r.t. α -closed sets, and from Proposition 5.3.4 that $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ is a sound α -abstraction of $\mathcal{T}_{\mathcal{A}}$. \square

And again as a consequence, like in Section 7.2, we get the following decidability and approximability results for single-clock STA.

Corollary 7.3.3. *Let \mathcal{A} be a one-clock labelled STA, and \mathbf{M} a DMA. Then:*

1. *we can decide whether \mathcal{A} satisfies almost-surely \mathbf{M} ;*
2. *for every initial distribution μ which is numerically amenable w.r.t. \mathcal{A} , we can compute arbitrarily close approximations of $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{M})$.*

Proof. Similarly to the proof of Corollary 7.2.3, this is an application of Theorem 6.1.25, Corollary 6.1.26 and of Sections 6.2.1 and 6.2.4. All the hypotheses are met:

- $\mathcal{T}_{\mathcal{A}}^{\text{tg}} \times M$ has a finite attractor, and
- $\mathcal{T}_{\mathcal{A}} \times M$ is decisive w.r.t. any α_M -closed sets.

Those can be deduced by similar arguments as in the proof of Corollary 7.2.3. We only observe that if $\mathcal{T}_{\mathcal{A}}$ is a single-clock STA, then so is $\mathcal{T}_{\mathcal{A}} \times M$ and that hypotheses (H2), (H3), (H4) and (H5) are preserved through the product with M as those only concern distributions over the STA which are not altered from the product with M . \square

Remark 7.3.4. The proof of the existence of an attractor is very similar to the one that was used for proving the fairness property in [BBB⁺14, Section 7.1]. However, for free, we get all the approximation results! It is worth noting that these results encompass the results of [BBBM08], where a strong assumption on cycles of the STA were made (but a closed-form for the probability could be computed). We remark here that the graph used in [BBBM08] is actually the graph of the attractor, as done in Section 6.1.2.

Conclusion and Future Work

We now conclude Part I with a summary of our different results on the qualitative and quantitative verification of general stochastic systems and with a quick word on some possibilities that are left for future work.

Inspired by [ABM07], in Part I we have been interested in the qualitative and quantitative verification of general stochastic systems, *i.e.* probabilistic systems defined with a Markovian kernel and with a possibly continuous set of states. The model considered is the STS model, which can be seen as labelled Markov processes with a single label [Pan09].

In [ABM07], the authors defined a notion of decisiveness for DMCs and showed that, roughly speaking, it allows to lift good properties from finite Markov chains to DMCs. They have strong results for the qualitative model-checking of reachability and repeated reachability properties, but also algorithms for the approximation of probabilities of the same properties. In Part I, we have been interested in an extension of those results to more general stochastic systems (STSs) and to a richer class of properties (properties given by a DMA). The ultimate objective was then to apply those results to STA.

The main contributions have been divided into four chapters. In Chapter 4, we have defined the notions of decisiveness, attractor and fairness in STSs and we have established links between those notions.

The verification of STSs can become very quickly difficult due to the possibly continuous set of states. It is thus important to be able to reduce the analysis of such a system to a smaller one. In Chapter 5, we have defined the notion of α -abstraction which is a STS that preserve the positive one-step behaviour of the initial STS. The main contributions are the following.

- The definition of the notion of α -abstraction. As α -abstractions are not sufficient to get strong results, we have also defined the notions of sound and complete α -abstraction, which add the preservation of the almost-sure behaviour of reachability properties (Section 5.1).
- The identification of properties that can be transferred through sound α -abstractions, *e.g.* decisiveness and attractors (Section 5.2.1).
- The identification of stronger hypotheses (than soundness) allowing the the transfer of decisiveness properties through α -abstractions (Section 5.2.2). The hypotheses require in particular for the α -abstraction to be a DMC or a finite Markov chain (which is much more easier to analyse than general STSs), and are easier to check than soundness.
- The identification of conditions that imply soundness or completeness of α -abstractions (Section 5.3).

Chapter 6 is then the essential part of our contributions. It provides decidability and approximability results of ω -regular properties directly for general STSs but also for STSs through denumerable α -abstractions which is the most interesting part. The main contributions are the following.

- The extension of the results of [ABM07] on the qualitative verification of (repeated) reachability properties to general STSs and the identification of conditions under which it can be reduced to an α -abstraction (Section 6.1.1). Mainly, the properties needed for all the results are decisiveness and soundness of the abstraction.
- The adaptation of techniques of [ABRS05] in order to get a procedure for the almost-sure verification of properties given by a DMA for DMCs, and the extension of this procedure to general STSs through a denumerable abstraction (Sections 6.1.2 and 6.1.3). The condition over the DMCs in order to get the procedure is to have a finite attractor, while the condition for general STSs is for the abstraction to be sound.
- The extension of the algorithms of [ABM07] in DMCs for the approximation of (repeated) reachability properties, into approximation schemes in general STSs for the approximation of the same properties (Section 6.2.1 and 6.2.2). Decisiveness is required for correctness and termination.
- The use of the procedures of Sections 6.1.2 and 6.1.3 in order to get approximation schemes for properties given by a DMA in DMCs and general

STSs through a denumerable abstraction (Sections 6.2.3 and 6.2.4). Those schemes require to approach finitely many reachability properties, using thus the approximation scheme of Section 6.2.1.

Section 6.3 gives a more precise summary of all the results!

Finally in Chapter 7, we have identified classes of STA (that can be seen as STSs) on which all the previous results can be applied. We have proven that the thick region graph viewed as a finite Markov chain (see Section 3.2) is an α -abstraction. We then identified two classes of STA (reactive STA and a subclass of one-clock STA [BBB⁺14]) in which the α -abstraction is moreover sound. This yields to the fact that all the decidability and approximability results of Chapter 6 can be applied to those classes of STA!

Perspectives for future work. We can list some perspectives for future work, the list is not exhaustive.

- **Open problems:** as stated in Remark 6.1.20, we do not know if the product of a STS with a DMA preserves the decisiveness of the STS or the soundness of an α -abstraction. If the answer is yes, it could simplify the hypotheses that need to be checked in order to prove the appliance of all the results.
- In [BBBC17], we showed also the application of the qualitative and quantitative results to a subclass of GSMPs. Perspectives for the future are obviously to find other classes where the results can be applied, whether it is in STA or GSMPs, but also in other stochastic models like stochastic (time) Petri nets or stochastic hybrid systems.
- One could also add prices in the considered models. The question would then be whether we can adapt the previous schemes in order to approach an expected reward or other measures of performance.

Part II

Composition of Stochastic Timed Automata

Interleaving Parallel Composition in STA

In this chapter, we are concerned with the definition of an operator of parallel composition in STA as defined in [BBCM16]. Inspired by the approach of [HZ11] (see Section 2.4.1 for a short discussion on the subject), we here define an operator that corresponds to the interleaving semantics for this model, as we will prove it. This is the first step towards a handshaking operator of composition.

In Section 9.1, we define the parallel composition that we will consider. It extends the interleaving operator in timed automata of Section 2.1.3 and it thus amounts to equip the product of timed automata with probability distributions over the delays and the edges. It is quite technical to get a definition. We illustrate the definition and we give several conditions on the probability distributions. Those conditions will allow us, in Section 9.2, to show that the parallel composition is well-defined interleaving. We exhibit problematic behaviours when those conditions are not satisfied. We then identify a large subclass of STA which is closed under parallel composition and under which the previous properties are met.

Finally in Section 9.3, we define a notion of bisimulation in STA and we prove that it is a congruence w.r.t. the defined parallel composition operator. As explained in Sections 2.1.3 and 2.4.2, this is an expected property for a proper compositional design in STA.

All the work of this chapter has been published in [BBCM16].

9.1. Definition of the parallel composition

In this section, we define a notion of parallel composition in STA as defined in [BBCM16], and we illustrate it. But first, we recall some notations and add some slight restrictions on the STA model.

In this chapter, we consider STA $\mathcal{A} = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X})$ ¹ as defined in Definition 3.1.2 with this little modification: we assume that $E \subseteq L \times \mathcal{G}_\times(X) \times 2^X \times L$ and $\text{Inv} : L \rightarrow \mathcal{G}_\times(X)$ where $\mathcal{G}_\times(X) \subseteq \mathcal{G}(X)$ is the subset of guards over X defined as follows: any finite conjunction of expressions of the form $x \sim c$ with $x \in X$, $c \in \mathbb{N}$ and $\sim \in \{<, \geq, >\}$. This restriction is made in order to avoid some technicalities. We recall that there is a transition system $T_{\mathcal{A}} = (Q, \mathbb{R}_+ \times E, \rightarrow)$ associated with \mathcal{A} and we refer to Chapter 3 for the probability distribution $\text{Prob}_q^{\mathcal{A}}$ defined over the set of infinite runs $(\text{Runs}(\mathcal{A}, q), \Omega_q^{\mathcal{A}})$.

We also need another hypothesis for the probability distributions over the edges: we need that for each edge e and each state q , the function $p_{q+\bullet}(e) : \mathbb{R}_+ \rightarrow [0, 1]$ that assigns to each $t \geq 0$ the value $p_{q+t}(e)$, is measurable. We write this condition (\star) .² This condition is needed for technical results of Section 9.3.

We finally remind that for each state $q \in Q$, and each edge $e \in E$, $I(q, e) = \{t \in \mathbb{R}_+ \mid \nu + t \models \text{Inv}(l) \text{ and } \exists q' \in Q \text{ s.t. } q \xrightarrow{t, e} q'\}$ and $I(q) = \bigcup_{e \in E} I(q, e)$. One can then easily show this technical lemma.

Lemma 9.1.1. Let \mathcal{A} be a STA. Then for each state q of \mathcal{A} , $I(q)$ is either the empty set or a finite disjoint union of open intervals or intervals of the form $[a, b[$ with $a, b \in \mathbb{R}_+$.

As already briefly argued in Remark 3.1.1, this is proven thanks to the form of the guards in $\mathcal{G}_\times(X)$. This lemma will prove to be useful in the sequel.

We recall also that it is assumed that the underlying timed automaton is non-blocking, *i.e.* $I(q) \neq \emptyset$ for any state q . Moreover, we assume here that all STA considered satisfy hypothesis (\ddagger) of page 65: writing Λ for the Lebesgue measure, for each state q , if $\Lambda(I(q)) > 0$ then μ_q is equivalent to the restriction of Λ on $I(q)$.

We are now ready to tackle the definition of the parallel composition in STA. We consider two STA $\mathcal{A}_i = (L_i, X_i, E_i, \text{Inv}_i, \text{AP}_i, \mathcal{L}_i, (\mu_q^{(i)}, p_q^{(i)})_{q \in L_i \times \mathbb{R}_+^{X_i}})$ for $i = 1, 2$ with $X_1 \cap X_2 = \emptyset$, and we first recall the standard (interleaving) parallel composition for the underlying timed automaton (see Definition 2.1.29). It is the timed automaton $(L, X, E, \text{Inv}, \text{AP}, \mathcal{L})$ where $L = L_1 \times L_2$, $X = X_1 \cup X_2$,

¹The labelling over the edges with actions is not required here.

²It should be noted that if STA \mathcal{A} satisfies hypothesis (H2) of page 70, then \mathcal{A} satisfies (\star) .

$\text{Inv} : L \rightarrow \mathcal{G}_\times(X)$ is such that $\text{Inv}((l_1, l_2)) = \text{Inv}_1(l_1) \wedge \text{Inv}_2(l_2)$, $\text{AP} = \text{AP}_1 \cup \text{AP}_2$, $\mathcal{L} : L \rightarrow 2^{\text{AP}}$ is such that $\mathcal{L}((l_1, l_2)) = \mathcal{L}_1(l_1) \cup \mathcal{L}_2(l_2)$ and where $E = E_{1,\bullet} \cup E_{\bullet,2}$ with $E_{1,\bullet} = \{((l_1, l_2), g, Y, (l'_1, l_2)) \mid (l_1, g, Y, l'_1) \in E_1, l_2 \in L_2\}$. Note that with this timed automaton, the following transition system is associated: $(Q, \mathbb{R}_+ \times E, \rightarrow)$ defined as in Definition 2.1.11, where $Q = Q_1 \times Q_2$ with $Q_1 = L_1 \times \mathbb{R}_+^{X_1}$ is the set of states of \mathcal{A}_1 and $Q_2 = L_2 \times \mathbb{R}_+^{X_2}$ is the set of states of \mathcal{A}_2 .

Back to the STA, the parallel composition $\mathcal{A}_1 \parallel \mathcal{A}_2$ has as underlying timed automaton the interleaving product of both underlying timed automata of \mathcal{A}_1 and \mathcal{A}_2 ; it remains to equip each state $q = (q_1, q_2) \in Q$ with probability distributions over both delays and edges, with the following constraints:

- distributions over delays from state (q_1, q_2) should reflect a *race* between the two components \mathcal{A}_1 and \mathcal{A}_2 from respectively states q_1 and q_2 (just like in CTMCs in Section 2.3);
- distributions over edges should be state-based (or memoryless), that is, should not depend on how long has been waited before taking that edge, or which other actions have been done meanwhile by other components;
- globally, the product-automaton should correspond to the interleaving of \mathcal{A}_1 and \mathcal{A}_2 , which we express as follows: given a state $q = (q_1, q_2) \in Q$ given a property φ_1 that only concerns³ \mathcal{A}_1 and a property φ_2 that only concerns \mathcal{A}_2 ,

$$\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\varphi_1 \wedge \varphi_2) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\varphi_1) \cdot \text{Prob}_{q_2}^{\mathcal{A}_2}(\varphi_2).$$

Before tackling those items, we first introduce a class of STA that will be of interest to us. In order to do so, we need some new notations.

Let \mathcal{A} be a STA and let $q = (l, \nu)$ be a state of \mathcal{A} . Since \mathcal{A} is assumed to satisfy (\ddagger) , Radon-Nikodym's theorem ensures us the existence of a density function of μ_q w.r.t. the Lebesgue measure; we write it f_q . It thus holds that for each Borel set $A \subseteq \mathbb{R}_+$, $\mu_q(A) = \int_A f_q(t) dt$. We write F_q for the cumulative function associated to f_q , i.e. $F_q : \mathbb{R}_+ \rightarrow [0, 1]$ is such that $F_q(t') = \int_0^{t'} f_q(t) dt = \mu_q([0, t'])$ for each $t' \in \mathbb{R}_+$.

Now we consider some random probability space $(Z, \mathcal{E}, \text{Prob})$. For each state q , we consider a random variable \mathbb{X}_q with density function f_q in the probability space $(Z, \mathcal{E}, \text{Prob})$, i.e. for each Borel set $A \subseteq \mathbb{R}_+$, $\text{Prob}(\mathbb{X}_q \in A) = \int_{t \in A} f_q(t) dt = \mu_q(A)$.

We now define a first class of STA, called CSTA (see [BBCM16]), which is suitable to define a parallel composition. We say that a STA \mathcal{A} is in CSTA if:

³We will make it clearer in the sequel.

- (A) for every state q of \mathcal{A} , the density function associated with μ_q , denoted by f_q , is continuous everywhere on \mathbb{R}_+ except in a finite number of points, and
- (B) the family of probability distributions $(\mu_q)_{q \in Q}$ is *weakly-memoryless*, i.e. for every $t, t' \geq 0$, $\text{Prob}(\mathbb{X}_q \geq t + t' \mid \mathbb{X}_q \geq t) = \text{Prob}(\mathbb{X}_{q+t} \geq t')$.

This second condition is a consistency condition between states which belong to the same ‘time-elapsing fiber’, that is, sets of the form $F = \{q + t \mid t \in \mathbb{R} \text{ and } q + t \in Q\}$. Indeed, \mathbb{X}_q (resp. \mathbb{X}_{q+t}) represents the delay after which state q (resp. $q + t$) is left via an edge. Hence if q_0 is the minimal (for time-elapsing) element of F , then for every $q = q_0 + t \in F$, the law of \mathbb{X}_q has to be equal to the law of \mathbb{X}_{q_0} conditioned by the fact that t time units have already passed. An arbitrary distribution can be taken in q_0 (satisfying condition (A)), and distributions for $q \in F$ can then be inferred.

Remark 9.1.2. Conditions above are not very restrictive, since they only impose to fix a law satisfying (A) at each initial element of a fiber, and since condition (A) is met by a large class of distributions. Let q_0 be an initial element of a fiber, we can check that for instance,

- if $I(q_0)$ is a bounded subset of \mathbb{R}_+ and if μ_{q_0} is a uniform distribution over $I(q_0)$, then for every $t \in \mathbb{R}_+$, (B) imposes that μ_{q_0+t} is also uniform over $I(q_0 + t)$;
- similarly, if $I(q_0) = \mathbb{R}_+$, and if μ_{q_0} is an exponential distribution with parameter λ (denoted $\text{Exp}(\lambda)$), then for every $t \in \mathbb{R}_+$, (B) imposes that μ_{q_0+t} is also an $\text{Exp}(\lambda)$ -distribution. This corresponds to the classical memoryless property satisfied in CTMCs (see Section 2.3).

We will illustrate this in Example 9.1.5.

Under condition (A), we have a useful characterisation of condition (B).

Lemma 9.1.3. The probability distribution μ_q is weakly-memoryless iff for every $t, t' \geq 0$,

$$f_q(t + t') = (1 - F_q(t))f_{q+t}(t') \quad (9.1)$$

except in a finite number of points.

Proof. First, it should be noted that $\text{Prob}(\mathbb{X}_q \geq t) = 0$ iff $I(q + t) = \emptyset$. This comes from Lemma 9.1.1. In that case, it is thus not possible to leave state q after $t+t'$ time units for each $t' \geq 0$, so that such t does not have to be considered

since it leads to blocking-states (and we assume that there is no blocking-state). Now let $t \geq 0$ be such that $\text{Prob}(\mathbb{X}_q \geq t) > 0$. Then,

$$\begin{aligned}
& \text{Prob}(\mathbb{X}_q \geq t + t' \mid \mathbb{X}_q \geq t) = \text{Prob}(\mathbb{X}_{q+t} \geq t') \\
& \iff \frac{\text{Prob}(\mathbb{X}_q \geq t + t', \mathbb{X}_q \geq t)}{\text{Prob}(\mathbb{X}_q \geq t)} = \text{Prob}(\mathbb{X}_{q+t} \geq t') \\
& \iff (1 - F_q(t + t')) = (1 - F_q(t))(1 - F_{q+t}(t')) \tag{9.2} \\
& \iff \partial_{t'}(1 - F_q(t + t')) = \partial_{t'}((1 - F_q(t))(1 - F_{q+t}(t'))) \\
& \iff f_q(t + t') = (1 - F_q(t))f_{q+t}(t')
\end{aligned}$$

for every $t' \geq 0$ in which $f_{q+t}(t')$ and $f_q(t + t')$ are continuous, *i.e.* everywhere except in a finite number of points from (A). \square

Remark 9.1.4. Note that since f_q and f_{q+t} are density functions, we can assume w.l.o.g. that (9.1) holds for every t and $t' \geq 0$.

We now show that uniform and exponential distributions are family of distributions that satisfy conditions (A) and (B).

Example 9.1.5. We fix some STA \mathcal{A} and some state q . Firstly, we assume that $I(q) = [0, a[$ for some $a > 0$ and that μ_q is the uniform distribution $\mathcal{U}([0, a[)$; *i.e.* its density function is given by

$$f_q(t) = \frac{1}{a} \mathbb{1}_{[0, a[}(t)$$

for any $t \geq 0$. Obviously, f_q satisfies condition (A). We now prove that condition (B) ensures that for each $t \in [0, a[$, μ_{q+t} is the uniform distribution over $I(q + t) = [0, a - t[$. We use the characterisation of Lemma 9.1.3. Observe that the cumulative function F_q is given by

$$F_q(t) = \frac{t}{a} \mathbb{1}_{[0, a[}(t) + \mathbb{1}_{[a, \infty[}(t)$$

for any $t \geq 0$. We fix $t \in [0, a[$. Characterisation (9.1) states that for any $t' \geq 0$,

$$f_{q+t}(t') = \frac{f_q(t + t')}{1 - F_q(t)}$$

Note that $1 - F_q(t) \neq 0$ since $t \in [0, a[$. We thus get that for any $t' \geq 0$,

$$\begin{aligned}
f_{q+t}(t') &= \frac{f_q(t + t')}{1 - F_q(t)} = \frac{\frac{1}{a} \mathbb{1}_{[0, a[}(t + t')}{1 - \frac{t}{a}} \\
&= \frac{\frac{1}{a} \mathbb{1}_{[0, a-t[}(t')}{\frac{a-t}{a}} = \frac{1}{a-t} \mathbb{1}_{[0, a-t[}(t')
\end{aligned}$$

which proves that μ_{q+t} is the uniform distribution over $I(q+t)$.

We now suppose that $I(q) = \mathbb{R}_+$ and that μ_q is the exponential distribution of parameter $\alpha > 0$ $\text{Exp}(\alpha)$, *i.e.* its density function is given by

$$f_q(t) = \alpha e^{-\alpha t}$$

for any $t \geq 0$. Obviously, f_q satisfies condition (A). We now prove that condition (B) ensures that for each $t \in \mathbb{R}_+$, μ_{q+t} is the same exponential distribution. We again use the characterisation of Lemma 9.1.3. Observe that the cumulative function F_q is given by

$$F_q(t) = (1 - e^{-\alpha t})$$

for any $t \geq 0$ and thus

$$1 - F_q(t) = e^{-\alpha t}$$

for any $t \geq 0$. We fix $t \in \mathbb{R}_+$, we still have $I(q+t) = \mathbb{R}_+$. Characterisation (9.1) states that for any $t' \geq 0$,

$$f_{q+t}(t') = \frac{f_q(t+t')}{1 - F_q(t)} = \frac{\alpha e^{-\alpha(t+t')}}{e^{-\alpha t}} = \alpha e^{-\alpha t'}$$

which proves that μ_{q+t} is the exponential distribution $\text{Exp}(\alpha)$.

Remark 9.1.6. Example 9.1.5 shows thus that any family of distributions that contains a uniform distribution for state q but some arbitrary other distribution for some state $q+t$, does not satisfy condition (B). The same holds true for exponential distributions. For the latter case, we will illustrate it in Example 9.1.13 and we will show that it yields some undesirable properties.

We can now explain how to build the probability distributions associated with a state $q = (q_1, q_2)$ of $\mathcal{A}_1 \parallel \mathcal{A}_2$. Firstly, recall that \mathbb{X}_{q_1} (resp. \mathbb{X}_{q_2}) represents the delay after which state q_1 (resp. q_2) is left via an edge for any states $q_1 \in Q_1$ and $q_2 \in Q_2$. Since we want to define here an interleaving operator, we may assume that \mathcal{A}_1 and \mathcal{A}_2 run independently. Hence for each state $q_1 \in Q_1$ and each state $q_2 \in Q_2$ we may assume that the random variables \mathbb{X}_{q_1} and \mathbb{X}_{q_2} are independent (the time to leave state q_1 does not depend on the time to leave state q_2).

Now fix a state $q = (q_1, q_2) \in Q = Q_1 \times Q_2$ of $\mathcal{A}_1 \parallel \mathcal{A}_2$. We have to choose a probability distribution μ_q over the delays and a probability distribution p_q over the set of enabled edges in q .

Since state (q_1, q_2) is left as soon as q_1 or q_2 are left, we naturally define $\mu_q(A) = \int_A f_q(t) dt$ for each Borel set A , where f_q is the density function of the random-variable $\min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2})$: as \mathbb{X}_{q_1} (resp. \mathbb{X}_{q_2}) corresponds to the time after which state q_1 (resp. q_2) is left in \mathcal{A}_1 (resp. \mathcal{A}_2), $\min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2})$ corresponds to

the time after which state q is left in $\mathcal{A}_1 \parallel \mathcal{A}_2$. Under hypothesis (A) for f_{q_1} and f_{q_2} , one can show that $f_q(t) = f_{q_1}(t)(1 - F_{q_2}(t)) + f_{q_2}(t)(1 - F_{q_1}(t))$ almost-surely for every $t \geq 0$ (w.r.t. the Lebesgue measure). This comes from the equality $1 - F_q(t) = \text{Prob}(\min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \geq t)$, from the independence of \mathbb{X}_{q_1} and \mathbb{X}_{q_2} and from a derivative computation.

Lemma 9.1.7. It holds that $\min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2})$ is a random variable of density function f_q defined by $f_q(t) = f_{q_1}(t)(1 - F_{q_2}(t)) + f_{q_2}(t)(1 - F_{q_1}(t))$ for almost every $t \geq 0$.

Proof. Let f_q be the density function of the random variable $\min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2})$. Then, it holds that the cumulative function of $\min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2})$ is defined by $F_q(t) = \text{Prob}(\min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \leq t)$ for every $t \geq 0$, and we therefore get that $1 - F_q(t) = \text{Prob}(\min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \geq t)$ for every $t \geq 0$. Now, since

$$\begin{aligned} \text{Prob}(\min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \geq t) &= \text{Prob}(\{\mathbb{X}_{q_1} \geq t\} \cap \{\mathbb{X}_{q_2} \geq t\}) \\ &= \text{Prob}(\mathbb{X}_{q_1} \geq t) \text{Prob}(\mathbb{X}_{q_2} \geq t) \\ &\quad \text{by independence of } \mathbb{X}_{q_1} \text{ and } \mathbb{X}_{q_2} \\ &= (1 - F_{q_1}(t))(1 - F_{q_2}(t)), \end{aligned}$$

we deduce that $(1 - F_q(t)) = (1 - F_{q_1}(t))(1 - F_{q_2}(t))$ for every $t \geq 0$ and thus, when we compute the derivative of this last equality we get that $f_q(t) = f_{q_1}(t)(1 - F_{q_2}(t)) + f_{q_2}(t)(1 - F_{q_1}(t))$ for every $t \geq 0$ such that f_{q_1} and f_{q_2} are continuous in t , *i.e.* for every $t \geq 0$ except a finite number of points. \square

In order to define the probability distribution p_q over the enabled edges in q , one could consider that from state q , both systems \mathcal{A}_1 and \mathcal{A}_2 are in a race to win the next edge, *i.e.* \mathcal{A}_1 wins the race if the first edge taken from q is in E_1 . Hence, given $t \in I(q)$, and an edge $e \in E_1$ enabled in $q + t$, one would like that $p_{q+t}(e) = w_q^1(t) p_{q_1+t}^{(1)}(e)$ where $w_q^1(t)$ is the probability that, starting from q , \mathcal{A}_1 wins the race knowing that it was won after a delay of t time units. Formally, we define $w_q^1(t) = \lim_{\varepsilon \rightarrow 0} \text{Prob}(\mathbb{X}_{q_1} = \min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \mid \min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \in [t, t + \varepsilon])$ for every $t \geq 0$ and, still under hypothesis (A) for f_{q_1} and f_{q_2} , we can show that if $f_q(t) \neq 0$, then $w_q^1(t) = \frac{f_{q_1}(t)(1 - F_{q_2}(t))}{f_q(t)}$ almost-surely.

Lemma 9.1.8. It holds that for every $t \in I(q)$ except a finite number of points, $w_q^1(t) = \frac{f_{q_1}(t)(1 - F_{q_2}(t))}{f_q(t)}$.

Proof. For any t in $I(q)$ we have that

$$\begin{aligned} w_q^1(t) &:= \lim_{\varepsilon \rightarrow 0} \text{Prob}(\mathbb{X}_{q_1} = \min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \mid \min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \in [t, t + \varepsilon]) \\ &= \lim_{\varepsilon \rightarrow 0} \frac{\frac{1}{\varepsilon} \text{Prob}(\mathbb{X}_{q_1} = \min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \wedge \min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \in [t, t + \varepsilon])}{\frac{1}{\varepsilon} \text{Prob}(\min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \in [t, t + \varepsilon])}. \end{aligned}$$

Since $t \in I(q)$, from Lemma 9.1.1 we get that there is $\delta > 0$ such that for every $\varepsilon < \delta$, $[t, t + \varepsilon] \subseteq I(q)$ and thus $\text{Prob}(\min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \in [t, t + \varepsilon]) \neq 0$. We can thus compute $w_q^1(t)$ as follows. We have

$$\begin{aligned} &\text{Prob}(\mathbb{X}_{q_1} = \min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \wedge \min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \in [t, t + \varepsilon]) \\ &= \text{Prob}(\mathbb{X}_{q_1} = \min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \wedge \mathbb{X}_{q_1} \in [t, t + \varepsilon]) \\ &= \int_{t_1=t}^{t+\varepsilon} \int_{t_2=t_1}^{+\infty} f_{q_2}(t_2) f_{q_1}(t_1) dt_2 dt_1 \\ &\quad \text{by classical results of probability theory, as } \mathbb{X}_{q_1} \text{ and } \mathbb{X}_{q_2} \text{ are independent} \\ &= \int_{t_1=t}^{t+\varepsilon} f_{q_1}(t_1)(1 - F_{q_2}(t_1)) dt_1 \quad \text{by definition of the cumulative function} \end{aligned} \tag{9.3}$$

for every t in $I(q)$. Let us denote $g(t_1) = f_{q_1}(t_1)(1 - F_{q_2}(t_1))$ for every $t_1 \geq 0$. Since f_{q_1} and f_{q_2} are continuous everywhere except in a finite number of points, it holds that g is also continuous on \mathbb{R}_+ except in a finite number of points. Let us assume that g is continuous in t . Then, it can be supposed that g is continuous on $[t, t + \varepsilon]$ and since it is a closed interval, we have that g reaches its bounds on $[t, t + \varepsilon]$. We have

$$\int_{t_1=t}^{t+\varepsilon} \min_{x \in [t, t+\varepsilon]} g(x) dt_1 \leq \int_{t_1=t}^{t+\varepsilon} g(t_1) dt_1 \leq \int_{t_1=t}^{t+\varepsilon} \max_{x \in [t, t+\varepsilon]} g(x) dt_1,$$

and thus,

$$\min_{x \in [t, t+\varepsilon]} g(x) \leq \frac{1}{\varepsilon} \int_{t_1=t}^{t+\varepsilon} g(t_1) dt_1 \leq \max_{x \in [t, t+\varepsilon]} g(x). \tag{9.4}$$

Now, since g is continuous on $[t, t + \varepsilon]$, we have

$$\min_{x \in [t, t+\varepsilon]} g(x) \xrightarrow{\varepsilon \rightarrow 0} g(t) \quad \text{and} \quad \max_{x \in [t, t+\varepsilon]} g(x) \xrightarrow{\varepsilon \rightarrow 0} g(t).$$

From (9.3) and (9.4), we thus deduce that

$$\frac{1}{\varepsilon} \text{Prob}(\mathbb{X}_{q_1} = \min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \wedge \min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \in [t, t + \varepsilon]) \xrightarrow{\varepsilon \rightarrow 0} f_{q_1}(t)(1 - F_{q_2}(t)) \tag{9.5}$$

almost-surely for every $t \in I(q)$ (it holds for every t such that f_{q_1} is continuous in t). Similarly, we have that

$$\begin{aligned} & \frac{1}{\varepsilon} \text{Prob}(\min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \in [t, t + \varepsilon]) \\ &= \frac{1}{\varepsilon} (\text{Prob}(\mathbb{X}_{q_1} \in [t, t + \varepsilon] \wedge \mathbb{X}_{q_1} \leq \mathbb{X}_{q_2}) + \text{Prob}(\mathbb{X}_{q_2} \in [t, t + \varepsilon] \wedge \mathbb{X}_{q_2} \leq \mathbb{X}_{q_1})) \\ &\xrightarrow{\varepsilon \rightarrow 0} f_{q_1}(t)(1 - F_{q_2}(t)) + f_{q_2}(t)(1 - F_{q_1}(t)) = f_q(t), \end{aligned} \quad (9.6)$$

for every $t \in I(q)$ such that f_{q_1} and f_{q_2} are continuous in t , *i.e.* for every $t \in I(q)$ except in a finite number of points. Finally, it should be noted that since f_q is equivalent to the restriction of the Lebesgue measure on $I(q)$, one can assume w.l.o.g. that $f_q(t) \neq 0$ for each $t \in I(q)$. Hence, from (9.5) and (9.6) we deduce that

$$w_q^1(t) = \frac{f_{q_1}(t)(1 - F_{q_2}(t))}{f_q(t)}$$

for every t in $I(q)$ except in a finite number of points. \square

We now formalize the definition of the parallel composition of two STA (see [BBCM16]).

Definition 9.1.9. Let $\mathcal{A}_i = (L_i, X_i, E_i, \text{Inv}_i, \text{AP}_i, \mathcal{L}_i, (\mu_q^{(i)}, p_q^{(i)})_{q \in L_i \times \mathbb{R}_+^{X_i}})$ for $i = 1, 2$ be two STA. We say that \mathcal{A}_1 and \mathcal{A}_2 are *composable* if \mathcal{A}_1 and \mathcal{A}_2 are in CSTA and $X_1 \cap X_2 = \emptyset$.⁴ In that case, we define the parallel composition of \mathcal{A}_1 and \mathcal{A}_2 as the STA $\mathcal{A}_1 \parallel \mathcal{A}_2 = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X})$, where for any state $q = (q_1, q_2)$ of $\mathcal{A}_1 \parallel \mathcal{A}_2$,

- (i) $(L, X, E, \text{Inv}, \text{AP}, \mathcal{L})$ is the composition of the underlying timed automata \mathcal{A}_1 and \mathcal{A}_2 ,
- (ii) μ_q is defined as follows:

$$\forall A \in \mathcal{B}(\mathbb{R}_+), \mu_q(A) = \int_A f_q(t) dt,$$

where $f_q(t) = f_{q_1}(t)(1 - F_{q_2}(t)) + f_{q_2}(t)(1 - F_{q_1}(t))$ for every $t \geq 0$, and

- (iii) for any $t \in I(q)$, p_{q+t} is defined as follows:

$$p_{q+t}(e) = \mathbf{1}_{E_1}(e)w_q^1(t)p_{q_1+t}^{(1)}(e) + \mathbf{1}_{E_2}(e)w_q^2(t)p_{q_2+t}^{(2)}(e)$$

⁴If this is not the case, we can rename the clocks and assume that the two sets of clocks are disjoint.

for every $e \in E$, where for any $t \in I(q)$,

$$w_q^1(t) := \frac{f_{q_1}(t)(1 - F_{q_2}(t))}{f_q(t)} \quad \text{and} \quad w_q^2(t) := \frac{f_{q_2}(t)(1 - F_{q_1}(t))}{f_q(t)}$$

if $f_q(t) \neq 0$, and $w_q^1(t) = w_q^2(t) = 0$ if $f_q(t) = 0$.

We illustrate the definition by composing two independent copies of the STA modelling the IPv4 Zeroconf protocol (see Example 3.1.10).

Example 9.1.10. In order to illustrate the notion of composition, we composed two independent copies of the STA modelling the IPv4 Zeroconf protocol (see Example 3.1.10). Part of the composed STA is depicted in Figure 9.1, in which we assume to begin from the location (IP_1, IP_2) .




Figure 9.1: The product of two STA modelling the IPv4 Zeroconf

We consider two STA \mathcal{A}_1 and \mathcal{A}_2 , modelling each the IPv4 Zeroconf protocol and that run completely independently. Observe that from the computations of Example 9.1.5, \mathcal{A}_1 and \mathcal{A}_2 are composable. We explain how we compute in $\mathcal{A}_1 \parallel \mathcal{A}_2$, the probability distributions over delays and edges in location (IP_1, IP_2) . We consider a state of the form $q = (q_1, q_2)$ with $q_1 = (IP_1, \nu_1)$ and $q_2 = (IP_2, \nu_2)$. We thus have that $I(q_1) = \mathbb{R}_+ \setminus \{T - \nu_1\}$ and $I(q_2) = \mathbb{R}_+ \setminus \{T - \nu_2\}$. We suppose that q_1 (resp. q_2) is equipped with $\text{Exp}(\mu_1)$ (resp. $\text{Exp}(\mu_2)$) for the probability over the delays and we refer to Example 9.1.5 for the density and cumulative functions of such distributions. Then the law of the minimum between these two probabilities is an exponential distribution of parameter $\mu_1 + \mu_2$, *i.e.* $\text{Exp}(\mu_1 + \mu_2)$. Indeed for any $t \geq 0$:

$$\begin{aligned} f_q(t) &= f_{q_1}(t)(1 - F_{q_2}(t)) + f_{q_2}(t)(1 - F_{q_1}(t)) \\ &= \mu_1 e^{-\mu_1 t} e^{-\mu_2 t} + \mu_2 e^{-\mu_2 t} e^{-\mu_1 t} \\ &= (\mu_1 + \mu_2) e^{-(\mu_1 + \mu_2)t}. \end{aligned}$$

Now one can see that for every $t \geq 0$, from $q + t$, there are two enabled edges: e_1 which is the only edge enabled from $q_1 + t$ in \mathcal{A}_1 and e_2 which is the only edge enabled from $q_2 + t$ in \mathcal{A}_2 . Hence, one can easily compute

$$p_{q+t}(e_1) = w_q^1(t) = \frac{\mu_1}{\mu_1 + \mu_2} \quad \text{and} \quad p_{q+t}(e_2) = w_q^2(t) = \frac{\mu_2}{\mu_1 + \mu_2}.$$

The rest of the automaton can be completed in a similar way.

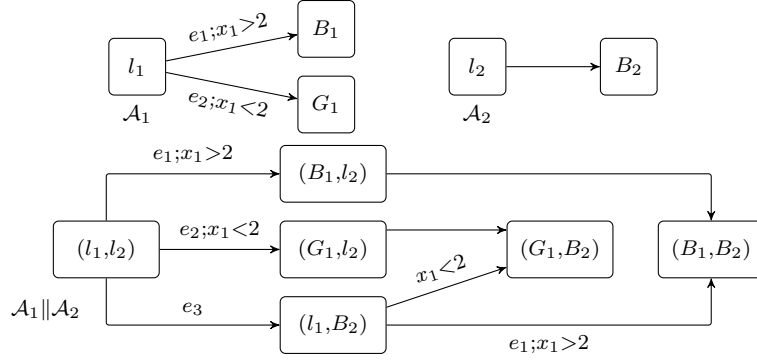
Remark 9.1.11. Observe that from Example 9.1.5, all STA with only uniform distributions or exponential distributions satisfying condition (B) can be composed between them.

Remark 9.1.12. It should be observed that if \mathcal{A} is a CTMC, then Definition 9.1.9 corresponds to the interleaving composition defined in Section 2.4.1 on page 48. Indeed, the minimum of two exponential distributions of parameters λ_1 and λ_2 gives an exponential distribution of parameter $\lambda_1 + \lambda_2$ which is what we get from the definition on 48 and from the semantics of CTMCs described in Section 2.3.

We now give an example of a family of probability measures that do not satisfy hypothesis (B), which yields undesirable properties in the parallel composition.

Example 9.1.13. We consider the single-clock STA \mathcal{A}_1 depicted in Figure 9.2. We assume $\mu_{q_1}^{(1)}$ is an exponential distribution of parameter λ_1 (resp. λ_2) if $q_1 = (l_1, \nu_1)$ with $\nu_1 < 1$ (resp. $\nu_2 \geq 1$), and with $\lambda_1 \neq \lambda_2$. Then for each $\nu_1 \in [0, 1[$, $\mu_{q_1}^{(1)}$ does not satisfy hypothesis (B) (as argued in Example 9.1.5). We assume that from l_1 , \mathcal{A}_1 can move to G_1 if $x_1 < 2$ or to B_1 if $x_1 > 2$, and from there, it stays in the same location with probability 1. We then compose \mathcal{A}_1 (using Definition 9.1.9) with the single-clock STA \mathcal{A}_2 , where \mathcal{A}_2 has l_2 as its initial location from which it can move at any time to B_2 where it stays with probability 1. We equip each state of the form $q_2 = (l_2, \nu_2)$ with an exponential distribution of parameter λ_2 over the delays. Then, assuming that we begin from state $q_1 = (l_1, 0)$ in \mathcal{A}_1 and from state $q_2 = (l_2, 0)$ in \mathcal{A}_2 , it can be shown that the probability to reach B_1 in \mathcal{A}_1 corresponds to the probability to reach (B_1, B_2) in $\mathcal{A}_1 \parallel \mathcal{A}_2$ iff $\ln(\lambda_1) - \ln(\lambda_2) = \lambda_1 - \lambda_2$ which is not true in general (in particular with $\lambda_1 = 1$ and $\lambda_2 = 2$). Observe that in order to have an interleaving semantic, we want this equality to hold true in any case.

We refer to Examples 9.1.5 and 9.1.10 for some notions on exponential distributions. The product $\mathcal{A}_1 \parallel \mathcal{A}_2$ is depicted below. From the initial state $q = (q_1, q_2) = ((l_1, 0), (l_2, 0))$, it should be noted that both clocks will always

Figure 9.2: $\mathcal{A}_1 \notin \text{CSTA}$.

have the same value since we never reset any clock to 0. Then, following Definition 9.1.9, one can see that

$$\mu_{((l_1, \nu), (l_2, \nu))} = \begin{cases} \text{Exp}(\lambda_1 + \lambda_2) & \text{if } \nu < 1, \\ \text{Exp}(2 \cdot \lambda_2) & \text{otherwise} \end{cases}$$

where $\text{Exp}(\lambda)$ denotes the exponential distribution of parameter λ , and that for each $i \in \{1, 2\}$ and for every $t \geq 0$,

$$w_{((l_1, \nu), (l_2, \nu))}^i(t) = \begin{cases} \frac{\lambda_i}{\lambda_1 + \lambda_2} & \text{if } \nu < 1, \\ \frac{1}{2} & \text{otherwise.} \end{cases}$$

In states of the form $((B_1, \nu), (l_2, \nu))$ and $((G_1, \nu), (l_2, \nu))$ we keep the distributions of (l_2, ν) , while in $((l_1, \nu), (B_2, \nu))$ we keep the distributions of (l_1, ν) .

Now, one can observe that the set of runs in \mathcal{A}_1 starting in q_1 that reach B_1 after 2 time units is given by $\text{Cyl}(\pi(q_1, e_1))$, while the set of runs in $\mathcal{A}_1 \parallel \mathcal{A}_2$ starting in q that reach B_1 is given by $\text{Cyl}(\pi(q, e_1)) \cup \text{Cyl}(\pi(q, e_3, e_1))$. Correspondingly to our interleaving semantics, we would like that

$$\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Cyl}(\pi(q, e_1)) \cup \text{Cyl}(\pi(q, e_3, e_1))) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi(q_1, e_1))). \quad (9.7)$$

It can easily be established that

$$\mathbb{P}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi(q_1, e_1))) = e^{-2\lambda_1}. \quad (9.8)$$

And we can compute

$$\begin{aligned}
& \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\text{Cyl}(\pi(q, e_1)) \cup \text{Cyl}(\pi(q, e_3, e_1)) \right) \\
&= \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\text{Cyl}(\pi(q, e_1)) \right) + \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\text{Cyl}(\pi(q, e_3, e_1)) \right) \\
&= \int_2^\infty \frac{\lambda_1}{\lambda_1 + \lambda_2} \cdot (\lambda_1 + \lambda_2) \cdot e^{-(\lambda_1 + \lambda_2)t} dt \\
&+ \int_{t_1=0}^1 \frac{\lambda_2}{\lambda_1 + \lambda_2} \cdot (\lambda_1 + \lambda_2) \cdot e^{-(\lambda_1 + \lambda_2)t_1} \int_{t_2=2-t_1}^\infty \lambda_1 \cdot e^{-\lambda_1 t_2} dt_2 dt_1 \\
&+ \int_{t_1=1}^2 \frac{\lambda_2}{\lambda_1 + \lambda_2} \cdot (\lambda_1 + \lambda_2) \cdot e^{-(\lambda_1 + \lambda_2)t_1} \int_{t_2=2-t_1}^\infty \lambda_2 \cdot e^{-\lambda_2 t_2} dt_2 dt_1 \\
&+ \int_2^\infty \frac{\lambda_2}{\lambda_1 + \lambda_2} \cdot (\lambda_1 + \lambda_2) \cdot e^{-(\lambda_1 + \lambda_2)t} dt \\
&= e^{-2(\lambda_1 + \lambda_2)} + e^{-2\lambda_1} \cdot (1 - e^{-\lambda_2}) + \frac{\lambda_2}{\lambda_1} \cdot e^{-2\lambda_2} \cdot (e^{-\lambda_1} - e^{-2\lambda_1}). \quad (9.9)
\end{aligned}$$

Hence, from (9.8) and (9.9), we have (9.7) iff

$$\begin{aligned}
& e^{-2(\lambda_1 + \lambda_2)} + e^{-2\lambda_1} - e^{-2\lambda_1 - \lambda_2} + \frac{\lambda_2}{\lambda_1} \cdot e^{-2\lambda_2} \cdot (e^{-\lambda_1} - e^{-2\lambda_1}) = e^{-2\lambda_1} \\
&\iff \frac{\lambda_2}{\lambda_1} \cdot e^{-2\lambda_2} \cdot (e^{-\lambda_1} - e^{-2\lambda_1}) = e^{-2\lambda_1 - \lambda_2} - e^{-2(\lambda_1 + \lambda_2)} \\
&\iff \frac{\lambda_2}{\lambda_1} \cdot e^{-2\lambda_2} \cdot (e^{-\lambda_1} - e^{-2\lambda_1}) = e^{-2\lambda_2} \cdot (e^{-2\lambda_1 + \lambda_2} - e^{-2\lambda_1}) \\
&\iff \frac{\lambda_2}{\lambda_1} \cdot e^{-\lambda_1} = e^{-2\lambda_1 + \lambda_2} \\
&\iff \lambda_2 e^{-\lambda_2} = \lambda_1 e^{-\lambda_1} \\
&\iff \ln(\lambda_1) - \ln(\lambda_2) = \lambda_1 - \lambda_2.
\end{aligned}$$

9.2. Properties of the parallel composition

We are now ready to prove that this parallel composition operator satisfies all the expected properties, in particular that the operator is interleaving. Again, all the results of the section can be found in [BBCM16]. In this section, we assume that \mathcal{A}_1 and \mathcal{A}_2 are two composable STA and we assume the notations of Definition 9.1.9. We first have this important result which states that parallel composition is well-defined in CSTA and it is internal within the class.

Lemma 9.2.1. The distributions μ_q and p_q are well-defined, and the STA $\mathcal{A}_1 \parallel \mathcal{A}_2$ belongs to the class CSTA. Moreover, $\mathcal{A}_1 \parallel \mathcal{A}_2$ satisfies conditions (\ddagger) and (\star) .

We first need a technical result that we do not give the proof here as it has no interest.

Lemma 9.2.2. Let $g_1, g_2, h : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be measurable functions such that there is $J \subseteq \mathbb{R}_+$ with $\Lambda(J) > 0$ and h almost-surely non null on J . If for every $I \in \mathcal{B}(\mathbb{R}_+)$,

$$\int_{t \in I} g_1(t)h(t) dt = \int_{t \in I} g_2(t)h(t) dt, \quad (9.10)$$

then $g_1 = g_2$ almost-surely on J .

Proof of Lemma 9.2.1. We first make clear what we mean by “well-defined”. Let \mathcal{A}_1 and $\mathcal{A}_2 \in \text{CSTA}$, in order to construct $\mathcal{A}_1 \parallel \mathcal{A}_2$, we have defined probability distributions over both delays and edges from each state q of the product. It can be easily established that μ_q and p_q are probability distributions over the delays and edges. For the delays, this is directly ensured from the construction of f_q . For the edges, we have to show that for each $t \in I(q)$, p_{q+t} defines a probability distribution over the set of enabled edges. Given a state q and $t \in I(q)$, we have that⁵

$$\begin{aligned} & \{((l_1, l_2), g, Y, (l'_1, l'_2)) \in E \mid (\nu_1 + t, \nu_2 + t) \models g\} \\ &= \{(l_1, g, Y, l'_1) \in E_1 \mid \nu_1 + t \models g\} \cup \{(l_2, g, Y, l'_2) \in E_2 \mid \nu_2 + t \models g\} \end{aligned} \quad (9.11)$$

where $q = ((l_1, l_2), (\nu_1, \nu_2))$, and for any $t \in I(q)$,

$$\begin{aligned} w_q^1(t) + w_q^2(t) &= \frac{f_{q_1}(t)(1 - F_{q_2}(t))}{f_q(t)} + \frac{f_{q_2}(t)(1 - F_{q_1}(t))}{f_q(t)} \\ &= \frac{f_{q_1}(t)(1 - F_{q_2}(t)) + f_{q_2}(t)(1 - F_{q_1}(t))}{f_q(t)} \\ &= 1 \quad \text{by definition of } f_q(t). \end{aligned} \quad (9.12)$$

Now if $t \in I(q_1)$ and $t \notin I(q_2)$, then $w_q^2(t) = 0$ and thus $w_q^1(t) = 1$ from (9.12). Thus,

$$\sum_{e \in E} p_{q+t}(e) = \sum_{e \in E_1} p_{q_1+t}^{(1)}(e) = 1,$$

since $p_{q_1+t}^{(1)}$ is a probability distribution over the set of (enabled) edges in E_1 , and thus p_{q+t} is a probability distribution over E . Similarly we have, if $t \notin I(q_1)$ and

⁵Recall the abusive correspondence between E_1 (resp. E_2) and $E_{1,\bullet}$ (resp. $E_{\bullet,2}$).

$t \in I(q_2)$, that $\sum_{e \in E} p_{q+t}(e) = 1$ and thus p_{q+t} is a probability distribution over E . Now, if $t \in I(q_1) \cap I(q_2)$, then from (9.11) and (9.12), we have that

$$\begin{aligned} \sum_{e \in E} p_{q+t}(e) &= \sum_{e \in E_1} p_{q+t}(e) + \sum_{e \in E_2} p_{q+t}(e) \\ &= w_q^1(t) \sum_{e \in E_1} p_{q_1+t}^{(1)}(e) + w_q^2(t) \sum_{e \in E_2} p_{q_2+t}^{(2)}(e) \\ &= w_q^1(t) + w_q^2(t) = 1. \end{aligned}$$

Hence, for any $t \in I(q)$, p_{q+t} is a probability distribution over the set of enabled edges in $q+t$.

Now, for the delays we have defined μ_q for each state q separately so that no problems could be encountered. However for the edges, from each state q we have defined the probability p_{q+t} for each $t \geq 0$. Hence, $\mathcal{A}_1 \parallel \mathcal{A}_2$ is well-defined only if for each state q and for each $t, t' \geq 0$, $p_{q+(t+t')} = p_{(q+t)+t'}$. This equality holds if $w_q^i(t+t') = w_{q+t}^i(t')$ for any $t, t' \geq 0$ and for any $i \in \{1, 2\}$, which is ensured thanks to hypothesis (B). Indeed, we have

$$\begin{aligned} w_q^1(t+t') &= w_{q+t}^1(t') \\ \iff \frac{f_{q_1}(t+t')(1-F_{q_2}(t+t'))}{f_q(t+t')} &= \frac{f_{q_1+t}(t')(1-F_{q+t}(t'))}{f_{q+t}(t')} \\ \iff f_{q_1}(t+t')(1-F_{q_2}(t+t'))f_{q_2+t}(t')(1-F_{q_1+t}(t')) & \\ &= f_{q_2}(t+t')(1-F_{q_1}(t+t'))f_{q_1+t}(t')(1-F_{q+t}(t')). \end{aligned}$$

We show that for each $i = 1, 2$, and for each t and $t' \geq 0$, $f_{q_i}(t+t')(1-F_{q_i+t}(t')) = f_{q_i+t}(t')(1-F_{q_i}(t+t'))$. Let $t, t' \geq 0$, from the equivalences in (9.2) we have that

$$\begin{aligned} f_{q_i}(t+t') &= (1-F_{q_i}(t))f_{q_i+t}(t') \\ \iff (1-F_{q_i}(t+t')) &= (1-F_{q_i}(t))(1-F_{q_i+t}(t')). \end{aligned} \quad (9.2)$$

First, let us notice that if $(1-F_{q_i+t}(t')) = 0$, then $(1-F_{q_i}(t+t')) = 0$ from the last equivalence and since $\mathcal{A}_i \in \text{CSTA}$. Thus $f_{q_i}(t+t')(1-F_{q_i+t}(t')) = f_{q_i+t}(t')(1-F_{q_i}(t+t'))$. Now if $(1-F_{q_i+t}(t')) \neq 0$, under hypothesis (B) and (9.2), we have that

$$\begin{aligned} f_{q_i}(t+t')(1-F_{q_i+t}(t')) &= f_{q_i+t}(t')(1-F_{q_i}(t+t')) \\ \iff f_{q_i}(t+t') &= f_{q_i+t}(t') \frac{(1-F_{q_i}(t+t'))}{(1-F_{q_i+t}(t'))} \\ \iff f_{q_i}(t+t') &= (1-F_{q_i}(t))f_{q_i+t}(t') \end{aligned}$$

which is true. We conclude that $w_q^1(t+t') = w_{q+t}^1(t')$ for every t and $t' \geq 0$, which ensures us that p_q is well-defined.

We prove now that $\mathcal{A}_1 \parallel \mathcal{A}_2 \in \text{CSTA}$: we need to show that for each state q of the product, f_q satisfies hypotheses (A) and (B). Let q be a state of $\mathcal{A}_1 \parallel \mathcal{A}_2$. Point (A) is easily established since f_{q_1} and f_{q_2} satisfy this point and since F_{q_1} and F_{q_2} are continuous. Now let $t, t' \geq 0$, from (9.2), we have that

$$\begin{aligned} f_q(t+t') &= (1 - F_q(t))f_{q+t}(t') \\ \iff (1 - F_q(t+t')) &= (1 - F_q(t))(1 - F_{q+t}(t')). \end{aligned}$$

Now by recalling that $(1 - F_q(t+t')) = \text{Prob}(\{\mathbb{X}_{q_1} \geq t+t'\} \cap \{\mathbb{X}_{q_2} \geq t+t'\})$ with \mathbb{X}_{q_1} and \mathbb{X}_{q_2} independent, we can compute $(1 - F_q(t+t'))$ as follows:

$$\begin{aligned} (1 - F_q(t+t')) &= (1 - F_{q_1}(t+t'))(1 - F_{q_2}(t+t')) \\ &= (1 - F_{q_1}(t))(1 - F_{q_1+t}(t'))(1 - F_{q_2}(t))(1 - F_{q_2+t}(t')) \\ &\quad \text{by hypotheses over } \mathcal{A}_1 \text{ and } \mathcal{A}_2 \\ &= ((1 - F_{q_1}(t))(1 - F_{q_2}(t)))((1 - F_{q_1+t}(t'))(1 - F_{q_2+t}(t'))) \\ &= (1 - F_q(t))(1 - F_{q+t}(t')) \end{aligned}$$

which is what we want.

It remains to show that $\mathcal{A}_1 \parallel \mathcal{A}_2$ satisfies hypotheses (\dagger) and (\star) . Condition (\star) is trivial: it requires to check that for each $e \in E$, $p_{q+\bullet}(e)$ is measurable which is the case since $p_{q_1+\bullet}^{(1)}(e_1)$ (resp. $p_{q_2+\bullet}^{(2)}(e_2)$) is measurable for each $e_1 \in E_1$ (resp. $e_2 \in E_2$) and since w_q^1 and w_q^2 are continuous on \mathbb{R}_+ except in a finite number of points (and thus measurable).

The satisfaction of condition (\dagger) is a little more tricky. It comes from the fact that $I((q_1, q_2)) = (I(q_1) \cap I_{\text{Inv}}(q_2)) \cup (I(q_2) \cap I_{\text{Inv}}(q_1))$ where $I_{\text{Inv}}(q_i) = \{t \in \mathbb{R}_+ \mid \nu_i + t \models \text{Inv}(l_i)\}$ with $q_i = (l_i, \nu_i)$ and from the fact that no blocking states are allowed in the model so that for each $t \in I(q_i)$, $I(q_i + t) \neq \emptyset$.

Fix a state $q = (q_1, q_2)$. First observe that thanks to similar arguments as for Lemma 9.1.1 and Remark 3.1.1, it holds that there is $c_i > 0$ such that for each $i = 1, 2$, $I_{\text{Inv}}(q_i) = [0, c_i[$ ($0 \in I_{\text{Inv}}(q_i)$ since q_i is a state). Secondly, it holds that for each i , $(1 - F_{q_i}(t)) > 0$ if and only if $t \in I_{\text{Inv}}(q_i)$: from Lemma 9.1.1 and from (\dagger) , $(1 - F_{q_i}(t)) > 0$ if and only if there are $t_2 > t_1 \geq t$ such that $[t_1, t_2[\subseteq I(q_i)$ and observe also that $I(q_i) \subseteq I_{\text{Inv}}(q_i)$. Since $I_{\text{Inv}}(q_i) = [0, c_i[$, we get there are $t_2 > t_1 \geq t$ such that $[t_1, t_2[\subseteq I_{\text{Inv}}(q_i)$ if and only if $t \in I_{\text{Inv}}(q_i)$.

We have to show that for each $A \in \mathcal{B}(\mathbb{R}_+)$, $\mu_q(A) = 0$ if and only if $\Lambda(A \cap I(q)) = 0$. Fix $A \in \mathcal{B}(\mathbb{R}_+)$. The fact that $\Lambda(A \cap I(q)) = 0$ implies $\mu_q(A) = 0$

is trivial from the definition of μ_q . Now for the other implication, assume that $\mu_q(A) = 0$, we want to show that $\Lambda(A \cap I(q)) = 0$. We will show that

$$\Lambda(A \cap I(q_1) \cap I_{\text{Inv}}(q_2)) + \Lambda(A \cap I(q_2) \cap I_{\text{Inv}}(q_1)) = 0,$$

or equivalently that $\Lambda(A \cap I(q_1) \cap I_{\text{Inv}}(q_2)) = 0$ and $\Lambda(A \cap I(q_2) \cap I_{\text{Inv}}(q_1)) = 0$.

W.l.o.g. and towards a contradiction, assume that $\Lambda(A \cap I(q_1) \cap I_{\text{Inv}}(q_2)) > 0$. It holds from the hypothesis that

$$\mu_q(A) = \int_{t \in A} f_{q_1}(t)(1 - F_{q_2}(t)) dt + \int_{t \in A} f_{q_2}(t)(1 - F_{q_1}(t)) dt = 0$$

and in particular,

$$\int_{t \in A} f_{q_1}(t)(1 - F_{q_2}(t)) dt = 0.$$

Observe that

$$\int_{t \in A} f_{q_1}(t)(1 - F_{q_2}(t)) dt = \int_{t \in A \cap I(q_1) \cap I_{\text{Inv}}(q_2)} f_{q_1}(t)(1 - F_{q_2}(t)) dt$$

from the previous observations for $(1 - F_{q_i})$ and from the fact that $\mu_{q_1}^{(1)}$ is a distribution over $I(q_1)$. Hence from Lemma 9.2.2 with $J = A \cap I(q_1) \cap I_{\text{Inv}}(q_2)$, $g_1 = f_{q_1}$, $g_2 = 0$ and $h = (1 - F_{q_2})$, we get that $f_{q_1} = 0$ almost-surely on $A \cap I(q_1) \cap I_{\text{Inv}}(q_2)$ and thus $\mu_{q_1}^{(1)}(A \cap I_{\text{Inv}}(q_2)) = 0$ which is a contradiction with $\Lambda((A \cap I_{\text{Inv}}(q_2)) \cap I(q_1)) > 0$ since \mathcal{A}_1 satisfies (\ddagger) . This concludes the proof. \square

We now make some technical observations that should be pointed out before going further into the section.

Remark 9.2.3. Allowing equalities or the inequality “ \leq ” in the guards would lead to major technicalities. Indeed, assume that \mathbb{X}_{q_1} is a uniform distribution over $[0, 2[$ and \mathbb{X}_{q_2} is a discrete distribution that charges only 1 (that is $\text{Prob}(\mathbb{X}_{q_2} = 1) = 1$, and \mathbb{X}_{q_2} does not admit a density function), with \mathbb{X}_{q_1} and \mathbb{X}_{q_2} independent. Then, one can easily compute $\text{Prob}(\min(\mathbb{X}_{q_1}, \mathbb{X}_{q_2}) \geq t) = 1 - \frac{t}{2}$ if $t \leq 1$, and 0 otherwise, which corresponds to the distribution assigning probability $\frac{1}{2}$ to $\{1\}$ and $\frac{1}{2}$ to $[0, 1[$. This is a distribution that is neither discrete, nor continuous and which does not admit a density function and moreover, that does not satisfy hypothesis (\ddagger) . Hence allowing equalities in constraints would significantly complicate the proofs, this is why we restrict to open guards in this paper.

Remark 9.2.4. It should be noted that this parallel composition is, in some sense, commutative. Indeed, if we compare $\mathcal{A}_1 \parallel \mathcal{A}_2$ and $\mathcal{A}_2 \parallel \mathcal{A}_1$, both automata will generate the same executions except that the states visited are in $Q_1 \times Q_2$ in the first case and in $Q_2 \times Q_1$ in the second case.

Also, Definition 9.1.9 can be extended to the composition of n stochastic timed automata. Let $\mathcal{A}_1, \dots, \mathcal{A}_n$ be n composable stochastic timed automata. Then, if we keep the same notations as before, we define $\mathcal{A}_1 \parallel \dots \parallel \mathcal{A}_n$ by letting

$$f_q(t) = \sum_{i=1}^n f_{q_i}(t) \prod_{j \neq i} (1 - F_{q_j}(t))$$

and

$$w_q^i(t) = \frac{f_{q_i}(t) \prod_{j \neq i} (1 - F_{q_j}(t))}{f_q(t)}$$

if $f_q(t) \neq 0$ and $w_q^i(t) = 0$ otherwise, for any $q = ((l_1, \dots, l_n), (\nu_1, \dots, \nu_n))$ with $\lambda(I(q_i)) > 0$ for each $i \in \{1, \dots, n\}$, for any $t \geq 0$ and for any $i \in \{1, \dots, n\}$.

It remains to identify when the parallel composition really coincides with an interleaving semantics. This is in general not true, as already shown in Example 9.1.13 (which does not satisfy Condition (B)), and witnessed further by Example 9.2.5 below (which satisfies both conditions (A) and (B)).

Example 9.2.5. We consider the STA \mathcal{A}_1 and \mathcal{A}_2 of Figure 9.3, equipped resp. with an $\text{Exp}(\lambda)$ -distribution and a uniform distribution. Let $q = (q_1, q_2)$ be a state of $\mathcal{A}_1 \parallel \mathcal{A}_2$, with $q_i = (l_i, 0)$. One can easily check that $\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(q \xrightarrow{*} \xrightarrow{e_1}) = 0^6$ while $\text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi(q_1, e_1))) = 1$ which contradicts the independence property we expect. One can notice that \mathcal{A}_2 is Zeno with probability 1 (see Sections 2.1.2 and Section 3.1 (page 64)).




Figure 9.3: \mathcal{A}_2 is Zeno

Hence we define a subclass CSTA^* of CSTA ; $\mathcal{A} \in \text{CSTA}^*$ will be in CSTA^* if:

(C) \mathcal{A} is almost-surely non-Zeno.

⁶ $q \xrightarrow{*} \xrightarrow{e_1}$ is a notation for the set of runs starting in q that will eventually take edge e_1 .

Remark 9.2.6. Hypothesis (C) is not too restrictive since Zeno runs can be seen as faulty behaviours as already argued in Section 2.1.2: they perform infinitely many actions in a finite amount of time, which is not realistic. We will see that hypothesis (C) is sufficient (together with (A) and (B)) to show that the parallel composition really coincides with an interleaving semantics. Note that condition (C) can be decided in various subclasses of STA [BBB⁺14], in particular the ones identified in Sections 7.2 and 7.3. For reactive STA, it is moreover shown that they are all almost-surely non-Zeno.

We give some more notations. Let \mathcal{A} be a STA and let φ be a property for \mathcal{A} . Given a state q , we say that φ is measurable from q if the set of runs starting from q satisfying φ is in $\Omega_{\mathcal{A}}^q$ ⁷; we write this set $\{q \models \varphi\}$ and recovering similar notations as in Chapter 4, we write $\text{Prob}_q^{\mathcal{A}}(\varphi)$ for the probability of this set. Now let \mathcal{A}_1 and \mathcal{A}_2 be two composable STA. We write ι_1 (resp. ι_2) for the natural projection of $\text{Runs}(\mathcal{A}_1 \parallel \mathcal{A}_2, q)$ onto $\text{Runs}(\mathcal{A}_1, q_1)$ (resp. $\text{Runs}(\mathcal{A}_2, q_2)$) for each state $q = (q_1, q_2) \in Q_1 \times Q_2$: we inductively define $\iota_1(q) = q_1$, $\iota_1(q \xrightarrow{t_1, e_1} q^{(1)}) = q_1 \xrightarrow{t_1, e_1} q_1^{(1)}$ if $e_1 \in E_1$; $\iota_1(q \xrightarrow{t_1, e_1} q^{(1)} \xrightarrow{t_2, e_2} q^{(2)}) = q_1 \xrightarrow{t_1+t_2, e_2} q_1^{(2)}$ if $e_1 \in E_2$ and $e_2 \in E_1$, ... For each $i = 1, 2$, given a measurable property φ_i in \mathcal{A}_i from q_i , we write $\{(q_1, q_2) \models \tilde{\varphi}_i\}$ for the set $\iota_i^{-1}(\{q_i \models \varphi_i\})$. The following theorem states that the defined parallel composition is indeed interleaving.

Theorem 9.2.7. *Let $\mathcal{A}_1, \mathcal{A}_2 \in \text{CSTA}^*$ be composable. Then $\mathcal{A}_1 \parallel \mathcal{A}_2 \in \text{CSTA}^*$. Moreover, for every state $q = (q_1, q_2)$ of $\mathcal{A}_1 \parallel \mathcal{A}_2$, for every properties φ_1 measurable in \mathcal{A}_1 from q_1 and φ_2 measurable in \mathcal{A}_2 from q_2 , we have*

$$\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\tilde{\varphi}_1 \wedge \tilde{\varphi}_2) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\varphi_1) \cdot \text{Prob}_{q_2}^{\mathcal{A}_2}(\varphi_2). \quad (9.13)$$

Proof. Given \mathcal{A}_1 and $\mathcal{A}_2 \in \text{CSTA}^*$, thanks to Lemma 9.2.1, in order to get that $\mathcal{A}_1 \parallel \mathcal{A}_2 \in \text{CSTA}^*$, it suffices to prove that $\mathcal{A}_1 \parallel \mathcal{A}_2$ is almost-surely non-Zeno. This will be ensured by (9.13). We thus first tackle the proof of (9.13).

Let $q = (q_1, q_2) = ((l_1, \nu_1), (l_2, \nu_2))$ be a state of $\mathcal{A}_1 \parallel \mathcal{A}_2$. The important first step to prove (9.13) consists in showing that, given an edge e_1 of \mathcal{A}_1 , the probability in $\mathcal{A}_1 \parallel \mathcal{A}_2$ that e_1 is the first edge performed from $q = (q_1, q_2)$ in a given set of delays \mathcal{C} corresponds to the probability in \mathcal{A}_1 that e_1 is the first edge performed from q_1 in the same set of delays \mathcal{C} , that is for every Borel set \mathcal{C} of \mathbb{R}_+ ,

$$\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Cyl}(\pi_{\mathcal{C}^*}(q, \mathcal{A}_2^*, e_1))) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi_{\mathcal{C}}(q_1, e_1))) \quad (9.14)$$

⁷Observe that in particular the grammar defined in Section 4.1.1 and DMA like in Section 4.1.2, give rise to measurable properties as quickly discussed in Remark 7.1.1.

where $\text{Cyl}(\pi_{\mathcal{C}^*}(q, \mathcal{A}_2^*, e_1)) = \iota_1^{-1}(\text{Cyl}(\pi_{\mathcal{C}}(q_1, e_1)))$. Formally, we have that

$$\text{Cyl}(\pi_{\mathcal{C}^*}(q, \mathcal{A}_2^*, e_1)) = \bigcup_{n \in \mathbb{N}} \bigcup_{(f_1, \dots, f_n) \in E_2^n} \text{Cyl}(\pi_{\mathcal{C}_n}(q, f_1, \dots, f_n, e_1))$$

where $\mathcal{C}_n = \{(\tau_1, \dots, \tau_{n+1}) \in \mathbb{R}_+^{n+1} \mid \tau_1 + \dots + \tau_{n+1} \in \mathcal{C}\}$ for every $n \in \mathbb{N}$, which is a countable union of disjoint cylinders. In order to show (9.14), hypothesis (B) is crucial. Indeed, if for instance $e_1 \in E_1$ and $e_2 \in E_2$, then the projection of $q \xrightarrow{\tau_1, e_2} \cdot \xrightarrow{\tau_2, e_1} q'$ in \mathcal{A}_1 is $q_1 \xrightarrow{\tau_1} q_1 + \tau_1 \xrightarrow{\tau_2, e_1} q'_1$ which is equivalent to $q_1 \xrightarrow{\tau_1 + \tau_2, e_1} q'_1$: the first movement $q \xrightarrow{\tau_1, e_2} \cdot$ is in \mathcal{A}_2 and has no impact over \mathcal{A}_1 , except the elapse of τ_1 time units. Hypothesis (B) ensures that the probability in \mathcal{A}_1 to leave q_1 after $\tau_1 + \tau_2$ time units knowing that we leave it after at least τ_1 time units coincides with the probability to leave $q_1 + \tau_1$ after τ_2 time units. It is formalized in the next proposition (see [BBCM16]).

Proposition 9.2.8. *Assuming the above notations, for every $e_1 \in E_1$ and for every Borel set \mathcal{C} of \mathbb{R}_+ , we have that*

$$\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Cyl}(\pi_{\mathcal{C}^*}(q, \mathcal{A}_2^*, e_1))) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi_{\mathcal{C}}(q_1, e_1))). \quad (9.14)$$

Proof. We first assume that $\mathcal{C} = \mathbb{R}_+$. We have to show that for every $e_1 \in E_1$,

$$\sum_{n \geq 0} p_n(q) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi(q_1, e_1))) \quad (9.15)$$

where $p_n(q)$ is the probability of the set of infinite runs in $\mathcal{A}_1 \parallel \mathcal{A}_2$ that start in q and that first perform n switch-transitions in E_2 and then choose e_1 as the first edge of E_1 , *i.e.*

$$p_n(q) = \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\bigcup_{(f_1, \dots, f_n) \in E_2^n} \text{Cyl}(\pi(q, f_1, \dots, f_n, e_1)) \right).$$

In order to prove (9.15), we first show that for each $n \geq 0$,

$$\sum_{i=0}^{n-1} p_i(q) + p'_n(q) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi(q_1, e_1))),$$

where $p'_n(q)$ is the probability of the set of infinite runs in $\mathcal{A}_1 \parallel \mathcal{A}_2$ that start in q and that first perform n switch-transitions in E_2 and then choose e_1 as the first edge of E_1 , knowing that the $n + 1$ th transition is won with probability 1 by \mathcal{A}_1 . This is proved in Lemma 9.2.11. The key point of this result lies in the fact that $p_n(q)$ corresponds to the probability that \mathcal{A}_1 chooses first e_1 , \mathcal{A}_2 performs its n

first transitions before \mathcal{A}_1 performs e_1 and its $n + 1$ th transition must be taken after e_1 , while $p'_n(q)$ corresponds to the probability that \mathcal{A}_1 chooses first e_1 and \mathcal{A}_2 performs its n first transitions before \mathcal{A}_1 . It is formalized in Lemma 9.2.9 and 9.2.10. These lemmas will lead to the fact that $p_n(q) + p'_{n+1}(q) = p'_n(q)$ for every $n \geq 0$. We then show that $p'_n(q) \xrightarrow{n \rightarrow +\infty} 0$ which will conclude the proof of Proposition 9.2.8.

In the sequel, we keep the same notations as before for the density and cumulative functions of the probability measures $\mu_{q_i}^{(i)}$, and we refer to Definition 9.1.9 for the probability measures considered in the automaton $\mathcal{A}_1 \parallel \mathcal{A}_2$. Note that given a state q' of $\mathcal{A}_1 \parallel \mathcal{A}_2$, we will sometimes write $f_{q',i}$ instead of $f_{q'_i}$ (when this latter notation will appear to be too heavy). Given a state q of $\mathcal{A}_1 \parallel \mathcal{A}_2$, given an edge $e \in E_1$ and an edge $f \in E_2$, we abusively write $I(q, e)$ for $I(q_1, e)$ and $I(q, f)$ for $I(q_2, f)$. Let us recall that given edges f_1, \dots, f_n we write $q_{t_1 \dots t_n}$ for the state such that

$$q \xrightarrow{t_1, f_1} \cdot \xrightarrow{t_2, f_2} \cdot \dots \xrightarrow{t_n, f_n} q_{t_1 \dots t_n}$$

(see Section 3.1). Let us notice that if f_1, \dots, f_n are all in E_2 then the projection of $q_{t_1 \dots t_n}$ in \mathcal{A}_1 is given by $q_1 + t_1 + \dots + t_n$.

Lemma 9.2.9. Assuming the above notations, for every $n \geq 0$, we have

$$\begin{aligned} p_n(q) = & \sum_{(f_1, \dots, f_n) \in E_2^n} \int_{t \in I(q, e_1)} f_{q_1}(t) p_{q_1+t}^{(1)}(e_1) \int_{t_1=0}^t f_{q_2}(t_1) p_{q_2+t_1}^{(2)}(f_1) \mathbb{1}_{I(q, f_1)}(t_1) \\ & \int_{t_2=0}^{t-t_1} f_{q_{t_1,2}}(t_2) p_{q_{t_1+t_2}}^{(2)}(f_2) \mathbb{1}_{I(q_{t_1}, f_2)}(t_2) \dots \\ & \int_{t_n=0}^{t-t_1 \dots - t_{n-1}} f_{q_{t_1 \dots t_{n-1}, 2}}(t_n) p_{q_{t_1 \dots t_{n-1} + t_n}}^{(2)}(f_n) \mathbb{1}_{I(q_{t_1 \dots t_{n-1}}, f_n)}(t_n) \\ & (1 - F_{q_{t_1 \dots t_n}, 2}(t - t_1 - \dots - t_n)) dt_n \dots dt_2 dt_1 dt. \end{aligned} \tag{9.16}$$

Proof. We prove the lemma by induction over n . If $n = 0$, then we have from

Definition 9.1.9

$$\begin{aligned}
p_0(q) &= \mathbf{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Cyl}(\pi(q, e_1))) \\
&= \int_{t \in I(q, e_1)} f_q(t) p_{q+t}(e_1) dt \\
&= \int_{t \in I(q, e_1)} f_q(t) w_q^1(t) p_{q_1+t}^{(1)}(e_1) dt \\
&= \int_{t \in I(q, e_1)} f_q(t) \frac{f_{q_1}(t)(1 - F_{q_2}(t))}{f_q(t)} p_{q_1+t}^{(1)}(e_1) dt \\
&= \int_{t \in I(q, e_1)} f_{q_1}(t) p_{q_1+t}^{(1)}(e_1) (1 - F_{q_2}(t)) dt,
\end{aligned}$$

and thus equality (9.16) is satisfied for $n = 0$, for every state q of $\mathcal{A}_1 \parallel \mathcal{A}_2$. Now, let $n \geq 0$ and let us assume that (9.16) is verified for every state q of $\mathcal{A}_1 \parallel \mathcal{A}_2$ and for every $0 \leq k \leq n$. We now show that it is still the case for $k = n + 1$. Let q be a state of $\mathcal{A}_1 \parallel \mathcal{A}_2$, we have

$$p_{n+1}(q) = \mathbf{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\bigcup_{(f_1, \dots, f_{n+1}) \in E_2^{n+1}} \text{Cyl}(\pi(q, f_1, \dots, f_{n+1}, e_1)) \right)$$

and thus, since $\text{Cyl}(\pi(q, f_1, \dots, f_{n+1}, e_1)) \cap \text{Cyl}(\pi(q, f'_1, \dots, f'_{n+1}, e_1)) = \emptyset$ whenever $(f_1, \dots, f_{n+1}) \neq (f'_1, \dots, f'_{n+1})$, we have from Definition 9.1.9 that

$$\begin{aligned}
p_{n+1}(q) &= \sum_{(f_1, \dots, f_{n+1}) \in E_2^{n+1}} \mathbf{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Cyl}(\pi(q, f_1, \dots, f_{n+1}, e_1))) \\
&= \sum_{(f_1, \dots, f_{n+1}) \in E_2^{n+1}} \int_{t_1 \in I(q, f_1)} f_{q_2}(t_1) (1 - F_{q_1}(t_1)) p_{q_2+t_1}^{(2)}(f_1) \\
&\quad \mathbf{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Cyl}(\pi(q_{t_1}, f_2, \dots, f_{n+1}, e_1))) dt_1.
\end{aligned}$$

Now, since the value $f_{q_2}(t_1)(1 - F_{q_1}(t_1)) p_{q_2+t_1}^{(2)}(f_1)$ only depends on f_1 and since E_2 is a finite set, we have

$$\begin{aligned}
p_{n+1}(q) &= \sum_{f_1 \in E_2} \int_{t_1 \in I(q, f_1)} f_{q_2}(t_1) (1 - F_{q_1}(t_1)) p_{q_2+t_1}^{(2)}(f_1) \\
&\quad \left(\sum_{(f_2, \dots, f_{n+1}) \in E_2^n} \mathbf{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Cyl}(\pi(q_{t_1}, f_2, \dots, f_{n+1}, e_1))) \right) dt_1 \\
&= \sum_{f_1 \in E_2} \int_{t_1 \in I(q, f_1)} f_{q_2}(t_1) (1 - F_{q_1}(t_1)) p_{q_2+t_1}^{(2)}(f_1) p_n(q_{t_1}) dt_1.
\end{aligned}$$

Now from the hypothesis of induction we can compute $p_n(q_{t_1})$ with (9.16), and thus

$$\begin{aligned}
p_{n+1}(q) = & \sum_{(f_1, \dots, f_{n+1}) \in E_2^{n+1}} \int_{t_1 \in I(q, f_1)} f_{q_2}(t_1) (1 - F_{q_1}(t_1)) p_{q_2+t_1}^{(2)}(f_1) \\
& \int_{u \in I(q+t_1, e_1)} f_{q_1+t_1}(u) p_{q+t_1+u}^{(1)}(e_1) \int_{t_2=0}^u f_{q_{t_1}, 2}(t_2) p_{q_{t_1}+t_2}^{(2)}(f_2) \mathbb{1}_{I(q_{t_1}, f_2)}(t_2) \\
& \int_{t_3=0}^{u-t_2} f_{q_{t_1 t_2}, 2}(t_3) p_{q_{t_1 t_2}+t_3}^{(2)}(f_3) \mathbb{1}_{I(q_{t_1 t_2}, f_3)}(t_3) \dots \\
& \int_{t_{n+1}=0}^{u-t_2 \dots - t_n} f_{q_{t_1 \dots t_n}, 2}(t_{n+1}) p_{q_{t_1 \dots t_n}+t_{n+1}}^{(2)}(f_{n+1}) \mathbb{1}_{I(q_{t_1 \dots t_n}, f_{n+1})}(t_{n+1}) \\
& (1 - F_{q_{t_1 \dots t_{n+1}}, 2}(u - t_2 - \dots - t_{n+1})) dt_{n+1} \dots dt_3 dt_2 du dt_1
\end{aligned}$$

since the projection of q_{t_1} in \mathcal{A}_1 is $q + t_1$ (as q_{t_1} is such that $q \xrightarrow{t_1, f_1} q_{t_1}$ and $f_1 \in E_2$) and since E_2 is a finite set. Under the hypotheses over \mathcal{A}_1 and \mathcal{A}_2 , we know that $f_{q_1+t_1}(u)(1 - F_{q_1}(t_1)) = f_{q_1}(u + t_1)$. Now, if we let $t = u + t_1$, then $u = t - t_1$ and we have that for every $t_1 \geq 0$,

$$\begin{aligned}
u \in I(q + t_1, e_1) & \iff \nu_1 + t_1 + u \models g_{e_1} \\
& \iff \nu_1 + t \models g_{e_1} \quad \text{and} \quad t \geq t_1 \\
& \iff t \in I(q, e_1) \cap [t_1, +\infty[
\end{aligned}$$

where g_{e_1} denotes the guard of edge e_1 . From classical results of integration by substitution, we obtain that

$$\begin{aligned}
p_{n+1}(q) = & \sum_{(f_1, \dots, f_{n+1}) \in E_2^{n+1}} \int_{t_1 \in I(q, f_1)} f_{q_2}(t_1) p_{q_2+t_1}^{(2)}(f_1) \\
& \int_{t \in I(q, e_1)} f_{q_1}(t) p_{q_1+t}^{(1)}(e_1) \mathbb{1}_{[t_1, +\infty[}(t) \int_{t_2=0}^{t-t_1} f_{q_{t_1}, 2}(t_2) p_{q_{t_1}+t_2}^{(2)}(f_2) \mathbb{1}_{I(q_{t_1}, f_2)}(t_2) \\
& \int_{t_3=0}^{t-t_1-t_2} f_{q_{t_1 t_2}, 2}(t_3) p_{q_{t_1 t_2}+t_3}^{(2)}(f_3) \mathbb{1}_{I(q_{t_1 t_2}, f_3)}(t_3) \dots \\
& \int_{t_{n+1}=0}^{t-t_1 \dots - t_n} f_{q_{t_1 \dots t_n}, 2}(t_{n+1}) p_{q_{t_1 \dots t_n}+t_{n+1}}^{(2)}(f_{n+1}) \mathbb{1}_{I(q_{t_1 \dots t_n}, f_{n+1})}(t_{n+1}) \\
& (1 - F_{q_{t_1 \dots t_{n+1}}, 2}(t - t_1 - \dots - t_{n+1})) dt_{n+1} \dots dt_3 dt_2 dt dt_1.
\end{aligned}$$

Now using the fact that $\mathbb{1}_{[t_1, +\infty[}(t) = \mathbb{1}_{[0, t]}(t_1)$ and using Fubini's theorem, we deduce that (9.16) is satisfied for $n + 1$ which concludes the proof of the lemma. \square

Lemma 9.2.10. Assuming the above notations, for every $n \geq 0$, we have

$$\begin{aligned}
p'_n(q) = & \sum_{(f_1, \dots, f_n) \in E_2^n} \int_{t \in I(q, e_1)} f_{q_1}(t) p_{q_1+t}^{(1)}(e_1) \int_{t_1=0}^t f_{q_2}(t_1) p_{q_2+t_1}^{(2)}(f_1) \mathbb{1}_{I(q, f_1)}(t_1) \\
& \int_{t_2=0}^{t-t_1} f_{q_{t_1, 2}}(t_2) p_{q_{t_1+t_2}}^{(2)}(f_2) \mathbb{1}_{I(q_{t_1}, f_2)}(t_2) \dots \\
& \int_{t_n=0}^{t-t_1 \dots - t_{n-1}} f_{q_{t_1 \dots t_{n-1}, 2}}(t_n) p_{q_{t_1 \dots t_{n-1} + t_n}}^{(2)}(f_n) \mathbb{1}_{I(q_{t_1 \dots t_{n-1}}, f_n)}(t_n) \\
& dt_n \dots dt_2 dt_1 dt. \tag{9.17}
\end{aligned}$$

Proof. We recall that, from the above notations, for every $n \geq 0$, $p'_n(q)$ is the probability of the set of infinite runs in $\mathcal{A}_1 \parallel \mathcal{A}_2$ that start in q and that first perform n switch-transitions in E_2 and then choose e_1 in the case where, after that \mathcal{A}_2 has performed n transitions, then the $n+1$ th transition is won with probability 1 by \mathcal{A}_1 . In other words in the case where, for every run $\rho = q \xrightarrow{t_1, f_1} \dots \xrightarrow{t_2, f_2} \dots \xrightarrow{t_n, f_n} q_{t_1 \dots t_n}$ with $f_1, \dots, f_n \in E_2$, we have $w_{q_{t_1 \dots t_n}}^1(t) = 1$ for every $t \in I(q_{t_1 \dots t_n})$. Now, let s be an arbitrary state of $\mathcal{A}_1 \parallel \mathcal{A}_2$. From Definition 9.1.9, for every $t \in I(s)$, we have

$$\begin{aligned}
w_s^1(t) = 1 & \iff \frac{f_{s_1}(t)(1 - F_{s_2}(t))}{f_s(t)} = 1 \\
& \iff f_{s_1}(t)(1 - F_{s_2}(t)) = f_{s_1}(t)(1 - F_{s_2}(t)) + f_{s_2}(t)(1 - F_{s_1}(t)) \\
& \iff f_{s_2}(t)(1 - F_{s_1}(t)) = 0.
\end{aligned}$$

Thus, if for every $t \in I(s)$, $w_s^1(t) = 1$ then for every $t \in I(s)$, $f_{s_2}(t)(1 - F_{s_1}(t)) = 0$. We can then prove that $f_s(t) = f_{s_1}(t)$ almost-surely. Let us prove (9.17) when $n = 0$. Let q be a state of $\mathcal{A}_1 \parallel \mathcal{A}_2$, we have that $w_q^1(t) = 1$ for every $t \in I(q)$. And thus,

$$\begin{aligned}
p'_0(q) &= \int_{t \in I(q, e_1)} f_q(t) w_q^1(t) p_{q_1+t}^{(1)}(e_1) dt \\
&= \int_{t \in I(q, e_1)} f_{q_1}(t) p_{q_1+t}^{(1)}(e_1) dt.
\end{aligned}$$

Hence, (9.17) holds for every state q when $n = 0$. We can then prove by induction over n that if for every $0 \leq k \leq n$ with $n \geq 0$, (9.17) is satisfied for every state q , then it is still satisfied for $k = n+1$, for every state q . It uses similar arguments as in the proof of Lemma 9.2.9. This concludes the proof. \square

Lemma 9.2.11. Assuming the above notations, for every $n \geq 0$, we have

$$\sum_{i=0}^{n-1} p_i(q) + p'_n(q) = \text{Prob}_{q_1}^{A_1}(\text{Cyl}(\pi(q_1, e_1))).$$

Proof. If $n = 0$, we have from Lemma 9.2.10 that $p'_0(q) = \text{Prob}_{q_1}^{A_1}(\text{Cyl}(\pi(q_1, e_1)))$. Now, let us assume that for every $0 \leq k \leq n$ with $n \geq 0$, we have

$$\sum_{i=0}^{k-1} p_i(q) + p'_k(q) = \text{Prob}_{q_1}^{A_1}(\text{Cyl}(\pi(q_1, e_1))),$$

and let us prove that it is still the case when $k = n + 1$. First, let us compute $p'_{n+1}(q)$. From Lemma 9.2.10, we have

$$\begin{aligned} p'_{n+1}(q) = & \sum_{(f_1, \dots, f_{n+1}) \in E_2^{n+1}} \int_{t \in I(q, e_1)} f_{q_1}(t) p_{q_1+t}^{(1)}(e_1) \int_{t_1=0}^t f_{q_2}(t_1) p_{q_2+t_1}^{(2)}(f_1) \mathbb{1}_{I(q, f_1)}(t_1) \\ & \int_{t_2=0}^{t-t_1} f_{q_{t_1,2}}(t_2) p_{q_{t_1+t_2}}^{(2)}(f_2) \mathbb{1}_{I(q_{t_1}, f_2)}(t_2) \dots \\ & \int_{t_n=0}^{t-t_1 \dots - t_{n-1}} f_{q_{t_1 \dots t_{n-1}, 2}}(t_n) p_{q_{t_1 \dots t_{n-1} + t_n}}^{(2)}(f_n) \mathbb{1}_{I(q_{t_1 \dots t_{n-1}}, f_n)}(t_n) \\ & \int_{t_{n+1}=0}^{t-t_1 \dots - t_n} f_{q_{t_1 \dots t_n, 2}}(t_{n+1}) p_{q_{t_1 \dots t_n + t_{n+1}}}^{(2)}(f_{n+1}) \mathbb{1}_{I(q_{t_1 \dots t_n}, f_{n+1})}(t_{n+1}) \\ & dt_{n+1} dt_n \dots dt_2 dt_1 dt. \end{aligned}$$

Now, one can observe that only $p_{q_{t_1 \dots t_n + t_{n+1}}}^{(2)}(f_{n+1}) \mathbb{1}_{I(q_{t_1 \dots t_n}, f_{n+1})}(t_{n+1})$ depends on f_{n+1} in the last integral. Thus, since E_2 is finite, we have

$$\begin{aligned} p'_{n+1}(q) = & \sum_{(f_1, \dots, f_n) \in E_2^n} \int_{t \in I(q, e_1)} f_{q_1}(t) p_{q_1+t}^{(1)}(e_1) \int_{t_1=0}^t f_{q_2}(t_1) p_{q_2+t_1}^{(2)}(f_1) \mathbb{1}_{I(q, f_1)}(t_1) \\ & \int_{t_2=0}^{t-t_1} f_{q_{t_1,2}}(t_2) p_{q_{t_1+t_2}}^{(2)}(f_2) \mathbb{1}_{I(q_{t_1}, f_2)}(t_2) \dots \\ & \int_{t_n=0}^{t-t_1 \dots - t_{n-1}} f_{q_{t_1 \dots t_{n-1}, 2}}(t_n) p_{q_{t_1 \dots t_{n-1} + t_n}}^{(2)}(f_n) \mathbb{1}_{I(q_{t_1 \dots t_{n-1}}, f_n)}(t_n) \\ & \sum_{f_{n+1} \in E_2} \int_{t_{n+1}=0}^{t-t_1 \dots - t_n} f_{q_{t_1 \dots t_n, 2}}(t_{n+1}) p_{q_{t_1 \dots t_n + t_{n+1}}}^{(2)}(f_{n+1}) \mathbb{1}_{I(q_{t_1 \dots t_n}, f_{n+1})}(t_{n+1}) \\ & dt_{n+1} dt_n \dots dt_2 dt_1 dt. \end{aligned}$$

Now, we have that

$$\begin{aligned}
& \sum_{f_{n+1} \in E_2} \int_{t_{n+1}=0}^{t-t_1 \dots -t_n} f_{q_{t_1 \dots t_n}, 2}(t_{n+1}) p_{q_{t_1 \dots t_n} + t_{n+1}}^{(2)}(f_{n+1}) \mathbb{1}_{I(q_{t_1 \dots t_n}, f_{n+1})}(t_{n+1}) \\
&= \int_{t_{n+1}=0}^{t-t_1 \dots -t_n} f_{q_{t_1 \dots t_n}, 2}(t_{n+1}) \\
&\quad \sum_{f_{n+1} \in E_2} \left(p_{q_{t_1 \dots t_n} + t_{n+1}}^{(2)}(f_{n+1}) \mathbb{1}_{I(q_{t_1 \dots t_n}, f_{n+1})}(t_{n+1}) \right) dt_{n+1} \\
&= \int_{t_{n+1}=0}^{t-t_1 \dots -t_n} f_{q_{t_1 \dots t_n}, 2}(t_{n+1}) \mathbb{1}_{I(q_{t_1 \dots t_n})}(t_{n+1}) \\
&= F_{q_{t_1 \dots t_n}, 2}(t - t_1 - \dots - t_n)
\end{aligned}$$

since $p_{q_{t_1 \dots t_n} + t_{n+1}}^{(2)}$ is a probability measure over the enabled edges in $q_{t_1 \dots t_n} + t_{n+1}$ when this set is not empty (otherwise, we assume that $p_{q_{t_1 \dots t_n} + t_{n+1}}^{(2)}$ is a function that assigns 0 to each edge of E_2). We deduce thus that

$$\begin{aligned}
p'_{n+1}(q) &= \sum_{(f_1, \dots, f_n) \in E_2^n} \int_{t \in I(q, e_1)} f_{q_1}(t) p_{q_1+t}^{(1)}(e_1) \int_{t_1=0}^t f_{q_2}(t_1) p_{q_2+t_1}^{(2)}(f_1) \mathbb{1}_{I(q, f_1)}(t_1) \\
&\quad \int_{t_2=0}^{t-t_1} f_{q_{t_1}, 2}(t_2) p_{q_{t_1}+t_2}^{(2)}(f_2) \mathbb{1}_{I(q_{t_1}, f_2)}(t_2) \dots \\
&\quad \dots \int_{t_n=0}^{t-t_1 \dots -t_{n-1}} f_{q_{t_1 \dots t_{n-1}}, 2}(t_n) p_{q_{t_1 \dots t_{n-1}} + t_n}^{(2)}(f_n) \mathbb{1}_{I(q_{t_1 \dots t_{n-1}}, f_n)}(t_n) \\
&\quad F_{q_{t_1 \dots t_n}, 2}(t - t_1 - \dots - t_n) dt_n \dots dt_2 dt_1 dt.
\end{aligned}$$

From this last equality and Lemmas 9.2.9 and 9.2.10, we can thus easily see that $p_n(q) + p'_{n+1}(q) = p'_n(q)$. From this last equality and the hypothesis of induction, we have

$$\begin{aligned}
\sum_{i=0}^n p_i(q) + p'_{n+1}(q) &= \sum_{i=0}^{n-1} p_i(q) + p_n(q) + p'_{n+1}(q) \\
&= \sum_{i=0}^{n-1} p_i(q) + p'_n(q) \\
&= \text{Prob}_{q_1}^{A_1}(\text{Cyl}(\pi(q_1, e_1)))
\end{aligned}$$

which concludes the proof. \square

Lemma 9.2.12. Assuming the above notations, we have that $p'_n(q) \xrightarrow{n \rightarrow +\infty} 0$.

Proof. Since E_2 is finite we have that

$$\begin{aligned} p'_n(q) = & \int_{t \in I(q, e_1)} f_{q_1}(t) p_{q_1+t}^{(1)}(e_1) \\ & \sum_{(f_1, \dots, f_n) \in E_2^n} \int_{t_1=0}^t f_{q_2}(t_1) p_{q_2+t_1}^{(2)}(f_1) \mathbb{1}_{I(q, f_1)}(t_1) \\ & \int_{t_2=0}^{t-t_1} f_{q_{t_1, 2}}(t_2) p_{q_{t_1}+t_2}^{(2)}(f_2) \mathbb{1}_{I(q_{t_1}, f_2)}(t_2) \dots \\ & \int_{t_n=0}^{t-t_1 \dots - t_{n-1}} f_{q_{t_1 \dots t_{n-1}, 2}}(t_n) p_{q_{t_1 \dots t_{n-1}}+t_n}^{(2)}(f_n) \mathbb{1}_{I(q_{t_1 \dots t_{n-1}}, f_n)}(t_n) \\ & dt_n \dots dt_2 dt_1 dt. \end{aligned}$$

We can then write $p'_n(q) = \int_{t \in I(q, e_1)} f(t) g_n(t) dt$, where $f(t) = f_{q_1}(t) p_{q_1+t}^{(1)}(e_1)$ and

$$\begin{aligned} g_n(t) = & \sum_{(f_1, \dots, f_n) \in E_2^n} \int_{t_1=0}^t f_{q_2}(t_1) p_{q_2+t_1}^{(2)}(f_1) \mathbb{1}_{I(q, f_1)}(t_1) \\ & \int_{t_2=0}^{t-t_1} f_{q_{t_1, 2}}(t_2) p_{q_{t_1}+t_2}^{(2)}(f_2) \mathbb{1}_{I(q_{t_1}, f_2)}(t_2) \dots \\ & \int_{t_n=0}^{t-t_1 \dots - t_{n-1}} f_{q_{t_1 \dots t_{n-1}, 2}}(t_n) p_{q_{t_1 \dots t_{n-1}}+t_n}^{(2)}(f_n) \mathbb{1}_{I(q_{t_1 \dots t_{n-1}}, f_n)}(t_n) dt_n \dots dt_1. \end{aligned}$$

Under the hypotheses over \mathcal{A}_1 and \mathcal{A}_2 (e.g. $\mathcal{A}_1, \mathcal{A}_2 \in \text{CSTA}^*$), we know that \mathcal{A}_2 is almost-surely non-Zeno. Let $t \geq 0$ and let $M \in \mathbb{N}$ such that $t \leq M$. From the discussion over the set of Zeno runs in Section 3.1 on page 64, we have that

$$g_n(t) \leq \sum_{(f_1, \dots, f_n) \in E_2^n} \text{Prob}_{q_2}^{\mathcal{A}_2}(\pi_{\mathcal{C}_{M,n}}(q, f_1, \dots, f_n)) \xrightarrow{n \rightarrow +\infty} 0 \quad (9.18)$$

and thus, $\lim_n g_n(t) = 0$ for every $t \geq 0$, since $g_n(t) \geq 0$ for every $t \geq 0$. Then, we have that $\lim_n f(t) g_n(t) = 0$ for every $t \geq 0$. Now, from inequality (9.18), we have that $g_n(t) \leq 1$ for every $t \geq 0$ and thus, $f(t) g_n(t) \leq f(t)$ for every $t \geq 0$. And since

$$\int_{t \in I(q, e_1)} f(t) = \int_{t \in I(q, e_1)} f_{q_1}(t) p_{q_1+t}^{(1)}(e_1) dt \leq \int_{\mathbb{R}_+} f_{q_1}(t) dt = 1,$$

we have, by dominated convergence, that

$$\lim_{n \rightarrow +\infty} p'_n(q) = \lim_{n \rightarrow +\infty} \int_{t \in I(q, e_1)} f(t) g_n(t) dt = \int_{t \in I(q, e_1)} \lim_{n \rightarrow +\infty} (f(t) g_n(t)) dt = 0$$

which concludes the proof. \square

We can now prove that (9.14) holds when $\mathcal{C} = \mathbb{R}_+$. We have

$$\begin{aligned} & \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Cyl}(\pi(q, \mathcal{A}_2^*, e_1))) \\ &= \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\bigcup_{n \geq 0} \bigcup_{(f_1, \dots, f_n) \in E_2^n} \text{Cyl}(\pi(f_1, \dots, f_n, e_1)) \right) \\ &= \sum_{n \geq 0} p_n(q) \end{aligned} \tag{9.19}$$

from the definition of $p_n(q)$. Now, from Lemma 9.2.11, we have that for every $n \geq 0$,

$$\sum_{i=0}^n p_i(q) + p'_{n+1}(q) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi(q_1, e_1)))$$

and thus,

$$\lim_{n \rightarrow +\infty} \left(\sum_{i=0}^n p_i(q) + p'_{n+1}(q) \right) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi(q_1, e_1))).$$

Hence, since $\lim_n p'_{n+1}(q) = 0$ from Lemma 9.2.12, we have $\sum_{n \geq 0} p_n(q) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi(q_1, e_1)))$ and we deduce from (9.19) that

$$\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Cyl}(\pi(q, \mathcal{A}_2^*, e_1))) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi(q_1, e_1))).$$

Now, we would like to get (9.14) for every Borel set \mathcal{C} of \mathbb{R}_+ . Given a Borel set \mathcal{C} , a reasoning similar to the ones in Lemmas 9.2.9, 9.2.10, 9.2.11 and 9.2.12 can be applied to $p_{n,\mathcal{C}}(q)$ and $p'_{n,\mathcal{C}}(q)$ where

$$p_{n,\mathcal{C}}(q) = \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\bigcup_{(f_1, \dots, f_n)} \text{Cyl}(\pi_{\mathcal{C}_n}(q, f_1, \dots, f_n, e_1)) \right),$$

with $\mathcal{C}_n = \{(t_1, \dots, t_{n+1}) \in \mathbb{R}_+^{n+1} \mid t_1 + \dots + t_{n+1} \in \mathcal{C}\}$, is the probability of the set of infinite runs in $\mathcal{A}_1 \parallel \mathcal{A}_2$ that start in q and that first perform n switch-transitions in E_2 , then choose e_1 and all these transitions are taken in a delay that is in \mathcal{C} , and $p'_n(q)$ is the probability of the same set in the case where, after that \mathcal{A}_2 has performed n transitions, then the $n + 1$ th transition is won with probability 1 by \mathcal{A}_1 . We can then prove, as in Lemmas 9.2.9 and 9.2.10, that for

every state q , for every Borel set \mathcal{C} and for every $n \geq 0$,

$$\begin{aligned}
p_{n,\mathcal{C}}(q) = & \sum_{(f_1, \dots, f_n) \in E_2^n} \int_{t \in I(q, e_1)} f_{q_1}(t) p_{q_1+t}^{(1)}(e_1) \mathbb{1}_{\mathcal{C}}(t) \\
& \int_{t_1=0}^t f_{q_2}(t_1) p_{q_2+t_1}^{(2)}(f_1) \mathbb{1}_{I(q, f_1)}(t_1) \\
& \int_{t_2=0}^{t-t_1} f_{q_{t_1,2}}(t_2) p_{q_{t_1+t_2}}^{(2)}(f_2) \mathbb{1}_{I(q_{t_1}, f_2)}(t_2) \dots \\
& \dots \int_{t_n=0}^{t-t_1 \dots - t_{n-1}} f_{q_{t_1 \dots t_{n-1}, 2}}(t_n) p_{q_{t_1 \dots t_{n-1} + t_n}}^{(2)}(f_n) \mathbb{1}_{I(q_{t_1 \dots t_{n-1}}, f_n)}(t_n) \\
& (1 - F_{q_{t_1 \dots t_n}, 2}(t - t_1 - \dots - t_n)) dt_n \dots dt_2 dt_1 dt
\end{aligned} \tag{9.20}$$

and

$$\begin{aligned}
p'_{n,\mathcal{C}}(q) = & \sum_{(f_1, \dots, f_n) \in E_2^n} \int_{t \in I(q, e_1)} f_{q_1}(t) p_{q_1+t}^{(1)}(e_1) \mathbb{1}_{\mathcal{C}}(t) \\
& \int_{t_1=0}^t f_{q_2}(t_1) p_{q_2+t_1}^{(2)}(f_1) \mathbb{1}_{I(q, f_1)}(t_1) \\
& \int_{t_2=0}^{t-t_1} f_{q_{t_1,2}}(t_2) p_{q_{t_1+t_2}}^{(2)}(f_2) \mathbb{1}_{I(q_{t_1}, f_2)}(t_2) \dots \\
& \int_{t_n=0}^{t-t_1 \dots - t_{n-1}} f_{q_{t_1 \dots t_{n-1}, 2}}(t_n) p_{q_{t_1 \dots t_{n-1} + t_n}}^{(2)}(f_n) \mathbb{1}_{I(q_{t_1 \dots t_{n-1}}, f_n)}(t_n) \\
& dt_n \dots dt_2 dt_1 dt.
\end{aligned} \tag{9.21}$$

It can be proved by induction over n as in the previous lemmas, by noticing that

$$p_{n+1,\mathcal{C}}(q) = \sum_{f_1 \in E_2} \int_{t_1 \in I(q, f_1)} f_{q_2}(t_1) (1 - F_{q_1}(t_1)) p_{q_2+t_1}^{(2)}(f_1) p_{n,\mathcal{C}_{t_1}}(q_{t_1}) dt_1 \tag{9.22}$$

where \mathcal{C}_{t_1} is a notation for $(\mathcal{C} - t_1) \cap \mathbb{R}_+$ and $(\mathcal{C} - t_1) = \{t - t_1 \mid t \in \mathcal{C}\}$, which is a Borel set. Indeed, we have that

$$\begin{aligned}
p_{n+1,\mathcal{C}}(q) = & \sum_{(f_1, \dots, f_{n+1}) \in E_2^{n+1}} \text{Prob}_q^{A_1 \| A_2}(\text{Cyl}(\pi_{\mathcal{C}_{n+1}}(q, f_1, \dots, f_{n+1}, e_1))) \\
= & \sum_{(f_1, \dots, f_{n+1}) \in E_2^{n+1}} \int_{t_1 \in I(q, f_1)} f_{q_2}(t_1) (1 - F_{q_1}(t_1)) p_{q_2+t_1}^{(2)}(f_1) \\
& \text{Prob}_q^{A_1 \| A_2}(\text{Cyl}(\pi_{\mathcal{C}_{t_1, n}}(q_{t_1}, f_2, \dots, f_{n+1}, e_1))) dt_1
\end{aligned} \tag{9.23}$$

where for every $t_1 \geq 0$, $\mathcal{C}_{t_1, n} = \{(t_2, \dots, t_{n+2}) \in \mathbb{R}_+^{n+1} \mid t_1 + t_2 + \dots + t_{n+2} \in \mathcal{C}\}$. Since we have that for every $t_1 \geq 0$, $(t_1, \dots, t_{n+2}) \in \mathcal{C}_{n+1}$ if and only if $(t_2, \dots, t_{n+2}) \in \mathcal{C}_{t_1, n}$ where $\mathcal{C}_{t_1, n} = \{(t_2, \dots, t_{n+2}) \in \mathbb{R}_+^{n+1} \mid (t_1, t_2, \dots, t_{n+2}) \in \mathcal{C}\}$, we obtain

$$\mathbb{1}_{\mathcal{C}_n}((t_1, \dots, t_{n+2})) = \mathbb{1}_{\mathcal{C}_{t_1, n}}((t_2, \dots, t_{n+2})).$$

In our case, we can see that $\mathcal{C}_{t_1, n} = \{(t_2, \dots, t_{n+2}) \in \mathbb{R}_+^{n+1} \mid t_2 + \dots + t_{n+2} \in (\mathcal{C} - t_1) \cap \mathbb{R}_+\}$. We thus have (9.23) and as in Lemma 9.2.9, we deduce (9.22). If we assume by induction that for every state q and for every Borel set \mathcal{C} , (9.20) holds, then we have

$$\begin{aligned} p_{n+1, \mathcal{C}}(q) &= \sum_{f_1 \in E_2} \int_{t_1 \in I(q, f_1)} f_{q_2}(t_1) (1 - F_{q_1}(t_1)) p_{q_2+t_1}^{(2)}(f_1) \\ &\sum_{(f_2, \dots, f_{n+1}) \in E_2^n} \int_{u \in I(q+t_1, e_1)} f_{q_1+t_1}(u) p_{q+t_1+u}^{(1)}(e_1) \mathbb{1}_{\mathcal{C}_{t_1}}(u) \\ &\int_{t_2=0}^u f_{q_{t_1, 2}}(t_2) p_{q_{t_1}+t_2}^{(2)}(f_2) \mathbb{1}_{I(q_{t_1}, f_2)}(t_2) \\ &\int_{t_3=0}^{u-t_2} f_{q_{t_1 t_2, 2}}(t_3) p_{q_{t_1 t_2}+t_3}^{(2)}(f_3) \mathbb{1}_{I(q_{t_1 t_2}, f_3)}(t_3) \dots \\ &\int_{t_{n+1}=0}^{u-t_2 \dots -t_n} f_{q_{t_1 \dots t_n, 2}}(t_{n+1}) p_{q_{t_1 \dots t_n}+t_{n+1}}^{(2)}(f_{n+1}) \mathbb{1}_{I(q_{t_1 \dots t_n}, f_{n+1})}(t_{n+1}) \\ &(1 - F_{q_{t_1 \dots t_{n+1}, 2}}(u - t_2 - \dots - t_{n+1})) dt_{n+1} \dots dt_2 du dt_1. \end{aligned}$$

Now, since $\mathcal{C}_{t_1} = (\mathcal{C} - t_1) \cap \mathbb{R}_+$ and since $u \geq 0$ for each $u \in I(q + t_1, e_1)$, we have that $\mathbb{1}_{\mathcal{C}_{t_1}}(u) = \mathbb{1}_{\mathcal{C}}(t_1 + u)$. Hence, by a substitution as in the proof of Lemma 9.2.9 and by Fubini's theorem, we obtain that (9.20) holds for $n + 1$. Now, following the same reasoning as in Lemmas 9.2.10, 9.2.11 and 9.2.12, it is easy to see that (9.21) holds, that for every $n \geq 0$,

$$\sum_{i=0}^{n-1} p_{i, \mathcal{C}}(q) + p'_{n, \mathcal{C}}(q) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi_{\mathcal{C}}(q_1, e_1))),$$

that $\lim_n p'_{n, \mathcal{C}}(q) = 0$ and thus that

$$\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Cyl}(\pi_{\mathcal{C}^*}(q, \mathcal{A}_2^*, e_1))) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi_{\mathcal{C}}(q_1, e_1))),$$

which concludes the proof of Proposition 9.2.8. \square

We can extend this result to the case where the n first movements in \mathcal{A}_1 are determined. For this, we write

$$\text{Cyl}(\pi_{\mathcal{C}^*}(q, \mathcal{A}_2^*, e_1, \dots, \mathcal{A}_2^*, e_n)) = \iota_1^{-1}(\text{Cyl}(\pi_{\mathcal{C}}(q_1, e_1, \dots, e_n)))$$

where \mathcal{C} is a Borel set of \mathbb{R}_+^n .

Proposition 9.2.13. *Assuming the previous notations, for every $n \geq 1$, for every $e_1, \dots, e_n \in E_1$ and for every Borel set \mathcal{C} of \mathbb{R}_+^n , we have*

$$\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Cyl}(\pi_{\mathcal{C}^*}(q, \mathcal{A}_2^*, e_1, \dots, \mathcal{A}_2^*, e_n))) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Cyl}(\pi_{\mathcal{C}}(q_1, e_1, \dots, e_n))).$$

We do not give the proof here as it is similar as the proof of Proposition 9.2.8. Now, this result can be extended to the elements of the σ -algebra:

Proposition 9.2.14. *Assuming the above notations, for every property φ_1 measurable in \mathcal{A}_1 from q_1 , we have*

$$\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\tilde{\varphi}_1) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\varphi_1).$$

Proof. The proof is immediate from Proposition 9.2.13 by noticing that, given a stochastic timed automaton \mathcal{A} , the complementary of a cylinder is a countable union of cylinders, that the union of two cylinders $\text{Cyl}(\pi_{\mathcal{C}_1}(q, e_1, \dots, e_n))$ and $\text{Cyl}(\pi_{\mathcal{C}_2}(q, e_1, \dots, e_m))$ with $n \leq m$ can be rewritten as

$$\text{Cyl}(\pi_{\mathcal{C}_1}(q, e_1, \dots, e_n)) \cup \text{Cyl}(\pi_{\mathcal{C}_2 \setminus \mathcal{C}_2^{(n)}}(q, e_1, \dots, e_m))$$

where $\mathcal{C}_2^{(n)} = \{(t_1, \dots, t_m) \in \mathcal{C}_2 \mid (t_1, \dots, t_n) \in \mathcal{C}_1\}$, which is the union of two disjoint cylinders, and by noticing that for every sequence $(A_n)_{n \geq 0} \subseteq \Omega_{\mathcal{A}}^q$ and for every $A \in \Omega_{\mathcal{A}}^q$, we have

$$\iota^{-1}\left(\bigcup_{n \geq 0} A_n\right) = \bigcup_{n \geq 0} \iota^{-1}(A_n) \quad \text{and} \quad \iota^{-1}(A^c) = \iota^{-1}(A)^c.$$

□

Similar results as Propositions 9.2.8, 9.2.13 and 9.2.14 hold when we alternate \mathcal{A}_1 and \mathcal{A}_2 . These propositions will lead to (9.13) but before getting to that, we need an extra notion.

Definition 9.2.15. Let e_1, \dots, e_n (resp. f_1, \dots, f_m) be edges of \mathcal{A}_1 (resp. \mathcal{A}_2) and let \mathcal{C}_1 (resp. \mathcal{C}_2) be Borel sets of \mathbb{R}_+^n (resp. \mathbb{R}_+^m). We define the *shuffle* of the cylinders $\text{Cyl}(\pi_{\mathcal{C}_1}(q_1, e_1, \dots, e_n))$ and $\text{Cyl}(\pi_{\mathcal{C}_2}(q_2, f_1, \dots, f_m))$ as the following set of runs of $\mathcal{A}_1 \parallel \mathcal{A}_2$:

$$\{\rho \in \text{Runs}(\mathcal{A}_1 \parallel \mathcal{A}_2, (q_1, q_2)) \mid \iota_1(\rho) \in \text{Cyl}(\pi_{\mathcal{C}_1}(q_1, e_1, \dots, e_n)) \\ \wedge \iota_2(\rho) \in \text{Cyl}(\pi_{\mathcal{C}_2}(q_2, f_1, \dots, f_m))\}.$$

We denote this set by $\text{Cyl}(\pi_{\mathcal{C}_1}(q_1, e_1, \dots, e_n)) \sqcup \text{Cyl}(\pi_{\mathcal{C}_2}(q_2, f_1, \dots, f_m))$.

Remark 9.2.16. The shuffle of two cylinders can be rewritten as a union of disjoint cylinders. As a simple example, assuming the same notations of Definition 9.2.15, we have that

$$\begin{aligned} \text{Cyl}(\pi(q_1, e_1, e_2)) \sqcup \text{Cyl}(\pi(q_2, f_1)) &= \text{Cyl}(\pi(q, f_1, \mathcal{A}_2^*, e_1, \mathcal{A}_2^*, e_2)) \\ &\cup \text{Cyl}(\pi(q, e_1, f_1, \mathcal{A}_2^*, e_2)) \cup \text{Cyl}(\pi(q, e_1, e_2, \mathcal{A}_1^*, f_1)). \end{aligned}$$

Let us also remark that Definition 9.2.15 trivially extends to sets of the σ -algebras $\Omega_{q_1}^{\mathcal{A}_1}$ and $\Omega_{q_2}^{\mathcal{A}_2}$. Hence we can notice that, given two properties φ_1 and φ_2 measurable in \mathcal{A}_1 and \mathcal{A}_2 ,

$$\{q \models \tilde{\varphi}_1\} \cap \{q \models \tilde{\varphi}_2\} = \{q_1 \models \varphi_1\} \sqcup \{q_2 \models \varphi_2\}.$$

We are now able to prove (9.13) that is

$$\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\tilde{\varphi}_1 \wedge \tilde{\varphi}_2) = \text{Prob}_{q_1}^{\mathcal{A}_1}(\varphi_1) \cdot \text{Prob}_{q_2}^{\mathcal{A}_2}(\varphi_2), \quad (9.13)$$

where φ_1 (resp. φ_2) is a property measurable in \mathcal{A}_1 (resp. \mathcal{A}_2) from q_1 (resp. q_2). As in Proposition 9.2.14, since for every properties φ_1 measurable in \mathcal{A}_1 from q_1 and φ_2 measurable in \mathcal{A}_2 from q_2 , we have

$$\{q \models \tilde{\varphi}_1\} \cap \{q \models \tilde{\varphi}_2\} = \{q_1 \models \varphi_1\} \sqcup \{q_2 \models \varphi_2\},$$

and $\{q_i \models \varphi_i\} \in \Omega_{q_i}^{\mathcal{A}_i}$ for $i \in \{1, 2\}$, it suffices to prove that for every $n, m \geq 0$, for every e_1, \dots, e_n edges of \mathcal{A}_1 , for every f_1, \dots, f_m edges of \mathcal{A}_2 , for every borel sets \mathcal{C}_1 of \mathbb{R}_+^n and for every borel sets \mathcal{C}_2 of \mathbb{R}_+^m ,

$$\begin{aligned} &\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_1}(q_1, e_1, \dots, e_n)) \sqcup \text{Cyl}(\pi_{\mathcal{C}_2}(q_2, f_1, \dots, f_m)) \right) = \\ &\text{Prob}_{q_1}^{\mathcal{A}_1} \left(\text{Cyl}(\pi_{\mathcal{C}_1}(q_1, e_1, \dots, e_n)) \right) \cdot \text{Prob}_{q_2}^{\mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_2}(q_2, f_1, \dots, f_m)) \right). \end{aligned} \quad (9.24)$$

We prove it by induction over (n, m) . We can first notice that if $n = 0$, then for every $m \geq 0$,

$$\text{Cyl}(\pi(q_1)) \sqcup \text{Cyl}(\pi_{\mathcal{C}_2}(q_2, f_1, \dots, f_m)) = \text{Cyl}(\pi_{\mathcal{C}_2^*}(q, \mathcal{A}_1^*, f_1, \dots, \mathcal{A}_1^*, f_m))$$

and thus (9.24) holds from Proposition 9.2.13. Hence, equality (9.24) holds for each $(0, m)$ with $m \geq 0$. It can similarly be shown that the property holds for each $(n, 0)$ with $n \geq 0$. Now, we fix $n, m \geq 0$ and we assume that (9.24) is satisfied for every (n', m') with $0 \leq n' \leq n$ and $0 \leq m' \leq m$, and we prove that

it is still verified for $(n + 1, m + 1)$. Let e_1, \dots, e_{n+1} and f_1, \dots, f_{m+1} be edges of \mathcal{A}_1 and \mathcal{A}_2 , and let \mathcal{C}_1 and \mathcal{C}_2 be Borel sets of \mathbb{R}_+^{n+1} and \mathbb{R}_+^{m+1} . Then

$$\begin{aligned}
& \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_1}(q_1, e_1, \dots, e_{n+1})) \sqcup \text{Cyl}(\pi_{\mathcal{C}_2}(q_2, f_1, \dots, f_{m+1})) \right) = \\
& \int_{t_1 \in I(q, e_1)} p_{q_1+t_1}^{(1)}(e_1) f_{q_1}(t_1) (1 - F_{q_2}(t_1)) \\
& \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_1^{t_1}}(q_{t_1}, e_2, \dots, e_{n+1})) \sqcup \text{Cyl}(\pi_{\mathcal{C}_2^{t_1,+}}(q_2 + t_1, f_1, \dots, f_{m+1})) \right) dt_1 \\
& + \int_{t_2 \in I(q, f_1)} p_{q_2+t_2}^{(2)}(f_1) f_{q_2}(t_2) (1 - F_{q_1}(t_2)) \\
& \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_1^{t_2,+}}(q_1 + t_2, e_1, \dots, e_{n+1})) \sqcup \text{Cyl}(\pi_{\mathcal{C}_2^{t_2}}(q_{t_2}, f_2, \dots, f_{m+1})) \right) dt_2,
\end{aligned} \tag{9.25}$$

where $\mathcal{C}_1^{t_1} = \{(\tau_2, \dots, \tau_{n+1}) \mid (t_1, \tau_2, \dots, \tau_{n+1}) \in \mathcal{C}_1\}$ which is a Borel set of \mathbb{R}_+^n and $\mathcal{C}_1^{t_2,+} = \{(\tau_1, \dots, \tau_{n+1}) \mid (t_2 + \tau_1, \tau_2, \dots, \tau_{n+1}) \in \mathcal{C}_1\}$ which is a Borel set of \mathbb{R}_+^{n+1} . This is similar with \mathcal{C}_2 . Now, by induction hypothesis, we obtain that

$$\begin{aligned}
& \int_{t_1 \in I(q, e_1)} p_{q_1+t_1}^{(1)}(e_1) f_{q_1}(t_1) (1 - F_{q_2}(t_1)) \\
& \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_1^{t_1}}(q_{t_1}, e_2, \dots, e_{n+1})) \sqcup \text{Cyl}(\pi_{\mathcal{C}_2^{t_1,+}}(q_2 + t_1, f_1, \dots, f_{m+1})) \right) dt_1 \\
& = \int_{t_1 \in I(q, e_1)} p_{q_1+t_1}^{(1)}(e_1) f_{q_1}(t_1) (1 - F_{q_2}(t_1)) \text{Prob}_{q_1}^{\mathcal{A}_1} \left(\text{Cyl}(\pi_{\mathcal{C}_1^{t_1}}(q_{t_1}, e_2, \dots, e_{n+1})) \right) \\
& \quad \text{Prob}_{q_2}^{\mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_2^{t_1,+}}(q_2 + t_1, f_1, \dots, f_{m+1})) \right) dt_1 \\
& = \int_{t_1 \in I(q, e_1)} p_{q_1+t_1}^{(1)}(e_1) f_{q_1}(t_1) (1 - F_{q_2}(t_1)) \text{Prob}_{q_1}^{\mathcal{A}_1} \left(\text{Cyl}(\pi_{\mathcal{C}_1^{t_1}}(q_{t_1}, e_2, \dots, e_{n+1})) \right) \\
& \quad \int_{u \in I(q_2+t_1, f_1)} p_{q_2+t_1+u}^{(2)}(f_1) f_{q_2+t_1}(u) \text{Prob}_{q_2}^{\mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_2^{t_1+u}}(q_{t_1+u}, f_2, \dots, f_{m+1})) \right) du dt_1.
\end{aligned}$$

Now, using similar substitution arguments as in the proof of Lemma 9.2.9 by letting $t_2 = t_1 + u$, it follows that

$$\begin{aligned}
& \int_{t_1 \in I(q, e_1)} p_{q_1+t_1}^{(1)}(e_1) f_{q_1}(t_1) (1 - F_{q_2}(t_1)) \\
& \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_1^{t_1}}(q_{t_1}, e_2, \dots, e_{n+1})) \sqcup \text{Cyl}(\pi_{\mathcal{C}_2^{t_1,+}}(q_2 + t_1, f_1, \dots, f_{m+1})) \right) dt_1 \\
& = \int_{t_1 \in I(q, e_1)} p_{q_1+t_1}^{(1)}(e_1) f_{q_1}(t_1) \text{Prob}_{q_1}^{\mathcal{A}_1} \left(\text{Cyl}(\pi_{\mathcal{C}_1^{t_1}}(q_{t_1}, e_2, \dots, e_{n+1})) \right) \\
& \quad \int_{t_2 \in I(q, f_1)} p_{q_2+t_2}^{(2)}(f_1) f_{q_2}(t_2) \mathbb{1}_{[t_1, +\infty[}(t_2) \\
& \quad \text{Prob}_{q_2}^{\mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_2^{t_2}}(q_{t_2}, f_2, \dots, f_{m+1})) \right) dt_2 dt_1.
\end{aligned} \tag{9.26}$$

Now, still by induction hypothesis and using similar arguments as before, we get that

$$\begin{aligned}
& \int_{t_2 \in I(q, f_1)} p_{q_2+t_2}^{(2)}(f_1) f_{q_2}(t_2) (1 - F_{q_1}(t_2)) \\
& \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_1^{t_2, +}}(q_1 + t_2, e_1, \dots, e_{n+1})) \sqcup \text{Cyl}(\pi_{\mathcal{C}_2^{t_2}}(q_2, f_2, \dots, f_{m+1})) \right) dt_2 \\
& = \int_{t_1 \in I(q, e_1)} p_{q_1+t_1}^{(1)}(e_1) f_{q_1}(t_1) \text{Prob}_{q_1}^{\mathcal{A}_1} \left(\text{Cyl}(\pi_{\mathcal{C}_1^{t_1}}(q_1, e_2, \dots, e_{n+1})) \right) \\
& \int_{t_2 \in I(q, f_1)} p_{q_2+t_2}^{(2)}(f_1) f_{q_2}(t_2) \mathbb{1}_{[0, t_1[}(t_2) \\
& \quad \text{Prob}_{q_2}^{\mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_2^{t_2}}(q_2, f_2, \dots, f_{m+1})) \right) dt_2 dt_1. \tag{9.27}
\end{aligned}$$

Finally, from (9.25), (9.26) and (9.27), we obtain that

$$\begin{aligned}
& \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_1}(q_1, e_1, \dots, e_{n+1})) \sqcup \text{Cyl}(\pi_{\mathcal{C}_2}(q_2, f_1, \dots, f_{m+1})) \right) = \\
& = \int_{t_1 \in I(q, e_1)} p_{q_1+t_1}^{(1)}(e_1) f_{q_1}(t_1) \text{Prob}_{q_1}^{\mathcal{A}_1} \left(\text{Cyl}(\pi_{\mathcal{C}_1^{t_1}}(q_1, e_2, \dots, e_{n+1})) \right) dt_1 \\
& \int_{t_2 \in I(q, f_1)} p_{q_2+t_2}^{(2)}(f_1) f_{q_2}(t_2) \text{Prob}_{q_2}^{\mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_2^{t_2}}(q_2, f_2, \dots, f_{m+1})) \right) dt_2 \\
& = \text{Prob}_{q_1}^{\mathcal{A}_1} \left(\text{Cyl}(\pi_{\mathcal{C}_1}(q_1, e_1, \dots, e_{n+1})) \right) \cdot \text{Prob}_{q_2}^{\mathcal{A}_2} \left(\text{Cyl}(\pi_{\mathcal{C}_2}(q_2, f_1, \dots, f_{m+1})) \right)
\end{aligned}$$

which concludes the proof of (9.13).

Now in order to complete the proof of Theorem 9.2.7, it remains to show that $\mathcal{A}_1 \parallel \mathcal{A}_2$ is almost-surely non-Zeno, *i.e.* we have to prove that for each state $q = (q_1, q_2)$, $\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Zeno}(q)) = 0$ ⁸. Fix a state $q = (q_1, q_2)$ of $\mathcal{A}_1 \parallel \mathcal{A}_2$. We recall that from the discussion in Section 3.1 on Zeno runs, $\text{Zeno}(q)$, $\text{Zeno}(q_1)$ and $\text{Zeno}(q_2)$ are all measurable in their respective σ -algebra. Assuming previous notations, we have that

$$\text{Zeno}(q) \subseteq \iota_1^{-1}(\text{Zeno}(q_1)) \cup \iota_2^{-1}(\text{Zeno}(q_2)).$$

Indeed, let $\rho \in \text{Runs}(\mathcal{A}_1 \parallel \mathcal{A}_2, q)$ be a Zeno run. Then we can write

$$\rho = q \xrightarrow{t_1, e_1} q_1 \xrightarrow{t_2, e_2} \dots$$

where $\sum_{k \geq 1} t_k < +\infty$. Since ρ is an infinite run, we have that there is $i \in \{1, 2\}$ such that $\{k \geq 1 \mid e_k \in E_i\}$ is an infinite set. It is then obvious that $\iota_i(\rho)$ is Zeno

⁸Recall that $\text{Zeno}(q)$ corresponds to the set of Zeno runs starting from q (see page 64).

and thus $\rho \in \iota_i^{-1}(\text{Zeno}(q_i))$. The above inclusion thus holds and therefore,

$$\begin{aligned} \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Zeno}(q)) &\leq \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\iota_1^{-1}(\text{Zeno}(q_1)) \cup \iota_2^{-1}(\text{Zeno}(q_2))) \\ &\leq \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\iota_1^{-1}(\text{Zeno}(q_1))) + \text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\iota_2^{-1}(\text{Zeno}(q_2))) \\ &= \text{Prob}_{q_1}^{\mathcal{A}_1}(\text{Zeno}(q_1)) + \text{Prob}_{q_2}^{\mathcal{A}_2}(\text{Zeno}(q_2)) \\ &\quad \text{from Proposition 9.2.14} \\ &= 0 \quad \text{since } \mathcal{A}_1 \text{ and } \mathcal{A}_2 \text{ are almost-surely non-Zeno.} \end{aligned}$$

Thus $\text{Prob}_q^{\mathcal{A}_1 \parallel \mathcal{A}_2}(\text{Zeno}(q)) = 0$, and since it holds true for each state q , we get that $\mathcal{A}_1 \parallel \mathcal{A}_2$ is almost-surely non-Zeno which concludes the proof of Theorem 9.2.7 \square

Remark 9.2.17. Note that given an almost-surely non-Zeno STA \mathcal{A} equipped with uniform or exponential distributions such that it satisfies conditions (A) and (B) (*i.e.* as in Remark 9.1.2), it holds that \mathcal{A} is in CSTA*. As said before, we have large classes of STA that are almost-surely fair. It holds that reactive STA [BBB⁺14] (see Section 7.2 here) are almost-surely non-Zeno. Equipping them with exponential distributions as in Remark 9.1.2) make them also composable. Hence reactive STA form a class in which the qualitative and quantitative model-checking problems of ω -regular properties are decidable (see Section 7.2), but that also has a compositional framework with an interleaving semantics. In [BBB⁺14], the authors also defined a class of weak-reactive STA that are also composable if we only equip them with uniform and exponential distributions.

As said in introduction, this interleaving semantics is a first for a compositional framework of STA. As stated in [HZ11], this is not the most interesting semantics for composition. An interest comes when the systems can interact between them, which yields to a handshaking operator. This will be the subject of Chapter 10. Before that, we need to introduce a notion of bisimulation in STA and to prove that it is a congruence w.r.t. parallel composition, which is an important property when dealing with composition.

9.3. Bisimulation and congruence

In this section, we define a notion of bisimulation for STA as done in [BBCM16], which naturally extends that for CTMCs [BHHK03, DP03] (see section 2.4.2). We importantly show that the defined bisimulation is a congruence w.r.t. parallel composition: as already briefly explained in Section 2.1.3, this means that, in a complex system, a component can be replaced by an equivalent one without

affecting the global behaviour of the system. This is crucial for a proper modular approach to system design.

9.3.1 Bisimulation

To define a bisimulation relation between STA, we are inspired by the approach of [DP03], in which the authors consider continuous-time Markov processes (CTMPs) – CTMPs generalize CTMCs (see Section 2.3) with general continuous state-spaces; this definition of bisimulation that is given for CTMPs can be adapted to our context (note however that STA cannot be seen as particular CTMPs).

We first define some notions. A subset $P \subseteq \mathbb{R}^n$ is a *polyhedral set* if it is defined by a (finite) boolean combination of constraints of the form $A_1x \leq b_1$ or $A_2x < b_2$, where $x = (x_1, \dots, x_n)$ is a variable, $A_1 \in \mathbb{R}^{m_1 \times n}$, $b_1 \in \mathbb{R}^{m_1}$, $A_2 \in \mathbb{R}^{m_2 \times n}$ and $b_2 \in \mathbb{R}^{m_2}$.

Let \mathcal{A} be a STA, Q be its set of states, and $P(Q) = \{\cup_{l \in L} \{l\} \times C_l \mid \forall l \in L, C_l \text{ polyhedral set of } \mathbb{R}_+^n\}$ where n is the number of clocks of \mathcal{A} . The set $P(Q)$ is a proper subset of the Borel σ -algebra over $L \times \mathbb{R}_+^n$, which is closed by projection (contrary to the Borel σ -algebra).

We then define the *closure of \mathcal{R} w.r.t. polyhedral sets*, that we write $\text{pcl}(\mathcal{R})$, as the following set: $\text{pcl}(\mathcal{R}) = \{A \in P(Q) \mid (a \in A \wedge a\mathcal{R}b) \implies b \in A\}$. One can notice that $\text{pcl}(\mathcal{R})$ corresponds to the set of all polyhedral unions of equivalence classes. Given two equivalence relations \mathcal{R} and \mathcal{R}' over S we say that \mathcal{R}' is *coarser* than \mathcal{R} or that \mathcal{R} is *finer* than \mathcal{R}' if $\mathcal{R} \subseteq \mathcal{R}'$.

We can now define a notion of bisimulation in STA (see [BBCM16]).

Definition 9.3.1. Let $\mathcal{A} = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X})$ be a STA. An equivalence relation \mathcal{R} over the set of states Q is a *bisimulation* for \mathcal{A} if for all $q, q' \in Q$ with $q\mathcal{R}q'$:

- (i) $\mathcal{L}(q) = \mathcal{L}(q')$, and
- (ii) for every $I \in \mathcal{B}(\mathbb{R}_+)$, for every $C \in \text{pcl}(\mathcal{R})$,

$$\text{Prob}_q^{\mathcal{A}}(q \xrightarrow{I, E} C) = \text{Prob}_{q'}^{\mathcal{A}}(q' \xrightarrow{I, E} C),$$

where $\{q \xrightarrow{I, E} C\}$ stands for $\{\rho \in \text{Runs}(\mathcal{A}, q) \mid \exists \tau \in I, \exists e \in E, \rho = q \xrightarrow{\tau, e} q_1 \rightarrow \dots \wedge q_1 \in C\}$.

States q and q' are *bisimilar* (written $q \sim q'$) if there is a bisimulation that contains (q, q') .

Given $q \in Q$, $I \in \mathcal{B}(\mathbb{R}_+)$ and $C \in \text{pcl}(\mathcal{R})$ the value $\text{Prob}_q^A(q \xrightarrow{I,E} C)$ can be expressed:

$$\text{Prob}_q^A(q \xrightarrow{I,E} C) = \int_{t \in I} P_{q+t}(C) f_q(t) dt$$

where the value $P_{q+t}(C)$ corresponds to the probability to reach instantaneously C from state $q + t$. Formally: if $C = \bigcup_{l \in L} \{l\} \times C_l$,

$$P_{q+t}(C) = \sum_{l' \in L} \sum_{e \in E_{l'}} p_{q+t}(e) \mathbb{1}_{C_{l'}(e,\nu)}(t)$$

for each $t \geq 0$ and each $C \in \text{pcl}(\mathcal{R})$, where, given $l' \in L$, $E_{l'}$ is the set of edges with target l' , and given $e = (l, g, Y, l')$, $C_{l'}(e, \nu) = \{t \in \mathbb{R}_+ \mid [Y \leftarrow 0](\nu + t) \in C_{l'}\}$. It can be shown that for every $t \geq 0$, P_{q+t} is a probability measure over Q .

Remark 9.3.2. In the case where $I = \mathbb{R}_+$, one can make a link with the Markov kernel $\kappa_{\mathcal{A}}(q, \cdot)$ defined in Section 7.1: for each $C \in \text{pcl}(\mathcal{R})$,

$$\text{Prob}_q^A(q \xrightarrow{\mathbb{R}_+,E} C) = \kappa_{\mathcal{A}}(q, C).$$

Also, given a STA \mathcal{A} , one can show that \sim is the coarsest bisimulation for \mathcal{A} .

Proposition 9.3.3. *For each STA \mathcal{A} , \sim is the coarsest bisimulation for \mathcal{A} .*

Definition 9.3.1 enjoys the following very nice characterisation which again shows a link with the definition of bisimulation over CTMCs [BHHK03].

Proposition 9.3.4. *Let \mathcal{A} be a STA and let \mathcal{R} be a bisimulation for \mathcal{A} . Then for all $q, q' \in Q$, $q \mathcal{R} q'$ if and only if*

- (i) $\mathcal{L}(q) = \mathcal{L}(q')$,
- (ii) $\mu_q = \mu_{q'}$, and
- (iii) for every $C \in \text{pcl}(\mathcal{R})$, $P_{q+t}(C) = P_{q'+t}(C)$ almost-surely for every $t \geq 0$.

Proof. In order to prove this characterisation, we will use Lemma 9.2.2.

The fact that if q and q' satisfy points (i), (ii) and (iii) then $q \mathcal{R} q'$ comes immediately from Definition 9.3.1. Indeed point (i) of Definition 9.3.1 is the same statement as point (i) of this proposition. Point (ii) of the same definition comes from point (ii) and (iii) of this proposition and from the fact that $\text{Prob}_q^A(q \xrightarrow{I,E} C) = \int_{t \in I} P_{q+t}(C) f_q(t) dt$ for each Borel sets $I \subseteq \mathbb{R}_+$ and each $C \in \text{cl}(\mathcal{R})$, where f_q is the density function associated with μ_q . Now let us assume that $q \mathcal{R} q'$. Point

(i) trivially holds from Definition 9.3.1. We also know that for every $I \in \mathcal{B}(\mathbb{R}_+)$ and for every $C \in \text{cl}(\mathcal{R})$,

$$\int_{t \in I} P_{q+t}(C) f_q(t) dt = \int_{t \in I} P_{q'+t}(C) f_{q'}(t) dt.$$

In particular, with $C = Q \in \text{cl}(\mathcal{R})$, we have $P_{q+t}(C) = 1 = P_{q'+t}(C)$ for every t and thus, for every $I \in \mathcal{B}(\mathbb{R}_+)$,

$$\int_{t \in I} f_q(t) dt = \int_{t \in I} f_{q'}(t) dt$$

which leads to the fact that $f_q = f_{q'}$ almost-surely and thus $\mu_q = \mu_{q'}$. It remains to show point (iii). If we fix $C \in \text{cl}(\mathcal{R})$ we can show that the function $P_{q+\bullet}(C)$, assigning the value $P_{q+t}(C)$ to each real positive number t , is measurable. Indeed, we have that

$$P_{q+t}(C) = \sum_{l' \in L} \sum_{e \in E_{l'}} p_{q+t}(e) \mathbb{1}_{\mathcal{C}_{l'}(e, \nu)}(t).$$

We already know from hypothesis (\star) (see page 176) that for each e , $p_{q+\bullet}(e)$ is measurable. If we assume that $q = (l, \nu)$ with $\nu \in \mathbb{R}_+^n$, then for each e and l' , the set $\mathcal{C}_{l'}(e, \nu)$ corresponds to the set $([Y \leftarrow 0] \circ g_\nu)^{-1}(C_l)$ where $g_\nu : \mathbb{R}_+ \rightarrow \mathbb{R}_+^n$ is the function that assigns the value $\nu + t$ for each $t \in \mathbb{R}_+$. Now since $[Y \leftarrow 0]$ and g_ν are two measurable functions and since C_l is a Borel set, we get that $\mathcal{C}_{l'}(e, \nu)$ is a Borel set. Hence $\mathbb{1}_{\mathcal{C}_{l'}(e, \nu)}(t)$ is a measurable function. We can thus conclude that $P_{q+\bullet}(C)$ is also measurable. The same reasoning can be applied to $P_{q'+\bullet}(C)$. Finally, from Lemma 9.2.2 applied to $g_1(t) = P_{q+t}(C)$, $g_2(t) = P_{q'+t}(C)$ and $h = f_q = f_{q'}$ almost-surely, we get that point (iii) is satisfied. \square

We now illustrate the notion of bisimulation on a simple example.

Example 9.3.5. Let us consider the simple STA \mathcal{A} with two clocks on Figure 9.4. We assume that each state of the form (l_1, ν) or (l_2, ν') with $\nu, \nu' \in \mathbb{R}_+^2$ is equipped with the same exponential distribution over delays, say $\text{Exp}(\lambda)$. Now, from a state of the form $q = (l_0, (\nu_1, \nu_2))$ with $\nu_1 < 1$ or $\nu_2 < 1$, $I(q) = [0, 1 - \min(\nu_1, \nu_2)[$ and so we can equip q with a uniform distribution on the interval $I(q)$ for the delays.

We now compute the equivalence classes for \sim . It can easily be established that the set of states $\{l_1, l_2\} \times \mathbb{R}_+^2$ is an equivalence class for \sim .

In order to find the equivalence classes associated with l_0 , we use the characterisation Proposition 9.3.4. It is obvious that the set $\{l_0\} \times [1, \infty[^2$ forms




Figure 9.4: A simple example for bisimulation.

an equivalence class for \sim as there are the only states from which no edges are enabled. We can then show that for each $\nu \in [0, 1[$,

$$A_\nu = \{l_0\} \times (\{(\nu_1, \nu) \mid \nu_1 \geq \nu\} \cup \{(\nu, \nu_2) \mid \nu_2 \geq \nu\})$$

is an equivalence class. First let us prove that given $q = (l_0, (\nu_1, \nu_2))$ and $q' = (l_0, (\nu'_1, \nu'_2)) \in A_\nu$, both states satisfy point (i), (ii) and (iii) of the characterisation. Point (i) is obvious. In order to get point (ii), it suffices to observe that $\min(\nu_1, \nu_2) = \min(\nu'_1, \nu'_2)$. Finally, for point (iii), from location l_0 , the only set $C \in \text{pcl}(\sim)$ that is reachable in one step is $C = \{l_1, l_2\} \times \mathbb{R}_+^2$, and we get that for each $t \geq 0$, $P_{q+t}(C) = \mathbb{1}_{I(q)}(t)$ and $P_{q'+t}(C) = \mathbb{1}_{I(q')}(t)$. But since $\min(\nu_1, \nu_2) = \min(\nu'_1, \nu'_2)$, we have $I(q) = I(q')$ and thus $P_{q+t}(C) = P_{q'+t}(C)$. Finally A_ν is an equivalence class since, if $q = (l_0, (\nu_1, \nu_2)) \in A_\nu$ and $q' = (l_0, (\nu'_1, \nu'_2)) \in A_{\nu'}$ with $\nu \neq \nu'$, then $\min(\nu_1, \nu_2) \neq \min(\nu'_1, \nu'_2)$ and thus point (ii) does not hold.

Construction of a bisimulation between two STA. We have defined a bisimulation as a relation between states of a STA. We would like now to define a bisimulation as a relation between STA with the same set of atomic propositions AP. A classical way to achieve this objective (see [BK08]), is to consider the disjoint union of two STA and to define a bisimulation between these two automata as a bisimulation for the disjoint union of both automata. We thus need to define the disjoint union of two STA. Let $\mathcal{A}_i = (L_i, X_i, E_i, \text{Inv}_i, \text{AP}, \mathcal{L}_i, (\mu_q^{(i)}, p_q^{(i)})_{q \in L_i \times \mathbb{R}_+^{X_i}})$ for $i \in \{1, 2\}$ be two STA with $L_1 \cap L_2 = \emptyset$, $X_1 \cap X_2 = \emptyset$ and $E_1 \cap E_2 = \emptyset$. The *disjoint union* of \mathcal{A}_1 and \mathcal{A}_2 is the stochastic timed automaton $\mathcal{A}_1 \cup \mathcal{A}_2$ defined by

$$\mathcal{A}_1 \cup \mathcal{A}_2 = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X}),$$

where $L = L_1 \cup L_2$, $X = X_1 \cup X_2$, $E = E_1 \cup E_2$, $\text{Inv}(l) = \text{Inv}_1(l)$ if $l \in L_1$ and $\text{Inv}(l) = \text{Inv}_2(l)$ if $l \in L_2$, μ_q and p_q are such that if $q = (l_1, \nu_1, \nu_2)$ with $l_1 \in L_1$,

$\nu_1 \in \mathbb{R}_+^{X_1}$ and $\nu_2 \in \mathbb{R}_+^{X_2}$ then $\mu_q = \mu_{(l_1, \nu_1)}^{(1)}$ and $p_q = p_{(l_1, \nu_1)}^{(1)}$ and similarly, if $q = (l_2, \nu_1, \nu_2)$ with $l_2 \in L_2$ then $\mu_q = \mu_{(l_2, \nu_2)}^{(2)}$ and $p_q = p_{(l_2, \nu_2)}^{(2)}$, and $\mathcal{L}(l) = \mathcal{L}_1(l)$ if $l \in L_1$ and $\mathcal{L}(l) = \mathcal{L}_2(l)$ if $l \in L_2$.

We can now define the notion of bisimulation between two STA with the same set of atomic propositions.

Definition 9.3.6. Let $\mathcal{A}_i = (L_i, X_i, E_i, \text{Inv}_i, \text{AP}, \mathcal{L}_i, (\mu_q^{(i)}, p_q^{(i)})_{q \in L_i \times \mathbb{R}_+^{X_i}})$ for $i \in \{1, 2\}$ be two STA with $L_1 \cap L_2 = \emptyset$, $X_1 \cap X_2 = \emptyset$ and $E_1 \cap E_2 = \emptyset$. An equivalence relation over $Q_1 \cup Q_2$ is a *bisimulation* between \mathcal{A}_1 and \mathcal{A}_2 if it is a bisimulation for $\mathcal{A}_1 \cup \mathcal{A}_2$. We say that \mathcal{A}_1 and \mathcal{A}_2 are *bisimilar from states* $q_1 \in Q_1$ and $q_2 \in Q_2$, and we write $\mathcal{A}_1 \sim \mathcal{A}_2$, if $q_1 \sim q_2$ in $\mathcal{A}_1 \cup \mathcal{A}_2$.

Remark 9.3.7. Let us note that given two STA \mathcal{A}_1 and \mathcal{A}_2 and given a state (l_1, ν_1, ν_2) of $\mathcal{A}_1 \cup \mathcal{A}_2$ with $l_1 \in L_1$, $\nu_1 \in \mathbb{R}_+^{X_1}$ and $\nu_2 \in \mathbb{R}_+^{X_2}$, it holds that $\forall \nu'_2 \in \mathbb{R}_+^{X_2}$, $(l_1, \nu_1, \nu'_2) \sim (l_1, \nu_1, \nu_2)$. This comes from the fact that from state (l_1, ν_1, ν_2) in the union, with $l_1 \in L_1$, only edges of \mathcal{A}_1 are enabled. Hence, the state (l_1, ν_1, ν_2) will behave in $\mathcal{A}_1 \cup \mathcal{A}_2$ exactly as state (l_1, ν_1) in \mathcal{A}_1 . Thus, the value of ν_2 has no impact on the behaviour of $\mathcal{A}_1 \cup \mathcal{A}_2$ from (l_1, ν_1, ν_2) . The same reasoning applies to states (l_2, ν_1, ν_2) with $l_2 \in L_2$. A state of the form (l_i, ν_1, ν_2) with $l_i \in L_i$ will then be abusively identified as the state (l_i, ν_i) .


We now illustrate this definition on a simple example, which considers again STA \mathcal{A} of Example 9.3.5.

Example 9.3.8. We consider again the STA \mathcal{A} depicted in Example 9.3.5 and on Figure 9.4.

Now, let us consider the single clock STA \mathcal{B} of Figure 9.5. Assuming that we have the same probability distributions as STA \mathcal{A} (*i.e.* uniform distribution for l'_0 and $\text{Exp}(\lambda)$ -distribution for l'_1), it can be easily established that $\mathcal{A} \sim \mathcal{B}$ from $q_0 = (l_0, (0, 0))$ and $q'_0 = (l'_0, 0)$. It is obvious that each state of $\{l_1, l_2\} \times \mathbb{R}_+^2$ is bisimilar to each state of $\{l'_1\} \times \mathbb{R}_+$ and that each state of $\{l_0\} \times [1, \infty[^2$ is bisimilar to each state of $\{l'_0\} \times [1, \infty[$. Finally, by a similar argument as in Example 9.3.5, it can be easily established that for each $\nu \in [0, 1[$, (l'_0, ν) is bisimilar to each state of A_ν .

9.3.2 Congruence

One of the main objectives of defining behavioural equivalences is to aim at modular design and proof of correctness. This is only possible if bisimulation is a *congruence w.r.t. parallel composition*, that is, if $\mathcal{A}_1 \sim \mathcal{A}_2$ from states $q_1 \in Q_1$ and $q_2 \in Q_2$, then for every \mathcal{B} , $\mathcal{A}_1 \parallel \mathcal{B} \sim \mathcal{A}_2 \parallel \mathcal{B}$ from states (q_1, q) and (q_2, q) ,

Figure 9.5: \mathcal{B} is bisimilar to \mathcal{A} .

for each state q of \mathcal{B} . We first prove the following natural lemma which is a key point for proving the congruence of the bisimulation w.r.t. parallel composition. Though very intuitive, the result is surprisingly quite technical to prove.

Lemma 9.3.9. Let $\mathcal{A}, \mathcal{B} \in \text{CSTA}^*$ with sets of states resp. Q_A and Q_B . If \mathcal{R} is a bisimulation for \mathcal{A} then the equivalence relation \mathcal{R}' over $Q_A \times Q_B$ defined by $\mathcal{R}' = \{((q_1, q), (q_2, q)) \mid q_1 \mathcal{R} q_2 \text{ and } q \in Q_B\}$, is a bisimulation for $\mathcal{A} \parallel \mathcal{B}$.

This result seems very intuitive, however the proof is quite technical. The tricky part comes when verifying if for each states q and q' of $\mathcal{A} \parallel \mathcal{B}$ with $q \mathcal{R}' q'$, it holds that for each polyhedral set $C \in \text{pcl}(\mathcal{R}')$ and for almost every $t \geq 0$, $P_{q+t}(C) = P_{q'+t}(C)$. The key part to prove this point, is to decompose and to project C in such a way that we can use the fact that \mathcal{R} is a bisimulation for \mathcal{A} . This is where we need polyhedral sets for C .

Lemma 9.3.10. Let \mathcal{A} be a stochastic timed automaton in CSTA^* and let \mathcal{R} be a bisimulation for \mathcal{A} . Let q and q' be two states of \mathcal{A} . If $q \mathcal{R} q'$ then $(q+t) \mathcal{R} (q'+t)$ for every $t \geq 0$.

Proof. Let q and q' be states of \mathcal{A} such that $q \mathcal{R} q'$. We have to prove that for every $t \geq 0$, $(q+t) \mathcal{R} (q'+t)$. Let $t \geq 0$, we have to show that $q+t$ and $q'+t$ satisfy points (i), (ii) and (iii) of Proposition 9.3.4, *i.e.* that (i) $\mathcal{L}(q+t) = \mathcal{L}(q'+t)$, that (ii) $\mu_{q+t} = \mu_{q'+t}$ and that (iii) for every $C \in \text{cl}(\mathcal{R})$ and for almost every $t' \geq 0$, $P_{(q+t)+t'}(C) = P_{(q'+t)+t'}(C)$. Point (i) is trivial as it only depends on the locations. Point (iii) comes from the fact that for every $C \in \text{cl}(\mathcal{R})$, $P_{q+t''}(C) = P_{q'+t''}(C)$ for almost every $t'' \geq 0$, as $q \mathcal{R} q'$. We thus get that for every $C \in \text{cl}(\mathcal{R})$, $P_{(q+t)+t'}(C) = P_{(q'+t)+t'}(C)$ almost-surely for every $t' \geq 0$. It remains to establish point (ii), *i.e.* that $\mu_{q+t} = \mu_{q'+t}$ or again $f_{q+t} = f_{q'+t}$ almost-surely. This immediate from condition (B) of CSTA^* : it holds that

$$f_{q_1+t}(t') = \frac{f_{q_1}(t+t')}{1 - F_{q_1}(t)} \quad \text{and} \quad f_{q_2+t}(t') = \frac{f_{q_2}(t+t')}{1 - F_{q_2}(t)}$$

for almost-surely each $t' \geq 0$. Then since $q_1 \mathcal{R} q_2$, we get that $f_{q_1}(t+t') = f_{q_2}(t+t')$ for almost-surely each $t' \geq 0$. It follows that $f_{q_1+t}(t') = f_{q_2+t}(t')$ for almost-surely each $t' \geq 0$ which concludes the proof. \square

We can now prove Lemma 9.3.9. We write $\mathcal{A} \parallel \mathcal{B} = (L_A \times L_B, X_A \cup X_B, E_A \cup E_B, \text{Inv}, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in Q_A \times Q_B})$. Given a state (q_i, q) of $\mathcal{A} \parallel \mathcal{B}$, we write $f_{q_i, A}$ for the density function associated with $\mu_{q_i}^A$, $f_{q, B}$ for the density function associated with μ_q^B and $f_{q_i, \min}$ for the density function of $\mu_{(q_i, q)}$ in the product.

Proof of Lemma 9.3.9. Let $q_1 = (l_1, \nu_1)$ and $q_2 = (l_2, \nu_2)$ be states of \mathcal{A} and let $q = (l, \nu)$ be a state of \mathcal{B} such that $(q_1, q) \mathcal{R}'(q_2, q)$, i.e. such that $q_1 \mathcal{R} q_2$. Then, from Proposition 9.3.4 it holds that $\mathcal{L}_A(q_1) = \mathcal{L}_A(q_2)$, that $f_{q_1, A} = f_{q_2, A}$ almost-surely and that for each $C \in \text{pcl}(\mathcal{R})$, $P_{q_1+\bullet}^A(C) = P_{q_2+\bullet}^A(C)$ almost-surely. From the definition of the composition we can immediately deduce that $\mathcal{L}((q_1, q)) = \mathcal{L}((q_2, q))$ and that $f_{q_1, \min} = f_{q_2, \min}$ almost-surely. It remains to show that for each $C \in \text{pcl}(\mathcal{R}')$, $P_{(q_1, q)+\bullet}(C) = P_{(q_2, q)+\bullet}(C)$ almost-surely. Let $C \in \text{pcl}(\mathcal{R}')$, we can write

$$C = \bigcup_{(l_A, l_B) \in L_A \times L_B} \{(l_A, l_B)\} \times C_{(l_A, l_B)}$$

where for each $(l_A, l_B) \in L_A \times L_B$, $C_{(l_A, l_B)}$ is a polyhedral set. Let us first compute the value $P_{(q_1, q)+t}(C)$. We have that

$$\begin{aligned} P_{(q_1, q)+t}(C) &= \frac{f_{q_1, A}(t)(1 - F_{q, B}(t))}{f_{q_1, \min}(t)} \sum_{l_A \in L_A} \sum_{e_A \in E_{l_A}} p_{q_1+t}^A(e_A) \mathbb{1}_{\mathcal{C}_{(l_A, l)}(e_A, \nu_1, \nu)}(t) \\ &\quad + \frac{f_{q, B}(t)(1 - F_{q_1, A}(t))}{f_{q_1, \min}(t)} \sum_{l_B \in L_B} \sum_{e_B \in E_{l_B}} p_{q+t}^B(e_B) \mathbb{1}_{\mathcal{C}_{(l_1, l_B)}(e_B, \nu_1, \nu)}(t), \end{aligned}$$

where, if $e_A = (l_1, g_A, Y_A, l_A)$ and $e_B = (l, g_B, Y_B, l_B)$ then

- $\mathcal{C}_{(l_A, l)}(e_A, \nu_1, \nu) = \{t \in \mathbb{R}_+ \mid ([Y_A \leftarrow 0](\nu_1 + t), \nu + t) \in C_{(l_A, l)}\}$, and
- $\mathcal{C}_{(l_1, l_B)}(e_B, \nu_1, \nu) = \{t \in \mathbb{R}_+ \mid (\nu_1 + t, [Y_B \leftarrow 0](\nu + t)) \in C_{(l_1, l_B)}\}$.

Computing the value $P_{(q_2, q)+t}(C)$ is similar. In order to show that for almost-surely every $t \geq 0$, $P_{(q_1, q)+t}(C) = P_{(q_2, q)+t}(C)$, it suffices to prove that for almost-surely every $t \geq 0$,

$$\begin{aligned} w_{(q_1, q)}^A(t) &\sum_{l_A \in L_A} \sum_{e_A \in E_{l_A}} p_{q_1+t}^A(e_A) \mathbb{1}_{\mathcal{C}_{(l_A, l)}(e_A, \nu_1, \nu)}(t) \\ &= w_{(q_2, q)}^A(t) \sum_{l_A \in L_A} \sum_{e_A \in E_{l_A}} p_{q_2+t}^A(e_A) \mathbb{1}_{\mathcal{C}_{(l_A, l)}(e_A, \nu_2, \nu)}(t) \end{aligned} \quad (9.28)$$

and that

$$\begin{aligned} w_{(q_1, q)}^B(t) & \sum_{l_B \in L_B} \sum_{e_B \in E_{l_B}} p_{q+t}^B(e_B) \mathbb{1}_{\mathcal{C}_{(l_1, l_B)}(e_B, \nu_1, \nu)}(t) \\ & = w_{(q_2, q)}^B(t) \sum_{l_B \in L_B} \sum_{e_B \in E_{l_B}} p_{q+t}^B(e_B) \mathbb{1}_{\mathcal{C}_{(l_2, l_B)}(e_B, \nu_2, \nu)}(t). \end{aligned} \quad (9.29)$$

First of all, let us observe that we already have that $w_{(q_1, q)}^A = w_{(q_2, q)}^A$ almost-surely and that $w_{(q_1, q)}^B = w_{(q_2, q)}^B$ since by hypothesis, $f_{q_1, A} = f_{q_2, A}$ almost-surely. Now, we can show that for each $l_B \in L_B$ and for each $e_B = (l, g_B, Y_B, l_B) \in E_{l_B}$, $\mathcal{C}_{(l_1, l_B)}(e_B, \nu_1, \nu) = \mathcal{C}_{(l_2, l_B)}(e_B, \nu_2, \nu)$ almost-surely. Indeed, for almost every $t \geq 0$, $(q_1 + t)\mathcal{R}(q_2 + t)$ from Lemma 9.3.10. Now if $t \in \mathcal{C}_{(l_1, l_B)}(e_B, \nu_1, \nu)$ and t is such that $(q_1 + t)\mathcal{R}(q_2 + t)$, then $(q_1 + t, [Y_B \leftarrow 0](q + t)) \in C$. Since $(q_1 + t, [Y_B \leftarrow 0](q + t))\mathcal{R}'(q_2 + t, [Y_B \leftarrow 0](q + t))$ and $C \in \text{pcl}(\mathcal{R}')$, we get that $(q_2 + t, [Y_B \leftarrow 0](q + t)) \in C$, *i.e.* $t \in \mathcal{C}_{(l_2, l_B)}(e_B, \nu_2, \nu)$. By a similar argument, we get the almost-sure equality between the two sets. This proves equality (9.29). In order to show equality (9.28) we first introduce some notation. We write $C_q^{(t)}$ for the following set of states in \mathcal{A} : $\{q_A \in Q_A \mid (q_A, q + t) \in C\}$. Equality (9.28) can thus be rewritten as

$$P_{q_1+t}^A(C_q^{(t)}) = P_{q_2+t}^A(C_q^{(t)}). \quad (9.30)$$

We will show that this equality holds almost-surely on each interval I of the form $[a, b]$ with $a, b \in \mathbb{R}_+$. Let $I = [a, b]$ be such an interval. For each $n \in \mathbb{N}$, we write

$$I = \bigcup_{k=0}^{2^n-1} I_k^{(n)}$$

where for each k , $I_k^{(n)} = [a + \frac{k(b-a)}{2^n}, a + \frac{(k+1)(b-a)}{2^n}]$. For each $n \in \mathbb{N}$ and each $0 \leq k \leq 2^n - 1$ we write

$$C_q^{(k, n)} = \{q_A \in Q_A \mid \exists t \in I_k^{(n)}, (q_A, q + t) \in C\}.$$

We can prove that $C_q^{(k, n)} \in \text{pcl}(\mathcal{R})$. Indeed if $q_A \in C_q^{(k, n)}$, then there is $t \in I_k^{(n)}$ such that $(q_A, q + t) \in C$. Let q'_A be such that $q_A \mathcal{R} q'_A$. Then, from definition of \mathcal{R}' , we have $(q_A, q + t)\mathcal{R}'(q'_A, q + t)$ and thus $(q'_A, q + t) \in C$ since $C \in \text{pcl}(\mathcal{R}')$. Hence, $q'_A \in C_q^{(k, n)}$. It remains to show that for each $l_A \in L_A$, there is a polyhedral set C_{l_A} such that

$$C_q^{(k, n)} = \bigcup_{l_A \in L_A} \{l_A\} \times C_{l_A}.$$

Such sets exist and are defined as follows. For each $l_A \in L_A$, we define C_{l_A} as follows

$$C_{l_A} = \text{Proj}_{X_A} \left(C_{(l_A, l)} \cap \left(\mathbb{R}_+^{X_A} \times \{\nu + t \mid t \in I_k^{(n)}\} \right) \right)$$

(where Proj_{X_A} denotes the projection of $\mathbb{R}_+^{X_A \cup X_B}$ over $\mathbb{R}_+^{X_A}$) which is a polyhedral set. Indeed, we have that for each n and k , $C_{(l_A, l)}$ and $\mathbb{R}_+^{X_A} \times \{\nu + t \mid t \in I_k^{(n)}\}$ are two polyhedral sets. Then the intersection is still a polyhedral set. And since the projection of a polyhedral set is a polyhedral set, we get that C_{l_A} is a polyhedral set. Which proves that $C_q^{(k, n)} \in \text{pcl}(\mathcal{R})$. It follows that for each $n \geq 0$, for each $0 \leq k \leq 2^n - 1$,

$$P_{q_1+t}^A(C_q^{(k, n)}) = P_{q_2+t}^A(C_q^{(k, n)}) \quad (9.31)$$

for almost every $t \geq 0$. We will now establish that for almost each $t \in I$,

$$\sum_{k=0}^{2^n-1} P_{q_1+t}^A(C_q^{(k, n)}) \mathbf{1}_{I_k^{(n)}}(t) \xrightarrow{n \rightarrow \infty} P_{q_1+t}^A(C_q^{(t)}). \quad (9.32)$$

We fix $t \in I$ such that t is not in $\{a + \frac{k(b-a)}{2^n} \mid n \in \mathbb{N} \wedge k \in \{0, \dots, 2^n\}\}$. It can be shown that $\Lambda(\{a + \frac{k(b-a)}{2^n} \mid n \in \mathbb{N} \wedge k \in \{0, \dots, 2^n\}\}) = 0$. Then, by construction, for each $n \geq 0$ there is a unique $k_{t, n} \in \{0, \dots, 2^n - 1\}$ such that $t \in I_{k_{t, n}}^{(n)}$. We thus have to show that

$$P_{q_1+t}^A(C_q^{(k_{t, n}, n)}) \xrightarrow{n \rightarrow \infty} P_{q_1+t}^A(C_q^{(t)}). \quad (9.33)$$

By construction, we have that for each $n \geq 0$, $I_{k_{t, n+1}}^{(n+1)} \subseteq I_{k_{t, n}}^{(n)}$ and thus also $C_q^{(k_{t, n+1}, n+1)} \subseteq C_q^{(k_{t, n}, n)}$. It follows that

$$C_q^{(k_{t, n}, n)} \xrightarrow{n \rightarrow \infty} \bigcap_{n \geq 0} C_q^{(k_{t, n}, n)} = C_q^{(t)}.$$

Since $P_{q_1+t}^A$ is a probability measure over Q_A we conclude, from a classical result, that equality (9.33) holds which proves (9.32). Similarly, we can prove that

$$\sum_{k=0}^{2^n-1} P_{q_2+t}^A(C_q^{(k, n)}) \mathbf{1}_{I_k^{(n)}}(t) \xrightarrow{n \rightarrow \infty} P_{q_2+t}^A(C_q^{(t)}). \quad (9.34)$$

for almost every $t \in I$. Now, from (9.31) we get that

$$\sum_{k=0}^{2^n-1} P_{q_1+t}^A(C_q^{(k, n)}) \mathbf{1}_{I_k^{(n)}}(t) = \sum_{k=0}^{2^n-1} P_{q_2+t}^A(C_q^{(k, n)}) \mathbf{1}_{I_k^{(n)}}(t),$$

for almost every $t \in I$, and thus (9.32) and (9.34) allow us to state that equality (9.30) holds almost-surely for each $t \in I$. Since it holds for each interval $I \subseteq \mathbb{R}_+$ of the form $[a, b]$, we get that (9.30) holds almost-surely for each $t \geq 0$ which terminates the proof of the lemma. \square

We can now state the main result of this section:

Theorem 9.3.11. *Bisimulation is a congruence w.r.t. parallel composition. That is: if \mathcal{A}_1 , \mathcal{A}_2 and \mathcal{B} are three STA in CSTA^* , for every states q_1 , q_2 and q of resp. \mathcal{A}_1 , \mathcal{A}_2 and \mathcal{B} , if $\mathcal{A}_1 \sim \mathcal{A}_2$ from q_1 and q_2 , then $\mathcal{A}_1 \parallel \mathcal{B} \sim \mathcal{A}_2 \parallel \mathcal{B}$ from (q_1, q) and (q_2, q) .*

Proof. We fix states q_1 , q_2 and q . From the hypothesis, there is a bisimulation \mathcal{R} for $\mathcal{A}_1 \cup \mathcal{A}_2$ such that $q_1 \mathcal{R} q_2$. We define

$$\mathcal{R}' = \{((q_1, q), (q_2, q)) \mid (q_1 \mathcal{R} q_2) \wedge q \in Q\}$$

and we show that it is a bisimulation between $\mathcal{A}_1 \parallel \mathcal{B}$ and $\mathcal{A}_2 \parallel \mathcal{B}$ from states (q_1, q) and (q_2, q) . It is trivial to get that $(q_1, q) \mathcal{R}' (q_2, q)$ and Lemma 9.3.9 states that \mathcal{R}' is a bisimulation for $(\mathcal{A}_1 \cup \mathcal{A}_2) \parallel \mathcal{B}$. Then Remark 9.3.7 allows us to conclude that \mathcal{R}' is a bisimulation for $(\mathcal{A}_1 \parallel \mathcal{B}) \cup (\mathcal{A}_2 \parallel \mathcal{B})$. \square

Interactive Stochastic Timed Automata and Handshaking Composition

Following [HZ11] (see Section 2.4.1), handshaking composition is a generalisation of both interleaving and synchronous composition. Hence the interleaving semantics introduced in Chapter 9 is not entirely satisfactory. However this was an important first step towards handshaking: the handshaking semantics obviously allows for synchronisations, however when no interactions are possible, the product has to be interleaving. Therefore, having already defined an interleaving operator for STA gives us a good basis.

As stated in [HZ11], in CTMCs the interleaving composition is the natural operator for the model. A way to extend handshaking composition to CTMCs, is to consider the model of IMCs ([Her02] and [HK09]) which adds non-determinism into the model (see Section 2.4) and allows for communication between the systems.

The objective of this chapter is thus to adapt the work of [Her02], [HK09] and [HZ11] to STA, *i.e.* to extend handshaking composition to STA by adding interactions in the model. In Section 10.1 we thus introduce the new model of *interactive STA* (ISTA for short) and we illustrate it. In Section 10.2, we define a handshaking parallel composition in ISTA. Similarly as in Chapter 9, we identify a class of ISTA for which parallel composition is well-defined. We also define a *hiding* operator and show how this operator along with parallel composition allow us to define the semantics of an ISTA as a STA. We also define a notion of bisimulation, that extends the one on IMCs ([Her02], [HK09]; see Section 2.4.2

here), and importantly show that it is a congruence w.r.t. parallel composition and hiding.

We end the chapter with a link with Part I. Roughly speaking, in Section 10.3, we identify a class of ISTA in which parallel composition is well-defined and whose semantics give rise to a class of STA that, we prove, enter the setting of Chapter 6. In other words, we identify a class of ISTA that are composable and on which the decidability and approximability results of Chapter 6 can be applied.

All the notions and the results of the chapter are neither published nor submitted work.

10.1. Syntax of ISTA

In this section, we introduce the new model of ISTA and illustrate it on an example. It is inspired from the IMC model ([Her02], [HK09] and Section 2.4.2) which extends the CTMC model.

We recall here a notation of Chapter 9. Given some set of clocks X , $\mathcal{G}_\times(X)$ denotes the set of guards of the following form: any finite conjunction of expressions of the form $x \sim c$ where $x \in X$, $c \in \mathbb{N}$ and $\sim \in \{<, \geq, >\}$. Similarly, we consider $\mathcal{G}_{\text{int}}(X)$ as the set of guards with $\sim \in \{<, \geq\}$. We now give immediately the definition of the ISTA model.

Definition 10.1.1. An *interactive stochastic timed automaton* (ISTA for short) is a tuple $\mathcal{A} = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X}, \Gamma^\tau, \dashrightarrow)$, where

- L is a finite set of locations partitioned into $L_{\text{int}} = \{l \in L \mid \forall (l', g, Y, l'') \in E, l' \neq l\}$, called the set of *interactive locations*, and L_{int}^c representing all other locations, *i.e.* the set of non-interactive locations;
- X is a finite set of clocks;
- $E \subseteq L \times \mathcal{G}_\times(X) \times 2^X \times L$ is a finite set of edges;
- $\Gamma^\tau = \Gamma \cup \{\tau\}$ where Γ is a finite set of actions and τ is an internal action;
- $\dashrightarrow \subseteq L \times \Gamma^\tau \times \mathcal{G}_{\text{int}}(X) \times 2^X \times L$ is a finite set of *interactive transitions*;
- $\mathcal{L} : L \rightarrow 2^{\text{AP}}$ is a labelling function;
- for each $q \in L_{\text{int}}^c \times \mathbb{R}_+^X$ (*i.e.* for each state with non-interactive location), μ_q is a probability distribution over $I(q)$, and p_q a probability distribution over E ;

- $\text{Inv} : L \rightarrow \mathcal{G}_\times(X)$ is an invariant function satisfying that for each $l \in L_{\text{int}}$, $\text{Inv}(l) = \text{true}$.

We immediately illustrate the model on Example 10.1.2. It represents a naive cooling system inspired from the one depicted in [BBJM12] as a STA.

Example 10.1.2. We illustrate a naive model for a cooling system as an ISTA. It is depicted in Figure 10.1. For this example, we omit the labels on the location. ISTA $\mathcal{A}_{\text{cool}} = (L, X, E, \text{Inv}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X}, \Gamma^\tau, \dashrightarrow)$ is defined with the following components: sets $L = \{\text{Up}, \text{Wait}, \text{Down}, \text{Failure}, \text{Burned Plant}\}$, $X = \{x\}$, E is described by the plain arrows on the figure just like in the STA model, $\Gamma = \{\text{down}, \text{fix}, \text{replace}, \text{burn}\}$, \dashrightarrow is also obviously described on Figure 10.1, all invariants are given by **true**, and the distributions are defined as follows: first observe that Wait, Down and Failure are interactive locations and thus, no distributions are defined in these locations. Then from Up and Burned Plant, the distributions over the delays could be exponential distributions, let us say $\text{Exp}(\lambda_1)$ and $\text{Exp}(\lambda_2)$. Finally we assume that, from Up, the probability over the edges gives 1/2 to both outgoing edges. In location Burned Plant, only one edge is enabled. One can check that all conditions of Definition 10.1.1 are met.

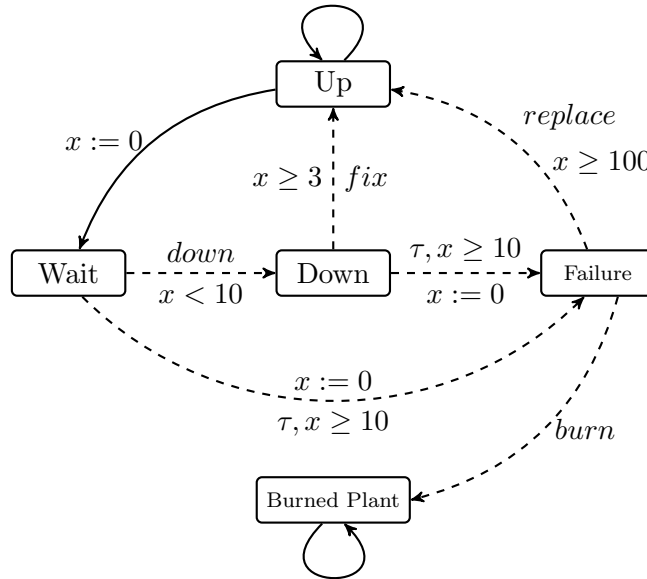


Figure 10.1: An ISTA $\mathcal{A}_{\text{cool}}$ describing a cooling system

We give now some light on the example. In location Up, the cooling system

works just fine. It can stay there forever $(\text{Up}, \mathbf{true}, \emptyset, \text{Up}) \in E$ or something wrong can happen at any time $(\text{Up}, \mathbf{true}, \{x\}, \text{Wait}) \in E$, *i.e.* a delay is randomly chosen according to $\text{Exp}(\lambda_1)$ and then, with probability 1/2 something wrong occurs and the system moves towards location Wait. In this location, the system is already down however it needs to wait for someone to fix it. We make the assumption that the system is better to be fixed within the next 10 time units, that is why clock x was reset to 0 and why guard $x < 10$ is used: once someone is available, a signal is sent through the interactive transition Wait $\xrightarrow{\text{down}, x < 10}$ Down. However if nobody is available to repair the cooling system within the 10 first time units, then an internal error can occur through the interactive transition Wait $\xrightarrow{\tau, x \geq 10, x := 0}$ Failure. In location Down, a worker is fixing the cooling system. It goes to Up as soon as the worker has fixed it. We assume that the system cannot be fixed in the 3 first time units since it was down, hence Down $\xrightarrow{\text{fix}, x \geq 3}$ Up. After the elapse of 10 time units, an internal error can occur like in location Wait, through the interactive transition Down $\xrightarrow{\tau, x \geq 10, x := 0}$ Failure. In location Failure, the system is beyond repair and has to be replaced but it takes much more time and it has to wait for someone to replace it: Failure $\xrightarrow{\text{replace}, x \geq 100}$ Up. But since the cooling system does not work, the power plant can burn at any time which is depicted through the interactive transition Failure $\xrightarrow{\text{burn}}$ Burned Plant. The systems stays then in this location forever $(\text{Burned Plant}, \mathbf{true}, \emptyset, \text{Burned Plant}) \in E$.

Observe that this system as depicted here, requires communication with other systems. For instance one could assume that two cooling systems are working in the power plant so that when one is down, the other one can still do the job. In that case, the plant would burn only if the two systems are beyond repair leading thus to a synchronisation on action *burn*. On the other hand, one could build an ISTA representing a worker who fixes the cooling system. Then the worker and the cooling system would need to communicate through actions *down*, *fix* and *replace*. We will come back to this in Example 10.2.5 when dealing with the handshaking composition.

As quickly observed in Example 10.1.2, and like in the IMC model, interactive transitions will allow for interaction between systems. The internal action τ has again a special role, and like in IMCs, we will make the *maximal progress assumption*: we suppose that in any ISTA, internal interactive transitions take precedence over the edges (*i.e.* the probabilistic transitions given by E) or, in other words, we assume that when a τ -labelled transition is enabled, it is taken immediately.

We are not yet ready to define a semantics for ISTA. Indeed, interactive

transitions (with external actions) can be seen as signals in the model. The system can then fire a given interactive transition once it has communicated on the corresponding signal with some required other system(s). In other words, interactive transitions need synchronisation(s). Therefore, we need to define an operator of composition with synchronisations. Then, once all interactions are done and thus the system does not need to synchronise, we can abstract all external actions into the internal action τ : due to the maximal progress assumption, we thus suppose that an enabled interaction occurs immediately. This abstraction is done through a *hiding* operator. Once all actions are hidden, we say that the system is closed. The semantics of the resulting ISTA will then be given through a STA.

Remark 10.1.3. Due to the presence of interactive locations, one cannot speak immediately of a STA induced by an ISTA. Edges (*i.e.* the probabilistic transitions given by E) can arrive in interactive locations. But such locations have no sense in a STA.

Before going further in the chapter, we need to establish some conditions on the ISTA that we will consider in the sequel. Firstly, like in STA, we assume that there are no-blocking states in ISTA. The set of states is defined as in Definition 2.1.5 by $Q = \{(l, \nu) \in L \times \mathbb{R}_+^X \mid \nu \models \text{Inv}(l)\}$ and a state q is blocking like in Definition 2.1.16, if $I(q) = \emptyset$, where $I(q)$ is defined as in the case of STA: it is the set of delays after which an edge is enabled. We do not consider interactive transitions in $I(q)$. For each $e \in E$ with $\text{source}(e) = l$, $I((l, \nu), e) = \{t \in \mathbb{R}_+ \mid \nu + t \models \text{Inv}(l) \wedge \nu + t \models g\}$ and $I((l, \nu)) = \bigcup_{e \in E_l} I((l, \nu), e)$ where $E_l = \{e \in E \mid \text{source}(e) = l\}$. We also assume that once an interactive transition is enabled, it is always enabled in the future for all states of the ISTA. For the case of an interactive location l , since $\text{Inv}(l) = \text{true}$, this implies that each interactive transition from l is guarded by a constraint of the form $\bigwedge_{x \in Y} x \geq c_x$ for some $Y \subseteq X$ and where for each $x \in Y$, $c_x \in \mathbb{N}$. We formalize those hypotheses below:

$$(I1) \text{ for each } l \in L_{\text{int}}^c \text{ and for each } \nu \in \mathbb{R}_+^X, \nu \models \text{Inv}(l) \implies I((l, \nu)) \neq \emptyset;$$

$$(I2) \text{ for each } l \in L, \text{ for each } \nu \in \mathbb{R}_+^X \text{ and for each } d = (l, a, g, Y, l') \in \dashrightarrow, \nu \models g \implies (\forall t \geq 0, (\nu + t \models \text{Inv}(l) \implies \nu + t \models g));$$

It should be noted that the ISTA $\mathcal{A}_{\text{cool}}$ does not satisfy hypothesis (I2) due to the interactive transition $\text{Wait} \xrightarrow{\text{down}, x \leq 10} \text{Down}$, although $\text{Inv}(\text{Wait}) = \text{true}$.

Like in Chapter 9, we also make the assumptions that the probability distributions over the delays satisfy hypothesis (‡) of page 65 and that the distributions over the edges satisfy condition (★) of page 176.

10.2. Semantics of ISTA

In this section, we define the semantics of ISTA. In order to do so, we first define a parallel composition operator that is handshaking, and a hiding operator. This is done in Section 10.2.1. Like in IMCs ([Her02] and [HK09]), the parallel composition operator is interleaving for non-interactive transitions, but it synchronises on the interactive transitions. The hiding operator abstracts some set of actions into the internal action τ . The idea is like in IMCs (see Section 2.4.2 for some details) that once an action does not need interactions anymore, it is hidden.

Once a system has performed all possible interactions, all actions are hidden. Then the semantics of the resulting ISTA can be given as a STA. This will be the subject of Section 10.2.2 and will be plainly explained.

10.2.1 Parallel composition and hiding operator

In this section, we define a handshaking operator of parallel composition and a hiding operator for ISTA. We then define a notion of bisimulation that extends the one on IMCs ([Her02] and [HK09]; see Section 2.4.2) and show that it is a congruence w.r.t. parallel composition and hiding.

We fix an ISTA \mathcal{A} for the rest of the section and we assume that it satisfies conditions (I1) and (I2) of page 233. As already said before the set of states Q is the set $\{(l, \nu) \in L \times \mathbb{R}_+^X \mid \nu \models \text{Inv}(l)\}$. Given a state q , we write $I_{\text{Inv}}(q) = \{t \in \mathbb{R}_+ \mid \nu + t \models \text{Inv}(l)\}$ recovering a notation used in the proof of Lemma 9.2.1. As stated in the same proof and similarly to arguments in Remark 3.1.1 or for Lemma 9.1.1, we can establish the following technical result.

Lemma 10.2.1. Given an ISTA \mathcal{A} and a state $q = (l, \nu)$, there is $t \in]0, \infty]$ such that $I_{\text{Inv}}(q) = [0, t[$.

Similarly to the class CSTA of Chapter 9 page 177 in STA, we define the class CISTA as the class of ISTA satisfying conditions (A) and (B) of CSTA. We recall that condition (A) asks that each distribution over the delays is given by a density function continuous everywhere, except in a finite number of points; condition (B) states that those distributions are *weakly-memoryless*. We now define when two ISTA are composable. We use the following notation: we write

$$\Gamma_l = \{a \in \Gamma \mid \exists g \in \mathcal{G}_{\text{int}}(X), \exists Y \subseteq X, \exists l' \in L, (l, a, g, Y, l') \in \dashrightarrow\}.$$

Given a location $l \in L$, Γ_l represents thus the set of actions a that are enabled in l , *i.e.* such that there is an interactive transition starting from l and labelled with a .

Definition 10.2.2. Consider for each $i \in \{1, 2\}$, ISTA $\mathcal{A}_i = (L_i, X_i, E_i, \text{Inv}_i, \text{AP}_i, \mathcal{L}_i, (\mu_q^{(i)}, p_q^{(i)})_{q \in L_i \times \mathbb{R}_+^{X_i}}, \Gamma_i^\tau, \dashrightarrow_i)$ and let $A \subseteq \Gamma_1 \cap \Gamma_2$. We say that \mathcal{A}_1 and \mathcal{A}_2 are *composable w.r.t. A* if \mathcal{A}_1 and \mathcal{A}_2 are in CISTA, $X_1 \cap X_2 = \emptyset$ and if for each $l_1 \in L_{1,\text{int}}$ and each $l_2 \in L_{2,\text{int}}$ one of the following statements holds true:

- $\Gamma_{l_1} \not\subseteq A$, or
- $\Gamma_{l_2} \not\subseteq A$, or
- $\Gamma_{l_1} \cap \Gamma_{l_2} \neq \emptyset$.

In this definition, A represents the set of actions on which the two systems will synchronise. The two first requirements imply that at in l_1 or in l_2 that are two interactive locations, at least one of the systems does not have to wait for the other system to synchronise: $\Gamma_{l_1} \not\subseteq A$ (resp. $\Gamma_{l_2} \not\subseteq A$) states that there is an interactive transition from l_1 (resp. l_2) that is labelled with some action b that is not in A . The last requirement is here in order to avoid a situation where only interactive transitions are possible but both systems are waiting for communications that are not possible and thus blocking the system. This will be made clearer in Remark 10.2.4 We can now define the handshaking composition operator.

Definition 10.2.3. Consider for each $i \in \{1, 2\}$, ISTA $\mathcal{A}_i = (L_i, X_i, E_i, \text{Inv}_i, \text{AP}_i, \mathcal{L}_i, (\mu_q^{(i)}, p_q^{(i)})_{q \in L_i \times \mathbb{R}_+^{X_i}}, \Gamma_i^\tau, \dashrightarrow_i)$ and let $A \subseteq \Gamma_1 \cap \Gamma_2$ be such that \mathcal{A}_1 and \mathcal{A}_2 are composable w.r.t. A . We define the parallel composition of \mathcal{A}_1 and \mathcal{A}_2 with synchronisations over the set A as the ISTA

$$\mathcal{A}_1 \parallel_A \mathcal{A}_2 = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X}, \Gamma^\tau, \dashrightarrow)$$

where

- $L = L_1 \times L_2$, $X = X_1 \cup X_2$, $E = E_{1,\bullet} \cup E_{\bullet,2}$ (see Definition 2.1.29: $E_{1,\bullet}$ represents the set of edges where the component of \mathcal{A}_1 performs an edge of E_1 while the component of \mathcal{A}_2 does not change), $\text{AP} = \text{AP}_1 \cup \text{AP}_2$;
- for each $(l_1, l_2) \in L$, $\mathcal{L}(l_1, l_2) = \mathcal{L}(l_1) \cup \mathcal{L}(l_2)$, $\text{Inv}(l_1, l_2) = \text{Inv}(l_1) \wedge \text{Inv}(l_2)$;
- $L_{\text{int}} = L_{1,\text{int}} \times L_{2,\text{int}}$;
- for each state $q = ((l_1, l_2), (\nu_1, \nu_2)) \in L_{\text{int}}^c \times \mathbb{R}_+^X$, writing $q_1 = (l_1, \nu_1)$ and $q_2 = (l_2, \nu_2)$,

- ▷ if $(l_1, l_2) \in L_{1,\text{int}}^c \times L_{2,\text{int}}^c$, μ_q and p_q are defined as in Definition 9.1.9: μ_q is the distribution of the minimum between $\mu_{q_1}^{(1)}$ and $\mu_{q_2}^{(2)}$; p_q is defined by a race between \mathcal{A}_1 and \mathcal{A}_2 to win the next edge;
- ▷ if $(l_1, l_2) \in L_{1,\text{int}}^c \times L_{2,\text{int}}$, $(\mu_q, p_q) = (\mu_{q_1}^{(1)}, p_{q_1}^{(1)})$,
- ▷ if $(l_1, l_2) \in L_{1,\text{int}} \times L_{2,\text{int}}^c$, $(\mu_q, p_q) = (\mu_{q_2}^{(2)}, p_{q_2}^{(2)})$;
- $\Gamma = \Gamma_1 \cup \Gamma_2$;
- \dashrightarrow is defined as follows:
 - ▷ for each $a \in A$, it holds that if $l_1 \xrightarrow{g_1, a, Y_1} l'_1$ and if $l_2 \xrightarrow{g_2, a, Y_2} l'_2$, then $(l_1, l_2) \xrightarrow{g_1 \wedge g_2, a, Y_1 \cup Y_2} (l'_1, l'_2)$;
 - ▷ for each $a \notin A$, it holds that if $l_1 \xrightarrow{g_1, a, Y_1} l'_1$, then for each $l_2 \in L_2$, $(l_1, l_2) \xrightarrow{g_1, a, Y_1} (l'_1, l_2)$;
 - ▷ for each $a \notin A$, it holds that if $l_2 \xrightarrow{g_2, a, Y_2} l'_2$, then for each $l_1 \in L_1$, $(l_1, l_2) \xrightarrow{g_2, a, Y_2} (l_1, l'_2)$.

Observe thus that like in IMCs with the rates (see Definition 2.4.3), the parallel composition is interleaving on the edges but also on the interactive transitions whose labels are not in A . Observe also that we cannot synchronise on the internal action τ .

Remark 10.2.4. As already said before, \mathcal{A}_1 and \mathcal{A}_2 are required to be composable w.r.t. A in order to avoid blocking-states: if $l_1 \in L_{1,\text{int}}$ and $l_2 \in L_{2,\text{int}}$ are such that $\Gamma_{l_1} \subseteq A$ and $\Gamma_{l_2} \subseteq A$ and $\Gamma_{l_1} \cap \Gamma_{l_2} = \emptyset$, then (l_1, l_2) has no outgoing (interactive) transitions and thus the system cannot evolve anymore.

We now illustrate the operator of composition on the cooling system of Example 10.1.2.

Example 10.2.5. We consider again the cooling system in the power plant of Example 10.1.2 described on Figure 10.1. In order to simplify the example, we make the assumption in the rest of the chapter that there is no interactive transition between Failure and Up. Moreover, in order to have the conditions (I1) and (I2) satisfied, we assume that guard of the interactive transition for Wait to Down to be true. We will see how we can force the τ -transition towards Failure once 10 time units have elapsed. This time we assume that there are two cooling systems in the power plant: $\mathcal{A}_{\text{cool}_1}$ and $\mathcal{A}_{\text{cool}_2}$. For each $i \in \{1, 2\}$, we thus consider the cooling system $\mathcal{A}_{\text{cool}_i}$ depicted on Figure 10.2. We therefore index each clock, location and action correspondingly to the ISTA it belongs to, except for

the action *burn* which is the only shared action: $\Gamma_1 \cap \Gamma_2 = \{burn\}$. Here we assume that the power plant is burned only when both cooling systems are beyond repair. We partially describe the ISTA $\mathcal{A}_{cool_1} \parallel_{\{burn\}} \mathcal{A}_{cool_2}$ on Figure 10.3.

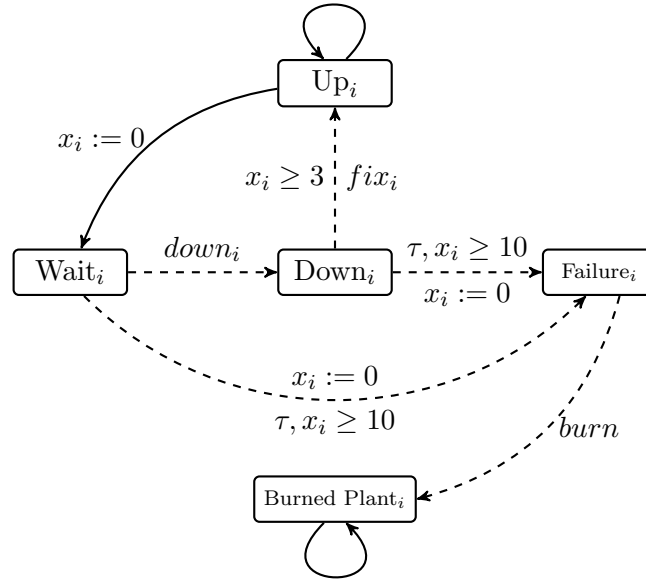


Figure 10.2: A slight variant of the cooling system; ISTA \mathcal{A}_{cool_i}

As said before, the ISTA described on Figure 10.3 is only a partial representation of $\mathcal{A}_{cool_1} \parallel_{\{burn\}} \mathcal{A}_{cool_2}$. Missing locations are (F_1, U_2) , (U_1, F_2) , (F_1, W_2) and (W_1, F_2) . Missing edges are all the self-loops when needed (for instance, two self-loops on (U_1, U_2) : one for \mathcal{A}_{cool_1} and the other one for \mathcal{A}_{cool_2}). Missing interactive transitions are for instance $(D_1, W_2) \xrightarrow{fix_{x_1}, x_1 \geq 3}$, (U_1, W_2) or $(W_1, D_2) \xrightarrow{fix_{x_2}, x_2 \geq 3}$, (W_1, U_2) . Observe that from (F_1, D_2) and (D_1, F_2) there should indeed not have interactive transitions with label *burn* (which are enabled from F_1 and F_2) as the synchronisation is not possible there. The invariant function is given trivially. The distributions over the delays and edges are not given as they are computed as in Chapter 9 (for instance Example 9.1.10). Observe however that, for instance, in location (W_1, U_2) , since W_1 is an interactive location, then (W_1, U_2) takes the distributions of location U_2 in \mathcal{A}_{cool_2} , namely $\text{Exp}(\lambda_1)$ and probability one half on the self-loop edge and one half on the edge toward Wait_2 .

We now briefly describe an ISTA representing a worker that should fix the cooling systems when they are down. It is given as ISTA \mathcal{A}_{worker} on Figure 10.4.

First observe that there is no interactive location, therefore each location

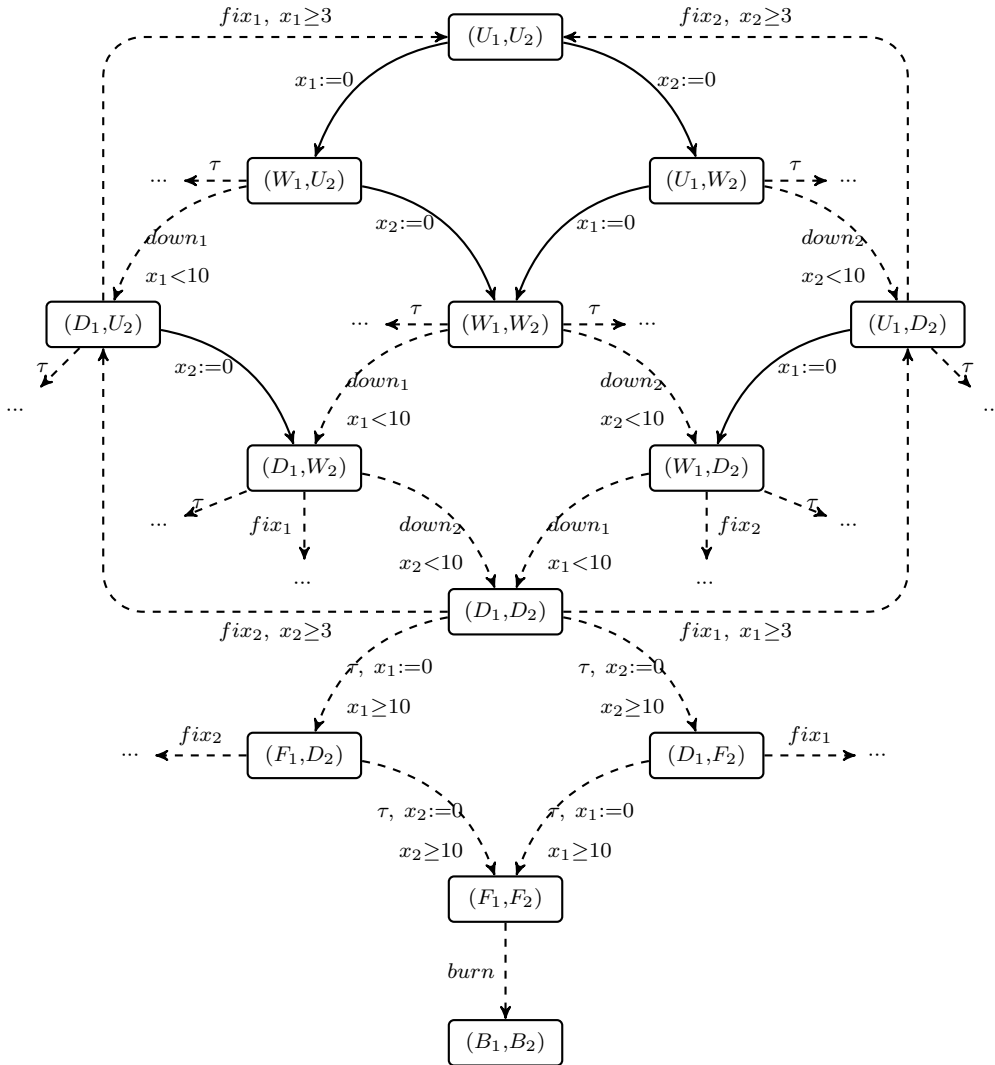


Figure 10.3: The handshaking composition $\mathcal{A}_{cool_1} \parallel_{\{burn\}} \mathcal{A}_{cool_2}$

should be equipped with adequate distributions. The invariant function should be given as follows: $\text{Inv}(\text{Other Tasks}) = y < 4$ and $\text{Inv}(l) = \text{true}$ for any other location. Now \mathcal{A}_{worker} represents a worker such that in location Idle, he is unoccupied. At any time, he can do some tasks in location Other Tasks and we make the assumption that he cannot spend more than 4 time units on his tasks. Once


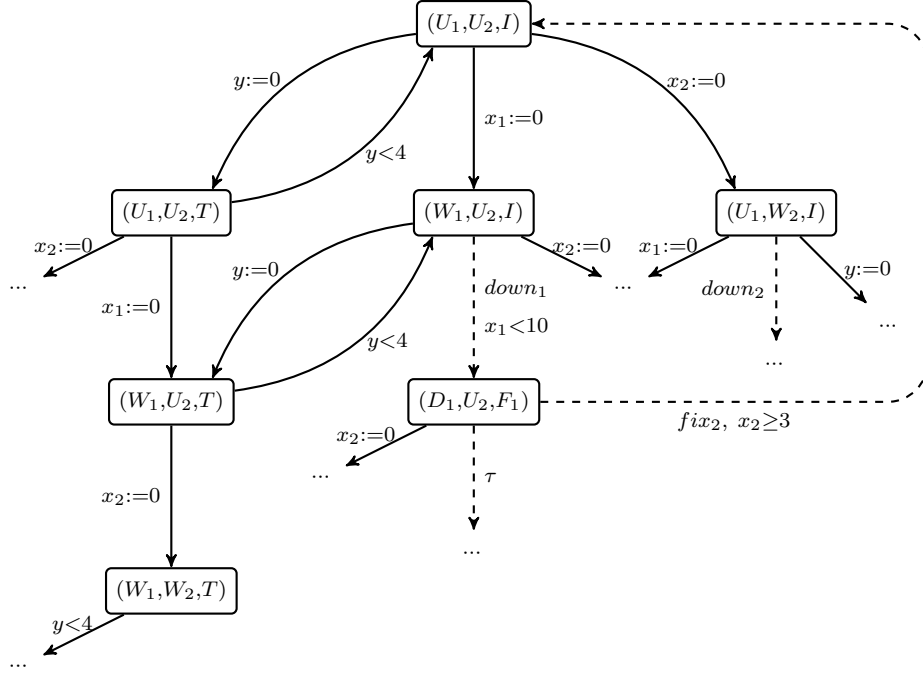


Figure 10.4: An ISTA $\mathcal{A}_{\text{worker}}$ representing a worker in a power plant, fixing cooling systems.

done, he goes back to location Idle where the same process happens. However, as soon as one of the cooling system is down, he has to go and repair it in Fixing₁ or Fixing₂ depending on which system is down or if both are down, depending on which one he chooses to fix first. Observe that if he is busy with fixing one of the system, he cannot take care of the other system. For each system, we assume that it takes at least 3 time units to repair it. At any time, if both systems are getting beyond repair while he is fixing one of them, the burning of the plant would lead to his death in location Dead.

Then, if we write $\Gamma' = \{down_1, down_2, fix_1, fix_2, burn\}$ for the set of actions of $\mathcal{A}_{\text{worker}}$, we are interested by the product $(\mathcal{A}_{\text{cool}_1} \parallel_{\{burn\}} \mathcal{A}_{\text{cool}_2}) \parallel_{\Gamma'} \mathcal{A}_{\text{worker}}$. Some of the first steps of this ISTA are given in Figure 10.5. Observe that F_1 and F_2 in the third component correspond to Fixing₁ and Fixing₂. Furthermore, T in the third component stands for Other Tasks.

The figure is quite self-explanatory, we just make a few comments. First we have again omitted the self-loops when it is required, and we will not be giving the invariant function nor the new distributions. We have also omitted the τ -labelled transition corresponding to Wait₁ or Wait₂ in each location. Observe that in location (W_1, U_2, T) and (W_1, W_2, T) , there is no interactive transition labelled with $down_1$ or $down_2$ while those are allowed in W_1 and W_2 . It comes from the


 Figure 10.5: The handshaking composition $(\mathcal{A}_{\text{cool}_1} \parallel_{\{\text{burn}\}} \mathcal{A}_{\text{cool}_2}) \parallel_{\Gamma'} \mathcal{A}_{\text{worker}}$

fact that we synchronise the worker and the two cooling systems on down_1 and down_2 in particular. In location T , the worker is doing some other task(s) and cannot receive the signal alerting that one or both cooling systems are down. On the other hand, as an example, transition $(W_1, U_2, I) \xrightarrow{\text{down}_1} (D_1, U_2, F_1)$ is the result of a synchronisation on down_1 between the product $\mathcal{A}_{\text{cool}_1} \parallel_{\{\text{burn}\}} \mathcal{A}_{\text{cool}_2}$ and $\mathcal{A}_{\text{worker}}$. Finally, one can also observe that the communication on burn can only occur in locations (F_1, F_2, F_1) or (F_1, F_2, F_2) , *i.e.* when the two cooling systems are beyond repair while the worker is trying to fix one of them.

Remark 10.2.6. Like in Chapter 9 with the interleaving composition (see Remark 9.2.4), the handshaking is in some sense commutative. One can also get associativity as follows. Fix three ISTA \mathcal{A}_1 , \mathcal{A}_2 and \mathcal{A}_3 with respectively Γ_1 , Γ_2 and Γ_3 for the set of actions. Fix $A_1 \subseteq \Gamma_1 \cap \Gamma_2$ and $A_2 \subseteq (\Gamma_1 \cup \Gamma_2) \cap \Gamma_3$. Observe that $A_2 = (A_2 \cap \Gamma_1) \cup (A_2 \cap \Gamma_2)$. We can show that

$$(\mathcal{A}_1 \parallel_{A_1} \mathcal{A}_2) \parallel_{A_2} \mathcal{A}_3 = \mathcal{A}_1 \parallel_{(A_2 \cap \Gamma_1) \cup A_1} (\mathcal{A}_2 \parallel_{A_2 \cap \Gamma_2} \mathcal{A}_3).$$

Similarly as in Chapter 9 with Lemma 9.2.1, we need to prove that parallel composition is well-defined, *i.e.* that the probability distributions are well-defined, but also that all the hypotheses made in Definition 10.1.1 are met. We assume that \mathcal{A}_1 and \mathcal{A}_2 are two ISTA, that $A \subseteq \Gamma_1 \cap \Gamma_2 \neq \emptyset$ and that \mathcal{A}_1 and \mathcal{A}_2 are composable w.r.t. A .

Lemma 10.2.7. The product $\mathcal{A}_1 \parallel_A \mathcal{A}_2$ is well-defined and is an ISTA in the class CISTA. Moreover, $\mathcal{A}_1 \parallel_A \mathcal{A}_2$ satisfies conditions (I1), (I2), (\ddagger) and (\star).

Proof. Given a state $q = ((l_1, l_2), (\nu_1, \nu_2))$ with $(l_1, l_2) \in L_{\text{int}}^c$, the fact that the distributions μ_q and p_q are well-defined and that the distributions μ_q are weakly-memoryless comes from the definition since either they are given as $\mu_q^{(1)}$ and $p_q^{(1)}$ (if $l_1 \in L_{1,\text{int}}^c$ and $l_2 \in L_{2,\text{int}}$), as $\mu_q^{(2)}$ and $p_q^{(2)}$ (if $l_1 \in L_{1,\text{int}}$ and $l_2 \in L_{2,\text{int}}^c$) or as defined in Definition 9.1.9 (if $l_1 \in L_{1,\text{int}}^c$ and $l_2 \in L_{2,\text{int}}^c$) and then Lemma 9.2.1 allows us to conclude. We get similarly that the distributions satisfy conditions (A) and (B) of CISTA and conditions (\ddagger) and (\star). It is trivial to get that $\mathcal{A}_1 \parallel_A \mathcal{A}_2$ is still an ISTA as defined in Definition 10.1.1. It remains to show that condition (I1) and (I2) are satisfied, *i.e.*

(I1) for each $l \in L_{\text{int}}^c$ and for each $\nu \in \mathbb{R}_+^X$, $\nu \models \text{Inv}(l) \implies I((l, \nu)) \neq \emptyset$;

(I2) for each $l \in L$, for each $\nu \in \mathbb{R}_+^X$ and for each $d = (l, a, g, Y, l') \in \dashrightarrow$, $\nu \models g \implies (\forall t \geq 0, (\nu + t \models \text{Inv}(l) \implies \nu + t \models g))$;

We begin with (I1). We fix $l = (l_1, l_2) \in L_{\text{int}}^c$ and $\nu = (\nu_1, \nu_2) \in \mathbb{R}_+^X$, and assume that $\nu \models \text{Inv}(l)$. We have to prove that $I(l, \nu) \neq \emptyset$. First it should be noted from the definition of $I(q)$ that

$$\begin{aligned} I(q) &= \{t \in \mathbb{R}_+ \mid (t \in I(q_1) \wedge \nu_2 + t \models \text{Inv}(l_2)) \vee (t \in I(q_2) \wedge \nu_1 + t \models \text{Inv}(l_1))\} \\ &= (I(q_1) \cap I_{\text{Inv}}(q_2)) \cup (I(q_2) \cap I_{\text{Inv}}(q_1)) \end{aligned}$$

where $q_1 = (l_1, \nu_1)$ and $q_2 = (l_2, \nu_2)$. Thus, $I(l, \nu) \neq \emptyset$ if and only if $I(q_1) \cap I_{\text{Inv}}(q_2) \neq \emptyset$ or $I(q_2) \cap I_{\text{Inv}}(q_1) \neq \emptyset$. First observe that if $l_1 \in L_{1,\text{int}}$, then $I(q_1) = \emptyset$, $\text{Inv}(l_1) = \text{true}$ and $l_2 \in L_{2,\text{int}}^c$. It follows that $I(q_2) \neq \emptyset$ (as \mathcal{A}_2 satisfies (I1)) and $I_{\text{Inv}}(q_1) = [0, \infty[$. We thus straightforwardly get that $I(q_2) \cap I_{\text{Inv}}(q_1) \neq \emptyset$. A similar argument holds if $l_2 \in L_{2,\text{int}}$. We can thus assume $l_1 \in L_{1,\text{int}}^c$ and $l_2 \in L_{2,\text{int}}^c$. Towards a contradiction, assume that

$$I(q_1) \cap I_{\text{Inv}}(q_2) = \emptyset \text{ and } I(q_2) \cap I_{\text{Inv}}(q_1) = \emptyset. \quad (10.1)$$

From Lemma 10.2.1, it holds that there are c_1 and $c_2 \in]0, \infty]$ such that $I_{\text{Inv}}(q_1) = [0, c_1[$ and $I_{\text{Inv}}(q_2) = [0, c_2[$. Since $l_1 \in L_{1,\text{int}}^c$ and $l_2 \in L_{2,\text{int}}^c$, and since $\nu \models$

$\text{Inv}(l) = \text{Inv}(l_1) \wedge \text{Inv}(l_2)$, we get that $I(q_1) \neq \emptyset$ and $I(q_2) \neq \emptyset$ (from hypothesis (I1) for \mathcal{A}_1 and \mathcal{A}_2). W.l.o.g., suppose that $c_2 \geq c_1$. From hypothesis (10.1) we get that

$$I(q_1) \subseteq [c_2, \infty[\subseteq [c_1, \infty[= I_{\text{Inv}}(q_1)^c$$

which contradicts the definitions of $I(q_1) = \{t \in \mathbb{R}_+ \mid \nu_1 + t \models \text{Inv}(l_1) \wedge \exists e_1 = (l_1, g_1, Y_1, l'_1) \in E_1, \nu_1 + t \models g\} \subseteq I_{\text{Inv}}(q_1)$. This proves point (I1).

Now in order to get (I2), fix $l = (l_1, l_2) \in L$, $\nu = (\nu_1, \nu_2) \in \mathbb{R}_+^X$ and $d = (l, a, g, Y, l') \in \dashrightarrow$ such that $\nu \models g$. First if $a \in A$, then $d = ((l_1, l_2), a, g_1 \wedge g_2, Y_1 \cup Y_2, (l'_1, l'_2))$ where $(l_1, a, g_1, Y_1, l'_1) \in \dashrightarrow_1$ and $(l_2, a, g_2, Y_2, l'_2) \in \dashrightarrow_2$. As $\nu \models g$, we have that $\nu_1 \models g_1$ and $\nu_2 \models g_2$. Now fix $t \geq 0$ and assume that $\nu + t \models \text{Inv}(l)$, *i.e.* $\nu_1 + t \models \text{Inv}(l_1)$ and $\nu_2 + t \models \text{Inv}(l_2)$. From condition (I2) for \mathcal{A}_1 and \mathcal{A}_2 , we thus get that $\nu_1 + t \models g_1$ and $\nu_2 + t \models g_2$ and thus $\nu + t \models g$. If $a \notin A$, then either $d = ((l_1, l_2), a, g_1, Y_1, (l'_1, l'_2))$ with $(l_1, a, g_1, Y_1, l'_1) \in \dashrightarrow_1$, or $d = ((l_1, l_2), a, g_2, Y_2, (l'_1, l'_2))$ with $(l_2, a, g_2, Y_2, l'_2) \in \dashrightarrow_2$. W.l.o.g., assume that we are in the first case, it thus holds that $\nu_1 \models g_1$. Finally fix $t \geq 0$ and assume that $\nu + t \models \text{Inv}(l)$, *i.e.* $\nu_1 + t \models \text{Inv}(l_1)$ and $\nu_2 + t \models \text{Inv}(l_2)$. Again we get that $\nu_1 + t \models g_1$ and therefore $\nu \models g$. This concludes the proof. \square

As already explained in Section 9.3 but also in Sections 2.1.3 and 2.4.2, bisimulation is an important notion when dealing with parallel composition. We thus define here a notion of bisimulation for ISTA that extends the one of IMCs (see [Her02], [HK09] and Section 2.4.2) and show that it is a congruence w.r.t. parallel composition. We fix an ISTA \mathcal{A} .

We first need a new notation. Given a state $q = (l, \nu)$, we write $I_\tau(q) = \{t \in \mathbb{R}_+ \mid \nu + t \models \text{Inv}(l) \text{ and } \exists (l, \tau, g, Y, l') \in \dashrightarrow, \nu + t \models g\}$, *i.e.* it corresponds to the set of delays after which, starting from q , a τ -action is enabled. We then have the following result.

Lemma 10.2.8. For each state q , it holds that $I_\tau(q)$ is either the empty set or a finite union of disjoint intervals of the form $[t_1, t_2[$ with $t_1, t_2 \in \mathbb{R}_+$.

The proof uses similar arguments as in Remark 3.1.1 and Lemmas 9.1.1 and 10.2.1. Here it comes from Lemma 10.2.1 and from the fact that guards on interactive transitions are in $\mathcal{G}_{\text{int}}(X)$. With this result, one can establish that for each state q , if $I_\tau(q) \neq \emptyset$, then there exists a minimal element in this set, *i.e.* $\min(I_\tau(q))$ exists. If $I_\tau(q) = \emptyset$, we will abusively write $\min(I_\tau(q)) = \infty$. This is why guards in $\mathcal{G}_{\text{int}}(X)$ are needed for interactive transitions.

We can now define a notion of bisimulation for ISTA. We refer to Section 9.3.1 for notations on equivalence relations and in particular on the closure of an

equivalence relation w.r.t. polyhedral sets. We recall that for each state $q = (l, \nu)$ with l non-interactive, f_q denotes the density function of μ_q and for an equivalence relation \mathcal{R} and $C = \bigcup_{l' \in L} \{l'\} \times C_{l'} \in \text{pcl}(\mathcal{R})$ (the closure of \mathcal{R} w.r.t. polyhedral sets), we write for each $t \geq 0$,

$$P_{q+t}(C) = \sum_{l' \in L} \sum_{e \in E_{l'}} p_{q+t}(e) \mathbb{1}_{C_{l'}(e, \nu)}(t)$$

where $E_{l'}$ is the subset of E with target l' and given $e = (l, g, Y, l')$, $C_{l'}(e, \nu) = \{t \in \mathbb{R}_+ \mid [Y \leftarrow 0](\nu + t) \in C_{l'}\}$.

Definition 10.2.9. Let $\mathcal{A} = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X}, \Gamma^\tau, \dashrightarrow)$ be an ISTA. An equivalence relation \mathcal{R} over the set of states Q is a *bisimulation for \mathcal{A}* if for each $q_1, q_2 \in Q$ with $q_1 \mathcal{R} q_2$ with $q_i = (l_i, \nu_i)$ for $i = 1, 2$,

- (1) $\mathcal{L}(q_1) = \mathcal{L}(q_2)$;
- (2) $l_1 \in L_{\text{int}} \iff l_2 \in L_{\text{int}}$;
- (3) $I_\tau(q_1) = \emptyset \iff I_\tau(q_2) = \emptyset$ and moreover if $I_\tau(q_1) \neq \emptyset$, then $\min(I_\tau(q_1)) = \min(I_\tau(q_2)) = t^*$;
- (4) if $l_1, l_2 \in L_{\text{int}}^c$, for each $C \in \text{pcl}(\mathcal{R})$, for almost-surely each $t < t^*$,
 - (a) $f_{q_1}(t) = f_{q_2}(t)$, and
 - (b) $P_{q_1+t}(C) = P_{q_2+t}(C)$,
- (5) (a) for each $a \in \Gamma$, for each $t < t^*$,

$$\exists q'_1, q_1 + t \xrightarrow{a} q'_1 \implies \exists q'_2, q_2 + t \xrightarrow{a} q'_2 \wedge q'_1 \mathcal{R} q'_2,$$

and vice versa,

- (b) if $t^* < \infty$,

$$\exists q'_1, q_1 + t^* \xrightarrow{\tau} q'_1 \implies \exists q'_2, q_2 + t^* \xrightarrow{\tau} q'_2 \wedge q'_1 \mathcal{R} q'_2,$$

and vice versa.

States q_1 and q_2 are *bisimilar* (written $q \sim q'$) if there is a bisimulation that contains (q, q') .

One can establish links with Definitions 2.4.7 and 9.3.1, and with Proposition 9.3.4. Firstly, observe that point (4) of Definition 10.2.9, is the counter-part of point (ii) of Definition 9.3.1 thanks to the characterisation of bisimulation in STA of Proposition 9.3.4. Observe that here, we require equivalence of the distributions only on $[0, t^*[$. Like in IMCs, this is due to the maximal progress assumption: internal actions take precedence over edges. Then, it has no importance how time progresses afterwards. Point (4) somewhat corresponds the second point of the definition of bisimulation in IMCs (Definitions 2.4.7 here). In [HK09], the authors require the distributions to be equivalent only when no internal action is allowed. This is due to the fact that in IMCs all guards would be **true** on each edge and interactive transitions. Which is not the case here in ISTA. Point (5) of Definition 10.2.9 corresponds to the first point of Definition 2.4.7, also keeping in mind that all guards are **true** in IMCs. Finally, point (1) is classical, points (2) and (3) are here to make sure that firstly, if time can progress in one state, then it must also be able to progress in the other state, and secondly, the next τ -action happens at the same time in both states.

We illustrate the notion on a simple example.

Example 10.2.10. We consider a simple variant of the STA of Example 9.3.5 (see Figure 9.4). First we add some location l_3 labelled by $\mathcal{L}(l_3) = \{b\}$ and we add an interactive transition from l_0 to l_3 with guard $x_1 \geq 1$ and labelled with τ , turning the STA into an ISTA. We assume that l_3 has only a self-loop with guard **true** and that the distribution over the delays in this location is always an exponential distribution of parameter $\mu > 0$. We recall that there is an edge from l_0 to l_1 and from l_0 to l_2 , but this time we assume that guards are given by respectively $0 < x_1 < 2$ and $0 < x_2 < 2$. From a state of the form $q = (l_0, (\nu_1, \nu_2))$ with $\nu_1 < 2$ or $\nu_2 < 2$, $I(q) = [0, 2 - \min(\nu_1, \nu_2)[$ and so we can again equip q with a uniform distribution on the interval $I(q)$ for the delays. Finally, we again consider that each state of the form (l_1, ν) or (l_2, ν) with $\nu \in \mathbb{R}_+^2$ is equipped with the same exponential distribution over delays, say $\text{Exp}(\lambda)$. We assume $\lambda \neq \mu$.

Like in Example 9.3.5, we compute the equivalence classes of \sim . First since $\lambda \neq \mu$, it is easily seen that $\{l_1, l_2\} \times \mathbb{R}_+^2$ and $\{l_3\} \times \mathbb{R}_+^2$ are two different classes.

We are now interested in location l_0 . It is again very similar to Example 9.3.5 except that here, once $x_1 \geq 1$, the τ -transition is always immediately enabled. It follows that $\{l_0\} \times ([1, \infty[\times \mathbb{R}_+)$ is an equivalence class of \sim . Finally, one can show that for each $\nu \in [0, 1[$, the set $\{l_0\} \times (\{(\nu_1, \nu) \mid \nu_1 \geq \nu\} \cup \{(\nu, \nu_2) \mid \nu_2 \geq \nu\})$. We thus get almost the same classes as in Example 9.3.5, the only difference comes from the τ -transition: once such a transition is enabled, we do not care about the distributions over the delays (due to the maximal progress assumption).




Figure 10.6: Bisimulation in ISTA

Like in Chapter 9, we will prove that bisimulation is a congruence w.r.t. parallel composition. Recall that in order to do so we need to define a bisimulation between ISTA. Like in STA (see Section 9.3.1 page 213) one can build a union ISTA from two ISTA \mathcal{A}_1 and \mathcal{A}_2 . We will then use similar terminology: ISTA \mathcal{A}_1 and \mathcal{A}_2 are said *bisimilar from states* q_1 and q_2 , and we write $\mathcal{A}_1 \sim \mathcal{A}_2$, if $q_1 \sim q_2$ in $\mathcal{A}_1 \cup \mathcal{A}_2$. Observe that for an interesting bisimulation, it should be the case that $\Gamma_1 = \Gamma_2$.

Before proving that \sim is a congruence w.r.t. parallel composition, we first need two technical lemmas. The first one is similar to Lemma 9.3.10.

Lemma 10.2.11. Let \mathcal{A} be an ISTA in CISTA and let \mathcal{R} be a bisimulation for \mathcal{A} . Let q_1 and q_2 be two states of \mathcal{A} . If $q_1 \mathcal{R} q_2$ and if $t^* = \min(I_\tau(q_1)) = \min(I_\tau(q_2))$, then $(q_1 + t) \mathcal{R} (q_2 + t)$ for each $t < t^*$. Moreover,

- (i) if $t^* = \infty$ (i.e. $I_\tau(q_1) = I_\tau(q_2) = \emptyset$), then $I_{\text{Inv}}(q_1) = I_{\text{Inv}}(q_2)$;
- (ii) if $t^* < \infty$, then $[0, t^*] \subseteq I_{\text{Inv}}(q_1) \cap I_{\text{Inv}}(q_2)$.

Proof. Fix $q_1 = (l_1, \nu_1)$ and $q_2 = (l_2, \nu_2)$ two states of \mathcal{A} such that $q_1 \mathcal{R} q_2$, and write $t^* = \min(I_\tau(q_1)) = \min(I_\tau(q_2))$ with the convention that $t^* = \infty$ if $I_\tau(q_1) = I_\tau(q_2) = \emptyset$. Let us prove that for each $t < t^*$, $(q_1 + t) \mathcal{R} (q_2 + t)$. Points (1) and (2) of Definition 10.2.9 are easily checked as they only depend on the locations and since $q_1 \mathcal{R} q_2$. For point (3), notice that if $t^* = \infty$, then $I_\tau(q_1) = I_\tau(q_2) = \emptyset$ and thus $I_\tau(q_1 + t) = I_\tau(q_2 + t) = \emptyset$, while if $t^* < \infty$, then $I_\tau(q_1 + t) \neq \emptyset$ and $I_\tau(q_2 + t) \neq \emptyset$ since $t < t^*$. Thus $I_\tau(q_1 + t) = \emptyset$ if and only if $I_\tau(q_2 + t) = \emptyset$. Finally for the second case, it is easily seen that $\min(I_\tau(q_1 + t)) = \min(I_\tau(q_2 + t)) = t^* - t$.

Establishing point (4) is very similar to the proof of Lemma 9.3.10 except that here, we want the equivalence of distributions only on $[0, t^* - t[$.

Finally, we have to check (5). First fix $a \in \Gamma$ and $t' < t^* - t$, and assume that there is q'_1 such that $(q_1 + t) + t' \xrightarrow{a} q'_1$, that is $q_1 + (t + t') \xrightarrow{a} q'_1$. As $q_1 \mathcal{R} q_2$, we know that there is q'_2 such that $q_2 + (t + t') \xrightarrow{a} q'_2$ and $q'_1 \mathcal{R} q'_2$. Hence there is q'_2 such that $(q_2 + t) + t' \xrightarrow{a} q'_2$ and $q'_1 \mathcal{R} q'_2$, which shows point (a). Suppose now that there is q'_1 such that $(q_1 + t) + (t^* - t) \xrightarrow{\tau} q'_1$, thus $q_1 + t^* \xrightarrow{\tau} q'_1$. By a similar argument as before, there exists q'_2 such that $(q_2 + t) + (t^* - t) \xrightarrow{\tau} q'_2$ which thus proves that $(q_1 + t) \mathcal{R} (q_2 + t)$.

It remains to show points (i) and (ii). Assume that $t^* = \infty$, i.e. $I_\tau(q_1) = I_\tau(q_2) = \emptyset$. Then from the hypothesis, we get that $f_{q_1} = f_{q_2}$ almost-surely on \mathbb{R}_+ . Thus $\mu_{q_1} = \mu_{q_2}$ and it follows that $I(q_1) = I(q_2)$. Towards a contradiction, assume that $I_{\text{Inv}}(q_1) \neq I_{\text{Inv}}(q_2)$. Then, w.l.o.g. we can assume that there is $t \geq 0$ such $\nu_1 + t \models \text{Inv}(l_1)$ and $\nu_2 + t \not\models \text{Inv}(l_2)$. From Lemma 10.2.1, we thus get that for each $t' \geq t$, $t' \notin I_{\text{Inv}}(q_2)$. And from hypothesis (I1) on \mathcal{A}_1 (there are no blocking-states), since $\nu_1 + t \models \text{Inv}(l_1)$, we have that $I(l_1, \nu_1 + t) \neq \emptyset$. Hence, there is $t' \geq t$ such that $t' \in I(q_1) = I(q_2)$. From the definition of $I(q_2)$, we get $t' \in I(q_2) \subseteq I_{\text{Inv}}(q_2)$ which leads to the contradiction.

Now assume that $t^* < \infty$. From the definitions of $I_\tau(q_1)$ and $I_\tau(q_2)$, it holds that $\nu_1 + t^* \models \text{Inv}(l_1)$ and $\nu_2 + t^* \models \text{Inv}(l_2)$. Hence $t^* \in I_{\text{Inv}}(q_1) \cap I_{\text{Inv}}(q_2)$. Since also $0 \in I_{\text{Inv}}(q_1) \cap I_{\text{Inv}}(q_2)$ because q_1 and q_2 are states, it follows from Lemma 10.2.1 that $[0, t^*] \subseteq I_{\text{Inv}}(q_1), I_{\text{Inv}}(q_2)$. This concludes the proof. \square

We can now prove the counter-part of Lemma 9.3.9 in ISTA.

Lemma 10.2.12. Let \mathcal{A} and $\mathcal{B} \in \text{CISTA}$ with sets of states resp. Q_A and Q_B , and with sets of actions resp. Γ_A and Γ_B . Assume \mathcal{A} and \mathcal{B} are composable w.r.t. any $A \subseteq \Gamma_A \cap \Gamma_B$ and let \mathcal{R} be a bisimulation for \mathcal{A} . Then, the equivalence relation $\mathcal{R}' = \{((q_1, q), (q_2, q)) \in Q_A \times Q_B \mid q_1 \mathcal{R} q_2\}$ is a bisimulation for $\mathcal{A} \parallel_A \mathcal{B}$ for each $A \subseteq \Gamma_A \cap \Gamma_B$.

We will use similar notations as in the proof of Lemma 9.3.9.

Proof. Fix $A \subseteq \Gamma_A \cap \Gamma_B$ and let $(q_1, q) \mathcal{R}' (q_2, q)$ with $q_1 = (l_1, \nu_1) \in Q_A$, $q_2 = (l_2, \nu_2) \in Q_A$ and $q = (l, \nu) \in Q_B$. We have to show that (q_1, q) and (q_2, q) satisfy the five points of Definition 10.2.9. Point (1) is immediate, because from Definition 10.2.3, $\mathcal{L}(q_1, q) = \mathcal{L}_A(q_1) \cup \mathcal{L}_B(q)$ and $\mathcal{L}(q_2, q) = \mathcal{L}_A(q_2) \cup \mathcal{L}_B(q)$ and since $q_1 \mathcal{R} q_2$, $\mathcal{L}_A(q_1) = \mathcal{L}_A(q_2)$. The second point can be established as follows:

$$\begin{aligned} (l_1, l) \in L_{\text{int}} &\iff (l_1 \in L_{A, \text{int}}) \wedge (l \in L_{B, \text{int}}) && \text{from Definition 10.2.3} \\ &\iff (l_2 \in L_{A, \text{int}}) \wedge (l \in L_{B, \text{int}}) && \text{since } q_1 \mathcal{R} q_2 \\ &\iff (l_2, l) \in L_{\text{int}}. \end{aligned}$$

In order to get point (3), assume that $I_\tau(q_1, q) \neq \emptyset$ and let us show that $I_\tau(q_2, q) \neq \emptyset$. We consider the first case where $I_\tau(q_1) = \emptyset$ (and thus $I_\tau(q_2) = \emptyset$). From Lemma 10.2.11, it holds that $I_{\text{Inv}}(q_1) = I_{\text{Inv}}(q_2)$. Observe that for each $i = 1, 2$, $I_\tau((q_i, q)) = (I_\tau(q_i) \cap I_{\text{Inv}}(q)) \cup (I_\tau(q) \cap I_{\text{Inv}}(q_i))$. From the previous observations, we easily get $I_\tau((q_1, q)) = I_\tau(q) \cap I_{\text{Inv}}(q_1) \neq \emptyset$ and therefore $I_\tau((q_2, q)) = I_\tau(q) \cap I_{\text{Inv}}(q_2) \neq \emptyset$.

Now assume that $I_\tau(q_1) \neq \emptyset$ (and thus $I_\tau(q_2) \neq \emptyset$). Recall that then, $\min(I_\tau(q_1)) = \min(I_\tau(q_2)) = t^*$ and from Lemma 10.2.11, $[0, t^*] \subseteq I_{\text{Inv}}(q_1) \cap I_{\text{Inv}}(q_2)$. Observe that if $t^* \in I_{\text{Inv}}(q)$, we obviously get that $t^* \in I_\tau((q_1, q)) \cap I_\tau((q_2, q))$. Otherwise if $t^* \notin I_{\text{Inv}}(q)$, then from Lemma 10.2.1, we get that $I_{\text{Inv}}(q) \subseteq [0, t^*[$. It follows that $I_\tau(q_1) \cap I_{\text{Inv}}(q) = \emptyset$ and thus if $t \in I_\tau(q_1, q)$, then $t \in I_\tau(q) \cap I_{\text{Inv}}(q_1)$. Since $I_\tau(q) \subseteq I_{\text{Inv}}(q) \subseteq [0, t^*[$ and $[0, t^*] \subseteq I_{\text{Inv}}(q_2)$, it follows that $t \in I_\tau(q) \cap I_{\text{Inv}}(q_2)$. This concludes that $I_\tau(q_2, q) \neq \emptyset$. Similarly, we can show that if $I_\tau(q_2, q) \neq \emptyset$ then $I_\tau(q_1, q) \neq \emptyset$.

Now, we have to prove that if $I_\tau(q_1, q) \neq \emptyset$ (and thus $I_\tau(q_2, q) \neq \emptyset$) then $\min(I_\tau(q_1, q)) = \min(I_\tau(q_2, q))$. We write $t^* = \min(I_\tau(q_1)) = \min(I_\tau(q_2))$ as usual, and we write $t' = \min(I_\tau(q_1, q))$. We first show that $\min(I_\tau(q_2, q)) \geq t'$. Towards a contradiction, suppose that $\min(I_\tau(q_2, q)) < t'$. Then, there is $t < t'$ such that there is a state (q'_2, q') such that $(q_2, q) + t \xrightarrow{\tau} (q'_2, q')$. There are two cases: first if $t \in I_\tau(q_2)$ (i.e. $q_2 + t \xrightarrow{\tau} q'_2$) and $\nu + t \models \text{Inv}(l)$. In that case, we get that $t^* < t < t'$, i.e. $\min(I_\tau(q_1)) < t < \min(I_\tau(q_1, q))$. Since $t^* < t$ and $\nu + t \models \text{Inv}(l)$, from Lemma 10.2.1 it holds that $\nu + t^* \models \text{Inv}(l)$. It follows that there is $q'_1 \in Q_A$ such that $(q_1, q) + t^* \xrightarrow{\tau} (q'_1, q + t^*)$ which contradicts that $t' = \min(I_\tau(q_1, q))$. Now consider the second case: $t \in I_\tau(q)$ (i.e. $q + t \xrightarrow{\tau} q'$) and $\nu_2 + t \models \text{Inv}(l_2)$. Since $t' = \min(I_\tau(q_1, q))$, we have that $\nu_1 + t' \models \text{Inv}(l_1)$. Hence from Lemma 10.2.1, as $t < t'$, $\nu_1 + t \models \text{Inv}(l_1)$. It follows that $(q_1, q) + t \xrightarrow{\tau} (q_1 + t, q')$ which contradicts that $t' = \min(I_\tau(q_1, q))$. Thus $\min(I_\tau(q_2, q)) \geq t'$. It remains to prove that there is a state (q'_2, q') such that $(q_2, q) + t' \xrightarrow{\tau} (q'_2, q')$.

In order to establish this, let us remark that as $t' = \min(I_\tau(q_1, q))$, there is a state (q'_1, q') such that $(q_1, q) + t' \xrightarrow{\tau} (q'_1, q')$. First assume that $q + t' \xrightarrow{\tau} q'$ and $q'_1 = q_1 + t'$ (i.e. $\nu_1 + t' \models \text{Inv}(l_1)$). Given that $\min(I_\tau(q_2, q)) \geq t'$ we deduce from Lemma 10.2.1 that $[0, t'] \subseteq I_{\text{Inv}}(l_2)$, hence $\nu_2 + t' \models \text{Inv}(l_2)$. It follows that $(q_2, q) + t' \xrightarrow{\tau} (q_2 + t', q')$. Now, if $q_1 + t' \xrightarrow{\tau} q'_1$ and $q' = q + t'$ (i.e. $\nu + t' \models \text{Inv}(l)$), then we get that $[0, t'] \subseteq I_{\text{Inv}}(q)$. Hence, $t' = t^*$ ($= \min(I_\tau(q_1))$). Indeed otherwise, we would have that $t^* < t'$ which then would lead to the fact that, from Lemma 10.2.1 for $\text{Inv}(l)$, $\nu + t^* \models \text{Inv}(l)$ and thus $t^* \in I_\tau(q_1, q)$ which contradicts that $t' = \min(I_\tau(q_1, q))$. We thus have that $t' = \min(I_\tau(q_1)) = \min(I_\tau(q_2))$. Hence there is $q'_2 \in Q_A$ such that $(q_2, q) + t' \xrightarrow{\tau} (q'_2, q + t')$.

We now prove points (5.a) and (5.b). In order to get this, first remark that, given the same notations as before, $t' \leq t^*$. Towards a contradiction, assume that $t' > t^*$. We know that there are $q'_1 \in Q_A$ and $q' \in Q_B$ such that $(q_1, q) + t' \xrightarrow{\tau} (q'_1, q')$. This implies that $\nu + t' \models \text{Inv}(l)$ (as either $t' \in I_\tau(q) \subseteq I_{\text{Inv}}(q)$ or $q' = q + t$ is a state) and from Lemma 10.2.1, it holds that $\nu + t^* \models \text{Inv}(l)$. By Definition 10.2.9, there is $q''_1 \in Q_A$ such that $q_1 + t^* \xrightarrow{\tau} q''_1$. Then $(q_1, q) + t^* \xrightarrow{\tau} (q''_1, q + t^*)$ which contradicts that $t' = \min(I_\tau(q_1, q))$.

Fix $a \in \Gamma = \Gamma_A \cup \Gamma_B$ and $t < t'$, and assume that there is $(q'_1, q') \in Q_A \times Q_B$ such that $(q_1, q) + t \xrightarrow{a} (q'_1, q')$. We have to show that there is $(q'_2, q'') \in Q_A \times Q_B$ such that $(q_2, q) + t \xrightarrow{a} (q'_2, q'')$ and $(q'_1, q') \mathcal{R}'(q'_2, q'')$. We have to consider several cases. First, if $a \in A$. Then $q_1 + t \xrightarrow{a} q'_1$ and $q + t \xrightarrow{a} q'$. As $q_1 \mathcal{R} q_2$ and $t < t' < t^*$, there is $q'_2 \in Q_A$ such that $q_2 + t \xrightarrow{a} q'_2$ and $q'_1 \mathcal{R} q'_2$. Then, $(q_2, q) + t \xrightarrow{a} (q'_2, q')$ and $(q'_1, q') \mathcal{R}'(q'_2, q')$ by definition of \mathcal{R}' . Secondly, if $a \notin A$. If $q + t \xrightarrow{a} q'$ and $q'_1 = q_1 + t$ (and thus $\nu_1 + t \models \text{Inv}(l_1)$), then since $\nu_2 + t \models \text{Inv}(l_2)$ (because $t < t' < t^*$ and from Lemma 10.2.11), we get that $(q_2, q) + t \xrightarrow{a} (q_2 + t, q')$. Finally as $t < t^*$, Lemma 10.2.11 implies that $(q_1 + t) \mathcal{R}(q_2 + t)$ and thus $(q_1 + t, q') \mathcal{R}'(q_2 + t, q')$. Otherwise, if $q_1 + t \xrightarrow{a} q'_1$ and $q' = q + t$ (and thus $\nu + t \models \text{Inv}(l)$), then since $q_1 \mathcal{R} q_2$ and $t < t^*$, it holds that there is $q'_2 \in Q_A$ such that $q_2 + t \xrightarrow{a} q'_2$ and $q'_1 \mathcal{R} q'_2$. We deduce that $(q_2, q) + t \xrightarrow{a} (q'_2, q + t)$ and $(q'_1, q + t) \mathcal{R}'(q'_2, q + t)$ which concludes the proof of point (5.a). We can get point (5.b) in a very similar way as the case where $a \notin A$.

Finally, point (4) can be shown in a similar way as in the proof of Lemma 9.3.9 in Section 9.3.2 by restricting the domain to $[0, t'[$ and by observing that $[0, t'[\subseteq [0, t^*]$. \square

Like in STA, Theorem 9.3.11, we can now extend this result to bisimulation between ISTA, leading to the congruence result.

Theorem 10.2.13. *Bisimulation is a congruence w.r.t. parallel composition. That is: if $\mathcal{A}_1, \mathcal{A}_2$ and \mathcal{B} are three ISTA in CISTA and if $A \subseteq (\Gamma_1 \cap \Gamma_B) \cap (\Gamma_2 \cap \Gamma_B)$ such that \mathcal{A}_1 and \mathcal{B} are composable w.r.t. A , and \mathcal{A}_2 and \mathcal{B} are composable w.r.t. B , then for every states q_1, q_2 and q of resp. $\mathcal{A}_1, \mathcal{A}_2$ and \mathcal{B} , if $\mathcal{A}_1 \sim \mathcal{A}_2$ from q_1 and q_2 , then $\mathcal{A}_1 \parallel_A \mathcal{B} \sim \mathcal{A}_2 \parallel_A \mathcal{B}$.*

The proof is similar to Theorem 9.3.11, the only differences come from Lemmas 10.2.11 and 10.2.12.

Hiding operator. Now that we have defined a handshaking parallel composition for ISTA, we can define a *hiding* operator like in IMCs (see Definition 2.4.5 or [HK09]). Basically given an ISTA \mathcal{A} , the hiding operator will transform some

of the actions in Γ of \mathcal{A} into the τ internal action. It is supposed to be done when some action a does not need any more synchronisations. The idea is that, once a system has established communication on all possible interactions, then all actions will be hidden. The resulting ISTA will then gives rise to a STA.

Definition 10.2.14. Let $\mathcal{A} = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X}, \Gamma^\tau, \dashrightarrow)$ be an ISTA and let $A \subseteq \Gamma$ be a set of actions. The *hiding of \mathcal{A} w.r.t. A* is the ISTA $\mathcal{A} \setminus A = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X}, \Gamma^\tau \setminus A, \dashrightarrow_*)$ where \dashrightarrow_* is defined as follows:

- if $a \in A$ and $l \xrightarrow{a, g, Y} l'$ then $l \xrightarrow{\tau, g, Y} l'$;
- if $a \notin A$ and $l \xrightarrow{a, g, Y} l'$ then $l \xrightarrow{a, g, Y} l'$.

Like in IMCs (Theorem 2.4.8 or [Her02] and [HK09]), it can be shown that bisimulation is a congruence w.r.t. hiding.

Theorem 10.2.15. *Bisimulation is a congruence w.r.t. hiding. That is: if \mathcal{A}_1 and \mathcal{A}_2 are two ISTA and if $A \subseteq \Gamma_1 \cap \Gamma_2 \neq \emptyset$, then for every states q_1 and q_2 of resp. \mathcal{A}_1 and \mathcal{A}_2 , if $\mathcal{A}_1 \sim \mathcal{A}_2$ from q_1 and q_2 , then $\mathcal{A}_1 \setminus A \sim \mathcal{A}_2 \setminus A$.*

The result comes from the fact that the same actions are hidden in both ISTA. Hence if two states of \mathcal{A}_1 and \mathcal{A}_2 are bisimilar, then those states will still be bisimilar after hiding some set of actions A : the key point is to prove that the first time before a τ -transition is enabled will decrease by the same amount of time or will remain the same in both states. Then all the points of Definition 10.2.9 will be satisfied using the fact that they were already satisfied for \mathcal{A}_1 and \mathcal{A}_2 .

As said before, an action a will be hidden when, after multiple compositions between systems, a will not need to synchronise again. We briefly illustrate it on Example 10.2.5.

Example 10.2.16. Consider the ISTA $(\mathcal{A}_{\text{cool}_1} \parallel_{\{\text{burn}\}} \mathcal{A}_{\text{cool}_2}) \parallel_{\Gamma'} \mathcal{A}_{\text{worker}}$ of Example 10.2.5 depicted in Figure 10.5. One could assume that the system does not need any more synchronisation and could thus hide all actions, *i.e.* consider the ISTA

$$\left((\mathcal{A}_{\text{cool}_1} \parallel_{\{\text{burn}\}} \mathcal{A}_{\text{cool}_2}) \parallel_{\Gamma'} \mathcal{A}_{\text{worker}} \right) \setminus \Gamma'.$$

On the other hand, one could also presume that when the plant is burning, some other system needs to intervene in order to minimize the damages. Thus, we would hide all actions except action *burn*.

10.2.2 Semantics through STA

We have now the main tools in hands, in order to define a semantics on ISTA. As already said previously, it will be given through a STA.

While some actions of an ISTA \mathcal{A} are not yet hidden, it is presumed that the system still waits to communicate with other systems. What will be of interest to us are *complete* ISTA, *i.e.* ISTA in which all actions are hidden.

Definition 10.2.17. An ISTA \mathcal{A} is said *complete* if all actions are hidden, *i.e.* if $\Gamma^\tau = \{\tau\}$.

Observe that for any ISTA \mathcal{A} , $\mathcal{A} \setminus \Gamma$ is closed. Like in IMCs, once we have composed several ISTA leading to a new ISTA that does not require to interact again with other systems, we can hide all actions leading to a complete ISTA.

We give a semantics on complete ISTA. Keep in mind that, as briefly explained in Section 10.1, we assume that once a τ -action is enabled, it is immediately performed: this is the maximal progress assumption, and this could correspond to synchronisation happening between some systems.

Fix a closed ISTA $\mathcal{A} = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X}, E^\tau)$. Note that since \mathcal{A} is closed, we can omit the set of actions Γ . And E^τ represents here the set of interactive transitions, as they are only labelled with τ . Intuitively, \mathcal{A} will behave as a STA with this difference: as soon as an interactive transition is enabled, the system will immediately take this interactive transition. Observe that several interactive transitions could be enabled at the same time. In order to deal with possible non-determinism, we equip \mathcal{A} with a weight function $w : E^\tau \rightarrow \mathbb{N}$ such that for each state $q = (l, \nu)$ if there exists $d = (l, \tau, g, Y, l') \in E^\tau$ such that $\nu \models g$, then there is $d' = (l, \tau, g', Y', l'') \in E^\tau$ such that $\nu \models g'$ and $w(d') > 0$.¹ We write this condition (I3). Recall that for any state $q = (l, \nu)$, we write $I_\tau(q) = \{t \geq 0 \mid \nu + t \models \text{Inv}(l) \text{ and } \exists (l, \tau, g, Y, l') \in E^\tau, \nu + t \models g\}$ and that whenever $I_\tau(q) \neq \emptyset$, $I_\tau(q)$ has a minimal element. Then each state of \mathcal{A} is equipped with a distribution over the delays and the edges as follows: for each $q = (l, \nu) \in Q$, writing $t^* = \min(I_\tau(q))$,

$$\mu'_q = \begin{cases} \mu_q & \text{if } I_\tau(q) = \emptyset \\ \min(\mu_q, \delta_{t^*}) & \text{if } l \in L_{\text{int}}^c \text{ and } I_\tau(q) \neq \emptyset \\ \delta_{t^*} & \text{otherwise,} \end{cases}$$

¹Observe that usually in IMCs, non-determinism is removed through a scheduler which explicitly chooses one τ -transition. We chose here this representation with a weight function because our model is defined with distributions over the edges.

and for each $t \geq 0$, if $t < t^*$ we define $p'_{q+t} = p_{q+t}$, otherwise for each $e \in E$, $p'_{q+t}(e) = 0$ and for each $d \in E^\tau$

$$p'_{q+t}(d) = \frac{\mathbf{w}(d)}{\sum_{d' \text{ enabled in } q} \mathbf{w}(d')}.$$

Observe that if l is an interactive location, then obviously $t^* < \infty$ and thus δ_{t^*} is well-defined. This induces a STA as defined in Definition 3.1.2. Notice however that hypothesis (†) (see page 65) that is needed to get all decidability and quantitative results of Chapter 7 is not necessarily satisfied.

Definition 10.2.18. Consider $\mathcal{A} = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X}, E^\tau)$ a closed ISTA and let $\mathbf{w} : E^\tau \rightarrow \mathbb{N}$ be a weight function satisfying condition (I3). The *induced STA of \mathcal{A} w.r.t. \mathbf{w}* is the tuple

$$\mathcal{A}_{\mathbf{w}} = (L, X, E', \text{Inv}, \text{AP}, \mathcal{L}, (\mu'_q, p'_q)_{q \in L \times \mathbb{R}_+^X})$$

where $E' = E \cup E^\tau$ and for each $q \in Q$, μ'_q and p'_q are defined as above.

We illustrate this on the ISTA of Example 10.2.16.

Example 10.2.19. One can observe that ISTA $\left((\mathcal{A}_{\text{cool}_1} \parallel_{\{\text{burn}\}} \mathcal{A}_{\text{cool}_2}) \parallel_{\Gamma'} \mathcal{A}_{\text{worker}} \right) \setminus \Gamma'$ of Example 10.2.16, is a complete ISTA. We write it \mathcal{A} . Considering the product before hiding of Example 10.2.5 on Figure 10.5, non-determinism can only occur when the worker is fixing one of the cooling systems and the corresponding system or both systems are in location Down. In that case, non-determinism comes from the τ -actions of $\mathcal{A}_{\text{cool}_1}$ and $\mathcal{A}_{\text{cool}_2}$, and from the synchronisation on fix_1 or fix_2 . In those cases, we put a weight of 1 on each interactive transitions, leading to STA $\mathcal{A}_{\mathbf{w}}$.

10.3. Decidability of ISTA

In Section 10.2.2, we have defined the semantics of complete ISTA as a STA. Recall that in Chapter 7, we have shown two classes of STA on which all decidability and approximability results of Chapter 6 can be applied. Thanks to the semantics as a STA, the question is now: can we find a class of complete ISTA on which we can apply the decidability results of Chapter 6?

The two classes of STA that were identified in Chapter 7 were the classes of reactive STA and single-clock STA. The latter case has no interest here: complete

ISTA come often from the composition of several systems leading, obviously, to several clocks. While reactive STA has no longer sense with complete ISTA: the definition of the probability distribution over the delays μ'_q is such that it is bounded from above by the first time at which an internal action is enabled, which is in contradiction with the definition of reactive STA where time can always progress everywhere.

Reactive STA, though, is the class that will influence our work here. The idea will be to consider ISTA in which all distributions only depend on the location and that are equivalent to the Lebesgue measure on \mathbb{R}_+ , just like in reactive STA. As stated above, this will not suffice. We will make also the assumption that it is not allowed to perform an infinite sequence of interactive transitions. This will yield to the existence of purely stochastic locations (*i.e.* without outgoing interactive transitions) that will be reached with probability 1 and that will allow to reach memoryless regions with probability 1, just like in reactive STA.

Given a STA \mathcal{A}' , we have seen that the thick region graph (see Definition 3.2.8) seen as a Markov chain, *i.e.* $\mathcal{T}_{\mathcal{A}'}^{\text{tg}}$, is an α -abstraction² of $\mathcal{T}_{\mathcal{A}'}$ (the STS induced by \mathcal{A}'), see Lemma 7.1.2. This is no longer the case here! This is due to the fact that given an ISTA \mathcal{A} and a weight function w , STA \mathcal{A}_w does not always satisfy hypothesis (‡) of page 65: indeed given a state q , the distribution μ'_q is over $I(q) \cap [0, t^*]$ with $t^* = \min(I_\tau(q))$ and thus we could get a subset of $I(q)$, I , such that $\Lambda(I) > 0$ but $\mu'_q(I) = 0$. Hence singular edges (see Definition 3.2.6), can no longer be removed in the same way as before.

We fix for the rest of the section an ISTA \mathcal{A} , a weight function w and the induced STA \mathcal{A}_w . Singular edges have now to be defined as follows. We refer to Sections 2.1.1 and 3.2 for words on the timed region automaton $R(\mathcal{A}_w)$.

Definition 10.3.1. An edge e of $R(\mathcal{A}_w)$ is said *singular* whenever, writing $(l, r) = \text{source}(e)$:

- for all $d \in E^\tau$, $\text{source}(d) \neq (l, r)$, and there exists $\nu \in r$ such that $I((l, \nu), e)$ is a single point but there is an edge e' of \mathcal{A}_w such that $I((l, \nu), e')$ is not a single point; or
- there is $d \in E^\tau$ such that $\text{source}(d) = (l, r)$, and $e \notin E^\tau$ or $e \in E^\tau$ and $w(e) = 0$.

We still importantly get that all singular edges are performed with probability 0. Indeed observe that the first point corresponds to Definition 3.2.6 while the second point comes from the definition of \mathcal{A}_w : whenever τ -transitions are

²Recall that $\alpha : L' \times \mathbb{R}_+^X \rightarrow L' \times R_{\mathcal{A}'}$ is such that $\alpha((l, \nu)) = (l, [\nu]_{\mathcal{A}'})$ where $R_{\mathcal{A}'}$ is the set of regions of \mathcal{A}' .

enabled, edges are taken with probability 0. Then, we can define the thick region graph just like in Definition 3.2.8, and we can construct a finite Markov chain from it, that we write $\mathcal{T}_{\mathcal{A}_w}^{\text{tg}}$ like in Section 7.1.1. Considering also the mapping $\alpha : L \times \mathbb{R}_+^X \rightarrow L \times R_{\mathcal{A}_w}$ such that $\alpha((l, \nu)) = (l, [\nu]_{\mathcal{A}_w})$, we still get that $\mathcal{T}_{\mathcal{A}_w}^{\text{tg}}$ is an α -abstraction of \mathcal{A}_w .

Lemma 10.3.2. It holds true that $\mathcal{T}_{\mathcal{A}_w}^{\text{tg}}$ is a finite α -abstraction of \mathcal{A}_w .

The proof is also immediate from the construction of $\mathcal{T}_{\mathcal{A}_w}^{\text{tg}}$ and from Definition 10.3.1.

We now would like to find to class of ISTA on which we could infer the whole class of decidability and approximability results of Chapter 6. As already discussed in Remark 7.1.4, since $\mathcal{T}_{\mathcal{A}_w}^{\text{tg}}$ is a finite α -abstraction of \mathcal{A}_w , we need to find a class of complete ISTA such that, considering STA \mathcal{A}_w :

- either we can get a finite-attractor (through an abstraction) which satisfies conditions (\dagger) of Proposition 5.2.6,
- or we can show that \mathcal{A}_w is fair w.r.t. α -closed sets, meeting thus the conditions of Proposition 5.2.7.

As we think that the first case is easier to prove, and inspired by how decidability was proven for reactive STA (see Section 7.2 or [BBB⁺14]), we define the class of *reactive* ISTA.

Definition 10.3.3. Let $\mathcal{A} = (L, X, E, \text{Inv}, \text{AP}, \mathcal{L}, (\mu_q, p_q)_{q \in L \times \mathbb{R}_+^X}, \Gamma^\tau, \dashrightarrow)$ be an ISTA. We say that \mathcal{A} is reactive if for each $l \in L_{\text{int}}^c$ and for each $\nu \in \mathbb{R}_+^X$, $I((l, \nu)) = \mathbb{R}_+$ and for each $l \in L_{\text{int}}^c$, there is a probability distribution μ_l over \mathbb{R}_+ , equivalent to the Lebesgue measure on \mathbb{R}_+ , such that for each $\nu \in \mathbb{R}_+^X$, $\mu_{(l, \nu)} = \mu_l$.

Basically, it is the same class as the class of reactive STA, except that obviously here interactive transitions and interactive locations are allowed. It follows that the induced STA from a reactive ISTA and a weight function is not necessarily a reactive STA due to the definitions of the probability distributions over the delays (see Definition 10.2.18): the support of such distributions will be in general upper bounded due to interactive transitions.

Remark 10.3.4. Note that if \mathcal{A} is a reactive ISTA, then it holds that for each $l \in L$, $\text{Inv}(l) = \text{true}$. In the sequel, we thus omit the component Inv when using a reactive ISTA.

Recall that the objective in Section 7.2 was to prove that reactive STA have a finite attractor (through the α -abstraction) that satisfies the hypotheses (\dagger) of Proposition 5.2.6. This attractor was the set of memoryless regions (see Propositions 7.2.1 and 7.2.2). It enjoyed the nice property that two states of same memoryless region, have the same probabilistic behaviours (from [BBB⁺14, Lemma F.4]). Observe that here with reactive ISTA, the set of memoryless regions may not be reached with probability 1 due to the definition of the distributions μ'_q of \mathcal{A}_w and to the presence of interactive transitions. In order to make sure that the set of memoryless regions forms an attractor of $\mathcal{T}_{\mathcal{A}_w}$, we need an extra condition: we cannot have an infinite sequence of interactive transitions within the model. This is formalised in the next definition.

Definition 10.3.5. Fix \mathcal{A} a closed ISTA, w a weight function and \mathcal{A}_w the induced STA. We say that \mathcal{A} is τ -cycle-free if there is $N \geq 0$ such that for each $q_0, \dots, q_{N+1} \in Q$ if there are $t_0, \dots, t_N \in \mathbb{R}_+$ and $e_0, \dots, e_N \in E'$ such that

$$q_0 \xrightarrow{t_0, e_0} q_1 \xrightarrow{t_1, e_1} \dots q_N \xrightarrow{t_N, e_N} q_{N+1}$$

in \mathcal{A}_w , then there is $i \geq 0$ such that $e_i \notin E^\tau$ (i.e. $e_i \in E$).

We write ISTA_τ^r for the class of τ -cycle-free reactive closed ISTA. For technicalities, we require moreover for conditions (I1) and (I2) of page 233 to be verified. For the rest of this section we assume that $\mathcal{A} \in \text{ISTA}_\tau^r$ and we write N for the constant corresponding to Definition 10.3.5. Moreover, we assume that for each $d \in E^\tau$, $w(d) > 0$. Like in Section 7.2, we write $\mathcal{R}_{\mathcal{A}_w}^{\text{mem}}$ for the set of memoryless regions of \mathcal{A}_w and we will again somehow prove that $\alpha^{-1}(L \times \mathcal{R}_{\mathcal{A}_w}^{\text{mem}})$ is an attractor for $\mathcal{T}_{\mathcal{A}_w}$ and that it satisfies hypothesis (\dagger) of Proposition 5.2.6.

Firstly, we can show the following result stating that \mathcal{A} possesses a location that has no outgoing interactive transitions.

Lemma 10.3.6. If $\mathcal{A} \in \text{ISTA}_\tau^r$, there is $l \in L$ that has no outgoing interactive transitions.

Proof. Towards a contradiction, assume that for each $l \in L$, there is an outgoing interactive transition. Then from the hypothesis (I2) and since \mathcal{A} is reactive, it holds that for each state $q \in L \times \mathbb{R}_+^X$ there is $t^* \geq 0$ such that for each $t \geq t^*$, there is an interactive transition enabled in $q + t$ (from the last point of Definition 10.1.1 as here, $\text{Inv}(l) = \text{true}$ for each l). It then follows that for each state q_0 we can construct an infinite run

$$q_0 \xrightarrow{t_0, e_0} q_1 \xrightarrow{t_1, e_1} q_2 \xrightarrow{t_2, e_2} \dots$$

where for each i , $t_i = \min(I_\tau(q_i))$ and $e_i \in E^\tau$ is an interactive transition enabled in $q_i + t_i$. This contradicts Definition 10.3.5. \square

We write L^r for the set of locations that do not have outgoing interactive transitions, which is thus non-empty from Lemma 10.3.6. We can now refine the attractor that we will use for $\mathcal{T}_{\mathcal{A}_w}$: $\alpha^{-1}(L^r \times \mathcal{R}_{\mathcal{A}_w}^{\text{mem}})$. Before showing that it is indeed an attractor and that it satisfies conditions (†) of Proposition 5.2.6, we need some technical lemmas.

The first one shows that in reactive closed ISTA, the probability to take an interactive transition from any state that can perform some interaction in the future, is lower bounded by a non-null value!

Lemma 10.3.7. Assume \mathcal{A} is reactive. There exist $\gamma > 0$ and $\gamma' > 0$ such that for each $q \in Q$ with $I_\tau(q) \neq \emptyset$, writing $t_q^* = \min(I_\tau(q))$, $\mu'_q(t_q^*) \geq \gamma$ and for each $t \in I_\tau(q)$ and each $d \in E^\tau$ enabled in $q + t$, $p'_{q+t}(d) \geq \gamma'$.

Proof. From the definition of p'_q and the weight function w that satisfies condition (I3) on page 242 (see Definition 10.2.18), it suffices to take

$$\gamma' = \min_{d \in E^\tau} \left\{ \frac{w(d)}{\sum_{d' \in E^\tau} w(d')} \right\} > 0$$

from the finiteness of E^τ .³ Now in order to get γ , we have to make the following observations. Fix $q = (l, \nu) \in Q$ such that $I_\tau(q) \neq \emptyset$. First if $l \in L_{\text{int}}$, then $\mu'_q(t_q^*) = 1 \geq \gamma$ no matter the choice of γ . Now if $l \notin L_{\text{int}}$, we write $q_0 = (l, \mathbf{0}_X)$ where $\mathbf{0}_X$ is the clock valuation that assigns 0 to each clock, and $t_l^* = \min(I_\tau(q_0))$. As $I_\tau(q) \neq \emptyset$, then $I_\tau(q_0) \neq \emptyset$ and thus t_l^* is well-defined. From the hypothesis (I2), since here $\text{Inv}(l) = \mathbf{true}$, it holds that each guard on interactive transitions are of the form $\bigwedge_{x \in Y} x \geq c_x$ for some $Y \subseteq X$, it follows that $t_l^* \geq t_q^*$.

We can now prove that $\mu'_q(t_q^*) \geq \mu'_{q_0}(t_l^*)$. Indeed from the definition of μ'_q and since \mathcal{A} is reactive, we get that

$$\begin{aligned} \mu'_q(t_q^*) \geq \mu'_{q_0}(t_l^*) &\iff \min(\mu_l, \delta_{t_q^*})(t_q^*) \geq \min(\mu_l, \delta_{t_l^*})(t_l^*) \\ &\iff \mu_l([t_q^*, \infty[) \geq \mu_l([t_l^*, \infty[) \end{aligned}$$

which is true since $t_q^* \leq t_l^*$. It then suffices to take $\gamma = \min_{l \in L_{\text{int}}^c} \mu_l([t_l^*, \infty[) > 0$ since for each $l \in L_{\text{int}}^c$, μ_l is a distribution on \mathbb{R}_+ equivalent to the Lebesgue measure on \mathbb{R}_+ . \square

Given $l \in L$, we write $E'_l = \{e \in E' \mid \text{target}(e) = l\}$. We refer to Section 4.1.1 for the event $\text{Ev}_{\mathcal{T}_{\mathcal{A}_w}}(\mathbf{F}_{\leq N+1} L^r \times \mathbb{R}_+^X)$ that we will abusively write $\text{Ev}_{\mathcal{T}_{\mathcal{A}_w}}(\mathbf{F}_{\leq N+1} L^r)$. Observe that for any state q ,

$$\text{Prob}_q^{\mathcal{T}_{\mathcal{A}_w}}(\mathbf{F}_{\leq N+1} L^r) = \text{Prob}_q^{\mathcal{A}_w} \left(\bigcup_{l \in L^r} \bigcup_{e_0, \dots, e_{N-1} \in E'} \bigcup_{e_N \in E'_l} \text{Cyl}(\pi(q, e_0, \dots, e_N)) \right).$$

³Observe that here, we use the assumption that $w(d) > 0$ for each $d \in E^\tau$.

The next lemma importantly shows that the probability to reach L^r with $N + 1$ steps from any state, is lower bounded by a non-null value.

Lemma 10.3.8. If $\mathcal{A} \in \text{ISTA}_\tau^r$, there is $\beta > 0$ such that for each $q \in Q$,

$$\text{Prob}_q^{\mathcal{T}^{\mathcal{A}_w}}(\mathbf{F}_{\leq N+1} L^r) \geq \beta.$$

Proof. Remark first that if $q = (l, \nu)$ with $l \in L^r$ then the inequality is obvious no matter the choice of $0 < \beta (\leq 1)$. Now assume that $q = (l, \nu)$ with $l \notin L^r$. We will then show the following statement: there exist $k \leq N$, $e_0, \dots, e_k \in E'$ with $\text{Prob}_q^{\mathcal{A}_w}(\text{Cyl}(\pi(q, e_0, \dots, e_k))) > 0$ and with for each $i < k$, $\text{target}(e_i) \notin L^r$, such that $\text{target}(e_k) \in L^r$ and $e_i \in E^\tau$ for each i . In other words, there exists a first passage to L^r that is made only through interactive transitions. Towards a contradiction, assume that it is not the case, that is each first passage to L^r has to be made through at least one edge of E . Then we can construct the following path:

$$q_0 = q \xrightarrow{t_0, e_0} q_1 \xrightarrow{t_1, e_1} q_2 \dots \xrightarrow{t_N, e_N} q_{N+1}$$

where for each i , $q_i = (l_i, \nu_i)$ with $l_i \notin L^r$, $t_i = \min(I_\tau(q_i))$ and $e_i \in E^\tau$. Indeed since $q_0 = (l, \nu)$ with $l \notin L^r$, we know that $\min(I_\tau(q_0)) < \infty$ from hypothesis (I2), and thus we can construct $q_0 \xrightarrow{t_0, e_0} q_1$ where e_0 is some enabled interactive transition in state $q_0 + t_0$. It should be observed that from the assumption, $q_1 = (l_1, \nu_1)$ with $l_1 \notin L^r$: each first passage to L^r has to be made through at least one edge. And you can repeat the argument. Assume that for some $0 \leq i < N$, for each $0 \leq j \leq i$ we have constructed

$$q_0 = q \xrightarrow{t_0, e_0} q_1 \xrightarrow{t_1, e_1} q_2 \dots \xrightarrow{t_j, e_j} q_{j+1}$$

satisfying the previous hypothesis. Let us show that we can extend the path for $j = i + 1$. We know that $q_{i+1} = (l_{i+1}, \nu_{i+1})$ with $l_{i+1} \notin L^r$. It follows that $\min(I_\tau(q_{i+1})) < \infty$ from hypothesis (I2) and thus there exists $e_{i+1} \in E^\tau$ such that $q_{i+1} \xrightarrow{t_{i+1}, e_{i+1}} q_{i+2}$ with $t_{i+1} = \min(I_\tau(q_{i+1}))$. We therefore have constructed

$$q_0 = q \xrightarrow{t_0, e_0} q_1 \xrightarrow{t_1, e_1} q_2 \dots \xrightarrow{t_{i+1}, e_{i+1}} q_{i+2} = (l_{i+2}, \nu_{i+2})$$

where $l_{i+2} \notin L^r$ since for each $0 \leq j \leq i + 1$, $l_j \notin L^r$ and $e_j \in E^\tau$. We thus have

$$q_0 = q \xrightarrow{t_0, e_0} q_1 \xrightarrow{t_1, e_1} q_2 \dots \xrightarrow{t_N, e_N} q_{N+1}$$

with $e_i \in E^\tau$ for each $0 \leq i \leq N$ which is a contradiction with the hypothesis and Definition 10.3.5. Thus there exist $k \leq N$, $e_0, \dots, e_k \in E'$ with

$\text{Prob}_q^{\mathcal{A}_w}(\text{Cyl}(\pi(q, e_0, \dots, e_k))) > 0$ and with for each $i < k$, $\text{target}(e_i) \notin L^r$, such that $\text{target}(e_k) \in L^r$ and $e_i \in E^\tau$ for each i . Observe that

$$\text{Prob}_q^{\mathcal{A}_w}(q \models \mathbf{F}_{\leq N+1} L^r) \geq \text{Prob}_q^{\mathcal{A}_w}(\text{Cyl}(\pi(q, e_0, \dots, e_k))).$$

We can compute:

$$\begin{aligned} \text{Prob}_q^{\mathcal{A}_w}(\text{Cyl}(\pi(q, e_0, \dots, e_k))) &= \mu'_q(t_q^*) \cdot p'_{q+t_q^*}(e_0) \cdot \mu'_{q_{e_0}}(t_{q_{e_0}}^*) p'_{q_{e_0}+t_{q_{e_0}}^*}(e_1) \cdots \\ &\quad \cdots \mu'_{q_{e_0 \dots e_{k-1}}}(t_{q_{e_0 \dots e_{k-1}}}^*) \cdot p'_{q_{e_0 \dots e_{k-1}}+t_{q_{e_0 \dots e_{k-1}}}^*}(e_k) \end{aligned}$$

where

- for each $q' \in Q$, $t_{q'}^* = \min(I_\tau(q'))$ which is always finite in the cases above, from hypothesis (I2) and from the hypothesis on the edges;
- for each $0 \leq i \leq k-1$, $q_{e_0 \dots e_{i-1}} \xrightarrow{e_i} q_{e_0 \dots e_i}$ with $q_{e_0 \dots e_{i-1}} = q$ if $i = 0$, $q_{e_0 \dots e_{i-1}} = q_{e_0}$ if $i = 1$.

From Lemma 10.3.7, we can hence deduce that

$$\text{Prob}_q^{\mathcal{A}_w}(\text{Cyl}(\pi(q, e_0, \dots, e_k))) \geq (\gamma\gamma')^{k+1},$$

since none of the states visited has a location in L^r (and thus all states visited can perform an interactive transition in the future). Now as k equals at most N , we take $\beta = (\gamma\gamma')^{N+1} > 0$ since $\gamma > 0$ and $\gamma' > 0$. It is now straightforward to get that

$$\text{Prob}_q^{\mathcal{T}_{\mathcal{A}_w}}(\mathbf{F}_{\leq N+1} L^r) \geq \beta$$

which concludes the proof. \square

We can now prove that $\alpha^{-1}(L^r \times \mathcal{R}_{\mathcal{A}_w}^{\text{mem}})$ is an attractor for $\mathcal{T}_{\mathcal{A}_w}$. The proof is very much inspired from the proof of [BBB⁺14, Lemma 14].

Proposition 10.3.9. *Fix $\mathcal{A}_w \in \text{ISTA}_\tau^r$. For each $q \in Q$, $\text{Prob}_q^{\mathcal{T}_{\mathcal{A}_w}}(\mathbf{F} \alpha^{-1}(L^r \times \mathcal{R}^{\text{mem}})) = 1$.*

Proof. Write M for the maximal constant appearing in a guard of an edge of E' . Given $q_0 \in Q$ and $n \in \mathbb{N}$, we write

- $\mathcal{D}_n^{>M}(q_0) = \{\rho = q_0 \xrightarrow{t_0, e_0} q_1 \xrightarrow{t_1, e_1} \dots \mid t_n > M\}$;
- $\mathcal{D}_n^r(q_0) = \{\rho = q_0 \xrightarrow{t_0, e_0} q_1 \xrightarrow{t_1, e_1} \dots \mid l_n \in L^r \wedge \forall i < n, l_i \notin L^r\}$;

$$\bullet \mathcal{D}_n^{>M,r}(q_0) = \mathcal{D}_n^r(q_0) \cap \mathcal{D}_n^{>M}.$$

Set $\mathcal{D}_n^{>M,r}(q_0)$ corresponds thus to the set of runs starting from q_0 that visit L^r for the first time at step n and that delay more than M time units in this state. Observe that from Lemma 10.3.8, it holds that for each $q \in Q$, $\text{Prob}_q^{A_w}(\bigcup_{n=0}^{N+1} \mathcal{D}_n^r(q)) \geq \beta$. We now establish that there is $\delta > 0$ such that for each $q \in Q$, for each $n \in \mathbb{N}$, $\text{Prob}_q^{A_w}(\mathcal{D}_n^{>M,r}(q)) \geq \delta \cdot \text{Prob}_q^{A_w}(\mathcal{D}_n^r(q))$. We prove this by induction on n . We write $E'_r = \{e \in E' \mid \text{target}(e) \in L^r\}$. Let $\delta = \min_{l \in L_{\text{int}}^c} \mu_l([M, \infty]) > 0$ as L_{int}^c is finite. Assume first that $n = 0$ and fix $q = (l, \nu) \in Q$. We want to show that $\text{Prob}_q^{A_w}(\mathcal{D}_0^{>M,r}(q)) \geq \delta \cdot \text{Prob}_q^{A_w}(\mathcal{D}_0^r(q))$. Remark that if $l \notin L^r$, then both sides of the inequality equal 0. Otherwise, $l \in L^r$ and thus

$$\text{Prob}_q^{A_w}(\mathcal{D}_0^{>M,r}(q)) = \mu_l([M, \infty]) \geq \delta = \delta \cdot \text{Prob}_q^{A_w}(\mathcal{D}_0^r(q))$$

since $\text{Prob}_q^{A_w}(\mathcal{D}_0^r(q)) = 1$. We also have to show the particular case where $n = 1$. Fix $q = (l, \nu) \in Q$. If $l \in L^r$ both sides of the inequality equal 0. Otherwise, we compute

$$\begin{aligned} \text{Prob}_q^{A_w}(\mathcal{D}_1^{>M,r}(q)) &= \sum_{e_0 \in E'_r} \int_{t_0 \in I(q, e_0)} p'_{q+t_0}(e_0) \cdot \text{Prob}_{q_{e_0}}^{A_w}(\mathcal{D}_0^{>M,r}(q_{e_0})) \, d\mu'_q(t_0) \\ &\quad \text{where } q \xrightarrow{t_0, e_0} q_{e_0} \\ &\geq \delta \cdot \sum_{e_0 \in E'_r} \int_{t_0 \in I(q, e_0)} p'_{q+t_0}(e_0) \, d\mu'_q(t_0) \quad \text{from previous case} \\ &= \delta \cdot \text{Prob}_q^{A_w}(\mathcal{D}_1^r(q)). \end{aligned}$$

Now assume that for each $0 \leq n \leq k$ with $k \geq 1$, the inequality holds true and let us prove that it is still the case for $n = k + 1$. Fix $q = (l, \nu) \in Q$. Since $n = k + 1 > 0$, it should be observed that if $l \in L^r$ then both sides of the inequality equal 0. Otherwise, $l \notin L^r$ and we thus get that

$$\begin{aligned} \text{Prob}_q^{A_w}(\mathcal{D}_{k+1}^{>M,r}(q)) &= \sum_{e_0 \in (E'_r)^c} \int_{t_0 \in I(q, e_0)} p'_{q+t_0}(e_0) \cdot \text{Prob}_{q_{e_0}}^{A_w}(\mathcal{D}_k^{>M,r}(q_{e_0})) \, d\mu'_q(t_0) \\ &\quad \text{where } q \xrightarrow{t_0, e_0} q_{e_0} \\ &\geq \delta \cdot \sum_{e_0 \in (E'_r)^c} \int_{t_0 \in I(q, e_0)} p'_{q+t_0}(e_0) \text{Prob}_{q_{e_0}}^{A_w}(\mathcal{D}_k^r(q_{e_0})) \, d\mu'_q(t_0) \\ &\quad \text{from induction hypothesis} \\ &= \delta \cdot \text{Prob}_q^{A_w}(\mathcal{D}_{k+1}^r(q)). \end{aligned}$$

Observe now that for each state q , $\bigcup_{n=0}^{N+1} \mathcal{D}_n^{>M,r}(q)$ and $\bigcup_{n=0}^{N+1} \mathcal{D}_n^r(q)$ are disjoint unions. It thus follows that

$$\begin{aligned} \text{Prob}_q^{\mathcal{A}_w} \left(\bigcup_{n=0}^{N+1} \mathcal{D}_n^{>M,r}(q) \right) &\geq \delta \cdot \text{Prob}_q^{\mathcal{A}_w} \left(\bigcup_{n=0}^{N+1} \mathcal{D}_n^r(q) \right) \\ &\geq \delta \cdot \beta > 0 \quad \text{from Lemma 10.3.8.} \end{aligned}$$

It is now trivial to see that for each $q \in Q$,

$$\text{Prob}_q^{\mathcal{T}_{\mathcal{A}_w}} (\mathbf{F} \alpha^{-1}(L^r \times \mathcal{R}_{\mathcal{A}_w}^{\text{mem}})) \geq \text{Prob}_q^{\mathcal{A}_w} \left(\bigcup_{n=0}^{N+1} \mathcal{D}_n^{>M,r}(q) \right),$$

we hence deduce that for each state q ,

$$\text{Prob}_q^{\mathcal{T}_{\mathcal{A}_w}} (\mathbf{F} \alpha^{-1}(L^r \times \mathcal{R}_{\mathcal{A}_w}^{\text{mem}})) \geq \delta \cdot \beta > 0$$

and using classical arguments like in the proof of Proposition 5.2.6, we get that for each state q ,

$$\text{Prob}_q^{\mathcal{T}_{\mathcal{A}_w}} (\mathbf{F} \alpha^{-1}(L^r \times \mathcal{R}_{\mathcal{A}_w}^{\text{mem}})) = 1.$$

□

It remains to prove that the hypotheses (†) of Proposition 5.2.6 are met. It is very similar to the proof of Proposition 7.2.2 and uses a very similar lemma as [BBB⁺14, Lemma F.4] that we formalise below.

Lemma 10.3.10. If $\mathcal{A} \in \text{ISTA}_\tau^r$, then for each pair of states $q = (l, \nu)$ and $q' = (l, \nu')$ such that for every $x \in X$, $\nu(x) = \nu'(x)$ or $\min(\nu(x), \nu'(x)) > M$, for each $e_1, \dots, e_n \in E$,

$$\text{Prob}_q^{\mathcal{A}_w} (\pi(q, e_1, \dots, e_n)) = \text{Prob}_{q'}^{\mathcal{A}_w} (\pi(q', e_1, \dots, e_n)).$$

Proof. The proof is very similar to the proof of Lemma F.4 in [BBB⁺14]. The difference resides in the fact that in Lemma F.4, we can state that $\mu_q = \mu_l = \mu_{q'}$. Here it is not always the case that $\mu'_q = \mu_l$ and $\mu'_{q'} = \mu_l$. However, we can show that $\mu'_q = \mu'_{q'}$ which will conclude the proof. We first show that $I_\tau(q) = I_\tau(q')$, or more precisely that for each $(l, g, Y, l') \in E^\tau$ and for each $t \geq 0$, we have $\nu + t \models g$ iff $\nu' + t \models g$. Since for each $l \in L$, $\text{Inv}(l) = \mathbf{true}$ and from hypothesis (I2), we can write:

$$g = \bigwedge_{x \in X_1} x \geq a_x$$

with X_1 a subset of X and $a_x \in \{0, \dots, M\}$ for each clock x . Notice that it is sufficient to show only one of the implications. W.l.o.g. assume that $\nu + t \models g$. Fix $x \in X$, we consider several cases:

- if $\nu(x) = \nu'(x)$, it is obvious that $\nu'(x) + t$ will satisfy the same constraints as $\nu(x) + t$;
- if $\nu(x) > M$ and $\nu'(x) > M$, then $\nu(x) + t > M$ and $\nu'(x) + t > M$. Since M is the maximal constant appearing in guards and $\nu + t \models g$, it is obvious that if $x \in X_1$, we get $\nu'(x) + t > M \geq a_x$.

This shows that $I_\tau(q) = I_\tau(q')$. Finally, write $t^* = \min(I_\tau(q)) = \min(I_\tau(q'))$, we then get that $\mu'_q = \mu'_{q'} = \min(\mu_l, \delta_{t^*})$. This terminates the proof. \square

We thus get the following result.

Proposition 10.3.11. *If $\mathcal{A} \in \text{ISTA}_\tau^+$, it holds that $\mathcal{T}_{\mathcal{A}_w}$ is decisive w.r.t. α -closed sets and that $\mathcal{T}_{\mathcal{A}_w}^{\text{tg}}$ is a sound α -abstraction of $\mathcal{T}_{\mathcal{A}}$.*

Proof. The proof is in all points similar to the proof of Proposition 7.2.2. We get that all the hypotheses (\dagger) of Proposition 5.2.6 are satisfied, in particular thanks to Proposition 10.3.9 and Lemma 10.3.10.

Thus Proposition 5.2.6 implies that $\mathcal{T}_{\mathcal{A}_w}$ is decisive w.r.t. α -closed sets and therefore, Proposition 5.3.4 implies that $\mathcal{T}_{\mathcal{A}_w}^{\text{tg}}$ is a sound α -abstraction of $\mathcal{T}_{\mathcal{A}}$. \square

Finally, we get a similar result as Corollary 7.2.3 for the qualitative and quantitative model-checking problems for τ -cycle-free reactive closed ISTA of ω -regular properties (see Definitions 4.1.19 and 4.1.20). We consider again a product between STA \mathcal{A}_w with a DMA \mathbb{M} . We refer to Lemma 6.1.19 for notation $\alpha_{\mathbb{M}}$ and the fact that $\mathcal{T}_{\mathcal{A}_w}^{\text{tg}} \times \mathbb{M}$ is an $\alpha_{\mathbb{M}}$ -abstraction of $\mathcal{T}_{\mathcal{A}_w} \times \mathbb{M}$. As consequences of Chapter 6, we get the following decidability and approximability results for reactive STA.

Corollary 10.3.12. *Let \mathcal{A} be a τ -cycle-free reactive closed ISTA, w be a weight function with $w(d) > 0$ for each $d \in E^\tau$, \mathcal{A}_w be the induced STA and \mathbb{M} be a DMA. Then:*

1. *we can decide whether \mathcal{A}_w satisfies almost-surely \mathbb{M} ;*
2. *for every initial distribution μ which is numerically amenable w.r.t. \mathcal{A}^4 , we can compute arbitrarily close approximations of $\text{Prob}_\mu^{\mathcal{T}_{\mathcal{A}_w}}(\mathbb{M})$.*

Proof. The proof is in all points similar to the proof of Corollary 7.2.3. This is an application of Theorem 6.1.25, Corollary 6.1.26 and of Sections 6.2.1 and 6.2.4. It should be noted that all the hypotheses are met:

⁴Recall that we say that a distribution μ is numerically amenable w.r.t. \mathcal{A} if, given $k \in \mathbb{N}$, given $\varepsilon > 0$ and given a sequence of locations and regions $(l_0, r_0), (l_1, r_1), \dots, (l_k, r_k)$, one can approximate $\text{Prob}_\mu^{\mathcal{A}}(\text{Cyl}((l_0, r_0), (l_1, r_1), \dots, (l_k, r_k)))$ up to ε .

- $\mathcal{T}_{\mathcal{A}_w}^{\text{tg}} \times M$ has a finite attractor: since $\mathcal{T}_{\mathcal{A}_w}^{\text{tg}}$ is a finite MC then so is $\mathcal{T}_{\mathcal{A}_w}^{\text{tg}} \times M$ and we get a trivial finite attractor;
- $\mathcal{T}_{\mathcal{A}_w} \times M$ is decisive w.r.t. any α_M -closed sets.

This second point is a little more tricky and comes from the fact that $\mathcal{T}_{\mathcal{A}_w} \times M$ can be obtained through a product between ISTA \mathcal{A} and DMA M and that this product is also in ISTA'_{τ} . \square

Conclusion and Future Work

We now conclude this thesis with a summary of our results of Part II for compositional verification in STA and with a quick word on some possibilities that are left for future work.

Following [HZ11], in Part II we have been interested in the definition of an operator of composition in STA. We have divided the work into two chapters: we have defined an interleaving operator in Chapter 9, we have introduced the model of ISTA and extended the interleaving operator into a handshaking one for ISTA in Chapter 10.

The interleaving semantics corresponds to the case where two systems run completely independently, while the handshaking semantics corresponds to the case where the two systems can communicate and have to synchronise. While a handshaking operator is obviously more interesting, we advocate that the interleaving operator of Chapter 9 was an important first step. Indeed, as stated in [HZ11], a handshaking operator encompasses in particular the interleaving operator. When the composed systems do not have to interact, then the semantics should be interleaving. To have an interleaving operator is thus a good start.

In Chapter 9, we thus have defined an operator of parallel composition that, we have proven, corresponds to the interleaving semantics. The main contributions are the following.

- The definition of a parallel composition operator for STA and the identification of a subclass of STA in which, we proved, parallel composition is well-defined and internal (Sections 9.1 and 9.2).
- The identification of another class of STA (a subclass of the previous one)

in which we proved that, moreover, the parallel composition corresponds to the interleaving semantics (Section 9.2).

- The definition of a notion of bisimulation in STA, which extends the one on CTMCs ([DP03]) and that is importantly a congruence w.r.t. parallel composition. This is an expected property for a proper compositional framework (Section 9.3).

Inspired by the IMC model [Her02] and [HK09], we have then introduced the new model of ISTA in Chapter 10. This extends the STA model with interactive transitions allowing for synchronisations. We thus have defined a handshaking operator of parallel composition. The main contributions are the following.

- The definition of the new model of ISTA which is inspired from the IMC model [HK09] (Section 10.1).
- The definition of a handshaking operator of parallel composition for ISTA and the identification of a subclass of ISTA in which parallel composition is well-defined and internal (Section 10.2.1).
- The definition of a notion of bisimulation in ISTA, which extends the one on IMCs ([HK09]) and that is importantly a congruence w.r.t. parallel composition (Section 10.2.1).
- The definition of a hiding operator and of a closed ISTA which can then induce a STA as we define it (Sections 10.2.1 and 10.2.2).
- The identification of a subclass of closed ISTA on which the qualitative and quantitative results of Chapter 6 can be applied; yielding to a framework with a compositional framework and strong decidability and approximability results. (Section 10.3).

Perspectives for future work. We can list several perspectives for future work, the list is not exhaustive.

- Following the approach of [DP03] on CTMCs, we would like to prove a logical characterisation of bisimulation in STA and ISTA using CSL properties (or a subset). Such characterisations are standard when dealing with bisimulation.
- Another standard technique is to quotient a system thanks to a bisimulation. For instance in timed automata [AD90] and [AD94], the quotient of

a timed automaton with the timed-abstract bismulation gives the region automaton. One would thus like to study the quotient of STA or ISTA with the bisimulations defined here, in order to build a smaller system that should be easier to analyse.

- Finally and obviously, with this compositional design, we are interested in the compositional verification. The qualitative and quantitative results of Chapter 7 in STA, although very nice, concern only a unique STA describing a (possibly) big system. Compositional verification should thus simplify the proposed solutions by designing big systems as the result of several smaller systems and by reducing the model-checking problems of the big system to the model-checking problems of the smaller systems. This is the main reason we have defined this compositional framework and this is thus an interesting perspective for the verification of (I)STA.

Bibliography

- [AAGT12] Manindra Agrawal, S. Akshay, Blaise Genest, and P. S. Thiagarajan. Approximate verification of the symbolic dynamics of markov chains. In *Proc. 27th Annual Symposium on Logic in Computer Science (LICS'12)*. IEEE Computer Society, 2012.
- [AB06] Rajeev Alur and Mikhail Bernadsky. Bounded model checking for GSMP models of stochastic real-time systems. In *Proc. 9th International Workshop on Hybrid Systems: Computation and Control (HSCC'06)*, volume 3927 of *Lecture Notes in Computer Science*, pages 19–33. Springer, 2006.
- [Aba07] Alessandro Abate. *Probabilistic reachability for stochastic hybrid systems: theory, computations, and applications*. PhD thesis, University of California, Berkeley, USA, 2007.
- [ABM07] Parosh Aziz Abdulla, Noomene Ben Henda, and Richard Mayr. Decisive Markov chains. *Logical Methods in Computer Science*, 3(4), 2007.
- [ABRS05] Parosh Aziz Abdulla, Nathalie Bertrand, Alexander Rabinovich, and Philippe Schnoebelen. Verification of probabilistic systems with faulty communication. *Information and Computation*, 202(2):141–165, 2005.
- [ACD93] Rajeev Alur, Costas Courcoubetis, and David L. Dill. Model-checking in dense real-time. *Information and Computation*, 104(1):2–34, 1993.
- [AD90] Rajeev Alur and David L. Dill. Automata for modeling real-time systems. In *Proc. 17th International Colloquium on Automata*,

- Languages and Programming (ICALP'90)*, volume 443 of *Lecture Notes in Computer Science*, pages 322–335. Springer, 1990.
- [AD94] Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
- [AKLP10] Alessandro Abate, Joost-Pieter Katoen, John Lygeros, and Maria Prandini. Approximate model checking of stochastic hybrid systems. *European Journal of Control*, 16(6):624–641, 2010.
- [BA07] Mikhail Bernadsky and Rajeev Alur. Symbolic analysis for GSMP models with one stateful clock. In *Proc. 10th International Workshop on Hybrid Systems: Computation and Control (HSCC'07)*, volume 4416 of *Lecture Notes in Computer Science*, pages 90–103. Springer, 2007.
- [BBB⁺07] Christel Baier, Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, and Marcus Größer. Probabilistic and topological semantics for timed automata. In *Proc. 27th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'07)*, volume 4855 of *Lecture Notes in Computer Science*, pages 179–191. Springer, 2007.
- [BBB⁺08] Christel Baier, Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, and Marcus Größer. Almost-sure model checking of infinite paths in one-clock timed automata. In *Proc. 23rd Annual Symposium on Logic in Computer Science (LICS'08)*, pages 217–226. IEEE Computer Society Press, 2008.
- [BBB⁺14] Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, Quentin Menet, Christel Baier, Marcus Größer, and Marcin Jurdziński. Stochastic timed automata. *Logical Methods in Computer Science*, 10(4):1–73, 2014.
- [BBBC16] Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, and Pierre Carlier. Analysing decisive stochastic processes. In *Proc. 43rd International Colloquium on Automata, Languages and Programming (ICALP'16)*, LIPIcs. Leibniz-Zentrum für Informatik, 2016.
- [BBBC17] Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, and Pierre Carlier. When are stochastic transition systems tameable? Submitted for publication, 2017.

-
- [BBBM08] Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, and Nicolas Markey. Quantitative model-checking of one-clock timed automata under probabilistic semantics. In *Proc. 5th International Conference on Quantitative Evaluation of Systems (QEST'08)*. IEEE Computer Society Press, 2008.
- [BBCM16] Patricia Bouyer, Thomas Brihaye, Pierre Carlier, and Quentin Menet. Compositional design of stochastic timed automata. In *Proc. 11th International Computer Science Symposium in Russia (CSR'16)*, volume 9691 of *Lecture Notes in Computer Science*. Springer, 2016.
- [BBJM12] Patricia Bouyer, Thomas Brihaye, Marcin Jurdzinski, and Quentin Menet. Almost-sure model-checking of reactive timed automata. In *Proc. 9th International Conference on Quantitative Evaluation of Systems (QEST'12)*, pages 138–147. IEEE Computer Society Press, 2012.
- [BDHK06] H. Bohnenkamp, P D'Argenio, H. Hermanns, and J.-P. Katoen. MODEST: A compositional modeling formalism for hard and softly timed systems. *IEEE Trans. Software Engineering*, 32(10):812–830, 2006.
- [BDL⁺06] Gerd Behrmann, Alexandre David, Kim G. Larsen, John Håkansson, Paul Pettersson, Wang Yi, and Martijn Hendriks. Up-paal 4.0. In *Proc. 3rd International Conference on Quantitative Evaluation of Systems (QEST'06)*, pages 125–126. IEEE Computer Society Press, 2006.
- [Ber06] Nathalie Bertrand. *Modèles stochastiques pour les pertes de messages dans les protocoles asynchrones et techniques de vérification automatique*. PhD thesis, École Normale Supérieure de Cachan, Cachan, France, 2006.
- [BHHK03] Christel Baier, Boudewijn Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Model-checking algorithms for continuous-time Markov chains. *IEEE Transactions on Software Engineering*, 29(7):524–541, 2003.
- [BK08] Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT Press, 2008.

- [BKKŘ11] Tomáš Brázdil, Jan Krčál, Jan Křetínský, and Vojtěch Řehák. Fixed-delay events in generalized semi-Markov processes revisited. In *Proc. 22nd International Conference on Concurrency Theory (CONCUR'11)*, volume 6901 of *Lecture Notes in Computer Science*, pages 140–155. Springer, 2011.
- [Bri06] Thomas Brihaye. *Verification and Control of O-Minimal Hybrid Systems and Weighted Timed Automata*. PhD thesis, Université de Mons-Hainaut, Belgium, 2006.
- [BvdSHV03] Henrik Bohnenkamp, Peter van der Stok, Holger Hermanns, and Frits Vaandrager. Cost-optimisation of the ipv4 zeroconf protocol. In *International conference on dependable systems and networks*. IEEE Computer Society Press, 2003.
- [CE81] Edmund M. Clarke and E. Allen Emerson. Design and synthesis of synchronous skeletons using branching-time temporal logic. In *Proc. 3rd Workshop on Logics of Programs (LOP'81)*, volume 131 of *Lecture Notes in Computer Science*, pages 52–71. Springer-Verlag, 1981.
- [CSKN05] Stefano Cattani, Roberto Segala, Marta Z Kwiatkowska, and Gethin Norman. Stochastic transition systems for continuous state spaces and non-determinism. In *Proc. 8th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'05)*, volume 3441 of *Lecture Notes in Computer Science*, pages 125–139. Springer, 2005.
- [CY88] Costas Courcoubetis and Mihalis Yannakakis. Verifying temporal properties of finite-state probabilistic programs. In *Proc. 29th Annual Symposium on Foundations of Computer Science (FOCS'88)*, pages 338–345. IEEE Computer Society Press, 1988.
- [DK05a] Pedro R. D'Argenio and Joost-Pieter Katoen. A theory of stochastic systems Part I: Stochastic automata. *Information and Computation*, 203(1):1–38, 2005.
- [DK05b] Pedro R. D'Argenio and Joost-Pieter Katoen. A theory of stochastic systems Part II: Process algebra. *Information and Computation*, 203(1):39–74, 2005.

-
- [DP03] Josée Desharnais and Prakash Panangaden. Continuous stochastic logic characterizes bisimulation of continuous-time Markov processes. *Journal of Logic and Algebraic Programming*, 56:99–115, 2003.
- [Fel66] William Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. John Wiley & Sons, 1966.
- [Fel69] William Feller. *An Introduction to Probability Theory and Its Applications*, volume 2. John Wiley & Sons, 1969.
- [FHH⁺11] Martin Fränzle, Ernst Moritz Hahn, Holger Hermanns, Nicolás Wolovick, and Lijun Zhang. Measurability and safety verification for stochastic hybrid systems. In *Proc. 14th ACM International Conference on Hybrid Systems: Computation and Control (HSCC'11)*, pages 43–52. ACM, 2011.
- [GBK16] Daniel Gburek, Christel Baier, and Sascha Klüppelholz. Composition of stochastic transition systems. In *Proc. 43rd International Colloquium on Automata, Languages and Programming (ICALP'16)*, LIPIcs. Leibniz-Zentrum für Informatik, 2016.
- [Gly89] Peter W. Glynn. A GSMP formalism for discrete event systems. *Proc. of the IEEE*, 77(1):14–23, 1989.
- [GTW02] Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002.
- [HBR84] Charles A.R. Hoare, Stephen D. Brookes, and Andrew W. Roscoe. A theory of communicating sequential processes. *Journal of the ACM*, 31(3):560–599, 1984.
- [Hen96] Thomas A. Henzinger. The theory of hybrid automata. In *Proc. 11th Annual Symposium on Logic in Computer Science (LICS'96)*, pages 278–292. IEEE Computer Society Press, 1996.
- [Her02] Holger Hermanns. *Interactive Markov Chains: The Quest for Quantified Quality*, volume 2428 of *Lecture Notes in Computer Science*. Springer, 2002.
- [HH14] A. Hartmanns and H. Hermanns. The Modest toolset: An integrated environment for quantitative modelling and verification. In

- Proc. 20th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, volume 8413 of *Lecture Notes in Computer Science*, pages 593–598. Springer, 2014.
- [HK09] Holger Hermanns and Joost-Pieter Katoen. The how and why of interactive Markov chains. volume 6286 of *Lecture Notes in Computer Science*, pages 311–337. Springer, 2009.
- [HLS00] Jianghai Hu, John Lygeros, and Shankar Sastry. Towards a theory of stochastic hybrid systems. In *Proc. 3rd International Conference on Hybrid Systems: Computation and Control (HSCC'00)*, pages 160–173. Springer, 2000.
- [HPRV12] András Horváth, Marco Paolieri, Lorenzo Ridi, and Enrico Vicario. Transient analysis of non-markovian models using stochastic state classes. *Performance Evaluation*, 69(7-8):315–335, 2012.
- [HZ11] Holger Hermanns and Lijun Zhang. From concurrency models to numbers – Performance and dependability. In *Software and Systems Safety – Specification and Verification*, volume 30 of *NATO Science for Peace and Security Series - D: Information and Communication Security*, pages 182–210. IOS Press, 2011.
- [KNP11] Marta Kwiatkowska, Gethin Norman, and David Parker. PRISM 4.0: verification of probabilistic real-time systems. In *Proc. 23rd International Conference on Computer Aided Verification (CAV'11)*, volume 6806 of *Lecture Notes in Computer Science*, pages 585–591. Springer, 2011.
- [KNSS00] Marta Z. Kwiatkowska, Gethin Norman, Roberto Segala, and Jeremy Sproston. Verifying quantitative properties of continuous probabilistic timed automata. In *Proc. 11th International Conference on Concurrency Theory (CONCUR'00)*, volume 1877 of *Lecture Notes in Computer Science*, pages 123–137. Springer, 2000.
- [KNSS02] Marta Z. Kwiatkowska, Gethin Norman, Roberto Segala, and Jeremy Sproston. Automatic verification of real-time systems with discrete probability distributions. *Theoretical Computer Science*, 282(1):101–150, 2002.
- [KSK76] John G. Kemeny, J. Laurie Snell, and Anthony W. Knapp. *Denumerable Markov Chains*, volume 40 of *Graduate Texts in Mathematics*. Springer, 1976.

-
- [KT05] Moez Krichen and Stavros Tripakis. State identification problems for timed automata. In *Proc. 17th IFIP International Conference on Testing of Communicating Systems (TESTCOM'05)*, volume 3502 of *Lecture Notes in Computer Science*, pages 175–191. Springer, 2005.
- [LMS04] François Laroussinie, Nicolas Markey, and Philippe Schnoebelen. Model checking timed automata with one or two clocks. In *Proc. 15th International Conference on Concurrency Theory (CONCUR'04)*, volume 3170 of *Lecture Notes in Computer Science*, pages 387–401. Springer, 2004.
- [LS91] Kim G. Larsen and Arne Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94:1–28, 1991.
- [LY93] Kim G. Larsen and Wang Yi. Time abstracted bisimulation: Implicit specifications and decidability. In *Proc. 9th Conf. Mathematical Foundations of Programming Semantics (MFPS IX)*, volume 802 of *Lecture Notes in Computer Science*, pages 160–176. Springer, 1993.
- [Mod] Webpage of Modest. <http://www.modestchecker.net/>.
- [Pan01] Prakash Panangaden. Measure and probability for concurrency theorists. *Theor. Comput. Sci.*, 253(2):287–309, 2001.
- [Pan09] Prakash Panangaden. *Labelled Markov Processes*. Imperial College Press, 2009.
- [Pet81] James Lyle Peterson. *Petri Net Theory and the Modeling of Systems*. Prentice Hall, 1981.
- [PHV16] Marco Paolieri, András Horváth, and Enrico Vicario. Probabilistic model checking of regenerative concurrent systems. *IEEE Transactions on Software Engineering*, 42(2):153–169, 2016.
- [Pnu77] Amir Pnueli. The temporal logic of programs. In *Proc. 18th Annual Symposium on Foundations of Computer Science (FOCS'77)*, pages 46–57. IEEE Computer Society Press, 1977.
- [Pnu83] Amir Pnueli. On the extremely fair treatment of probabilistic algorithms. In *Proc. 15th Ann. Symp. Theory of Computing (STOC'83)*, pages 278–290. ACM Press, 1983.

- [Pri] Webpage of PRISM. <http://www.prismmodelchecker.org/>.
- [QS82] Jean-Pierre Queille and Joseph Sifakis. Specification and verification of concurrent systems in CESAR. In *Proc. 5th International Symposium on Programming*, volume 137 of *Lecture Notes in Computer Science*, pages 337–351. Springer, 1982.
- [RS59] Michael O. Rabin and D. Scott. Finite automata and their decision problems. *IBM Journal of Research and Developments*, 3:115–125, 1959.
- [SA13] Sadegh Esmail Zadeh Soudjani and Alessandro Abate. Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes. *Journal on Applied Dynamical Systems*, 12(2):921–956, 2013.
- [Upp] Webpage of UPPAAL. <http://www.uppaal.org/>.
- [Var85] Moshe Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proc. 26th Annual Symposium on Foundations of Computer Science (FOCS'85)*, pages 327–338. IEEE Computer Society Press, 1985.
- [VW94] Moshe Y. Vardi and Pierre Wolper. Reasoning about infinite computations. *Information and Computation*, 115(1):1–37, 1994.

Résumé en français

Sujet de la thèse. Dans cette thèse, nous nous intéressons à la vérification d'automates temporisés et stochastiques (écrits STA pour la suite). La vérification de systèmes informatiques est devenue courante de nos jours : certains de ces systèmes pouvant être critiques (avions, centrales nucléaires, ...), il est important de pouvoir vérifier si ces systèmes fonctionnent correctement. Cela se fait par le biais d'une modélisation mathématique du système. Ces modèles mathématiques nous fournissent des outils qui permettent de décider si un système est correct. En particulier, étant donné un modèle (représentant un système informatique) et une propriété (exprimée sur le modèle mathématique), on veut pouvoir vérifier si le modèle satisfait la propriété. On appelle cela le *model-checking*.

Afin de considérer ce problème, de nombreux modèles ont été considérés. Nous nous intéressons ici à des systèmes informatiques avec des contraintes de temps (on veut que certaines actions d'un programme soient uniquement exécutées dans un certain délai). L'un des modèles les plus répandus pour étudier le model-checking de systèmes informatiques temporisés sont les automates temporisés [AD94]. Nous nous intéressons à une extension probabiliste de ce modèle : les STA.

Automates temporisés. Un automate temporisé est un automate fini (donné par un ensemble fini de locations et un ensemble fini de transitions entre les états) enrichi d'un ensemble de variables à valeurs réelles et positives, appelées horloges. Les valeurs des horloges changent avec l'écoulement du temps et les actions faites. Le model-checking sur un automate temporisé s'exprime de cette façon : étant donné un automate temporisé \mathcal{A} et une propriété φ , on veut savoir si \mathcal{A} satisfait φ .

Le succès du modèle vient du fait que de nombreuses propriétés sont vérifiables pour les automates temporisés, et que de nombreux outils ont été développés pour

vérifier ces systèmes (par exemple Uppaal [BDL⁺06])

Automates temporisés et stochastiques. On considère une extension probabiliste des automates temporisés : une des motivations est que de nombreuses applications ont besoin de gérer à la fois des contraintes temporelles et des aspects probabilistes (des protocoles de communication par exemple [?]). Nous nous intéressons aux STA, présentés dans [BBB⁺14]. C'est une extension naturelle des automates temporisés où les délais et les choix discrets sont probabilisés.

De nombreux résultats sont déjà connus sur les STA. Notamment, on sait que le model-checking presque sûr (c'est-à-dire vérifier si oui ou non, un STA satisfait avec probabilité 1 (ou 0) une propriété donnée) pour les formules LTL est décidable pour une certaine classe de STA (toujours dans [BBB⁺14]). Cependant, plusieurs aspects manquent encore pour une étude plus complète de la vérification de STA, entre autres :

- une approche compositionnelle du modèle, afin d'étudier la vérification compositionnelle ;
- la vérification quantitative : étant donné un STA \mathcal{A} et une propriété φ , fournir un algorithme approchant la probabilité que \mathcal{A} vérifie φ .

Ce sont les deux aspects auxquels nous nous intéressons dans ma thèse.

Vérification presque sûre des systèmes probabilistes. Afin de considérer la vérification des STA, nous nous intéressons d'abord aux systèmes de transitions stochastiques (écrits STS). Ces systèmes peuvent être vus comme des chaînes de Markov avec un ensemble d'états qui peut-être non-dénombrable et munis donc de distributions qui peuvent être continues. Nous nous intéressons à deux types de problèmes dans ces STS :

- la vérification qualitative déjà citée plus haut ;
- la vérification quantitative : étant donné un STS \mathcal{T} et une propriété φ nous voudrions pouvoir fournir un algorithme qui approche la probabilité que \mathcal{T} vérifie φ .

Pour cela, nous nous sommes intéressé à ce papier : [ABM07].

Ce papier s'intéresse aux chaînes de Markov à temps discret (MC) et introduit une notion de *decisiveness* qui permet de faire passer les bonnes propriétés des MC finies aux MC infinies. Cette notion de decisiveness établit ceci : une MC \mathcal{M} est décisive par rapport à un certain ensemble d'états B si la probabilité d'atteindre un état dans B ou un état duquel on ne peut pas atteindre B , vaut 1. Il est montré que sous cette notion :

- le model-checking presque sûr de propriétés d'atteignabilité et d'atteignabilité répétée, se réduit à la vérification d'une sorte de propriété d'atteignabilité sur le graphe sous-jacent et est décidable sous certaines hypothèses supplémentaires raisonnables ;
- les algorithmes fournis pour approcher les probabilités d'une propriété d'atteignabilité et d'atteignabilité répétée se terminent bien et sont décidables sous certaines hypothèses supplémentaires raisonnables.

STS et decisiveness. Nous avons voulu alors adapter cette notion de decisiveness aux STS. Le problème pour cette extension est que les MC ont plusieurs hypothèses non-vérifiées par les STS : un ensemble d'états et un branchement au plus dénombrables. La présence de distributions continues pour les STS fait que l'on n'entre plus dans ce cadre.

Dans cette thèse, nous définissons une notion de decisiveness pour les STS quelconques : [BBBC16]. Cette notion de decisiveness permet d'obtenir des résultats similaires à ceux sur les MC :

- le model-checking presque sûr de propriétés d'atteignabilité et d'atteignabilité répétée, se réduit au 0-model-checking d'une sorte de propriété d'atteignabilité ;
- les mêmes algorithmes pour approcher les probabilités d'une propriété d'atteignabilité et d'atteignabilité répétée se terminent bien.

De plus, nous avons utilisé une procédure introduite dans [ABRS05] sur les *lossy channel systems* probabilistes (écrits LCSP) qui permet la vérification qualitative des propriétés données par un automate déterministe de Muller (écrit DMA) dans les MC. La première contribution a été d'adapter cette procédure afin d'obtenir un schéma d'approximation pour les propriétés données par un DMA. Cette procédure joue sur l'importante notion d'attracteur fini. Le but est ensuite d'étendre cette procédure aux STS.

Il reste alors deux grandes difficultés : comment prouver qu'un processus stochastique complexe (les STS) est décisif ? Et la notion d'attracteur fini n'étant pas adaptée ni raisonnable pour les STS, comment l'adapter ?

Pour répondre à ces problèmes, nous sommes passés par la notion d'abstraction. On montre que si un STS \mathcal{T}_1 possède une *bonne* abstraction \mathcal{T}_2 (c'est-à-dire \mathcal{T}_2 est une MC qui préserve le model-checking presque sûr de propriétés d'atteignabilité), alors \mathcal{T}_2 est décisif implique que \mathcal{T}_1 est également décisif. Si \mathcal{T}_1 possède une bonne abstraction, il suffit alors de montrer que son abstraction est décisive pour obtenir les résultats précédents. De plus, étant donné que la bonne

abstraction préserve l'atteignabilité presque sûre, il s'ensuit qu'un attracteur fini de \mathcal{T}_2 donne un attracteur dans \mathcal{T}_1 que nous pouvons utiliser afin d'obtenir une procédure similaire à celle dans les MC et qui nous permet de vérifier qualitativement et quantitativement les propriétés données par un DMA dans les STS.

Application aux STA. On peut montrer qu'un STA peut être exprimé par un STS. On travaille alors sur une abstraction bien connue des STA : le graphe des régions vu comme une MC ([BBB⁺14]). Cette MC est finie, [ABM07] nous permet donc de conclure que l'abstraction est décisive! Malheureusement cette abstraction n'est pas toujours *bonne* au sens dont nous avons besoin... Mais nous identifions plusieurs classes de STA dans lesquelles cette abstraction est bonne. Pour ces classes nous obtenons donc tous les résultats qualitatifs (déjà connus dans [BBB⁺14] mais l'approche proposée ici est plus uniforme) et surtout de nouveaux schémas d'approximations!

Vérification compositionnelle. En général, un système informatique est le résultat de plusieurs plus petits systèmes qui peuvent fonctionner indépendamment les uns des autres, ou en synchronisation avec les autres. Il est alors souvent plus simple de modéliser un à un les différents modèles et d'instaurer un opérateur de composition qui mettra en parallèle les différents systèmes et décrira le comportement du système informatique initial. C'est ce qu'on appelle la composition parallèle. La composition parallèle peut être avec ou sans synchronisation sur un certain ensemble d'actions.

Etant donné $\mathcal{A}_1, \dots, \mathcal{A}_n$ n systèmes modélisés pour un certain modèle, on écrit $\mathcal{A}_1 \parallel \dots \parallel \mathcal{A}_n$ pour la composition parallèle des n systèmes, avec ou sans ensemble de synchronisation. La vérification compositionnelle s'exprime de la façon suivante : étant donné une propriété φ , existe-t-il des propriétés $\varphi_1, \dots, \varphi_n$ telles que si pour chaque i compris entre 1 et n , \mathcal{A}_i vérifie φ_i , alors la composition des n systèmes (c'est-à-dire $\mathcal{A}_1 \parallel \dots \parallel \mathcal{A}_n$) vérifie la propriété φ ? L'objectif est d'étudier cet aspect sur les STA.

STA et composition. La notion de composition n'était pas encore définie sur les STA. Dans un premier temps, nous définissons un opérateur de composition parallèle pour les STA. Dans cette thèse, Nous définissons un tel opérateur dans le cas où l'on suppose que les automates que l'on compose sont indépendants. Nous donnons une classe de STA dans laquelle l'opérateur est bien défini et interne. Cette classe n'est définie que par des restrictions sur les lois de distributions, restrictions qui nous semble raisonnables. Nous montrons qu'avec l'hypothèse supplémentaire que l'ensemble des exécutions Zeno (c'est-à-dire des exécutions

infinies mais bornées en temps) a une probabilité nulle, l'opérateur exprime bien l'indépendance des automates composés : en restant peu précis, étant donné deux STA \mathcal{A}_1 et \mathcal{A}_2 , et deux propriétés φ_1 et φ_2 , on montre que la probabilité que $\mathcal{A}_1 \parallel \mathcal{A}_2$ vérifie φ_1 et φ_2 correspond à la multiplication des probabilités que d'une part, \mathcal{A}_1 vérifie φ_1 et d'autre part, \mathcal{A}_2 vérifie φ_2 .

Nous avons alors défini une notion de bisimulation sur les STA. Cette notion étend naturellement celle sur les chaînes de Markov à temps continu (CTMC) [DP03]. Nous montrons que la bisimulation est une congruence par rapport à la composition parallèle. Une telle propriété est attendue à chaque fois que l'on parle de composition parallèle dans différents modèles. Elle est requise afin de pouvoir étudier proprement une approche compositionnelle des modèles.

Afin de considérer la composition avec synchronisation, nous définissons le modèle des STA interactifs (écrits ISTA). Les ISTA étendent les STA dans la même façon que les MC interactives ([Her02]) étendent les MC : en plus des transitions probabilistes, nous avons des transitions interactives étiquetées par une action. Elles permettent d'établir des interactions entre différents ISTA. Nous définissons un opérateur de composition avec synchronisation dans les ISTA. Il se base sur le précédent dans les STA, il ajoute une condition de synchronisation étant donné un certain ensemble d'action A (pour faire court, les ISTA se synchronisent sur les actions présentes dans A). De manière similaire à précédemment, nous identifions une classe de ISTA dans laquelle la composition est bien définie et interne. Nous définissons également une notion de bisimulation qui est une congruence par rapport à la composition.

Nous terminons la thèse par un lien entre les deux grandes parties de la thèse. Nous identifions une classe de ISTA dans laquelle la composition est bien définie et dans laquelle, les STA sous-jacents ont une bonne abstraction, ce qui implique que tous les résultats qualitatifs et quantitatifs précédents peuvent être appliqués.

Conclusion. Nous avons fait quelques progrès dans la vérification des STA, qui peuvent se résumer de la façon suivante :

- l'apport d'une approche unifiée pour la vérification qualitative des STA et de nouveaux schémas d'approximation pour des propriétés d'atteignabilité, d'atteignabilité répétée et de DMA, qui se terminent sous une hypothèse de decisiveness ou d'attracteur ou encore sous l'hypothèse d'avoir une bonne abstraction ;
- l'apport d'un cadre compositionnel pour les STA que ce soit via l'opérateur de composition indépendant dans les STA ou bien via les ISTA et leur

opérateur de composition avec synchronisations, les premières étapes importantes pour se diriger vers l'étude de la vérification compositionnelle des STA.

Titre : Vérification des automates temporisés et stochastiques

Mots clés : vérification, automates, temporisé, stochastique

Résumé : Dans cette thèse, nous nous intéressons à la vérification formelle. On considère le modèle des automates temporisés et stochastiques (STA) qui est une extension probabiliste des automates temporisés, très connus. Les contributions de cette thèse se distinguent en deux parties : on étudie le *model-checking* qualitatif et quantitatif des STA, et la vérification compositionnelle des STA.

Dans la première partie, nous abordons l'analyse qualitative et quantitative des STA par les mêmes analyses de systèmes de transition stochastiques généraux (STS) qui peuvent être vus comme des chaînes de Markov (MC) générales avec un ensemble d'états continu.

Dans la deuxième partie, on définit un opérateur de composition dans les STA. D'abord, nous étudions le cas où les STA composés fonctionnent indépendamment, aboutissant à un opérateur *interleaving*. Pour permettre des interactions entre les systèmes, on définit le modèle des STA interactifs (ISTA) basé sur le modèle des MC interactives. Nous définissons alors un opérateur de composition dans les ISTA, qui est *handshaking*.

Nous terminons cette thèse par un lien entre les deux parties. Nous identifions une classe de ISTA dans laquelle la composition est bien définie et les résultats qualitatifs et quantitatifs précédents peuvent être appliqués.

Title : Verification of Stochastic Timed Automata

Keywords : verification, automaton, timed, stochastic

Abstract : In this thesis, we are interested in formal verification. We consider the stochastic timed automaton model (STA) which is a probabilistic extension of the well-known timed automaton model. The contributions of the thesis are twofold : we study the qualitative and quantitative model-checking problems of STA, and the compositional verification of STA.

In the first part, we tackle the qualitative and quantitative analysis of STA through the same analyses on general stochastic transition systems (STS), which can be seen as general Markov chains (MC) with a continuous set of states.

In the second part, we define an operator of composition in STA. We first study the case where the composed STA run independently, leading to an interleaving operator. In order to allow interactions between the systems, we define the interactive STA model (ISTA) based on the interactive MC model. We then define an operator of composition in ISTA, which is handshaking.

We end up the thesis with a link between the two parts. We identify a class of ISTA in which parallel composition is well-defined and in which the previous qualitative and quantitative results can be applied.

