MASARYK UNIVERSITY
FACULTY OF INFORMATICS

Formal Analysis of Discrete-Event Systems with Hard Real-Time Bounds

PH.D. THESIS

Jan Krčál

Brno, 2013

Declaration

I declare that this thesis is my own work and has not been submitted in any form for another degree or diploma at any university or other institution of tertiary education. Information derived from the published or unpublished work of others has been acknowledged in the text and a list of references is given.

Advisor: prof. RNDr. Antonín Kučera, Ph.D.

Acknowledgement

To my advisor
Tony Kučera
for teaching me that I know nothing and for his hamble attitude to excellence

To my colleague
Vojta Řehák
for his time for long discussions and for raising the spirits so many times

To my consultant
Tomáš Brázdil
for his generous scientific care and for his day-to-day patience with my ignorance

To the previous Ph.D. generation
Vojta and Vašek
for serving me as an attractor to the group (and being great friends as well)

To many other
friends
for support & love

To colleagues & friends
Petr, Ľuboš, Sebastian, Marek, Ondrej, and Reshma
for scientific debates as well as for making the life social in our great departement

To my brothers
Pavel, Ondra, and Marek
for inspiration for doing research and great things in general

To my colleague and friend
Jan Křetínský
for being a great sparing-partner for research and for deep discussions about life and its beauties

To my
Father
for teaching me to think

To my
Mother
for teaching me to thank

To my proof-readers
Vojtěch Forejt, Jan Křetínský, Barbora Krčálová
for valuable feedback and for spotting so many mistakes in the manuscript

Abstract

Discrete-event systems (DES) are widely used as a modelling formalism in probabilistic verification and performance evaluation. The behaviour of these models is driven by discrete events that occur randomly in continuous time. We study the impact of hard real-time bounds within DES such as time-outs changing the state of the model or deadlines in the specification of the desired behaviour of the model. Previously, there was no rigorous foundational material on such hard real-time bounds despite their presence in numerous practically oriented papers.

We show that DES with hard real-time bounds can exhibit an unstable behaviour and in general do not have a long-run average distribution. This is rather surprising as it contradicts several previous results. We also provide sufficient conditions upon which DES with hard real-time bounds are guaranteed to be stable. Furthermore, we study what changes if the hard real-time bounds are not part of the DES but a part of the specification of its desired behaviour. We show that such systems do not suffer from the instability observed in the previous case. Lastly, we define a 2-player game extension of DES with hard real-time bounds in the specification. We make use of the previous insight in the structure of such DES and provide a quantitative solution of the game extension.

Contents

Chapter 1

Introduction

For decades we have built more complicated systems than we are able to understand properly. We try to construct them in the most efficient way and without any errors in behaviour. Mathematical modelling, especially using software tools, helps us to cope with the increasing complexity of the systems we create. The general outline of the process called *formal analysis* is as follows. In the first step, the real system is abstracted using an appropriate modelling formalism. In the second step, we express formally the specification the system should satisfy or the performance aspects we want to analyse. In the third step, the actual analysis is performed automatically by a computer. The crucial assumption for successful utilization of formal analysis tools is their *correctness*. The playground for a theoretical computer scientist in this field is therefore finding powerful enough formalisms for step one and step two in the process above – with the nice attribute that there are *correct* and *efficient* algorithms to perform step three.

The key part of formal modelling is finding of an appropriate abstraction. One powerful approach to abstraction is *randomness*, as it allows us to simplify the model by declaring complex and intertwined phenomena as independent random events. Another approach is *non-determinism* allowing us to claim that we *do not know some parts of the system*. Such a non-determinism is later resolved during the analysis. Out of the possible options, either the most favourable behaviour is chosen if the unknown part is in our control, or the least favourable behaviour is chosen if the unknown part is out of our control and we want to find the formal guarantees the system provides.

Some systems, such as execution of a computer program, can be specified using *discrete-time* models that evolve in discrete steps. Other systems, such as water flow in a pipe, are continuous in nature and inappropriate for such simplification. We focus on *discrete-event systems (DES)* that change their state only in *discrete* points in time when *events* occur. However, the events occur *randomly* in *continuous-time* and their precise timing is important. A well-known example of such systems are the *continuous-time Markov chains (CTMC)*. The characteristic property of a CTMC, called also the memoryless or the Markov property, is that at any point in time its future behaviour depends only on its current state.

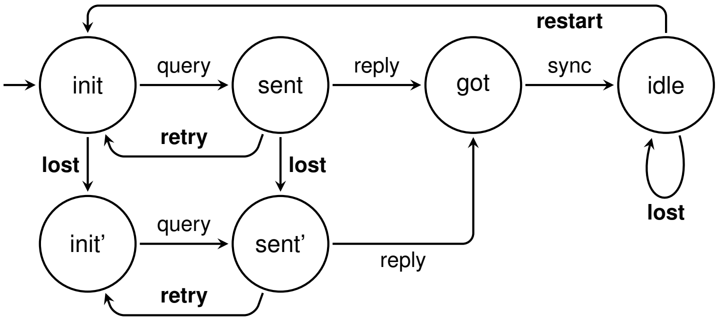

Figure 1.1: A GSMP model of a clock synchronization protocol. Fixed-delay events are printed in boldface. The delays of lost, restart, and retry are 100, 80, and 5, respectively. The remaining events are distributed continuously. (For example, the events query and reply may have Erlang distribution with shape 2 and rate 1, and the event sync may be distributed exponentially with rate 2.)

This greatly simplifies the analysis but also limits the modelling power. Namely it implies that the waiting times between the individual events are distributed according to an exponential distribution. On the one hand, the exponential distribution has various occurrences in nature as well as in man-made systems. Namely, it models the inter-arrival time of an event that can be triggered by a large amount of independent agents (e.g. a customer coming to a gas-station, a request reaching a web server, or a free electron causing isomerisation of a molecule). On the other hand, this distribution does not model faithfully many other natural phenomena. This leads to *non-Markovian* discrete-event systems that have arbitrarily distributed waiting times such as *semi-Markov processes* or *generalized semi-Markov processes (GSMP)* [Mat62].

A GSMP is basically a transition system over a finite set S of states where each transition is labelled by an event from a finite set $\mathscr{E}$ of events. A transition labelled with an event e leading from a state s denotes that e is *scheduled* to occur in s. All events scheduled to occur in s are awaited in parallel, each such event e occurs after a delay chosen randomly according to a fixed probability distribution F_e. When first of the events occurs, the corresponding transition leading to some state s' is taken. In state s', another event e' that was scheduled previously in s may stop being scheduled. When e' later becomes scheduled again due to another change of state, it is awaited anew regardless the amount of time it was scheduled previously. All the delay distributions F_e are often assumed to be continuous.

A GSMP may be further extended by inclusion of both the controllable and the non-controllable non-determinism into a two player turn-based *generalized semi-Markov game* (GSMG). In such a game, there are then two players and each player

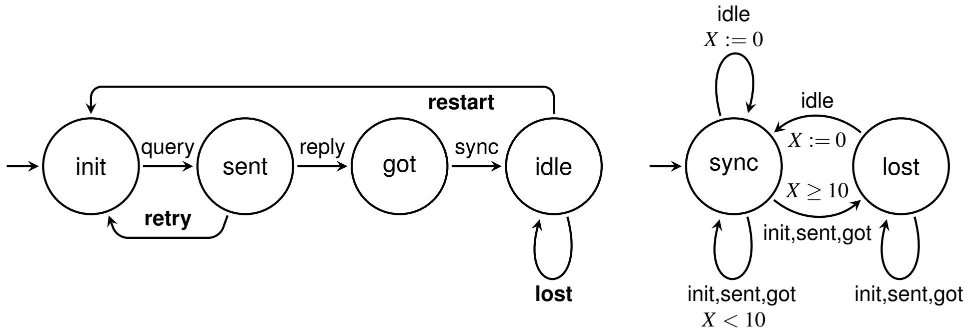

Figure 1.2: A GSMP model of a clock synchronization protocol observed by a deterministic timed automaton with one clock X measuring the time since the last synchronization.

controls her set of decision states. The player $\Box$ tries to guarantee with her decisions that the specification is satisfied, the player $\Diamond$ tries the opposite.

In the thesis, we study the impact of adding hard real-time bounds in non-Markovian discrete-event systems such as GSMP or GSMG. The bounds may be

1. a *part of the model* in the form of events that occur exactly after a fixed delay (the distribution F_e is then allowed to be concentrated on one point); or

2. a *part of the specification* such as time bounds in a logic formula or guards in a timed automaton that observes the behaviour of the process.

Let us illustrate these two approaches with an example of a simplified protocol for time synchronization. Via message exchange, the protocol sets and keeps a client clock sufficiently close to a server clock. Each message exchange is initialized by the client querying the server for the current time. The server replies back its current timestamp. This message exchange provides a reliable data for synchronization if it is realized within a specified round-trip delay. Otherwise, the client has to retry the procedure. After a success, the client is considered to be synchronized until a given delay elapses and the synchronization is lost due to potential clock skew. Since the aim is to keep the clocks synchronized all the time, the client restarts the synchronization process sooner, i.e. after a given delay that runs out before the synchronization is lost. Notice that the client gets desynchronized whenever several unsuccessful synchronizations occur in a row.

As regards the first approach, Figure 1.1 shows a GSMP model of this protocol where the fixed delays are modelled using fixed-delay events whereas the communication is modelled by continuously distributed events. As regards the second approach, observe that the time-out lost does not influence the synchronization.

Hence, we can make it a part of the specification as depicted in Figure 1.2. On the right, there is a *deterministic timed automaton (DTA)* that observes the process. It means that it evolves synchronously with the process and when the process takes a transition into a new state s', the DTA reads the input letter s'. Let us briefly recall that a DTA is basically a finite automaton enriched with real-valued clocks. As the time flows, the value of clocks increases. When an input letter is read, exactly one edge has its constraints (such as $X < 10$) over the current values of clocks satisfied. An edge may further prescribe some of the clocks to be reset (such as $X := 0$) when it is taken. Observe that Figure 1.2 actually combines both the approaches. In some situations it may be convenient to have the hard real-time bounds only a part of the DTA specification yielding a simpler model.

On the state space of the GSMP (or of the DTA if present) we study:

- probabilistic verification properties: the *reachability* (and the *Büchi*) property, i.e. the probability that a given target state is reached (infinitely often);

- performance evaluation measures: the discrete frequency $\mathbf{d}_s$ of visits to a given target state s and the timed frequency $\mathbf{c}_s$, i.e. the ratio of time spent in the given target state s. Formally, these random variables are defined by limits of partial ratios

  $$\mathbf{d}_s = \lim_{n\to\infty} \frac{\text{\# visits to } s \text{ in first } n \text{ steps}}{n} \qquad \mathbf{c}_s = \lim_{t\to\infty} \frac{\text{time spent in } s \text{ up to } t}{t}$$

  We say that $\mathbf{d}_s$ and $\mathbf{c}_s$ are *well-defined* if the limits exist with probability one. If they are well-defined, we are ideally interested in the probability distributions of their values.

In the examples above, we may study for example the probability that sync *is visited infinitely often in the DTA observer*, the frequency $\mathbf{c}_{\mathsf{lost}}$ in the DTA observer expressing the ratio of time the clocks are not synchronized, or the frequency $\mathbf{c}_{\mathsf{init}'} + \mathbf{c}_{\mathsf{sent}'}$ expresses the same in the GSMP in Figure 1.1.

After introducing the models and properties of interest, let us state that our primary concern is the *stochastic stability* [MT09] of these models with hard real-time bounds. Stability is a loosely defined concept related to the recurring patterns in the infinite behaviour of the process. Specifically, we ask questions such as:

- Under what conditions is a state visited infinitely often in a process with strongly connected transition graph?

- Under what conditions the long-run average behaviour stabilises so that the frequency measures $\mathbf{d}_s$ and $\mathbf{c}_s$ are well-defined?

Observe that rather than studying detailed *quantitative* questions, stochastic stability tackles the *qualitative* questions helping to understand the structure of the systems. For further clarification of this notion, see e.g. the discussion in [MT09, Section 1.3].

1.1 Contribution of the thesis

The main goal of the thesis is a fundamental research on stability of DES with hard real-time bounds.

Generalized semi-Markov processes We show that contrarily to previous results, GSMP with fixed-delay events may exhibit unstable behaviour. In particular,

- we show that the *region graph* previously applied to the qualitative analysis of GSMP [ACD91; ACD92] does not fully capture the qualitative behaviour of the system. Surprisingly, it does not hold that with probability one a bottom strongly connected component of the region graph is reached and all its nodes are visited infinitely often.
- Furthermore, we show that the discrete and timed frequencies do not have to be well-defined. As a result, the *steady-state distribution* may not exist for these systems. The previous literature presented various approximation algorithms for this quantity [Lin93; GL94; LS96; LRT99; BPS+98; HTT00; ZFG+00; ZFH01; Hor02; SDP03; HMM05; CGV09] without questioning its existence.

We also provide conditions on stability of the systems.

- As regards the first type of hard real-time bounds we consider, we define the class of *single-ticking* GSMP with fixed-delay events where the region graph characterizes the qualitative behaviour and where the discrete and timed frequencies are guaranteed to be well-defined. The conditions are easy to check algorithmically.
- As regards the second type of hard real-time bounds, we show that any DTA observer does not add any instability in the system. The proof is by reduction to a single-ticking GSMP that simulates the DTA observer on-the-fly.
- Furthermore, we address the closely related formalism of *deterministic and stochastic Petri nets (DSPN)*. We define the class of almost-monotone DSPN that we show to be stable again by reduction to single-ticking GSMP.

- Lastly, we show that the discrete and timed frequencies in a single-ticking GSMP can be effectively approximated.

Generalized semi-Markov games The first contribution is the definition of the formalism. Then, according to our goal, we focus on fundamental qualitative questions. We deal with both the reachability and the Büchi specification in the DTA observer. We show that

- player □ does not need to have an optimal strategy (recall that player □ tries to *satisfy* the specification).
- However if player □ has a strategy that guarantees winning with probability one, then □ also has a *structurally simple* such strategy that can be finitely represented using a DTA. This finite strategy suffices due to the stability of the observed system.
- Furthermore, we provide an algorithm that decides existence of such a strategy and constructs it if it exists.

1.2 Author's contribution

In this section, we summarize the author's contribution to the research in theoretical computer science.

Journals

[BFK+13] T. Brázdil, V. Forejt, J. Krčál, J. Křetínský, and A. Kučera. "Continuous-time stochastic games with time-bounded reachability". In: *Information and Computation* 224 (2013), pp. 46–70.

My contribution: Participated on the proceedings version of the paper (see below). Provided proofreading and detailed feedback. 15 %

International conference proceedings

[BFK+09] T. Brázdil, V. Forejt, J. Krčál, J. Křetínský, and A. Kučera. "Continuous-Time Stochastic Games with Time-Bounded Reachability". In: *Proceedings of the 29th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*. LIPIcs. Schloss Dagstuhl, 2009, pp. 61–72.

My contribution: Participated in discussions, formulated the algorithm and some of the appendix proofs. 20 %

[BKK+10b] T. Brázdil, J. Krčál, J. Křetínský, A. Kučera, and V. Řehák. "Stochastic real-time games with qualitative timed automata objectives". In: *Proceedings of 21st International Conference on Concurrency Theory (CONCUR)* (2010), pp. 207–221.

My contribution: Participated in discussions, provided some crucial insights, written various technical proofs and some parts of the main body. 30 %

[BKK+11a] T. Brázdil, J. Krčál, J. Křetínský, A. Kučera, and V. Řehák. "Measuring Performance of Continuous-Time Stochastic Processes using Timed Automata". In: *Proceedings of 14th International Conference on Hybrid Systems: Computation and Control (HSCC'11)*. ACM Press, 2011, pp. 33–42.

My contribution: Participated in discussions, devised and written major part of the technical proofs and some parts of the main body of the paper. 35 %

[BKK+11b] T. Brázdil, J. Krčál, J. Křetínský, and V. Řehák. "Fixed-delay events in generalized semi-Markov processes revisited". In: *Proceedings of 22nd International Conference on Concurrency Theory (CONCUR)*. Springer, 2011, pp. 140–155.

My contribution: Participated in discussions. Devised and written major part of the technical proofs of the crucial Theorems 4 and 5. Written various parts of the main body of the paper. 35 %

[BHK+12] T. Brázdil, H. Hermanns, J. Krčál, J. Křetínský, and V. Řehák. "Verification of Open Interactive Markov Chains". In: *Proceedings of 32th International Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*. LIPIcs. Schloss Dagstuhl, 2012, pp. 474–485.

My contribution: Participated in discussions, heavily influenced the problem formulation. Together with Jan Křetínský devised and written all the proofs. Participated on writing of the main body. 30 %

[BKK+13] T. Brázdil, L'. Korenčiak, J. Krčál, J. Křetínský, and V. Řehák. "On time-average limits in deterministic and stochastic Petri nets". In: *Proceedings of the ACM/SPEC International conference on performance engineering (ICPE)*. Poster paper. ACM. 2013, pp. 421–422.

My contribution: Participated in discussions, written the abstract of the paper. 35 %

[HKK13] H. Hermanns, J. Krčál, and J. Křetínský. "Compositional Verification and Optimization of Interactive Markov Chains". In: *Proceedings of 24th International Conference on Concurrency Theory (CONCUR)*. 2013, pp. 364–379.

My contribution: Heavily participated in discussions, devised and written major part of the proofs. Participated on various parts of the main body of the paper. 45 %

The thesis is based on the conference papers [BKK+10b; BKK+11a; BKK+11b; BKK+13]. However, most of the material was completely rewritten. In particular,

- compared to [BKK+10b], all the definitions have been rewritten resulting in a cleaner presentation. The proofs have been rewritten (and various mistakes have been corrected). Furthermore, the Büchi specification has been added.
- Results from [BKK+11a] have been proven by a novel reduction to the problem considered in the follow-up paper [BKK+11b].
- Compared to [BKK+11b], most of the proofs have been revised (and various mistakes have been corrected).
- The material from [BKK+13] has been enhanced with formal proofs that were omitted in the proceedings due to space constraints.

1.3 Outline of the thesis

Let us outline the structure of the thesis.

Chapter 2 provides the preliminaries and definitions used throughout the thesis. We first define the modelling formalism, mainly the GSMP and GSMP with fixed-delay events. Then we formalize the specification formalisms that we use later.

Chapter 3 provides an overview of related research in the area. We review related modelling formalisms, mainly those impacted by our results. Then we comment on related specification formalisms. In the third section, we provide an overview of results on stability of non-Markovian DES. In the final section, we survey the solution methods for GSMP and related formalisms and comment on the impact of our results on the methods.

Chapter 4 explains the instability results of GSMP with fixed-delay events. First, we prove the insufficiency of the region graph for the qualitative verification of the Büchi specifications. Then we show that the frequency measures are not well-defined in general.

Chapter 5 then provides conditions for stability. First, we address the stability of GSMP with fixed-delay events. Second, we analyse the stability of GSMP observed by a DTA. Third, we study the stability of the related formalism of deterministic and stochastic Petri nets. Finally, the chapter is concluded by showing that the frequency measures can be effectively approximated in stable GSMP.

Chapter 6 deals with the non-deterministic extension of GSMP, namely the generalized semi-Markov games. First, the formalism is defined. In the following section, the qualitative reachability in a DTA observer is addressed. We give an algorithm for constructing a simple optimal strategy (if any optimal strategy exists). In the last section, the reachability results are applied to solving the Büchi specification.

Chapter 7 concludes the thesis and discusses ideas for future work.

Chapter 2

Discrete-event systems (DES)

In this chapter, we formally introduce the area of the discrete-event systems. We avoid the unnecessary technical details; when appropriate, we refer the reader to more rigorous sources. Instead, we illustrate the notions by numerous examples. Let us briefly define the basic notation.

In this text, the sets of all positive integers, non-negative integers, real numbers, positive real numbers, and non-negative real numbers are denoted by $\mathbb{N}$, $\mathbb{N}_0$, $\mathbb{R}$, $\mathbb{R}_{>0}$, and $\mathbb{R}_{\geq 0}$, respectively. For a non-negative real number $r \in \mathbb{R}_{\geq 0}$, $\lfloor r \rfloor$ denotes its integral part, i.e. the largest integer smaller than r, and $\langle r \rangle$ denotes its fractional part, i.e. $r - \lfloor r \rfloor$. Let A be a finite or countably infinite set. A *probability distribution* on A is a function $f : A \to \mathbb{R}_{\geq 0}$ such that $\sum_{a \in A} f(a) = 1$. The set of all distributions on A is denoted by $\mathscr{D}(A)$.

A *σ-field* over a set Ω is a set $\mathscr{F} \subseteq 2^{\Omega}$ that includes Ω and is closed under complement and countable union. A *measurable space* is a pair $(\Omega, \mathscr{F})$ where Ω is a set called *sample space* and $\mathscr{F}$ is a σ-field over Ω whose elements are called *measurable sets*. Given a measurable space $(\Omega, \mathscr{F})$, we say that a function $f : \Omega \to \mathbb{R}$ is a random variable if the inverse image of any real interval is a measurable set. A *probability measure* over a measurable space $(\Omega, \mathscr{F})$ is a function $\mathscr{P} : \mathscr{F} \to \mathbb{R}_{\geq 0}$ such that $\mathscr{P}(\Omega) = 1$ and we have $\mathscr{P}(\bigcup_{i \in I} X_i) = \sum_{i \in I} \mathscr{P}(X_i)$ for each countable collection $\{X_i\}_{i \in I}$ of pairwise disjoint elements of $\mathscr{F}$. A *probability space* is a triple $(\Omega, \mathscr{F}, \mathscr{P})$, where $(\Omega, \mathscr{F})$ is a measurable space and $\mathscr{P}$ is a probability measure over $(\Omega, \mathscr{F})$. We say that a property $A \subseteq \Omega$ holds for *almost all* elements of a measurable set Y if $\mathscr{P}(Y) > 0$, $A \cap Y \in \mathscr{F}$, and $\mathscr{P}(A \cap Y \mid Y) = 1$.

A function $f : \mathbb{R} \to \mathbb{R}_{\geq 0}$ is called a *density* if $\int_0^\infty f(x)dx = 1$. For a density f and a value $t \in \mathbb{R}_{\geq 0}$ we define a "shifted" density function $f_{|t}$ as

$$f_{|t}(x) = \frac{f(x+t)}{\int_t^\infty f(y)\,dy}.$$

Notice the denominator scales the function up so that it is a density again. A density f is called *Markovian* if $f(x) = f_{|t}(x)$ for any $t, x \in \mathbb{R}_{\geq 0}$.

2.1 Modelling formalisms

The field of discrete-event systems is broad [CL08], we focus on generalized semi-Markov processes (with deterministic events). As this formalism is an extension of the widely used continuous-time Markov chains, we start with CTMCs. All the models we define share the same concepts: discrete state space, event-driven behaviour, and random arrival time of events.

Example 2.1.1. *Consider an example of a public photo booth. It has two states –* free *and* occupied*. With the flow of time it remains in its current state until an event comes –* a person enters the booth *or* a person leaves the booth*. Notice the current state determines which events can come – if the booth is free, no one can leave it. Furthermore, we can easily measure or estimate the times we wait for the individual events. It can happen that a person leaves the booth* 2 *seconds or* 2 *hours after entering it but it is rare. Such observation may lead us to model this waiting time using, e.g., the log-normal distribution.*

The formalism of CTMCs offers the most limited choice of probability distributions on events' arrival: only the exponential distributions are allowed.

2.1.1 Continuous-time Markov chains (CTMC)

A *continuous-time Markov chain (CTMC)* is a tuple $\mathscr{C} = (S, \mathbf{R}, \alpha_0)$, where S is a finite set of states, $\mathbf{R} : S \times S \to \mathbb{R}_{\geq 0}$ is a rate matrix, and $\alpha_0 \in \mathscr{D}(S)$ is an initial distribution.

The set S is the discrete state space. The system starts in a state randomly chosen according to the discrete distribution α_0. Then it moves from state to state whenever an (exponentially distributed) event occurs. An exponential distribution is fully specified by a single parameter, its *rate*. As its name suggests, it corresponds to the rate of occurrence of the event per time unit. For example, if on average 3 people try to enter the photo booth per hour, we use the rate 3. For any two states s and s', the rate matrix $\mathbf{R}(s,s')$ specifies the rate of such an event that is awaited in state s and upon whose arrival the system moves to state s'. If no such event can take place, it is expressed by a rate 0. CTMCs are often graphically represented as directed graphs (allowing self-loops) where each edge corresponds to an event. For the photo booth example, see Figure 2.1.

A *run* of a CTMC is an infinite sequence of the form $\sigma = s_0\, t_0\, s_1\, t_1\, s_2\, \cdots$ where $s_i \in S$ and $t_i \in \mathbb{R}_{>0}$ for each $i \in \mathbb{N}_0$. This means that the system starts in state s_0, stays there for t_0 time units until an event occurs and the system moves to state s_1. It waits in state s_1 for time t_1 and moves to s_2 and so forth. Since there are uncountably many runs, each individual run has probability 0. But we can assign

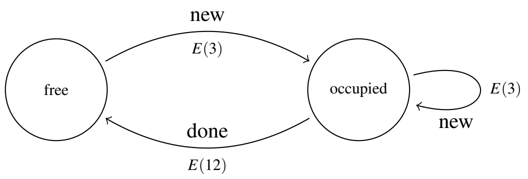

Figure 2.1: A photo booth modelled as a CTMC. The notation $E(\lambda)$ denotes the exponential distribution with rate λ. Notice that if there was no delay between the customers' visits so that the booth was occupied all the time, 12 people would leave it on average per hour. The initial distribution assigns probability 0.5 to the state *free* and 0.5 to the state *occupied* and is not depicted.

positive probabilities to (measurable) sets of runs. An example of such a set of runs is *the first* 5 *people enter the booth in the first hour*. Formally these probabilities are specified using a probability space $(\Omega, \mathscr{F}, \mathscr{P})$ over the set of all runs Ω. For the case of CTMCs, see e.g. [BH03]. We will get into more details when defining GSMP with fixed-delay events.

2.1.2 Generalized semi-Markov processes (GSMP)

For some systems, the exponential distribution is an appropriate abstraction. Yet, it is a strong limitation for many other systems; various natural phenomena can be more faithfully modelled using non-Markovian distributions. The generalized semi-Markov processes are able to model parallel systems with arbitrarily distributed arrival times.[1]

Definition 2.1.2 (GSMP)**.** *A* generalized semi-Markov process (GSMP) *is a tuple* $\mathscr{M} = (S, \mathscr{E}, \mathbf{E}, \mathbf{P}, \alpha_0)$ *where*

- *S is a finite set of states,*
- *$\mathscr{E}$ is a finite set of independent continuously distributed events where we associate to each event e its* density function *f_e,*
- *$\mathbf{E} : S \to 2^{\mathscr{E}}$ assigns to each state the set of events scheduled to occur in this state,*
- *$\mathbf{P} : S \times \mathscr{E} \to \mathscr{D}(S)$ is a successor function, and*
- *$\alpha_0 \in \mathscr{D}(S)$ is an initial distribution.*

We illustrate with an example how the events are scheduled to occur in parallel.

[1] The earlier formalism of *semi-Markov processes*, the sequential extension of CTMCs, allow only for one non-Markovian action to be performed at a time.

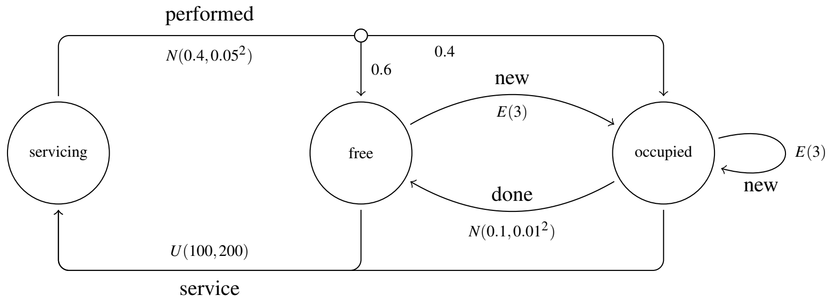

Figure 2.2: Photo booth with servicing, an example of a GSMP. The *service* event can occur in both *free* and *occupied* states and its distribution depends on the elapsed time.

Example 2.1.3. *In Figure 2.2 we combine the customers of the photo booth with a (simplified) servicing procedure. Imagine that the process starts in the state* free*. There are two parallel events, the event* new*, exponentially distributed with rate* 3*, and the event* service*, distributed uniformly on* $[100,200]$*. As the time flows, many new customers come and the process alternates between the states* free *and* occupied*. But the fact that new customers arrive does not postpone the servicing. After the first* 50 *hours, it is distributed uniformly between* 50 *and* 150 *hours. The process is not Markovian any more: knowing the current state does not give you the complete information. We must also remember what time ago each event was scheduled - the* elapsed time *of the event.*

To process starts in a state s_0 randomly chosen according to α_0. At this moment, the *elapsed time* of each event $e \in \mathbf{E}(s_0)$ is zero. The scheduled events occur randomly in continuous-time causing the process to change state. When is state s an event e occurs, the process moves into a random state according to the fixed probability distribution given by $\mathbf{P}(s,e)$. For example, when a service is *performed* the process moves with probability 0.6 to the state *free* and with probability 0.4 to the state *occupied*. As long as the process moves between the states where an event e is scheduled to occur, its elapsed time accumulates. When this event is triggered or the process enters a state where e is not scheduled, its elapsed time is reset. When the process returns into a state where the event e is scheduled to occur, its elapsed time starts accumulating from 0 again. The elapsed time t of an event e influences the probability that is happens in the future; it occurs with density $f_{e|t}$ (provided it does not stop being scheduled by an occurrence of another event before it occurs itself).

More formally, a *configuration* is a triple (s, ξ, t) where $s \in S$, ξ is a vector of *elapsed time* which assigns to every event $e \in \mathbf{E}(s)$ the amount of time that elapsed since the event e was scheduled,[2] and t is the time spent in the previous configuration. We define $\xi(e) = \bot$ whenever $e \notin \mathbf{E}(s)$. The process starts in an initial configuration $(s_0, \xi_0, 0)$ where s_0 is chosen randomly according to the initial distribution α_0 and ξ_0 assigns zero to all events enabled in s_0. Then the process moves from configuration to configuration forming a run $(s_0, \xi_0, t_0)(s_1, \xi_1, t_1)(s_2, \xi_2, t_2 s)\cdots$. In configuration (s_i, ξ_i, t_i), some time $t_{i+1} > 0$ is spent and then some event e occurs with density

$$f_{e|\xi(e)}(t_{i+1}) \cdot \prod_{c \in \mathbf{E}(s_i) \setminus \{e\}} \int_{t_{i+1}}^{\infty} f_{c|\xi(c)}(y)\, dy.$$

Observe that the first term corresponds to event e occurring at time t_{i+1} and the second term corresponds to all other scheduled events occurring after a delay longer than t_{i+1}. Then the process moves to $(s_{i+1}, \xi_{i+1}, t_{i+1})$ where s_{i+1} is chosen randomly according to $\mathbf{P}(s_i, e)$ and ξ_{i+1} is obtained from ξ_i as follows. The elapsed time of *old* events of $\mathbf{E}(s_i) \smallsetminus \mathbf{E}(s_{i+1})$ is discarded to $\bot$, the elapsed time of each *inherited* event of $(\mathbf{E}(s_{i+1}) \cap \mathbf{E}(s_i)) \smallsetminus \{e\}$ is increased by t_{i+1}, and the elapsed time of each *new* event of $(\mathbf{E}(s_{i+1}) \smallsetminus \mathbf{E}(s_i)) \cup (\mathbf{E}(s_{i+1}) \cap \{e\})$ is set to 0. The described behaviour again induces a probability measure $\mathscr{P}$ over the measurable space $(\Omega, \mathscr{F})$. This process was introduced by Matthes [Mat62], for a more recent formal definition including the probability space see e.g. [Whi80b]. Notice that we define the probability space rigorously in the following subsection for a more general model.

2.1.3 GSMP with fixed-delay events

Not every distribution has a density function. In particular, discrete distributions cannot be expressed this way. In the area of electronic systems, discrete distributions play a special role since they allow expressing deterministic events such as time-outs. The definition needs to be changed as multiple (fixed-delay) events can occur at the same moment of time.

Definition 2.1.4 (GSMP with fixed-delay events)**.** *A* generalized semi-Markov process with fixed-delay events *is a tuple* $(S, \mathscr{E}, \mathbf{E}, \mathrm{Succ}, \alpha_0)$ *where*

- S *is a finite set of* states*,*
- $\mathscr{E}$ *is a finite set of* events *where to every* $e \in \mathscr{E}$ *we associate the* lower bound $\ell_e \in \mathbb{N}_0$ *and the* upper bound $u_e \in \mathbb{N} \cup \{\infty\}$ *of its delay. We say that* e *is a* fixed-delay *event if* $\ell_e = u_e$*, and a* variable-delay *event if* $\ell_e < u_e$*. To*

[2] Alternatively, the configuration is defined to store the time left before the event appears. Our definition (similar to [Gly89; YS04]), is equivalent and (a) more convenient for our proof techniques in the general setting where both bounded and unbounded events appear and (b) necessary for the game extension of GSMP.

each variable-delay event e we assign a density function $f_e : \mathbb{R} \to \mathbb{R}$ *such that* $\int_{\ell_e}^{u_e} f_e(x)\,dx = 1$.

- $\mathbf{E} : S \to 2^{\mathscr{E}}$ *assigns to each state s a set of events* $\mathbf{E}(s) \neq \emptyset$ scheduled *to occur in s,*
- $\mathrm{Succ} : S \times 2^{\mathscr{E}} \to \mathscr{D}(S)$ *is the* successor *function, i.e. assigns a probability distribution specifying the successor state to each state and set of events that occur simultaneously in this state, and*
- $\alpha_0 \in \mathscr{D}(S)$ *is the* initial distribution.

Furthermore, let e be any variable-delay event. We say that e is *bounded* if $u_e \neq \infty$, and *unbounded*, otherwise. We assume f_e to be positive and continuous on the whole $[\ell_e, u_e]$ or $[\ell_e, \infty)$ if e is bounded or unbounded, respectively, and zero elsewhere. Finally, we require that there is a real number $r \in \mathbb{R}$ such that for any unbounded event e it holds $\sup_{t \geq \ell_e} \mathrm{E}[e \,|\, t] \leq r$ where $\mathrm{E}[e \,|\, t]$ denotes $\int_0^\infty x \cdot f_{e|t}(x)\,dx$. Intuitively, we require that the expected waiting for an event is finite and does not increase beyond any bounds when we increase the time we have already waited. Observe that most of the standard distributions satisfy this assumption.

The intuitive dynamics of the GSMP with fixed-delay events is the same as of GSMP, except for the fact that several events can occur simultaneously as illustrated by the following example.

Example 2.1.5. *In Figure 2.3, there is a photo-booth that can close itself automatically, when some servicing is necessary. Every* 100 *hours,* cleaning *is* needed*; every* 300 *hours paper and toner also* need *to be* refilled. *Notice that for the servicing events to occur simultaneously in the state* in use*, they both need to be scheduled also in* out of order.

A formal semantics of GSMP is usually defined in terms of general state-space Markov chains (GSSMC, see, e.g., [MT09]). A GSSMC is a stochastic process Φ over a measurable state-space $(\Gamma, \mathscr{F})$ whose dynamics is determined by an initial measure μ on $(\Gamma, \mathscr{F})$ and a *transition kernel P* which specifies one-step transition probabilities.[3] A given GSMP induces a GSSMC whose state-space consists of all configurations, the initial measure μ is induced by α_0 in a natural way, and the transition kernel is determined by the dynamics of GSMP described above. Formally,

- Γ is the set of all configurations $\{(s, \xi, t) \mid s \in S, \xi : \mathscr{E} \to \mathbb{R}_{\geq 0}, t \in \mathbb{R}_{\geq 0}\}$, and $\mathscr{F}$ is a σ-field over Γ induced by the discrete topology over S and the Borel σ-fields over vectors of elapsed times and over $\mathbb{R}_{\geq 0}$;

[3] Precisely, transition kernel is a function $P : \Gamma \times \mathscr{F} \to [0,1]$ such that $P(z, \cdot)$ is a probability measure over $(\Gamma, \mathscr{F})$ for each $z \in \Gamma$; and $P(\cdot, A)$ is a measurable function for each $A \in \mathscr{F}$.

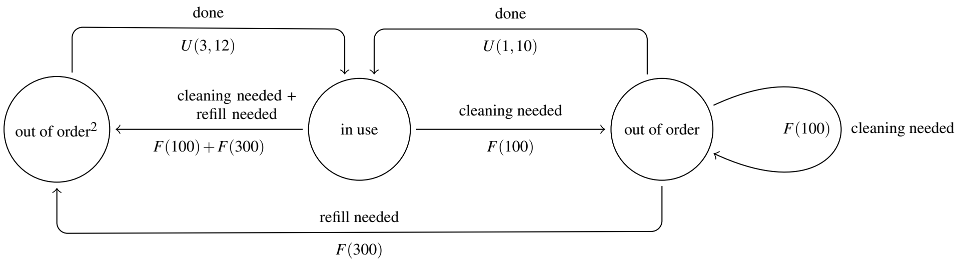

Figure 2.3: Servicing of a photo booth, an example of a generalized semi-Markov process with fixed-delay events. The distribution with fixed delay t is denoted by $F(t)$. For better clarity, we omit transitions for subsets of events that occur with probability 0 such as only $\{\text{refill needed}\}$ in the states in use and out of order.

- the initial measure μ allows starting only in configurations with zero elapsed time, i.e. for each measurable set $A \in \mathscr{F}$ we have $\mu(A) = \sum_{s \in S, (s,\mathbf{0},0) \in A} \alpha_0(s)$;
- the transition kernel $P(z,A)$ describing the probability to move in one step from a configuration $z = (s, \xi, t)$ to any configuration in a measurable set A is defined as follows. It suffices to consider A of the form $\{s'\} \times X \times I$ where X is a measurable set of vectors of elapsed time and I is a measurable subset of $\mathbb{R}_{\geq 0}$. Let V and F be the sets of variable-delay and fixed-delay events, respectively, that are scheduled in s. Let $F' \subseteq F$ be the set of fixed-delay events that can occur as first among the fixed-delay event enabled in z, i.e. that have in ξ the minimal remaining time u. Note that two variable-delay events occur simultaneously with probability zero. Hence, along with F after time u we consider all combinations of $e \in V$ after time $t < u$:

$$P(z,A) = \begin{cases} \sum_{e \in V} \int_0^\infty \mathrm{Hit}(\{e\},t) \cdot \mathrm{Wins}(\{e\},t)\, dt & \text{if } F = \emptyset \\ \sum_{e \in V} \int_0^u \ \mathrm{Hit}(\{e\},t) \cdot \mathrm{Wins}(\{e\},t)\, dt & \\ \quad + \mathrm{Hit}(F',u) \cdot \mathrm{Wins}(F',u) & \text{otherwise,} \end{cases}$$

where the term $\mathrm{Hit}(E,t)$ denotes the conditional probability of hitting A under the condition that E occurs at time t and the term $\mathrm{Wins}(E,t)$ denotes the probability (density) of E occurring at time t. Formally,

$$\mathrm{Hit}(E,t) = \mathrm{Succ}(s,E)(s') \cdot \mathbf{1}[\xi' \in X \wedge t \in I]$$

where $\mathbf{1}[\xi' \in X]$ is the indicator function and ξ' is the elapsed time after the transition, i.e. $\xi'(e)$ is $\perp$, or $\xi(e) + t$, or 0 for each old, or inherited, or new event e, respectively. The most complicated part is the definition of $\mathrm{Wins}(E,t)$ which intuitively corresponds to the probability that E is the set of events "winning" the competition among the events scheduled in s at time t. Recall that $f_{e|\xi(e)}$ denotes the "shifted" density function that takes into account that the time $\xi(e)$ has already elapsed. We have

$$\mathrm{Wins}(E,t) = \begin{cases} f_{e|\xi(e)}(t)\cdot\prod_{c\in V\setminus E}\int_t^\infty f_{c|\xi(c)}(y)\,dy & \text{if } E=\{e\}\subseteq V,\\ \prod_{c\in V}\int_t^\infty f_{c|\xi(c)}(y)\,dy & \text{if } E=F'\subseteq F,\\ 0 & \text{otherwise.}\end{cases}$$

A *run* of the Markov chain is an infinite sequence $\sigma = z_0\, z_1\, z_2 \cdots$ of configurations. The Markov chain is defined on the probability space $(\Omega,\mathscr{F},\mathscr{P})$ where Ω is the set of all runs, $\mathscr{F}$ is the product σ-field $\bigotimes_{i=0}^{\infty}\mathscr{F}$, and $\mathscr{P}$ is the unique probability measure [MT09] such that for every finite sequence $A_0,\cdots,A_n\in\mathscr{F}$ we have that

$$\mathscr{P}(\Phi_0{\in}A_0,\cdots,\Phi_n{\in}A_n) = \int\limits_{z_0\in A_0}\cdots\int\limits_{z_{n-1}\in A_{n-1}} \mu(dz_0)\cdot P(z_0,dz_1)\cdots P(z_{n-1},A_n)$$

where each Φ_i is the i-th projection of an element in Ω (the i-th configuration of a run). Finally, we define an m-step transition kernel P^m inductively as $P^1(z,A) = P(z,A)$ and $P^{i+1}(z,A) = \int_\Gamma P(z,dy)\cdot P^i(y,A)$.

Remark 2.1.6. *Observe that, for simplicity, we allow only events with either continuous or Dirac distribution. This is not a huge restriction as every discrete distribution with finite support (or arbitrary support) can be simulated (or approximated) by a chain of fixed-delay events. Furthermore, any distribution can be uniquely decomposed to its discrete part and its continuous part and hence simulated by a continuous event and a sequence of fixed-delay events likewise.*

For a given DES we denote by $\mathscr{P}_z$ the probability measure of the system when it starts in the configuration z, i.e. for any measurable set A of configurations the initial measure is $\mu(A)=1$ if $z\in A$ and $\mu(A)=0$, otherwise.

2.2 Specification formalisms

In this section, we address the ways to specify the desired behaviour of DES. We are mainly concerned with the long-run behaviour. Namely with the ω-regular properties and performance measures in both the GSMP and its DTA observer as outlined in the following table.

	GSMP $\mathscr{M}$	DTA $\mathscr{A}$ observing $\mathscr{M}$
ω-regular properties	$\text{Büchi}_{\mathscr{M}}(T)$	$\text{Reach}_{\mathscr{A}}(T)$, $\text{Büchi}_{\mathscr{A}}(T)$
performance measures	$\mathbf{d}_s^{\mathscr{M}}$, $\mathbf{c}_s^{\mathscr{M}}$	$\mathbf{d}_q^{\mathscr{A}}$, $\mathbf{c}_q^{\mathscr{A}}$

The notions from the table are defined in the following subsections. For an overview of related specification formalisms studied in the literature, see Chapter 3.

2.2.1 ω-regular properties in a GSMP

The Büchi properties are a widely studied class of ω-regular properties.

Definition 2.2.1 (Büchi specification)**.** *For a given GSMP $\mathscr{M}$, the Büchi specification over a subset of target states $T \subseteq S$ is the set of runs* $\text{Büchi}_{\mathscr{M}}(T)$ *that visit some state in T infinitely often.*

Observe that the Büchi specification does not involve any explicit time constraints. Nevertheless, the analysis of this specification is far from straightforward, as shown in Section 4.1.

Example 2.2.2. *If we set $G = \{\mathsf{servicing}\}$ for the system from Figure 2.2, we specify that servicing must occur infinitely many times. It is easy to see that for this system we have $\mathscr{P}(\text{Büchi}_{\mathscr{A}}(G)) = 1$.*

2.2.2 Performance measures in a GSMP

Though the Büchi specification requires some behaviour to occur infinitely often, it cannot quantify *how often* it actually occurs. These performance concerns are important for many applications [BHH+05] and can be captured by various performance measures. A typical type of performance measures are the time-average limits of *reward functions* or more complicated *reward structures*. Another type of performance measures quantifies various delays [Haa10] in the discrete-event systems. The basic approach that we deal with in the thesis is the *discrete frequency* of visits to a given state and the *timed frequency*, i.e. the ratio of time spent in a give state,

Definition 2.2.3 (Discrete and timed frequencies)**.** *For a given GSMP $\mathscr{M}$ with set of states S, the* discrete *and* timed *frequency of visits to $s \in S$ along a run $\sigma = (s_0, \xi_0, t_0)(s_1, \xi_1, t_1)\cdots$ of $\mathscr{M}$, denoted by $\mathbf{d}_s^{\mathscr{M}}(\sigma)$ and $\mathbf{c}_s^{\mathscr{M}}(\sigma)$, are defined by*

$$\begin{aligned} \mathbf{d}_s^{\mathscr{M}}(\sigma) &= \lim_{n\to\infty} \frac{\sum_{i=0}^{n} 1_{s_i=s}}{n} \\ \mathbf{c}_s^{\mathscr{M}}(\sigma) &= \lim_{n\to\infty} \frac{\sum_{i=0}^{n} t_{i+1} \cdot 1_{s_i=s}}{\sum_{i=1}^{n} t_{i+1}} \end{aligned}$$

where $1_{s_i=s}$ equals 1 if $s_i = s$ and $1_{s_i=s}$ equals 0, otherwise. We write $\mathbf{d}_s$ and $\mathbf{c}_s$ if $\mathscr{M}$ is clear from context. We say that the measure $\mathbf{d}_s$ or $\mathbf{c}_s$ is well-defined *for a run σ if the corresponding limit exists.*

Observe that these performance measures make sense only if they are (almost surely) well-defined. This, roughly speaking, corresponds to *stability* of the system under study.

Example 2.2.4. *For the system from Figure 2.3, the measure*

$$\mathbf{d}_{\textsf{out of order}^2} / (\mathbf{d}_{\textsf{out of order}} + \mathbf{d}_{\textsf{out of order}^2})$$

specifies the ratio of servicing where also refill is needed which is equal with probability 1 *to* $1/3$*. Similarly, the measure* $\mathbf{c}_{\textsf{out of order}} + \mathbf{c}_{\textsf{out of order}^2}$ *specifies the ratio of time when the booth is out of order. A straightforward computation reveals that with probability* 1 *it is* ≈ 0.06.

2.2.3 DTA specifications

Timed automata have been introduced [AD94] as a modelling formalism for worst-case analysis of real-time systems. Since then a huge community grew up around this model extending it and using it in many other directions. One such direction is using deterministic timed automata (DTA) as a specification formalism for stochastic real-time systems, originally proposed by Alur et al. [ACD92] for the qualitative analysis of GSMP. Almost two decades later, the stochastic logic CSL^{TA} incorporating one-clock timed automata was defined [DHS07] and timed automata were also used [CHK+09] for quantitative analysis of CMTC. We use this specification formalism to study the long-run behaviour of GSMP (and its 2-player game extension).

A timed automaton is basically a finite automaton endowed with a finite set of real-valued clocks. The timed automaton reads timed words, i.e. sequences of input letters and time stamps. Notice that a run of a stochastic system can be easily translated into an infinite time word (a sequence of states and times spent there).

Let $\mathscr{X}$ be a finite set of *clocks*. A *valuation* is a function $\nu : \mathscr{X} \to \mathbb{R}_{\geq 0}$. For a valuation ν and a subset $X \subseteq \mathscr{X}$ of clocks, we denote by $\nu[X := \mathbf{0}]$ the valuation assigning 0 to $x \in X$ and $\nu(x)$ to $x \notin X$. Further, for a valuation ν and $t \in \mathbb{R}_{\geq 0}$, we denote by $\nu + t$ the valuation assigning $\nu(x) + t$ to each $x \in \mathscr{X}$.

A *clock constraint* (or *guard*) is a finite conjunction of basic constraints of the form $x \bowtie c$, where $x \in \mathscr{X}$, $\bowtie \in \{<, \leq, \geq, >\}$, and $c \in \mathbb{N}_0$. For a valuation ν and a clock constraint g, let $\nu \models g$ denote that ν satisfies g (the satisfaction relation is defined in the expected way). Further, by $[g]$ we denote the set of valuations that satisfy g. Finally, the set of all guards over $\mathscr{X}$ is denoted by $\mathscr{B}(\mathscr{X})$.

Definition 2.2.5 (DTA)**.** *A* deterministic timed automaton (DTA) *is a tuple* $\mathscr{A} = (Q, \Sigma, \mathscr{X}, \longrightarrow, q_{init})$*, where* Q *is a finite set of* locations*,* Σ *is a finite* alphabet*,* $\mathscr{X}$

is a finite set of clocks*,* $q_{init} \in Q$ *is an* initial location*, and* $\longrightarrow \subseteq Q \times \Sigma \times \mathscr{B}(\mathscr{X}) \times 2^{\mathscr{X}} \times Q$ *is an* edge relation *such that for all* $q \in Q$ *and* $a \in \Sigma$ *we have the following:*

1. *the guards are deterministic, i.e., for all pairs of distinct edges of the form* (q,a,g_1,X_1,q_1) *and* (q,a,g_2,X_2,q_2) *we have that* $[g_1] \cap [g_2] = \emptyset$*;*

2. *the guards are total, i.e., for every* $q \in Q$*,* $a \in \Sigma$*, and valuation* ν *there is an edge* (q,a,g,X,q') *such that* $\nu \models g$*.*

A *configuration* of $\mathscr{A}$ is a pair (q,ν), where $q \in Q$ and ν is a valuation. An *infinite timed word* is an infinite sequence $\omega = c_0c_1c_2\cdots$ where each c_i is either a letter of Σ or a non-negative real number denoting a time stamp. The *run* of $\mathscr{A}$ on ω is the unique infinite sequence $\mathscr{A}(\omega) = (q_0,\nu_0)\,c_0\,(q_1,\nu_1)\,c_1\,\cdots$ such that $q_0 = q_{init}$, $\nu_0 = \mathbf{0}$, and for each $i \in \mathbb{N}_0$, slightly abusing the notation, we have that $(q_i,\nu_i) \xrightarrow{c_i} (q_{i+1},\nu_{i+1})$ which holds under the following conditions:

- if c_i is a time stamp $t \in \mathbb{R}_{\geq 0}$, then we require that $q_{i+1} = q_i$ and $\nu_{i+1} = \nu_i + t$,

- if c_i is an input letter $a \in \Sigma$, there is a unique edge (q_i,a,g,X,q) such that $\nu_i \models g$, and we require that $q_{i+1} = q$ and $\nu_{i+1} = \nu_i[X := \mathbf{0}]$.[4]

Lastly, we define the standard *region relation* $\sim$. It partitions the configurations into finitely many equivalence classes called *regions*. For $a,b \in \mathbb{R}_{\geq 0}$, we say that a and b *agree on integral part* if $\lfloor a \rfloor = \lfloor b \rfloor$ and neither or both a, b are integers. Let B be the maximal constant in the guards of $\mathscr{A}$. We put $(q,\nu) \sim (q',\nu')$ if $q = q'$ and

- for all $x \in \mathscr{X}$, $\nu(x)$ and $\nu'(x)$ agree on integral part or are both $> B$,

- for all $x,y \in \mathscr{X}$ lower than B in ν, $\langle \nu(x) \rangle \leq \langle \nu(y) \rangle$ iff $\langle \nu'(x) \rangle \leq \langle \nu'(y) \rangle$.

Definition 2.2.6 (DTA observer)**.** *Let* $\mathscr{M}$ *be a GSMP with a set of states S. We say that a DTA* $\mathscr{A}$ *is an* observer *of* $\mathscr{M}$ *if its alphabet is S.*[5] *We define a random variable W that assigns to each run* $\sigma = (s_0,\xi_0,t_0)(s_1,\xi_1,t_1)\cdots$ *of* $\mathscr{M}$ *its induced timed word* $W(\sigma) = s_0\,t_1\,s_1\,t_2\cdots$*. Finally, the* observation *of* $\mathscr{A}$ *over* $W(\sigma)$ *is the computation of* $\mathscr{A}$*, i.e.* $\mathscr{A}(W(\sigma))$*.*

Now, let us discuss several forms of specifications based on DTA observers. We start with basic ones.

[4] Sometimes, we use the relation for finite words $w = c_0\cdots c_n$. We write $(q_0,\nu_0) \xrightarrow{w} (q_{n+1},\nu_{n+1})$ if there are configurations $(q_1,\nu_1),\ldots,(q_n,\nu_n)$ such that $(q_i,\nu_i) \xrightarrow{c_i} (q_{i+1},\nu_{i+1})$ for all $0 \leq i \leq n$.

[5] Another option is to have a labelling function for the states of $\mathscr{M}$ and define the timed automaton over the labels. For the sake of simplicity, we define it directly this way.

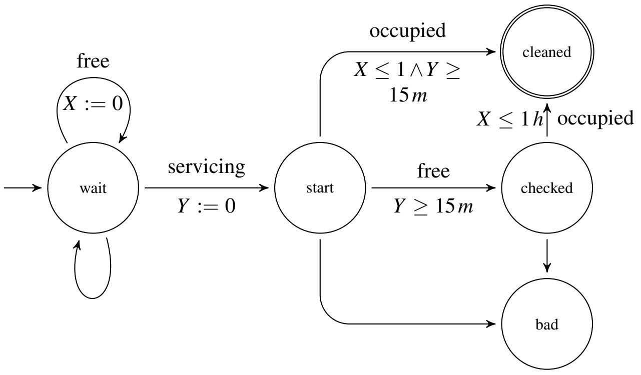

Figure 2.4: Thorough servicing. An example of a reachability specification in a DTA observer. Notice that edges are labelled with states of the servicing GSMP in Figure 2.2. The edges without labels represent the remaining behaviour (remaining input letters and clock valuations not covered by the other edges).

Definition 2.2.7 (Reachability and Büchi properties in the DTA observer)**.** *Let us fix a GSMP $\mathscr{M}$, a DTA observer $\mathscr{A}$, and a set of target locations $T \subseteq Q$. The* reachability specification *over T is the set of runs* $\text{Reach}_{\mathscr{A}}(T) = \{\mathscr{A}(W) \text{ visits } T\}$*, i.e. runs for which the observer visits a location in T. Furthermore, the* Büchi specification *over T is the set of runs* $\text{Büchi}_{\mathscr{A}}(T) = \{\mathscr{A}(W) \text{ visits } T \text{ infinitely often}\}$.

Notice that reachability in the observer TA can easily express reachability and time-bounded reachability in the original system as well as various more elaborate properties. Let us illustrate the concept by an example.

Example 2.2.8. *Figure 2.4 shows an example of a reachability DTA specification. It specifies that a servicing takes at least quarter of an hour and around the servicing there is at most one hour without a customer. Observe that this property cannot be specified by a formula in the logic CSL because of the simultaneous testing of two deadlines.*

A DTA observer is also suitable for expressing more complex performance measures. We basically transfer the discrete and timed frequencies from the run of the stochastic system to the computation of the timed automaton. The definition is analogous.

Definition 2.2.9 (Frequencies in the observer)**.** *Let $\mathscr{M}$ be a GSMP, $\mathscr{A}$ its DTA observer, and q a location of $\mathscr{A}$. We define the* discrete frequency *and* timed frequency

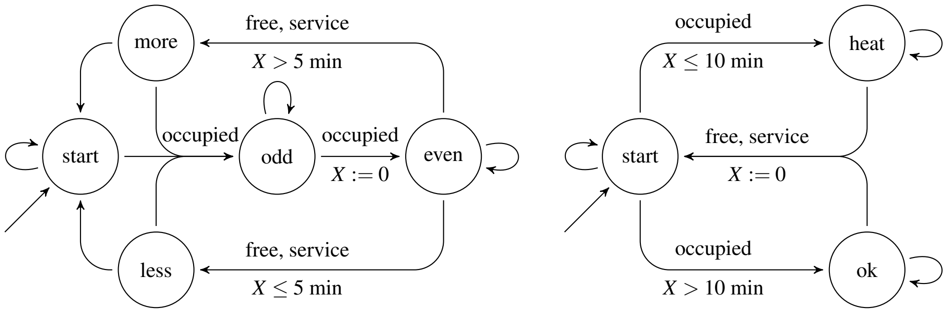

Figure 2.5: Performance of the serviced photo booth, an example of DTA performance measures. The automaton on the left measures *out of the even customers, the ratio of those who spend in the booth more than 5 minutes*. It corresponds to the number $d_{\text{more}}/(d_{\text{more}} + d_{\text{less}})$. The automaton on the right measures the *ratio of time the printer in the booth is overheated* (for this measure we assume that the printer remains overheated if there are two customers in the booth closer than 10 minutes to each other). It corresponds to the number c_{heat}.

of q in a run σ with $\mathscr{A}(W(\sigma)) = (q_0, v_0)\, s_0\, (q_1, v_1)\, t_1\, (q'_1, v'_1)\, s_1\, (q_2, v_2) \cdots$ by

$$\begin{aligned} \mathbf{d}^{\mathscr{A}}_q(\sigma) &= \lim_{n \to \infty} \frac{\sum_{i=0}^{n} 1_{q_i = q}}{n} \\ \mathbf{c}^{\mathscr{A}}_q(\sigma) &= \lim_{n \to \infty} \frac{\sum_{i=0}^{n} t_i \cdot 1_{q_i = q}}{\sum_{i=1}^{n} t_i} \end{aligned}$$

where, likewise, $1_{q_i=q}$ equals 1 if $q_i = q$ and $1_{q_i=q}$ equals 0, otherwise.

Example 2.2.10. *Figure 2.5 demonstrates that these specifications allow measuring performance w.r.t. behaviour that satisfies complicated time or structural constraints. For example,* out of the even customers, what is the ratio of those that spend in the booth more than 5 minutes? *The time constraints stem from the clocks and guards of the timed automaton, the structural constraints stem from its transition structure of a finite automaton. None of these aspects are possible to express within the classical* reward structures *or* delays*. Only some of these aspects are expressible within other performance frameworks.*

Chapter 3

State of the Art

In this chapter, we review the state of the art related to hard real-time bounds in discrete-event systems and how the related work has been impacted by our results.

3.1 Related modelling formalism

As many related papers study different formalisms describing similar concepts, we introduce these formalisms briefly in this section. The main dichotomy in the literature stems from the stochastic Petri nets, an area extensively studied in the field of performance evaluation. Stochastic Petri nets have arisen from queuing networks (apart from standard Petri nets). One of the reasons for the performance evaluation community to study stochastic Petri nets is its appeal to people in practice; the models are easy to understand without any knowledge of statistics, automata theory, or formal methods.

3.1.1 Stochastic Petri Nets

Stochastic Petri nets are based on standard Petri nets. Let us briefly recall the Petri nets formalism. It consists of a finite set of places and a finite set of transitions. For each transition t, there are sets of input places, inhibitor places, and output places. *Marking*, the current configuration of the net, is a function that assigns to each place a non-negative number of *tokens*. A transition is said to be *enabled* in a marking m if each its input place has at least one token in the marking m and each its inhibitor place has no token in the marking m. When a transition is *fired*, one token is removed from each its input place and one token is added into each its output place.

The net starts in the initial marking and then changes its marking by firing enabled transitions. If there are multiple transitions enabled at a time, one of them is chosen non-deterministically. A marking m is *reachable* if there is a sequence of markings $m_0 \cdots m_n$ where m_0 is the initial marking, $m_n = m$, and for each $i \in \mathbb{N}_0$, m_{i+1} is obtained from m_i by firing one of the transitions enabled in m_i. A Petri net is called *bounded* if the number of reachable markings is finite.

Stochastic Petri net is a Petri net where the transitions are fired randomly in continuous time and is usually restricted to be bounded. We define a general class of stochastic Petri nets that allows us to easily identify the studied subclasses.

Definition 3.1.1. *A* Stochastic Petri Net (SPN) *is a tuple* [1]

$$\mathscr{N} = (P, T, T', I, H, O, F, \mathbf{p}, m_0) \quad \textit{where}$$

- *P is a finite set of places, T is a finite set of transitions, and $T' \subseteq T$ are the immediate transitions (the transitions $T \setminus T'$ are called* timed*);*
- *$I(t)$, $H(t)$, and $O(t)$ are the sets of input places, inhibitor input places, and output places for each $t \in T$;*
- *$F(\cdot, t)$ is a clock-setting distribution function for each transition $t \in T \setminus T'$;*
- *$\mathbf{p}(\cdot, E)$ is a distribution on $\{-1, 0, 1\}^{|P|}$ for each set $E \subseteq T$ of transitions that fire concurrently;*
- *m_0 is the initial marking.*

A *configuration* of a SPN in a triple (m, τ, x) where m is a marking, τ assigns to each transition its current time to fire, and x is the time spent in the previous marking. A transition t is *firable* in a configuration if it is enabled and its time to fire $\tau(t)$ is not higher than the time to fire of any other enabled transition. In a configuration (m, τ, x), the next configuration (m', τ', x') is obtained by firing the set of firable transition E as follows.

- A vector $v \in \{-1, 0, 1\}^{|P|}$ is chosen randomly according to $\mathbf{p}(\cdot, E)$ and we set $m' = m + v$.
- Let t' be a transition enabled in m'. We say that t' is *newly enabled* if t' was not enabled in m or $t' \in E$. If t' is immediate, we set $\tau'(t')$ to 0. If t' is timed and newly enabled, $\tau'(t')$ is randomly sampled according to $F(\cdot, t')$. If t' is timed and not newly enabled, we set $\tau'(t') = \tau(t') - \delta$ where $\delta = \tau(t)$ for a $t \in E$.
- The time x' equals $\tau(t)$ for a $t \in E$.

This stochastic extension resolves all the non-determinism that was originally present in a Petri net. Hence, $\mathscr{N}$ induces a jump marking process $(\Xi_n)_{n \in \mathbb{N}_0}$ where each Ξ_i is the configuration of $\mathscr{N}$ after i firings.

Example 3.1.2. *In Figure 3.1, there is a stochastic Petri net corresponding to the GSMP model from Figure 2.3. On the left there is the marking at time* 0 *and on the right there is the marking that is reached with probability* 1 *at time* 300.

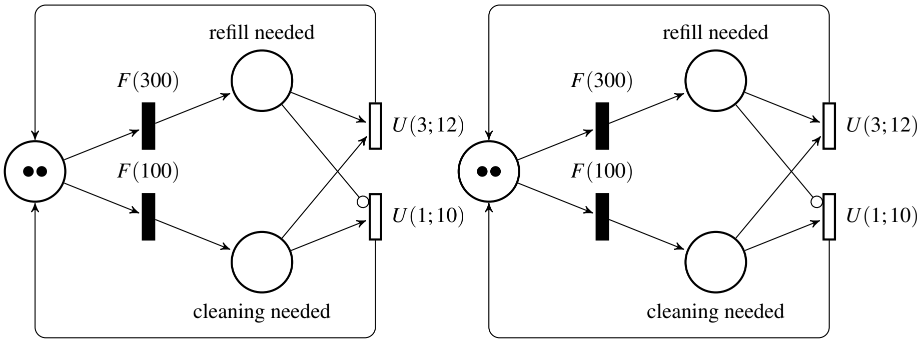

Figure 3.1: Servicing of a photo booth, an example of a stochastic Petri net. Places are drawn as circles, fixed-delay transitions as thick bars, and variable-delay transitions as empty bars. Input and output places of a transition are marked with normal arrows, inhibitor places with a circled arrow. When both the fixed-delay transitions occur at once, both of them move the tokens concurrently.

Out of the variety of classes of SPN defined in the literature, there are several classes relevant to the topic of the thesis:

- *Generalized stochastic Petri nets (GSPN)*[ACB84] restrict $F(\cdot,t)$ to be exponential for each timed transition t. This class closely corresponds to the continuous-time Markov chains.
- *Deterministic and stochastic Petri nets (DSPN)*[MC87], restricts $F(\cdot,t)$ for each timed transition t to be exponential or *deterministic*. A distribution function $F(\cdot,t)$ is called deterministic if there is a $b \in \mathbb{N}$ such that $F(x,t) = 0$ for each $x < b$ and $F(x,t) = 1$ for each $x \geq b$. This class roughly corresponds to GSMP with fixed-delay events where each variable-delay event is exponentially distributed. It is also similar to recently introduced delayed CTMC [GGH+12] from the area of computational biology.
- *Stochastic timed Petri nets (sTPN)*[CGV09] restricts $F(\cdot,t)$ for each timed transition t to be either deterministic, or to have positive density on an interval $[\ell_t, u_t)$ for $\ell_t \in \mathbb{N}_0$ and $u_t \in \mathbb{N} \cup \{\infty\}$. This class is even closer to our definition of GSMP with fixed-delay events.

Modelling power of SPN and GSMP Now we restate in our notation the result of [Haa10] that SPN and GSMP have the same modelling power. As our definition

[1] There are numerous different definitions; ours is based on [Haa10]

of GSMP is more restricted compared to [Haa10], we also restrict in the theorem below to the subclass of sTPN.

Definition 3.1.3 (Simulation of SPN and GSMP)**.** *We say that*

- *a GSMP $\mathscr{M}$ with its associated stochastic process Φ* simulates *a SPN with the marking process Ξ, if there is a mapping ϕ from the configurations of the GSMP to the configurations of the SPN such that Ξ_n and $\phi(\Phi_n)$ have the same distributions for each $n \in \mathbb{N}_0$;*
- *a SPN with the marking process Ξ* simulates *a GSMP $\mathscr{M}$ with its associated stochastic process Φ if there is a mapping ξ from the configurations of the SPN to the configurations of the GSMP such that Φ_n and $\xi(\Xi_n)$ have the same distributions for each $n \in \mathbb{N}_0$.*

Theorem 3.1.4. *For any sTPN $\mathscr{N}$, there is a GSMP with fixed-delay events that simulates $\mathscr{N}$. Furthermore, for any GSMP (with fixed-delay events) $\mathscr{M}$, there is a sTPN that simulates $\mathscr{M}$.*

The result follows from [Haa10, Theorems 3.4 and 4.6]. Note that [Haa10] defines the configurations of the underlying general state space Markov chain of a GSMP to store the *remaining* time to occurrence of events instead of the *elapsed* time. However, both the ways are used interchangeably in the literature and their equivalence is easy to show. Let us now briefly review other related modelling formalisms.

3.1.2 Stochastic extensions of timed automata

An alternative approach to DTA observers is to combine DES and TA within one model. One such formalism is the extension of probabilistic timed automata (PTA), the *continuous PTA* [KNS+00]. In a PTA, the discrete transitions when reading letters are probabilistic. The continuous randomness is then added by the means of clock resets: clocks are not reset to 0 but to a random point chosen according to a continuous distribution with bounded density. First difference to our approach is that the continuous PTA does not allow unbounded continuous events. Second, there is still a non-determinism in the flow of time in the continuous PTA similar to the non-determinism of TA. For this model, verification of a logic PTCTL is addressed by discretization.

A different approach is taken in [BBB+07; BBB+08a; BBB+08b; BBJ+12]. A stochastic semantics of TA is defined by assigning to every configuration (q,ν) a measure $\mu_{(q,\nu)}$ on the delay of the next step. The non-determinism among edges is also resolved randomly according to weights assigned to edges. This definition yields a broad class of models, various restrictions are assumed for practical solutions. *Qualitative* model-checking algorithms are given in [BBB+08a] for one-clock stochastic timed automata and in [BBJ+12] for *reactive* stochastic timed automata with arbitrarily many clocks. This class of reactive models roughly corresponds to GSMP observed by DTA where all events of the GSMP have variable delay supported on the whole $\mathbb{R}_{\geq 0}$. Note that it is a subclass of GSMG that we deal with in Chapter 6 and the algorithm of [BBJ+12] is very similar to our algorithm from [BKK+10a]. However, their proof technique allows a wider class of specifications than considered in Chapter 6. A *quantitative* model checking algorithm is given in [BBB+08b] for one-clock *reactive* stochastic timed automata where all delays are exponentially distributed. Note that the problem is similar to the analysis of CTMC observed by one-clock DTA [CHK+09].

This general definition of stochastic timed automata is further extended with non-deterministic flow of time in the 2-player game model of *stochastic timed games* [BF09]. Due to its expressiveness, the quantitative reachability is shown to be undecidable. The only fragment that is shown to be decidable is the qualitative reachability in a one-player game with only one event.

Stochastic automata [DKB97] is another formalism introducing stochastic flow of time into timed automata. It however rather corresponds to GSMP (with fixed delay events) as it does not allow the non-deterministic flow of time. Compared to GSMP, this model is better suited for process algebraic extensions (see below) as it still contains the non-determinism when taking discrete transitions. Most of the research in this area focuses on process algebraic questions directed at efficient compositional modelling. The analysis is performed by simulation or by automatically abstracting the models to simpler formalisms where the standard techniques can be applied [BdH+06]. None of this related work considers stability of systems combining DES with TA.

3.1.3 Process algebras and game extensions of DES

A lot of attention has been addressed to various game extensions of CTMC. Some previous literature also deals with (one-player) game extension of GSMP. One-player non-determinism also arises from the interleaving of synchronization in *process algebras* defined over DES. Process algebras are languages for compositional modelling of parallel systems. Using a process algebra, models can be built bottom-up; larger components are obtained by parallel composition of multiple smaller components that synchronize using message passing.

CTMC The game extension of CTMC are covered in closer detail by the recent Ph.D. thesis [Kre13]. Various authors considered one player game extension, the *continuous-time Markov decision processes* [BHK+05; NSK09; BS11; BHH+11]. In this formalism, an action, chosen by a scheduler in each state of the system, influences the rates of events that are awaited. A similar model of *interactive Markov chains* [Her02; ZN10; HH13] builds upon CTMC a process algebra in order to allow compositional modelling. In this formalism, instead of actions there are internal transitions in some states that are taken in zero time. Another process algebra stemming from CTMC is PEPA [Hil96].

Two-player turn-based games over CTMC, the *continuous-time stochastic games* were defined in [BFK+09; BFK+13] and further studied by [RS11; RS13]. In this setting, the set of states is partitioned among two players where the action is always chosen by the player in charge. Recently a limited form of concurrency appeared in [BHK+12; HKK13] where two-player games are studied as a solution technique to compositional verification of interactive Markov chains.

The work in this area addresses synthesis of optimal (or ε-optimal) schedulers and computation of guarantees provided by such schedulers. The specification is usually the time-bounded reachability or a formula in the logic CSL (see below).

GSMP The *generalized semi-Markov decision processes* were introduced in the community of applied statistics by [Dos79] and in the area of probabilistic verification by [YS04]. Another approach is discussed in [GY94] where the current rates of events are subject to continuous control. The scheduler can thus within some bounds continuously choose the distribution functions of events. A process algebraic extension of GSMP, called *interactive generalized semi-Markov process*, was introduced by [BG02]. Similarly to the research on stochastic automata [DKB97], this research focuses more on the process-algebraic questions than on the analysis. A one-player non-determinism is also present in the continuous PTA [KNS+00] discussed above.

To the best of our knowledge, the only two-player game model with stochastic non-Markovian flow of time are the *stochastic timed games* [BF09]. As discussed above, this model however combines the stochastic flow of time with the non-deterministic flow of time. Furthermore, the distributions of clocks' delays are configuration-dependent and only a very restricted problem is shown to be decidable for this model.

3.2 Related specification formalisms

Let us fix a discrete-event system with a state space S. A standard property of the system to be specified and analysed is the *transient distribution* $\pi(t)$ at time t. It is the discrete distribution of a random variable that to each run of the system assigns the state at time t. In the following text, we focus on properties related to long-run analysis of DES.

Steady-state analysis The *stationary distribution* π is the limit $\lim_{t\to\infty}\pi(t)$ of the transient distributions. This limit always exists for finite CTMC [Kul95] and coincides with the timed frequency $\mathbf{c}$. However, for semi-Markov processes and hence also for generalized semi-Markov processes, this limit may not exist [LHK01]. Instead, the timed frequency is studied [LHK01; Alf97]. The notion *steady-state* distribution is loosely used to denote both the stationary distribution for CTMC and the timed frequency for non-Markovian models.

Logics The most prominent logic for the stochastic continuous-time systems is the *continuous stochastic logic (CSL)* introduced by Aziz et al. [ASS+00] and later extended by Baier and Haverkort [BH03]. It is a branching-time temporal logic that also allows specifying long-run properties. The syntax is as follows:

$$\varphi ::= \mathtt{tt} \mid a \mid \neg\varphi \mid \varphi\wedge\varphi \mid P_{\trianglelefteq p}(X_I\varphi) \mid P_{\trianglelefteq p}(\varphi U_I \varphi) \mid S_{\trianglelefteq p}(\varphi)$$

where a is an atomic proposition, $\trianglelefteq \in \{\leq,<,>,\geq\}$, $p\in[0,1]$ is a real number, and $I\subseteq\mathbb{R}_{>0}$ is a non-empty interval. For the sake of simplicity we assume here that the set of atomic proposition coincides with the set of states of the model.

Let us first define the semantics for CTMC. For a state s we set $s\models tt$ and $s\models a_s$ where a_s is the atomic proposition of the state s. The negation and conjunction is defined in the expected way. Further, we set $s\models P_{\trianglelefteq p}(X_I\varphi)$ if the set of runs $A=\{s_0t_0s_1t_1\cdots \mid t_0\in I, s_1\models\varphi\}$ has probability $\mathscr{P}_s(A)\trianglelefteq p$ when starting from s. Similarly, $s\models P_{\trianglelefteq p}(\varphi U_I\psi)$ if the set of runs $B=\{s_0t_0s_1t_1\cdots \mid \exists i: t_0+\cdots+t_{i-1}\in I, s_i\models\psi, \forall j<i: s_j\models\varphi\}$ satisfies $\mathscr{P}_s(B)\trianglelefteq p$. Finally, we set $s\models S_{\trianglelefteq p}(\varphi)$ if $\sum_{s'\models\varphi}\pi_s(s')\trianglelefteq p$ where π_s is the stationary distribution when the chain starts in s.

Example 3.2.1. *We illustrate the logic on several formulae.*

- $P_{\geq 0.9}(\mathtt{tt}\,U_{[0,100]}\,(a_{\text{servicing}}\wedge P_{\geq 0.4}(a_{\text{servicing}}\,U_{[0,1]}\,a_{\text{free}})))$ *specifies that* with probability at least 0.9 the next servicing of the photo booth will start in the following 100 hours and when it starts, with probability at least 0.4 the booth will be free again within one hour.
- $S_{\geq 0.1}(\text{occupied})$ *specifies that* the photo booth is occupied at least 10% of time in the long-run.
- $S_{\geq 0.1}(a_{\text{occupied}}\wedge P_{\geq 0.9}(\mathtt{tt}\,U_{[0,0.1]}\,a_{\text{free}}))$ *specifies that* the photo booth is at least 10% of time in such a state that a person is inside the booth and the person leaves the booth in up to 6 minutes with probability at least 0.9.

The semantics has been similarly defined for semi-Markov processes [LHK01]; instead of the stationary distribution the timed frequency is used for the semantics of the steady-state operator S. Notice that for GSMP the semantics needs to be defined on configurations instead of states [YS02] because the satisfaction of a formula when entering a state depends also on the elapsed times of the events. Apart from that it goes along the same lines.

The logic CSL is similar to the logic PTCTL [KNS+00]. In the recent years, several follow-up logics appeared, such as asCSL [BCH+07] and CSL^{TA} [DHS07]. They differ from CSL in the way they specify the path restrictions. The logic asCSL defines labelling for both states and events and specifies paths by regular expressions over the language of pairs of state and event labels. The logic CSL^{TA} specifies paths by deterministic one-clock timed automata. A similar approach focusing rather on the structure are the *experiments* proposed by de Alfaro [Alf98] for the long-run average behaviour analysis. Experiments are basically trees of the possible behaviour with reward in leaves. A set of variables is defined. Each variable has a fixed value in each state of the discrete-event system. Every node in the tree then contains a constraint over variables. The experiment tree is traversed so that the constraints are kept satisfied. The experiments can be used to measure the long-run average reward or the long-run average time of traversal from the root to the leaves.

3.3 Stability of discrete-event systems

Research on stability of GSMP mainly addresses the existence of *invariant measure*. A probability measure π on the measurable space $(\Gamma, \mathscr{F})$ of configurations of the GSMP is called *invariant* if for all $A \in \mathscr{F}$ it holds $\pi(A) = \int_{\Gamma} \pi(dx)P(x,A)$, i.e. the measure does not change when a step is taken. As we discuss later in Chapter 5, the existence of invariant measure implies that the frequency measures **d** and **c** are almost-surely well-defined. Hence, this research is closely related to our topic.

Work from early 80's [Whi80a; Gly83] shows that a GSMP has a unique invariant measure if all events have continuous densities and finite moments. This is further extended by [HG02; Haa10] by proving the existence of the invariant measure for GSMP where the events have continuous part with positive density on $[0,x]$ for some $x > 0$ (again excluding the fixed-delay events). A different approach [GY94] proves stability in GSMP with arbitrary distributions that satisfy strict structural *monotonicity* conditions. The require the GSMP to be (1) *non-interruptive*, i.e. no scheduled event may get switched off, and (2) *permutable*, i.e. if a sequence of events can occur in two different permuted orders, the set of events scheduled in the respective target states coincide. Stability is also implied by the strong notion of *insensitivity* that has been extensively studied on GSMP, see for example [BS81; Tay89; CT92]. A DES is called insensitive if its steady-state distribution depends on the events' distributions only via their means. None of these papers targets GSMP with fixed-delay events. Another related result [HG01] studies the regenerative structure of GSMP for the sake of efficient simulation but only minor results are relevant to GSMP with fixed-delay events. Note that *regeneration* roughly speaking means visiting a configuration with all elapsed times ξ equal to zero.

The regenerative behaviour is further stressed in *regenerative GSMP* [HS87] or *Markov regenerative stochastic Petri nets* [CKT94] pointing out regenerative subclasses that are stable but not discussing the stability outside their subclass. Furthermore the subclass is quite restrictive: at most one non-Markovian event enabled in each state. Most of the work in this area (including ours) bases its proofs on the abstract treatment of stability for Markov chains with general state space [MT09].

3.4 Quantitative solution methods

The main result of this thesis, the study of stability of GSMP with fixed-delay events addresses *qualitative* questions. There are various solution methods for *quantitative* analysis of GSMP (or SPN). In this section, we review the various proposed methods and put them into perspective with our instability result.

3.4.1 Analytical solutions

One approach is to formulate the solution precisely using a mathematical expression that preferably admits an efficient numerical solution. Due to the rigorous nature of this approach, these algorithms provide a correct answer. As discussed below, they either a priori restrict to a stable subclass of models or do not terminate when applied to an unstable model.

Uniformization of CTMC We start with *uniformization* [Jen53], the key technique for analysis of CTMC as further methods for more complicated formalisms build upon this technique. Observe that the transient probability of state s at time t can be expressed as

$$\pi(t)(s) = \sum_{i=1}^{\infty} \sum_{s_1 \cdots s_i \in S^{i-1}s} D(s_1 \cdots s_i) \cdot T_t(s_1 \cdots s_i) \tag{3.1}$$

where $D(s_1 \cdots s_i) = \alpha_0(s_1) \prod_{j=1}^{i-1} P(s_j, s_{j+1})$ is the time-abstract probability of taking path $s_1 \cdots s_i$ with $P(s_j, s_{j+1}) = Q(s_j, s_{j+1}) / -Q(s_j, s_j)$ and $T_t(s_1 \cdots s_i)$ is the probability that at time t this path is traversed and no further step is made. In a CTMC, the term T_t depends on a sum of exponentially distributed random variables each having a different rate. The probability that such a sum is lower that t

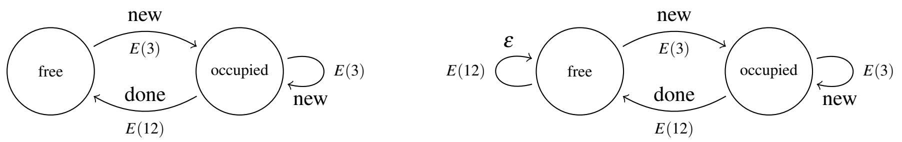

Figure 3.2: Uniformization. The original CTMC is on the left. By adding self loops we change the exit rate of all states to be the same. The resulting CTMC on the right is equivalent to the original one up to *stuttering*, i.e. taking self-loop transitions. Note that stuttering does not alter the transient distributions. The matrix P is then the transition matrix of the *embedded* Markov chain of the CTMC on the right. An embedded Markov chain of a DES is a discrete time process such that its discrete steps correspond to the state changes in the continuous-time DES.

has a closed form expression [AM97] but the computation of (3.1) [BFK+13] is inefficient in practice [BHH+11]. By creating a discrete-time Markov chain with transition matrix $P = Q/\mu + I$ where $\mu = \max_{s \in S} -Q(s,s)$ is the maximal exit rate in the system, we get

$$\pi(t)(s) = \sum_{i=1}^{\infty} (\alpha_0 P^i)(s) \cdot Poiss(i, \mu t)$$

where $Poiss(i, \mu t)$ is the probability that a Poisson-distributed random variable with parameter μt equals to i. This way, the discrete-time and continuous-time behaviour is totally decomposed. The method is illustrated in Figure 3.2.

This method has many uses in the area of probabilistic verification. By making a target state s absorbing, the transient probability $\pi_t(s)$ expresses the probability that s is reached *within* time t. Similarly, by making all states s with $s \models \neg\varphi \vee \psi$ absorbing, the probability $\pi_t(s)$ can be used to decide whether the CSL formula $P_{\trianglelefteq p}(\varphi U^{[0,t]} \psi)$ is satisfied. Similar approach can be taken for other forms of the interval in the operator U yielding the CSL logic decidable [ASS+00] and efficiently (approximately) verifiable in practice [BH03; KNP11; KZH+09].

Method of subordinated Markov chains for DSPN Recall that DSPN is a class of stochastic Petri nets with exponential or deterministic distribution on the firing times of the transitions. In this class the nets are restricted so that always at most one deterministic transition is enabled at a time. This restriction was imposed when the class was introduced [MC87]. The method of *subordinated Markov chains* applies uniformization to the steady-state analysis of DSPN. In this method, a discrete-time Markov chain is build such that its states correspond to a subset M' of markings of the Petri net where either (a) no deterministic transition is enabled or (b) a deterministic transition newly became enabled.

- The transition probabilities from a state s of type (a) are as in the embedded chain as there is no deterministic transition.

- The transition probabilities from a state s of type (b) are computed as follows. Let $m \in M'$ be the marking corresponding to s. Either the deterministic transition fires after its delay d (after firing some exponential transitions) and we end up in some $m' \in M'$; or this transition gets disabled before time d and we again reach some marking $m' \in M'$. The subordinated Markov chain is a CTMC obtained from the marking process of the Petri net by ignoring the deterministic transitions. The CTMC starts in m and all other markings from M' are made absorbing. The transition probabilities from s are finally simply obtained from the transient distribution π_d.

Finally, the steady-state distribution is computed from the steady-state distribution of the DTMC, roughly speaking, by weighting it by the expected times one spends in the states of the DTMC. This method has been later extended to the transient analysis [CKT93]. The idea of subordinated Markov chains is closely related to the method used for quantitative reachability in CTMC observed by one-clock timed automaton [CHK+09], in verifying the logic CSL^{TA} [DHS07], or for analysing DTA properties of population models [BL13].

Symbolical integration There is another analytical approach for models where the events have either fixed delay or have *expolynomial* distribution, i.e. its density is piecewise defined by expolynomials of the form $\sum_{i=0}^{n} \sum_{j=0}^{m} a_{ij} x^i e^{-\lambda_{ij} x}$ with each $a_{ij} \in \mathbb{R}$ and each $\lambda_{ij} \in \mathbb{R}_{\geq 0}$.

First, [CGL94] showed for semi-Markov models that transition probabilities of the embedded DTMC can be computed by symbolical integration instead of computing them numerically [CKT94; LHK01]. This symbolical approach was implemented in the tool TimeNET [ZFG+00].

Later, the symbolical integration of expolynomials was extended [SV07] to generalized semi-Markov models using the concept of zone graph [Dil90] from the theory of timed automata. By endowing the zones with densities over events' delays [BPS+05], the zone graph then actually corresponds to the embedded Markov chain and allows efficient steady-state analysis. A similar approach to the transient analysis of π_t was proposed by [AB06] but it only restricted to at most k occurrences of events up to time t. Later, the unrestricted transient analysis was addressed [HRV10b; HPR+12] by introducing into the zone graph an artificial event that occurs at time t. Both steady-state and transient analysis using the zone graph are implemented in a tool Oris [BCR+10].

Note that conceptually, the symbolical integration extends the algorithm for transient analysis of uniformized models to the wider class of expolynomial distributions. Similar such extensions of Jensen's method were used in various other papers, such as [GL94; BA07; BFK+13].

Remark 3.4.1. *Notice that the method of symbolical integration is the first method that we mention that allows for multiple concurrent fixed-delay events. If the algorithm terminates, it is guaranteed to return a correct answer. However, for some (non-regenerative) models, the zone graph may be infinite and hence, the algorithm does not terminate. Later, approximation of the densities using Berstein expolynomials has been introduced [HRV10a] without any proof of correctness, that can thus provide an incorrect answer.*

3.4.2 Approximation techniques

Similarly to Remark 3.4.1, there are many methods that promise to approximate the steady-state or transient distributions. However, rigorous proofs of correctness for these methods are rare. Thus, for models with multiple fixed-delay events, one can compute a distribution [Lin93; GL94; LS96; LRT99; BPS+98; HTT00; ZFG+00; ZFH01; Hor02; SDP03; HMM05; CGV09] that should approximate the timed frequency. In fact, the timed frequency may not exist and the notion "approximation" may make no sense. In the following, we sketch the various methods and comment how they cope with their correctness.

Method of supplementary variables We start with the simple model of DSPN extended to support multiple concurrent deterministic transitions where the method of *supplementary variables* [Cox55b] was proposed [GL94]. In the initial work, the state space is extended with real-valued variables of remaining time to fire of the events.[2] Thus the continuous-time is transformed into continuous-space yielding a discrete-time Markov chain with general state space [MT09]. They then express the steady-state distribution using a system of PDE that can be solved by discretization.

A follow-up paper [LS96] instead proposes an explicit discretization by time steps Δ. The steps of the defined general state space Markov chain hence do not correspond to occurrences of events but to lapse of time Δ. Again, by extending the method of Jensen, the (infinite) transition kernel of the state space with supplementary variables can be analytically characterized. However, a numerical solution is

[2] Observe that in our formal treatment we build on this method, even though we encode into the state space the elapsed times of events instead.

applied (solving the Voltera integral equations of the second type). This method was later extended to the transient analysis [LT99] and implemented in a tool DSPN-Express 2.0 [LRT99]. None of the approaches discusses the existence of the timed frequency and as such may provide incorrect answers.

Approximation by a DSPN with single deterministic transition A completely different approach to solving DSPN with multiple concurrent deterministic transitions has been proposed ca. 10 years later [HMM06]. One new deterministic transition t_Δ is added such that its delay Δ divides the delay of any other deterministic transition. This transition fires repeatedly like ticking of a clock and simulates the remaining transition. As the deterministic transition can become enabled at any time but fire only when t_Δ fires, their delay is only approximate. Again, the correctness of this approximation is not proven (and in fact, it may be incorrect due to Theorems 4.2.1 and 5.1.2). Later, this method was independently reinvented by [GGH+12] for the analysis of delayed CTMC in the area of computational biology.

Phase-type approximation The nice analytical properties of the CTMCs can be exploited for the GSMPs by approximating them using a CTMC. Each arbitrarily distributed event is replaced with a similar event with the *phase-type (PH) distribution* [Cox55a]. F is a PH distribution if there exists a CTMC $\mathscr{C}$ with designated absorbing target state s such that for all $x \in \mathbb{R}_{\geq 0}$ we have $F(x) = \pi_x(s)$. Here, π_x is the transient distribution of $\mathscr{C}$ at time x. The overall CTMC that approximates the GSMP can be built by composing the CTMC gadgets $\mathscr{C}_e$ for each non-exponential event e, see Figure 3.3. Note that any continuous distribution can be approximated up to arbitrary precision by a PH distribution [Neu81]. However the impact on the transient or steady-state distributions has not been rigorously studied; not even in the context of probabilistic verification of GSMDP [YS04]. Furthermore, for models with fixed-delay events, this approach may be incorrect as the resulting CTMC always has the steady-state distribution unlike the original GSMP.

Most of the research on PH approximation of non-Markovian models has been conducted in the area of stochastic Petri nets. The continuous PH approximation was introduced here by [BC84]. Later a different approach of discrete PH approximation was studied [Mol85]. The concept is the same, a discrete-time Markov chain is created such that its transient and steady-state distribution approximates the distributions in the original model. Each step of the DTMC now corresponds to lapse of time Δ in the original model for some appropriate $\Delta > 0$. Observe, that the potential instability of the model is again not preserved by this approximation as every finite-space DTMC almost-surely attains the discrete frequencies. A

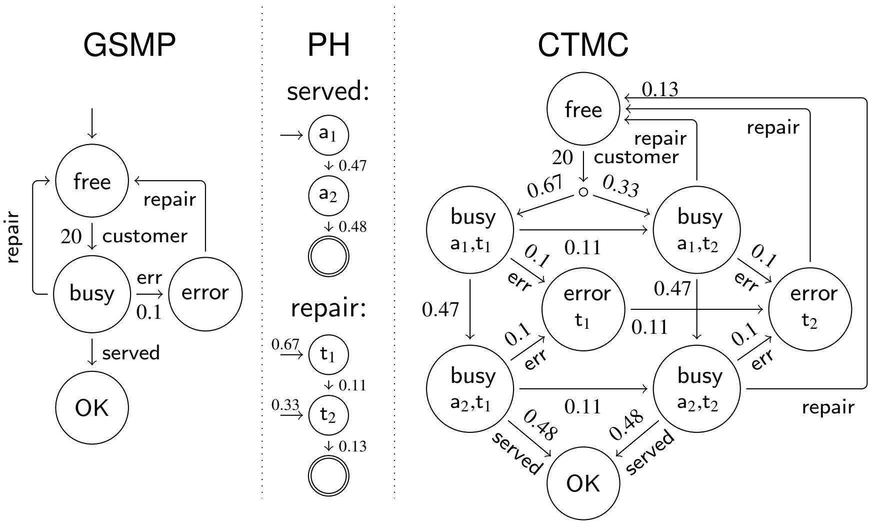

Figure 3.3: Phase type approximation. On the left, there is a GSMP model of one customer being served in a photo-booth. A repair comes after a fixed delay of 10 days, customer comes with an exponential distribution with rate 20, err is exponential with rate 0.1 corresponding to an error occurring each 10 days of the use of the booth on average, and served is a non-exponential variable-delay distribution. In the middle, there are 3-state PH approximations of events served and repair. On the right, there is the CTMC PH approximation obtained roughly speaking as a product of the GSMP and the two PH gadgets.

hybrid approach combining the discrete and continuous PH has been also considered [JC01] which creates an approximating model similar to DSPN. This method preserves the (in)stability of the model, which may however get lost in the further solution method.

As regards the efficiency of this method, observe that the size of the resulting model is exponential in the sizes of the gadgets and in the parallelism of the model. This method therefore performs better for models where the events are easy to fit (such as various continuous distributions supported on the whole $\mathbb{R}_{\geq 0}$) and the parallelism is in some sense local [HMC97]. Furthermore, expressing the product state space using the Kronecker algebra [Neu81; SB98] permits the models to be analysed without explicitly building the whole rate or transition matrix.

Discretization Discretization as a solution method is prevalent in the transient analysis of the game extensions of discrete-event systems, mainly in continuous-time Markov decision processes (CTMDP). In this solution method, the process is transformed into a discrete-time Markov decision process that “models” the continuous-time process at times that are multiples of the discretization step Δ. An assumption that at most one event occurs within each interval of size Δ is used and an error bound derived [Neu10]. Follow-up methods allow for multiple events to occur within each interval [HH13], for adaptive sizes of the intervals [BS11], or for two-player games where one player controls the flow of time non-deterministically [BHK+12; HKK13]. Discretization was also used for solving a continuous-time stochastic extension of timed automata [KNS+00].

As regards models without non-determinism, this approach was proposed for non-Markovian stochastic Petri nets [ZFH01] and is closely related to discrete PH approximation. The main difference is that phase-type approximation may produce smaller model (with additional loss of precision) by adding self-loops or feedbacks in the DTMC gadgets; discretization rather corresponds to gadgets of the form of a chain with exit transitions into the target state from each state of the chain. A related research studies the *proxel based simulation* [Hor02; Laz05; KH09] which is in fact discretization where the state space is generated on the fly and states with insignificant probability are disposed. In contrast to discretizing CTMDPs, none of the work on non-Markovian models discusses the correctness of the discretization with respect to the stability of the models.

Simulation Another simple approximation method for transient and steady-state analysis of DES is simulation, also called Monte Carlo method [MU49]. Basically, the result is obtained by averaging a huge amount of randomly generated samples of the behaviour. This method is often used as an engineering tool that *usually works*. However, there is a vast literature on how to guarantee that the result lies within a small enough *confidence interval* around the actual value with high enough *confidence level*, see e.g. [Gly89; Haa10] and [JS89; HTT00] for the tool support. To the best of our knowledge, there is no research on such rigorous guarantees for GSMP with fixed-delay events. As a side effect of our theoretical research, we provide such guarantees for the restricted class of single-ticking GSMP in Section 5.4. More importantly, we show that in the unrestricted class, simulation may not provide correct answers to steady-state analysis as the steady-state distribution may not even exist.

The simulation method currently draws a lot of attention in the area in probabilistic verification as the so called *statistical model checking* [SVA04]. Supported by a number of tools [You05; KZH+09; KNP11], this approach attracts research in verification [YKN+06; LDB10] as well as computational biology [HLG+09; JCL+09]. The most related to the topic of this thesis is the work of Younes [YS02] on statistical model checking of GSMP (without fixed-delay events) against a time-bounded fragment of the logic CSL. Due to the time-boundedness, formal guarantees are obtained in a rather straightforward way using a variant of the Wald's sequential probability ratio test. Not much insight into the inner workings of the system is needed nor provided. Due to its simplicity, the statistical model checking is claimed [DLL+11] to allow us to solve many instances of undecidable problems. For DES with hard real-time bounds, this thesis also aims at not loosing correctness along the way.

Chapter 4

Unstable Behaviour in GSMP with Fixed-delay Events

In this chapter, we show the unstable behaviour of GSMP caused by fixed-delay events. First, we address the simpler Büchi specification and show that previous algorithms [ACD91; ACD92] are not correct. Second, we address the more intricate frequency measures. We show that GSMP models with fixed-delay events can also be unstable with respect to the discrete and timed frequencies, i.e. the variables **d** and **c** may not be well-defined. In other words, the steady-state distribution may not exist. This observation refines various previous results, for example [GL94; LS96; HMM05; ZFH01; Hor02], where the possibility that the steady-state distribution may not exist was not considered at all. This chapter is based on [BKK+11b].

4.1 Büchi specifications

We first consider the qualitative model-checking problem whether for a given GSMP $\mathscr{M}$ and a given set of target states T, we have $\mathscr{P}(\text{Büchi}_{\mathscr{M}}(T)) = 1$. In [ACD91; ACD92] there are algorithms for this model-checking problem based on the *region graph*. They rely on two crucial statements of the papers:

1. Almost all runs end in some of the bottom strongly connected components (BSCC) of the region graph.

2. Almost all runs entering a BSCC visit all regions of the component infinitely often.

The qualitative model checking then (according to these statements) reduces to SCC decomposition. Both of these statements are true for finite-state Markov chains. In this section, we show that neither of them has to be valid for region graphs of GSMP.

Region graph Let us first define the region graph of a GSMP. Let $C = \max(\{\ell_e, u_e \mid e \in \mathscr{E}\} \setminus \{\infty\}) + 1$ be greater than the maximal finite bound of the events of $\mathscr{M}$. Analogously to the region relation of timed automata, we put $(s, \xi, t) \sim (s', \xi', t')$ if $s = s'$ and

- for all $e \in \mathscr{E}$, $\xi(e)$ and $\xi'(e)$ agree on integral part or are both $> C$,
- for all $e, f \in \mathscr{E}$ lower than C in ξ, $\langle\xi(e)\rangle \leq \langle\xi(f)\rangle$ iff $\langle\xi'(e)\rangle \leq \langle\xi'(f)\rangle$,

Note that $\sim$ is an equivalence with finite index. The equivalence classes of $\sim$ are called *regions*.

Definition 4.1.1. Region graph *of $\mathscr{M}$ is a finite graph $G = (V, E)$ where the set of vertices V is the set of regions of $\mathscr{M}$ and for every pair of regions r, r' there is an edge $(r, r') \in E$ iff $P(z, r') > 0$ for all $z \in r$.*

The construction is correct because all configurations in the same region have the same one-step qualitative behaviour as formalized by the following lemma.

Lemma 4.1.2. *Let $z \sim z'$ be configurations and r be a region. We have $P(z, r) > 0$ iff $P(z', r) > 0$.*

Proof. For the sake of contradiction, let us fix a region r and a pair of configurations $z \sim z'$ such that $P(z, r) > 0$ and $P(z', r) = 0$. Let $z = (s, \xi, t)$ and $z' = (s, \xi', t')$.

First, let us assume that the part of $P(z, r)$ contributed by the variable-delay events V is zero, i.e. $\sum_{e \in V} \int_0^\infty \mathrm{Hit}(\{e\}, t) \cdot \mathrm{Wins}(\{e\}, t)\, dt = 0$. Then the set E of fixed-delay events scheduled with the minimal remaining time in z must be non-empty, i.e. some $e \in E$. We have

$$P(z, r) = \mathrm{Succ}(s, E)(s') \cdot \mathbf{1}[\bar{\xi} \in r] \cdot \prod_{c \in V} \int_{\xi(c)}^{\infty} f_{c|\xi(c)}(y)\, dy > 0 \tag{4.1}$$

$$P(z', r) = \mathrm{Succ}(s, E)(s') \cdot \mathbf{1}[\bar{\xi}' \in r] \cdot \prod_{c \in V} \int_{\xi(c)}^{\infty} f_{c|\xi'(c)}(y)\, dy = 0 \tag{4.2}$$

where s' is the control state of the region r and $\bar{\xi}$ and $\bar{\xi}'$ are the vectors of elapsed time after the transitions from z and z', respectively. It is easy to see that from $z \sim z'$ we get that $\bar{\xi} \in r$ iff $\bar{\xi}' \in r$. Hence, $P(z, r)$ and $P(z', r)$ can only differ in the big product in (4.1) and (4.2). For any fixed $c \in V$, we show that $\int_{\xi(e)}^{\infty} f_{c|\xi'(c)}(y)\, dy$ is positive. Recall that the density function f_c can qualitatively change only on integral values. Both z and z' have the same order of events' values. Hence, the integral is positive for ξ' iff it is positive for ξ. We get $P(z', r) > 0$ which is a contradiction.

Second, let us assume that there is a variable-delay event $e \in V$ such that

$$\int_0^\infty \mathrm{Succ}(s, \{e\})(s') \cdot \mathbf{1}[\xi_t \in r] \cdot f_{e|\xi(e)}(t) \cdot \prod_{c \in V \setminus \{e\}} \int_t^\infty f_{c|\xi(c)}(y)\, dy\, dt > 0$$

where ξ_t is the vector of elapsed time after the transition from z with waiting time t. There must be an interval I such that for every $t \in I$ we have that $f_{e|\xi(e)}(t)$ is positive, $\mathbf{1}[\xi_t \in r] = 1$, and $\int_t^\infty f_{c|\xi(c)}(y)\, dy > 0$ for any $c \in V \setminus \{e\}$. From the definition

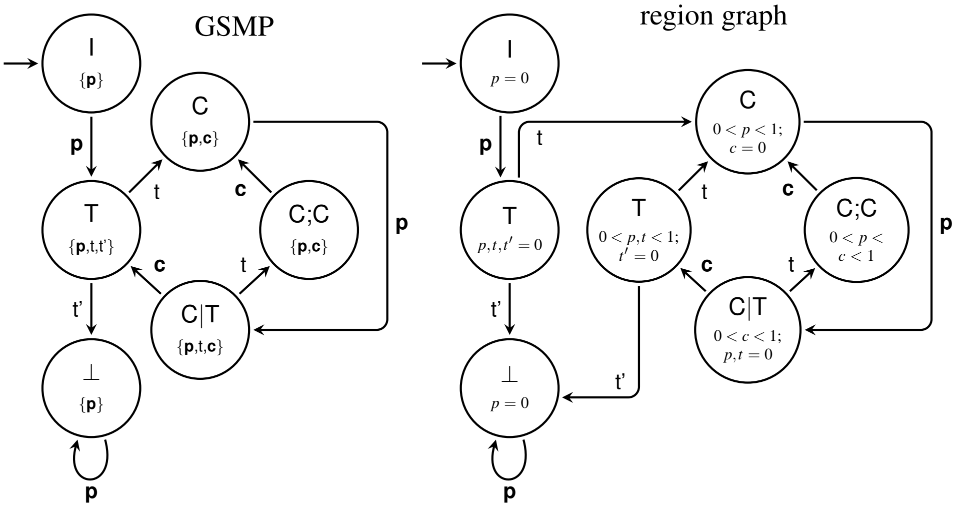

Figure 4.1: A GSMP model of a producer-consumer system on the left and its region graph on the right. The events p, t, and c model that a packet production, transport, and consumption is finished, respectively. Below each state label, there is the set of scheduled events (or the constraints describing the region). There are two fixed-delay events p and c (with $l_{\mathsf{p}} = u_{\mathsf{p}} = l_{\mathsf{c}} = u_{\mathsf{c}} = 1$), and uniformly distributed variable-delay events t, t' (with $l_{\mathsf{t}} = l_{\mathsf{t'}} = 0$ and $u_{\mathsf{t}} = u_{\mathsf{t'}} = 1$). In the region graph, only regions reachable with non-zero probability are depicted. Furthermore, the edges of the region graph are labelled with event names only for convenience.

of the region relation, this interval I corresponds to an interval between two adjacent events in ξ. Since $z \sim z'$, there must be also an interval I' such that for every $t \in I'$ we have that $f_{e|\xi'(e)}(t)$ is positive, $\mathbf{1}[\xi'_t \in r] = 1$, and $\int_t^\infty f_{c|\xi'(c)}(y)\,dy > 0$ for any $c \in V \setminus \{e\}$. Hence, $P(z', r) > 0$, contradiction. □

Counterexamples After defining the region graph, let us state the examples contradicting statements 1. and 2. above. Let us consider the GSMP depicted in Figure 4.1. It models a producer-consumer system with three components – a producer, a transporter and a consumer of packets. The components work in parallel but each component can process (i.e. produce, transport, or consume) at most one packet at a time.

Every 1 time unit, a packet is produced. The first production, at time 1.0, moves the system from the initial state I into the state T. Then, the packet is transported to the consumer; the transport takes between 0 and 1 time unit (and is uniformly distributed). Say it takes 0.9 time unit. Hence, at time 1.9 the consumption starts in the state C. As the consumption also takes exactly 1 time unit, it will finish at time 2.9. Before that a new packet is produced at time 2.0 and the system moves to the state C|T. Then, the ongoing consumption and the transport of the new packet interleave. If the transport is finished sooner (i.e. before 2.9), the new packet is placed in a buffer in the state C;C and its consumption will start at time 2.9. If the transport takes longer, the consumer idles from time 2.9 in the state T until the new packet is transported, say at time 2.95, and the new consumption begins. Then at time 3.0 a new packet is produced and the whole procedure is repeated. Notice that the state T is visited whenever the current transport takes more time than it has ever taken. Since the transport time is uniform and bounded by 1, the state T is visited less and less frequently and a smaller and smaller amount of time is spent there on every visit.

In state T, the whole system can break down and move into $\bot$. It is caused by the event t' which occurs uniformly in the interval $(0,1)$. Due to the shorter stays in state T, the probability of t' occurring decreases during the run resulting in the following lemma. Its formal proof is in Section 4.3.

Lemma 4.1.3. *The probability to reach $\bot$ from the state* I *is strictly less than* 1.

Observe that the region graph of this simple GSMP coincides with the transition graph of the GSMP as shown in Figure 4.1 on the right. Hence, in the region graph, the state $\bot$ forms the only BSCC. Thus, Lemma 4.1.3 disproves statement 1 as summarized by the following theorem.

Theorem 4.1.4. *There is a GSMP (with two fixed-delay and two variable delay events) where the probability to reach any BSCC of the region graph is strictly less than 1.*

Now consider in Figure 4.1 a transition under the event p from the state $\bot$ to the state I instead of the self-loop. This turns the whole region graph into a single BSCC. We prove that the state $\bot$ is almost surely visited only finitely often. Indeed, let $p < 1$ be the probability to reach $\bot$ guaranteed by Lemma 4.1.3. The probability to reach $\bot$ from $\bot$ again is also p as the only transition leading from $\bot$ enters the initial configuration. Therefore, the probability to reach $\bot$ infinitely often is $\lim_{n\to\infty} p^n = 0$. Hence, the statement 2 of [ACD91; ACD92] is disproved, as well:

Theorem 4.1.5. *There is a GSMP (with two fixed-delay and two variable delay events) with strongly connected region graph and with a region that is reached infinitely often with probability* 0.

Observe that these results not only show incorrectness of some previous algorithms but also the insufficiency of the region construction for qualitative behaviour

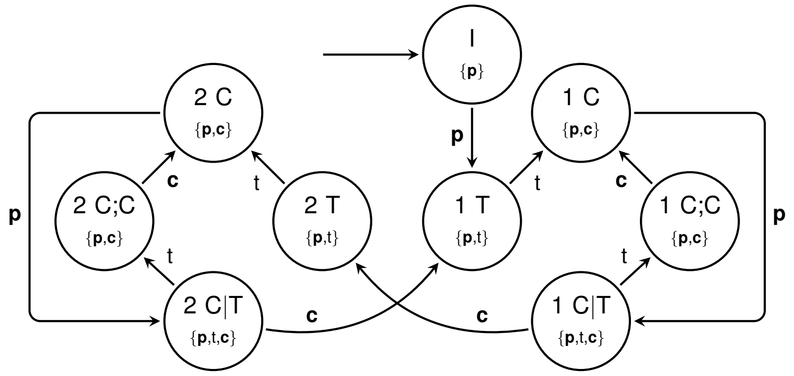

Figure 4.2: A variant of a producer-consumer system without the error state $\perp$ and with two consumer modules. The fixed-delay events **p** and **c** have again $l_{\mathsf{p}} = u_{\mathsf{p}} = l_{\mathsf{c}} = u_{\mathsf{c}} = 1$ and the uniformly distributed variable-delay event t has $l_{\mathsf{t}} = 0$ and $u_{\mathsf{t}} = 1$.

of general GSMP. In Chapter 5, we show that under some restrictions on the GSMP the region relation captures the qualitative behaviour well enough.

4.2 Performance measures

Let us now turn our attention to the frequency measures in GSMP. In Figure 4.2, we show an example of a GSMP with two fixed-delay events and one variable-delay event for which it is not true that the variables **d** and **c** are well-defined for almost all runs. It is a variant of the model from Figure 4.1, but there is no error state $\perp$ and the consumer has two modules – one is in operation and the other idles at a time. When the state T is entered, the consumer switches the modules. The labels 1 and 2 denote which module of the consumer is in operation.

Similarly to the previous example, the state 1 T or 2 T is entered (and the modules are switched) if and only if the current transport takes more time than it has ever taken. As the transport time is bounded by 1, it gets harder and harder to break the record. As a result, the system stays in the current module on average for longer time than in the previous module. Therefore, the frequencies for 1-states and 2-states oscillate. Precise computations are in the following subsection. We conclude the above observation by the following theorem.

Theorem 4.2.1. *There is a GSMP (with two fixed-delay events and one variable-delay event) for which it is* not *true that the variables* **c** *and* **d** *are almost surely well-defined.*

4.3 Formal proofs

Since the models from Figures 4.1 and 4.2 are closely related, we treat them similarly. By saying for example "state T" in the context of the model from Figure 4.2, we mean any of the states 1 T and 2 T. The key observation is that the state T is entered whenever the current transport takes more time than it has ever taken. As the transport time is bounded by 1, the difference of the maximal transport time from 1 gets smaller and smaller. We call this difference the *distance* of a configuration (s, ξ, t) and define it by $\langle \xi(c) - \xi(p) \rangle$. Using this notion, we first build up the tools needed for the proofs.

At the beginning of the run, the distance it initially set (uniformly on $(0,1)$) and the state C|T is reached. The time before the state C|T is reached again is called an *attempt* (to lessen the distance). This way, the whole run can be divided into a sequence of *attempts* (possibly a finite sequence, if the state $\perp$ is reached at the end) where each attempt takes exactly 1 time unit. An attempt is either successful or failed. A failed attempt means traversing the cycle C|T – C;C – C – C|T where the distance stays the same; a successful attempt means traversing the cycle C|T – T – C – C|T where the distance gets smaller. In each attempt, the probability of success is d (where d is the current distance) and the probability of failure is $1-d$ since t is distributed uniformly.

A *phase* is a maximum continuous sequence of attempts where exactly the last one is successful. Observe that in the model from Figure 4.2, the first phase is spent in 1-state, the second phase in 2-states, the third phase in 1-states, etc.

For the proofs, we need to characterize phases that lessen the distance substantially. A phase is called *strong* if the newly generated distance is at most half of the old one. Further, we define a *half-life* to be a maximum continuous sequence of phases where exactly the last one is strong. Every run can thus be uniquely decomposed into a sequence of half-lives. The random variable $D_{i,j}$ denotes the distance at the beginning of the j-th phase of the i-th half-life. Denoting the number of phases in the i-th half-life by $L(i)$ we get $D_{n-1,L(i)} \geq 2D_{n,1}$.

Proof of Lemma 4.1.3. In the following, we prove that the probability to reach the state $\perp$ is strictly less than 1. Firstly, we prove that we can restrict to *quick* runs that have few phases in each half-life. A run is *quick* if for any $i \leq H$ it has at most $2i$ phases in the i-th half-life where H is the number of half-lives on the run. The probability that $2i$ consecutive phases are not strong, i.e. $L(i) > 2i$, is $1/2^{2i}$ as t is distributed uniformly. Therefore, the probability that there is $i \in \mathbb{N}$ with $L(i) > 2i$ is less than $\sum_{i=1}^{\infty} 1/2^{2i} = 1/3$. Hence, at least $2/3$ of runs are quick.

Secondly, we show that quick runs that never visit $\perp$ have positive probability. From now on, we restrict to quick runs and talk about probability conditioned by the set of quick runs. By definition (and simple induction), for every run with at least i half-lives, it holds $D_{i,1} \leq 1/2^i$. As the distance can only get smaller during the half-life, it also holds $D_{i,j} \leq 1/2^i$ for any $j \leq L(i)$. Since t' is distributed uniformly, the probability that $\perp$ is not reached during one phase of the i-th half-life is at least $(1-1/2^i)$. The probability that $\perp$ is not reached during the whole i-th half-life is at least $(1-1/2^i)^{2i}$. Hence, the probability that $\perp$ is never reached is greater than

$$m := \prod_{i=1}^{\infty}(1-1/2^i)^{2i}.$$

It remains to show that $m > 0$. This is equivalent to $\sum_{i=1}^{\infty} \ln(1-1/2^i)^{2i} > -\infty$ which in turn can be rewritten as $\sum_{i=1}^{\infty} 2i \ln\left(2^i/(2^i-1)\right) < \infty$. As regards the logarithm,

$$\ln\left(\frac{2^i}{2^i-1}\right) \;=\; \frac{\ln(2^i)-\ln(2^i-1)}{1} \;\leq\; \ln'(2^i-1) \;=\; \frac{1}{2^i-1}$$

where the inequality is obtained by concavity of ln: the derivative of ln can be substituted for its approximation in 2^i-1. In total, we get

$$\sum_{i=1}^{\infty} 2i \cdot \ln\left(\frac{2^i}{2^i-1}\right) \;\leq\; 2\sum_{i=1}^{\infty}\frac{i}{2^i-1} \;\leq\; 2\sum_{i=1}^{\infty}\frac{i}{2^{i-1}} \;=\; 2\cdot 4 \;<\; \infty \qquad \square$$

In the following, we prove that in our example there are runs of positive measure where **d** and **c** are not well-defined. Namely, for these runs the partial sums oscillate. Let $S_{i,j}$ be the number of attempts in the j-th phase of the i-th half-life, i.e. a *length* of this phase. Roughly speaking, we show that there are runs (of overall positive measure) where some phase is longer than the overall length of all phases up to that point (causing the frequencies to oscillate). Note that the precise statement of the lemma implies moreover that this happens even infinitely often on runs of overall positive measure.

Lemma 4.3.1. *There are $\alpha > 0$ and $m > 0$, such that for every $n > 1$ there is a set $\mathscr{R}_n$ of runs of measure at least m satisfying*

$$S_{n,1} \geq \alpha \sum_{\substack{i=1..n-1 \\ j=1..L(i)}} S_{i,j}$$

Proof. We set $\alpha = 1/18$ and $m = 1/8$ and let $n > 1$ be arbitrary. We define the set $\mathscr{R}_n$ to be the set of all runs σ such that the following conditions hold:

1. $S_{n,1} > 1/(2D_{n,1})$,
   (the length of the "last" phase is above its expected value),

2. for all $1 \leq i < n$, $L(i) \leq (n-i)+3$,
   (the half-lives so far have no more phases than $4, 5, \ldots, n+1, n+2$, respectively),

3. for all $1 \leq i < n$ and $1 \leq j \leq L(i)$, $S_{i,j} \leq 3(n-i)/D_{i,j}$,
   (all phases in previous half-lives are "short" w.r.t their expectations).

Denote $D := D_{n,1}$. We firstly prove that $S_{n,1} \geq \alpha \sum_{i=1..n-1, j=1..L(i)} S_{i,j}$ for all runs in $\mathscr{R}_n$. Observe that by simple induction, we have for all natural $i < n$ and $j \leq L(i)$,

$$D_{i,j} \geq 2^{n-i} \cdot D. \tag{4.3}$$

Due to this inequality and requirements 2. and 3., we can bound the overall length of all previous phases by

$$\begin{aligned}\sum_{\substack{i=1..n-1\\ j=1..L(i)}} S_{i,j} &\leq \sum_{\substack{i=1..n-1\\ j=1..L(i)}} \frac{3(n-i)}{D_{i,j}} \leq \sum_{i=1}^{n-1} \frac{((n-i)+3)\cdot 3(n-i)}{2^{n-i}\cdot D} \\ &= \sum_{i=1}^{n-1} \frac{(i+3)\cdot 3i}{2^i \cdot D} \leq \sum_{i=1}^{\infty} \frac{3(i^2+3i)}{2^i \cdot D} = \frac{3(6+3\cdot 2)}{D} = \frac{1}{2\alpha D}\end{aligned}$$

and conclude by the requirement 1.

It remains to prove that measure of $\mathscr{R}_n$ is at least m. We investigate the measures of the runs described by requirements 1.–3. Firstly, the probability that $S_{n,1} > \frac{1}{2D_{n,1}}$ is $(1-D_{n,1})^{1/2D_{n,1}}$, which is greater than $1/2$ for $D_{n,1} \leq 1/2$, i.e. for $n \geq 2$. Out of this set of runs of measure $1/2$ we need to cut off all runs that do not satisfy requirements 2. or 3. As for 2., the probability of i-th half-life failing to satisfy 2. is $(1/2)^{(n-i)+3}$ corresponding to at least $(n-i)+3$ successive non-strong phases. Therefore, 2. cuts off less than $\sum_{i=1}^{n-1} 1/2^{(n-i)+3} = \sum_{i=1}^{n-1} 1/2^{i+3} \leq \sum_{i=1}^{\infty} 1/2^{i+3} = 1/2^3$. From the remaining runs we need to cut off all runs violating 3. Since the probability of each $S_{i,j}$ failing is $(1-D_{i,j})^{3(n-i)/D_{i,j}}$, the overall probability of all violating runs is due to (4.3) at most

$$\begin{aligned}\sum_{i=1}^{n-1}\sum_{j=1}^{L(i)} (1-D_{i,j})^{3(n-i)/D_{i,j}} &= \sum_{i=1}^{n-1}\sum_{j=1}^{L(i)} (1-2^{n-i}D)^{3(n-i)/2^{n-i}D} \\ &\leq \sum_{i=1}^{n-1} (i+3)(1-2^iD)^{3i/2^iD} \leq \sum_{i=1}^{\infty} (i+3)(1/e)^{3i} \\ &= \frac{4e^3-3}{(e^3-1)^2} < 1/4\end{aligned}$$

Altogether the measure of $\mathscr{R}_n$ is at least $1/2 - 1/8 - 1/4 = 1/8 = m$. □

We can now easily finish the proof of Theorem 4.2.1.

Proof of Theorem 4.2.1. Due to the previous lemma, there is a set $\mathscr{R}$ of runs of positive measure such that each run of $\mathscr{R}$ is contained in infinitely many $\mathscr{R}_n$'s. We prove that neither $\mathbf{d}(\sigma)$ nor $\mathbf{c}(\sigma)$ is well-defined on any $\sigma \in \mathscr{R}$ where (slightly abusing the notation) $\mathbf{d}(\sigma)$ and $\mathbf{c}(\sigma)$ denotes the sum of frequencies of all 1-states instead of one single state s. Since attempts last for one time unit, non-existence of $\mathbf{d}(\sigma)$ implies non-existence of $\mathbf{c}(\sigma)$.

Assume for a contradiction that $\mathbf{d}(\sigma)$ is well-defined. Let $\mathbf{d}(\sigma) \leq 1/2$, the other case is handled symmetrically. Because 1-states are visited exactly in odd phases, we have $\mathbf{d}(\sigma) = \lim_{n\to\infty} (\sum_{i=1}^{n} s_i \cdot odd(i)) \,/\, (\sum_{i=1}^{n} s_i)$ where s_i is the number of attempts in the i-th phase and $odd(i) = 1$ if i is odd and 0 otherwise. By the definition of limit, for every $\varepsilon > 0$ there is n_0 such that for all $n > n_0$

$$\left| \frac{\sum_{i=1}^{n} s_i \cdot odd(i)}{\sum_{i=1}^{n} s_i} - \frac{\sum_{i=1}^{n-1} s_i \cdot odd(i)}{\sum_{i=1}^{n-1} s_i} \right| < \varepsilon \tag{4.4}$$

Let $\varepsilon > 0$ be such that $\alpha \geq \varepsilon/(1-2\varepsilon-\mathbf{d}(\sigma))$. Due to Lemma 4.3.1, there is an odd $n > n_0$ satisfying $s_n \geq \alpha \sum_{i=1}^{n-1} s_i$. Denoting $A = \sum_{i=1}^{n-1} s_i$ and $O = \sum_{i=1}^{n-1} s_i \cdot odd(i)$ we get from (4.4)

$$\frac{O+s_n}{A+s_n} - \frac{O}{A} \geq \frac{O+\alpha A}{A+\alpha A} - \frac{O}{A} \overset{(*)}{\geq} \frac{\varepsilon \cdot \left(1-\frac{O}{A}\right)}{1-\mathbf{d}(\sigma)-\varepsilon} \overset{(**)}{\geq} \frac{\varepsilon \cdot (1-\mathbf{d}(\sigma)-\varepsilon)}{1-\mathbf{d}(\sigma)-\varepsilon} = \varepsilon$$

which is a contradiction with (4.4). Notice that we omitted the absolute value from (4.4) because for an odd n the term is non-negative. The inequality $(*)$ is obtained by substituting $\varepsilon/(1-2\varepsilon-\mathbf{d}(\sigma))$ for α and by straightforward manipulation. In $(**)$ we use, similarly to (4.4), that $|\frac{O}{A} - \mathbf{d}(\sigma)| < \varepsilon$. □

Chapter 5

Conditions on Stability

After showing instability in GSMP models with fixed-delay events, we turn our attention to positive results. First, we identify a stable subclass of GSMP with fixed-delay events that we call *single-ticking* GSMP. Second, we address GSMP observed by DTA and show that the DTA observer does *not* introduce any instability in the model. In contrast to the proof in [BKK+11a], we show it by reduction to single-ticking GSMP. Third, we provide conditions on stability for a closely related formalism of deterministic and stochastic Petri nets. We use an existing reduction [Haa10] to GSMP and again show that the resulting GSMP with fixed-delay events is single-ticking. Fourth, we conclude the chapter by showing that the frequency measures of single-ticking GSMP not only exist but also can be effectively approximated. This chapter is based on [BKK+11a; BKK+11b; BKK+13].

5.1 GSMP with Fixed-delay Events

First of all, motivated by the previous counterexamples, we identify the behaviour of the fixed-delay events that may cause **d** and **c** to be undefined. The problem lies in fixed-delay events that can immediately schedule themselves whenever they occur; such an event can occur periodically like ticking of clocks. In the example of Figure 4.1, there are two such events p and c. The phase difference of their ticking gets smaller and smaller, causing the unstable behaviour. For the following definition we generalize the *ticking* of one event to cyclic periodic occurrence of multiple fixed-delay events. Furthermore, the definition is based on the region graph, see Definition 4.1.1.

We say that a set of fixed-delay events E is *ticking* in a region r if there is a cycle $r_1,\ldots,r_n$ in the region graph and sets of fixed-delay events $E_1\ldots,E_n$ with $r_1 = r_n = r$ and $E_1 = E_n = E$ such that each E_i is a maximal set of events with the same fractional value in r_i and upon traversing the cycle,

- each event scheduled to occur in r either occurs or stops being scheduled;
- the sets of fixed-delay events schedule each other in the following sense. For each $1 \leq i < n$ either a set of other events $E \cap E_i = \emptyset$ occurs in r_i and $E_{i+1} \subseteq E_i$

(some events from E_i may stop being scheduled in r_{i+1}); or a subset $E \subseteq E_i$ of the events occur and in r_{i+1} the set $F \subseteq E_i$ stops being scheduled and the set of fixed-delay events G is newly scheduled such that $E_{i+1} = ((E_i \setminus E) \setminus F) \cup G$. Notice that E are the events of E_i with their elapsed time closest to their delay in r_i.

When traversing this cycle forever, due to the first condition the time diverges, due to the second condition the fixed-delay events occur periodically like ticking of the clock.

Definition 5.1.1. *A GSMP is called* single-ticking *if in every region of its region graph there is at most one set of ticking events.*

From now on we restrict to single-ticking GSMP and prove our main positive result.

Theorem 5.1.2. *In single-ticking GSMP, the random variables $\mathbf{d}_s$ and $\mathbf{c}_s$ are well-defined for almost every run and any $s \in S$. Furthermore, almost every run reaches a BSCC of the region graph and for each BSCC B and $s \in S$ there are $d_s, c_s \in [0,1]$ such that $\mathbf{d}_s(\sigma) = d_s$ and $\mathbf{c}_s(\sigma) = c_s$ for almost every run σ reaching the BSCC B.*

The rest of this section is devoted to the proof of Theorem 5.1.2. First, we show that almost all runs end up trapped in some BSCC of the region graph. Second, we solve the problem while restricting to runs that *start* in a BSCC (as the initial part of a run outside of any BSCC is not relevant for the long run average behaviour). We show that in a BSCC, the variables $\mathbf{d}_s$ and $\mathbf{c}_s$ are almost surely constant. The second part of the proof relies on several standard results from the theory of general state space Markov chains. Formally, the proof follows from Propositions 5.1.4 and 5.1.8 stated below.

Remark 5.1.3. *Observe that the definition of single-ticking GSMP is slightly more complicated than in [BKK+11b]. This is necessary for Section 5.2 where we show by a simple reduction to single-ticking GSMP the similarly complicated result of [BKK+11a] that GSMP observed by a DTA are stable as well.*

5.1.1 Reaching a BSCC

Proposition 5.1.4. *In single-ticking GSMP, almost all runs reach a BSCC of the region graph.*

The rest of this subsection forms the proof of Proposition 5.1.4. The proof uses methods similar to [ACD92]. By definition, the process moves along the edges of the region graph. From every region, there is a path of minimal length through the

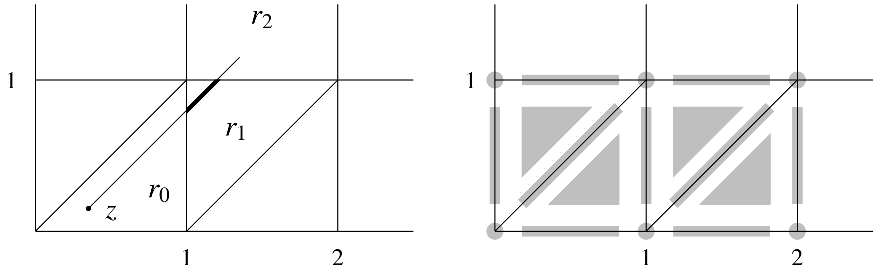

Figure 5.1: Intuition for δ-separated parts of regions.

region graph into a BSCC, let n be the maximal length of all such paths. Hence, in at most n steps the process reaches a BSCC with positive probability from any configuration. Observe that if this probability was bounded from below, we would eventually reach a BSCC from any configuration almost surely.

However, such a bound does not exist in general; the transition probabilities between two regions approach zero as the starting configuration approaches the boundary of its region (i.e. as the fractional parts of two clocks approach each other). The left part of Figure 5.1 shows a part of the region graph of a system with a single state and two fixed-delay events with delay 3. There is also a single variable-delay event, which is positive on $(0,1)$ and its elapsed time is not depicted. Now observe that if the starting configuration z comes closer and closer to the diagonal, the probability that the (only) event happens in the region r_1 is smaller and smaller. Similarly, if z comes closer and closer to the bottom or right boundary of the region, respectively, the probability that the event happens in the region r_2 or r_0 is smaller and smaller.

Notice that the transition probabilities depend on the difference of the fractional values of the clocks, we call this difference *separation*. The δ-separated parts of regions are depicted in grey in the right part of Figure 5.1. Here, we are either at least δ-away from the boundary of the region or exactly at the boundary.

Definition 5.1.5. *Let $\delta > 0$. We say that a configuration (s,ξ,t) is* δ-separated *if for every $x,y \in \{0\} \cup \{\xi(e) \mid e \in \mathbf{E}(s)\}$, we have either $|\langle x\rangle - \langle y\rangle| > \delta$ or $\langle x\rangle = \langle y\rangle$.*

We fix a $\delta > 0$. To finish the proof using the concept of δ-separation, we need two observations. First, from *any* configuration we reach in m steps a δ-separated configuration with probability at least $q > 0$. Second, the probability to reach a fixed region from *any* δ-separated configuration is bounded from below by some $p > 0$.

The second observation only holds if we can bound from below the events' densities; this is satisfied in *inner* configurations. We say that a configuration (s,ξ,t) is *inner* if $\xi(e) \leq 2r + B + 1$ for every $e \in \mathbf{E}(s)$. (Recall that r is the bound on the conditional expectations of all events and $B = \max\left(\{\ell_e, u_e \mid e \in \mathscr{E}\} \setminus \{\infty\}\right)$. The reason why the bound is $2r + B + 1$ is technical, the main idea is that the events' values are bounded by *some* constant.) The observations are formulated in Lemmata 5.1.6 and 5.1.7.

Lemma 5.1.6. *There is $\delta > 0$, $m \in \mathbb{N}$ and $q > 0$ such that from every configuration we reach a δ-separated inner configuration in m steps with probability at least q.*

Proof. In the following, by *value* of an event e we mean the fractional part $\langle \xi(e) \rangle$ when the vector of elapsed time ξ is clear from context. Let us fix a configuration z. In the following proof, we restrict the probabilistic behaviour in such a way that all the runs satisfying this restriction starting in z have probability at least q and visit a δ-separated inner configuration in m steps.

There are two sources of stochastic behaviour: (i) occurrence of events and (ii) choice of successor states. As regards the choice of the successor for a given set of events, we do not restrict the behaviour at all. Furthermore, we cannot restrict when the fixed-delay events occur; we only restrict the timing of variable-delay events in two phases as follows.

1. In the first phase, we aim at reaching an inner configuration. It lasts $m_1 := 2r \cdot |\mathscr{E}| \cdot 2 + |\mathscr{E}|$ steps. For each variable-delay event already scheduled in z, we (a) restrict that within the first $2r$ time units of the first phase it either occurs or stops being scheduled. Furthermore, (b) the elapsed time of each variable-delay event newly scheduled within the first phase never exceeds $2r$ during the first phase and (c) any variable-delay event newly scheduled within the first phase never occurs before its elapsed time reaches $1/2$. We show that the probability of runs satisfying this restriction is bounded from below. As regards (a), the bound is obtained by the Markov's inequality stating that for a random variable X with expected value $\mathrm{E}(X)$ and $a \in \mathbb{R}_{>0}$ we have $\mathscr{P}(X \geq a) \leq \mathrm{E}(X)/a$. Indeed the expected time to occurrence of each scheduled event is $\leq r$ and the probability that such an event occurs or stops being scheduled before $2r$ is $\geq 1/2$. As the events are independent, we get an overall bound of $1/2^{|\mathscr{E}|}$. As regards (b) and (c), the bound is obtained as there is $p > 0$ such that for each event e, we have $\int_{1/2}^{2r} f_e(x)dx > p$. Due to (c), the first phase takes at least $2r$ time units. Hence, due to (a) and (b), each scheduled event has elapsed time at most $2r$ at the end of the first phase.

2. In the second phase, we aim at reaching a δ-separated configuration for separation $\delta = 1/(4 \cdot |\mathscr{E}|)$. We restrict that (1) each variable-delay occurs within the first time unit of its support (i.e. at latest one time unit after it exceeds its lower bound). If the event is already above its lower bound at the beginning

of the second phase, we restrict it to occur in the first time unit of the second phase.

In addition, we (2) restrict the times where each variable-delay event occurs within its allowed unit interval so that after each occurrence of a variable-delay event, no other variable-delay event occurs within time δ. Now we explain in greater detail the restriction (2) and why the runs satisfying (1) and (2) have probability bounded from below.

The restrictions are based on the *values* of the events, i.e. on the fractional parts of the events' elapsed times. Hence, we focus on the $[0,1]$ line segment modulo one. For each scheduled variable-delay event, we place on the line segment a *reservation bar* of length δ that we shift along with the events' values, as the time flows.[1] During the second phase, we restrict each variable-delay event to occur when its reservation bar passes over 0 and when (1) is satisfied.

At last, we explain the placement of the reservation bars. We divide the $[0,1]$ line segment into $4 \cdot |\mathscr{E}|$ *slots* of length δ. We say that a slot is occupied if it contains value of some fixed-delay event or a piece of a reservation bar; or empty, otherwise. Whenever a variable-delay event e gets newly scheduled during the first or the second phase, we place its reservation bar of length δ on an empty slot with its left and right neighbouring slots (modulo one) also empty. Such a slot must always exist since there are $4 \cdot |\mathscr{E}|$ slots and each event of $\mathscr{E} \setminus \{e\}$ occupies at most two slots at a time (and two neighbouring slots that must stay empty).

The second phase lasts for $m_2 := |V|$ steps where $|V|$ is the size of the region graph. The probability is bounded from below since for each variable-delay event e, the probability that it occurs with its elapsed time within any δ-long subinterval of the interval $[l_e, \min\{u_e, 2r + B + 1\}]$ is bounded from below.

The resulting configuration is δ-separated from two reasons. First, variable-delay events are separated from each other and from fixed-delay events because of the way we place the reservation bars. Second, each fixed-delay event scheduled in the resulting configuration is separated from other fixed-delay events because during the second phase at least one cycle in the region graph is traversed. Each fixed-delay event has been scheduled either by the ticking events on that cycle (observe that the definition of δ-separation allows two events to have equal value); or by an occurrence of some variable-delay event separated from other fixed-delay events

[1] We treat the reservation bars also modulo one, i.e. when a piece of a reservation bar exceeds 1, it actually overlaps to the beginning of the line segment.

(possibly followed by occurrences of fixed-delay events not changing the value and thus not changing the separation).

The resulting configuration is inner because after the first phase, all the variable-delay events have elapsed time below $2r$, and in the second phase we do not allow any event to overpass the value $2r+B+1$. □

Notice that Lemma 5.1.6 does not hold for general GSMP. As in the example of Figure 4.1, the separation may be non-increasing for all runs. On the contrary, Lemma 5.1.7 does hold even for unrestricted GSMP.

Lemma 5.1.7. *For every $\delta > 0$ and $k \in \mathbb{N}$ there is $p > 0$ such that for any regions r, r' connected by a path of length k and for any δ-separated inner configuration $z \in r$, we have $P^k(z,r') > p$.*

Proof. Let $z \in r = r_0$ and $r_0, r_1, \dots, r_k$ be a shortest path in the region graph from the region $r = r_0$ to the region $r' = r_k$. We restrict to runs that follow this path so that in each step they lose two thirds of the separation. At last, a $(\delta/3^k)$-separated configuration in the target region r_k is reached. We show how we obtain the overall bound on probabilities from bounds on every step.

In each step either a variable-delay event or a set of fixed-delay events occur. Let δ_i be the separation in the current step. To follow the region path, a specified event must occur in an interval between two specified values which are δ_i-separated. A fixed-delay event occurs in this interval for sure because it has been scheduled this way. For a variable-delay event, we divide this interval into thirds and restrict the event to occur in the middle subinterval of length δ_{i+1}. This happens with a probability bounded from below if the events' densities are bounded from below. Furthermore, to follow the path in the region graph, no other event can occur sooner. Every other event has at least δ_{i+1} to its upper bound; the probability that it does *not* occur is again bounded from below if the events' densities are bounded from below.

It remains to show that the events' densities are bounded from below on this region path. Notice that not all the configurations on the path have to be inner. Still, the time the path must take is bounded by $|V|(B+1)$ since the shortest path in the region graph must be shorter than the size of the region graph $|V|$; and each step takes at most $(B+1)$ time because waiting more time does not change the target region of that step (recall that $B = \max\left(\{\ell_e, u_e \mid e \in \mathcal{E}\} \setminus \{\infty\}\right)$). All the (conditional) events' densities are bounded from below in the interval $[0, 2r+B+1+|V|(B+1)]$. □

By repeating the two observations ad infinitum, we reach some BSCC almost surely concluding the proof of Proposition 5.1.4.

5.1.2 Frequency in a BSCC

From now on, we deal with the bottom strongly connected components that are reached almost surely. Hence, we assume that the region graph G is strongly connected. We have to allow an arbitrary initial configuration $z_0 = (s, \xi, t)$; in particular, ξ does not have to be a zero vector.[2]

Proposition 5.1.8. *In a single-ticking GSMP with strongly connected region graph, there are values $d_s, c_s \in [0,1]$ for each $s \in S$ such that for any initial configuration z_0 and for almost all runs σ starting from z_0, we have that $\mathbf{d}_s$ and $\mathbf{c}_s$ are well-defined and $\mathbf{d}_s(\sigma) = d_s$ and $\mathbf{c}_s(\sigma) = c_s$.*

Let us fix $s \in S$. We assume that the region graph is aperiodic in the following sense. A *period* p of a graph G is the greatest common divisor of lengths of all cycles in G. The graph G is *aperiodic* if $p = 1$. First, we assume that G is aperiodic, later we discuss the opposite case. We show that the aperiodic chain Φ is in some sense stable. Namely that (i) Φ has a unique invariant measure that is independent of the initial measure and (ii) the strong law of large numbers (SLLN) holds for Φ.

Standard results for general state space Markov chains yield (i) and (ii) for chains where the whole set of configurations is *small*. Roughly speaking, it means that the transient distribution after n steps is similar (up to a factor of ε) for all starting configurations z. Intuitively, the process "regenerates" to some extent each n steps.

Definition 5.1.9. *Let $n \in \mathbb{N}$, $\varepsilon > 0$, and κ be a probability measure on $(\Gamma, \mathscr{F})$. The set Γ is (n, ε, κ)-*small *if for all $z \in \Gamma$ and $A \in \mathscr{F}$ we have that $P^n(z,A) \geq \varepsilon \cdot \kappa(A)$.*

The proof of the following lemma is the most demanding part of this chapter and we deal with it in the next subsection.

Lemma 5.1.10. *There is $n \in \mathbb{N}$, $\varepsilon > 0$, and κ such that Γ is (n, ε, κ)-small.*

As a further step, we show how smallness of the state space implies (i) and (ii).

Lemma 5.1.11. *If the set of configurations Γ of a GSMP is (n, ε, κ)-small,*

(i) there is a unique probability measure π on $(\Gamma, \mathscr{F})$ that is invariant, *i.e.*

$$\pi(A) \quad = \quad \int_\Gamma \pi(dx) P(x,A) \qquad\qquad \text{for all } A \in \mathscr{F},$$

(ii) Γ satisfies the strong law of large numbers, i.e. for each function $h : \Gamma \to \mathbb{R}$ such that $E_\pi[h] < \infty$, we almost surely have

$$\lim_{n\to\infty} \frac{\sum_{i=0}^{n-1} h(\Phi_i)}{n} \quad = \quad E_\pi[h], \tag{5.1}$$

where $E_\pi[h]$ is the expected value of h according to the invariant measure π.

[2] Technically, the initial measure is $\mu(A) = 1$ if $z_0 \in A$ and $\mu(A) = 0$, otherwise.

Proof. The theorem is a consequence of standard results for GSSMCs; we only give pointers to appropriate places in [MT09]. Since Γ is (m,ε,ν)-small, observe that

- the chain is by definition φ-irreducible for $\varphi = \nu$; thus also ψ-irreducible by [MT09, Proposition 4.2.2];
- Γ is by definition also (a,ε,ν)-petite (see [MT09, Section 5.5.2]), where a is the Dirac distribution on $\mathbb{N}_0$ with $a(m) = 1$, $a(n) = 0$ for $n \neq m$.

As Γ is trivially not uniformly transient, the chain Φ is recurrent by [MT09, Theorem 8.0.1], Thus by [MT09, Theorem 10.0.1], there exists a unique invariant probability measure π. Since π is trivially also subinvariant, the chain Φ is positive Harris by [MT09, Theorem 10.4.10 (ii)]. Therefore, (ii) is obtained by applying [MT09, Theorem 17.0.1 (i)]. □

Next, we show that (i) and (ii) imply the proposition. We set in the equality (5.1) the function $h((s',\xi,t))$ to the delta function $\delta(s')$ where $\delta(s') = 1$ if $s' = s$, and 0, otherwise. We have $E_\pi[h] < \infty$ since $h \leq 1$. From (5.1) we obtain that almost surely

$$\mathbf{d}_s \quad = \quad \lim_{n\to\infty} \frac{\sum_{i=0}^{n-1} h(\Phi_i)}{n} \quad = \quad E_\pi[h].$$

As a result, $\mathbf{d}_s$ is well-defined and equals the constant value $E_\pi[h]$ for almost all runs. We treat the variable $\mathbf{c}_s$ similarly. Here, we assume w.l.o.g. that each state s' has an unique predecessor state $\overline{s'}$. Indeed, we can extend the state space with the information about the previous state (the initial state having itself as the predecessor). Let $W((s',\xi,t)) = t$ and $\tau((s',\xi,t)) = W((s',\xi,t)) \cdot \delta(\overline{s'})$. Since all the events have finite expectation, we have $E_\pi[W] < \infty$ and $E_\pi[\tau] < \infty$. Now we show that

$$\mathbf{c}_s \quad = \quad \frac{E_\pi[\tau]}{E_\pi[W]},$$

yielding that $\mathbf{c}_s$ is well-defined and equals the constant $E_\pi[\tau]/E_\pi[W]$ for almost all runs. Let us consider a run $\sigma = (s_0,\xi_0,t_0)\ (s_1,\xi_1,t_1)\cdots$. We have that

$$\begin{aligned}\mathbf{c}_s(\sigma) &= \lim_{n\to\infty} \frac{\sum_{i=0}^{n-1} \delta(s_i)\cdot t_{i+1}}{\sum_{i=0}^{n-1} t_{i+1}} = \lim_{n\to\infty} \frac{\sum_{i=0}^{n} \delta(\overline{s_i})\cdot t_i}{\sum_{i=0}^{n} t_i} = \lim_{n\to\infty} \frac{\sum_{i=0}^{n-1} \delta(\overline{s_i})\cdot t_i}{n} \cdot \frac{n}{\sum_{i=0}^{n-1} t_i} \\ &= \frac{\lim_{n\to\infty}(\sum_{i=0}^{n-1} \delta(\overline{s_i})\cdot t_i)/n}{\lim_{n\to\infty}(\sum_{i=0}^{n-1} t_i)/n} = \frac{\lim_{n\to\infty}(\sum_{i=0}^{n-1} \tau((s_i,\xi_i,t_i)))/n}{\lim_{n\to\infty}(\sum_{i=0}^{n-1} W((s_i,\xi_i,t_i)))/n} = \frac{E_\pi[\tau]}{E_\pi[W]}\end{aligned}$$

The last equality follows from (5.1) and by proving the existence of the limit it also justifies the previous manipulations with the limits. This concludes the proof of Proposition 5.1.8 under the aperiodicity assumption.

If the region graph has period $p > 1$, we can employ the standard technique and decompose the region graph (and the Markov chain) into p aperiodic components [MT09, Proposition 5.4.6]. Namely, we decompose Φ into p stochastic processes $\Phi_0, \dots, \Phi_{p-1}$ where each Φ_k makes steps corresponding to p steps of the original process Φ and the initial measure of Φ_k is $\alpha_k(A) = \int_{x \in \Gamma} \mu(dx) \cdot P^k(x, A)$ for each $A \in \mathscr{F}$. Each Φ_k is aperiodic and hence small (this follows by slightly generalizing the arguments of Lemma 5.1.10). By the same arguments as above, each Φ_k has frequencies $\mathbf{c}_s$ and $\mathbf{d}_s$ almost surely well-defined and equal to some constants $d_{s,k}$ and $c_{s,k} = \tau_{s,k}/W_{s,k}$. Finally, by a straightforward manipulation, we obtain the frequency of visits to s in $\mathscr{M}$ as an average of the corresponding frequencies in $\Phi_0, \dots, \Phi_{p-1}$, i.e. $\mathbf{d}_s(\rho) = 1/p \cdot (d_{s,0} + \dots + d_{s,p-1})$ and $\mathbf{c}_s(\rho) = (\tau_{s,0} + \dots + \tau_{s,p-1})/(W_{s,0} + \dots + W_{s,p-1})$ for almost all runs ρ.

5.1.3 Formal proofs

In the whole subsection, we prove Lemma 5.1.10. We show that there is a set of configurations C such that

(I) there is $o \in \mathbb{N}$ and $\zeta > 0$ such that for every $z \in \Gamma$ we have $P^o(z, C) \geq \zeta$;

(II) there is $p \in \mathbb{N}$, $\eta > 0$, and a probability measure κ such that for every $z \in C$ and $A \in \mathscr{F}$ we have $P^p(z, A) \geq \eta \cdot \kappa(A)$.

The lemma is then obtained by setting $n := o + p$ and $\varepsilon := \zeta \cdot \eta$. We choose some reachable inner region r either with no set of ticking events or with its set of ticking events having the greatest value among all events scheduled in r. There clearly is such a region. We choose C to be the set of δ-separated configurations in r where δ is fixed below.

The delicate part about (I) is that this set has to be reached with positive probability from any configuration in *exactly* o steps. First, from Lemma 5.1.6, there is δ', m, and q such that we reach from any $z \in \Gamma$ in m steps some δ'-separated configuration z' with probability at least q. From z', we need to get to C. We use a standard result from the theory of Markov chains, see for example [Ros06, Lemma 8.3.9], that in every ergodic Markov chain there is m' such that between any two states there is a path of length exactly m'. The same result holds for the aperiodic region graph G. Hence, from z' we have a path of length m' to the region r. From Lemma 5.1.7, we have $q' > 0$ such that we reach r from z' in m' steps with probability at least q'. Furthermore, we end up in a $(\delta'/3^{m'})$-separated configuration of the region r. Hence, we set $\delta = \delta'/3^{m'}$, $\zeta := q \cdot q'$, $o := m + m'$, and we obtain the first property.

As regards the property (II), we show that there is $p \in \mathbb{N}$, $\delta' > 0$, and a configuration z^* such that from any $z \in C$, there is a δ'-*wide* path of length p from z to z^*.

Furthermore, we require that these paths have the same *trace*. We then show that such paths guarantee reaching a neighbourhood of z^* with "uniform" probability yielding the measure κ and the constant $\eta > 0$. Let us define the notions.

Definition 5.1.12. *For $\delta' > 0$, we call a path $(s_0, \xi_0, t_0) \cdots (s_p, \xi_p, t_p)$ δ'-*wide *if (1) all its configurations are δ'-separated and inner and (2) no bounded variable-delay event gets fired δ'-close to its upper bound, i.e. for any $0 \leq i < p$ and bounded $e \in \mathbf{E}(s_i)$ we have $\xi_i(e) + t_{i+1} < u_e - \delta'$.*

We say that a path $(s_0, \xi_0, t_0) \cdots (s_p, \xi_p, t_p)$ has a trace *$r_0 E_1 r_1 E_1 \cdots E_p r_p$ if for every $0 \leq i \leq p$ we have $s_i \in r_i$ and for every $0 < i \leq p$ we can get from $(s_{i-1}, \xi_{i-1}, t_{i-1})$ to (s_i, ξ_i, t_i) via occurrence of the set of events E_i after time t_i.*

Further, *total time* of a path $(s_0, \xi_0, t_0) \cdots (s_p, \xi_p, t_p)$ is $\sum_{i=1}^{p} t_i$ and by M we denote the sum of u_e of all fixed-delay events. We first show the existence of δ-wide paths.

Lemma 5.1.13. *There is $p \in \mathbb{N}$, $\delta' > 0$, trace $T = r_0 E_1 \cdots E_p r_p$, and a configuration z^* such that from any $z \in C$, there is a δ'-wide path to z^* with trace T and total time $t \geq M$.*

Proof. In the proofs of this subsection, by saying *value* of an event e, we again mean the fractional part $\langle \xi(e) \rangle$ when the vector of elapsed time ξ is clear from context. We use a similar concept as in the proof of Lemma 5.1.6. Let us fix a $z \in C$ and recall that z is δ-separated. Let a be the greatest value of all event scheduled in z. Observe, that no value is in the interval $(a, a + \delta)$. When we build the δ'-wide path step by step, we use a variable s denoting start of this interval of interest which flows with time. Before the first step, we have $s := a$. After each step, which takes t time, we set $s := \langle s + t \rangle$.

In the interval $[s, s + \delta]$ we make a grid of $2 \cdot |\mathscr{E}|$ points that we shift along with s, and set $\delta' = \delta / (2 \cdot |\mathscr{E}|)$. On this grid, a procedure similar to the δ-separation takes place. We build the δ'-wide path by choosing sets of events E_i to occur, waiting times t_i of the individual transitions, and target states z_i after each transition so that

- every variable-delay event occurs at a soonest possible moment such that it is exactly at an empty point of the grid (i.e. at a time when an empty point has value 0), and
- the built path is "feasible", i.e. all the specified events can occur after the specified waiting time, and upon each occurrence of a specified event we move to the specified target state with positive probability,

These rules guarantee that the path we create is δ'-wide. Indeed, the initial configuration is inner and δ'-separated since $\delta' < \delta$; upon every new transition, the δ'-neighbourhood of 0 is empty; and every variable-delay event occurs at a point different from its current point, whence it occurs at least δ' prior to its upper bound. Furthermore, no event occurs after a delay greater than $2r+B+1$ because it could occur also at least one time unit sooner (contradicting the first rule). It is easy to see that such choices are possible since there are only $\mathscr{E}$ events, but $2\cdot|\mathscr{E}|$ points (thus in the course of every time unit there are enough empty points for all the events to occur).

Now we show that this procedure lasts only a fixed amount of steps before all the scheduled events lie on the grid. Notice that if there is a set of ticking events in r, their value lies at a point of the grid from the very beginning because we define the grid adjacent to their value. Value of any other scheduled fixed-delay event gets eventually placed at a point of a grid. Indeed, every such event is scheduled by a variable-delay event after traversing one cycle, i.e. after $|V|$ steps, since we assume a single-ticking GSMP. In total, after $p := |V|+1$ steps with trace $r_0E_1\cdots E_pr_p$, we can set $z^* := z_p$.

It remains to show that from any other δ'-separated configuration $z' \in r$, we can build a δ'-wide path of length p, with trace $r_0E_1\cdots E_pr_p$ that ends in z^*. We start in the same region. From the definition of the region relation and from the fact that all events occur in the empty interval $(a, a+\delta)$ we get the following. By appropriately adjusting the waiting times so that the events occur at the same points of the grid as before, we can follow the same trace going through the same regions and build a path $z'_0\ldots z'_p$ such that $z'_p = z^*$. Indeed, all scheduled events have the same value in z'_p as in z_p because they lie on the same points of the grid. In fact, this holds for z'_{p-1} and z_{p-1} as well (because in the first $p-1$ steps at least one cycle is traversed) but $t'_{p-1} \neq t_{p-1}$. Finally, also t'_p and t_p have the same value in z'_p as in z_p because there is no need to alter the waiting time in the last step. By the same arguments as before, the built path is also δ'-wide.

□

Now we show how δ'-wide paths guarantee reaching the neighbourhood of z^* with "uniform" probability. Hence, the property (II) is directly yielded by connecting Lemmata 5.1.13 and 5.1.14 concluding the proof of Lemma 5.1.10.

Lemma 5.1.14. *Let $\delta' > 0$, $p \in \mathbb{N}$, $T = r_0E_1\cdots E_pr_p$ be a trace, and (s_p, ξ_p, t_p) be a configuration. There is a probability measure κ and $\eta > 0$ such that the probability satisfies $P^p((s_0,\xi_0,t_0),Y) \geq \eta\cdot\kappa(Y)$ for any $Y \in \mathscr{F}$ and any (s_0,ξ_0,t_0) such that there is a δ'-wide path $(s_0,\xi_0,t_0)\cdots(s_p,\xi_p,t_p)$ with trace T and total time $t \geq M$.*

Proof. Notice that all delays' densities are bounded from below by some $c_{\mathfrak{D}} > 0$ on the interval $[0,2r]$. Since we assume inner paths, we can make use of this bound.

We will find a set of configurations Z "around" the state $z_p = (s_p,\xi_p,t_p)$ and define the probability measure κ on this set Z such that $\kappa(Z) = 1$. Then we show for each measurable $Y \subseteq Z$ the desired property.

Intuitively, configurations around z_p are of the form (s_p, ξ'_p, t'_p) where each $\xi'_p(e)$ is either exactly $\xi_p(e)$ or in a small interval around $\xi_p(e)$ (and the same for t'_p and t_p). We now discuss which case applies to which event e for a fixed trace T. All the following notions are defined with respect to T. Let $G_0, \dots, G_p$ be the sets of events such that each G_i is either ticking in r_i or empty. Furthermore, we require that G_i is empty iff either there are no ticking events in r_i or the ticking events "died" earlier, i.e. some G_j is empty for $j < i$. We say that an event $e \in \mathbf{E}(s_p)$ is *originally scheduled in the i-th step by f* if

- either $f \in G_{i-1}$ or f is a variable-delay event; and
- there is $k \geq 1$ and a chain of events $e_1 \in E_{c_1}, \dots, e_k \in E_{c_k}$ such that
  - $e_1 = f$, $c_1 = i$ and for $2 \leq i \leq k$ we have that each e_i is a fixed-delay event such that $e_i \notin G_i$;
  - for $1 \leq i < k$ we have that occurrence of each E_{c_i} newly schedules e_{i+1} and occurrence of E_{c_k} newly schedules e; and
  - e stays scheduled since, i.e. $e \in \mathbf{E}(s_{c_k}) \cap \dots \cap \mathbf{E}(s_{p-2}) \cap \mathbf{E}(s_{p-1})$.

Notice that the length of the last step t_p is also a part of the state space. Therefore, it is important to study how we can affect it. We say that *the last step is variable* if E_p is either a singleton of a variable-delay event or all the events in E_p are originally scheduled by a variable-delay event. We say that *the last step is fixed* otherwise, i.e. if $E_p \subseteq G_p$ or all events from E_p are originally scheduled by some $e \in G_i$ for some $i \in \mathbb{N}_0$.

Intuitively, we cannot alter the value of an event e on the trace T (i.e., $\xi'_p(e) = \xi_p(e)$) if the last step is fixed and e is originally scheduled by the ticking event. In all other cases, the value of e can be altered such that $\xi'_p(e)$ lies in a small interval around $\xi_p(e)$. The rest of the proof is divided in two cases.

The last step is fixed We divide the events $\mathbf{E}(s_p)$ into sets A, B, and C as follows

$e \in A$	if e is originally scheduled by a variable-delay event and $\langle \xi_p(e) \rangle \neq 0$;
$e \in B$	if e is originally scheduled by a variable-delay event and $\langle \xi_p(e) \rangle = 0$;
$e \in C$	if e is originally scheduled by a ticking event from some G_i.

Let $a_1, \dots, a_d$ be the distinct fractional values of the events A in the vector of elapsed time ξ_p ordered increasingly by the step in which the corresponding events were originally scheduled. This definition is correct because two events with the same fractional value must be originally scheduled by the same event in the same step. Furthermore, let $F_1, \dots, F_d$ be the corresponding sets of events, i.e. $\langle \xi_p(e_i) \rangle =$

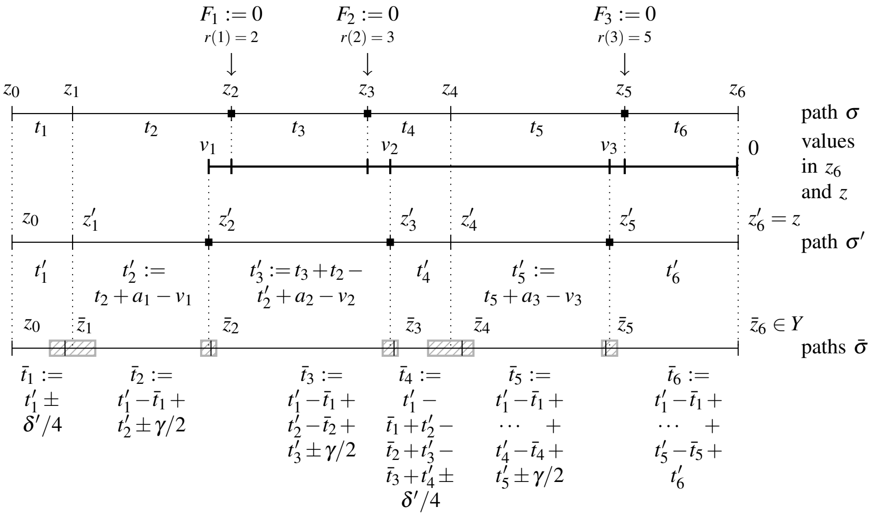
$F_1 := 0$
$r(1) = 2$
$F_2 := 0$
$r(2) = 3$
$F_3 := 0$
$r(3) = 5$
z_0
z_1
z_2
z_3
z_4
z_5
z_6
path σ
t_1
t_2
t_3
t_4
t_5
t_6
v_1
v_2
v_3
0
values in z_6 and z
z_0
z'_1
z'_2
z'_3
z'_4
z'_5
$z'_6 = z$
path σ'
t'_1
$t'_2 :=$ $t_2 + a_1 - v_1$
$t'_3 := t_3 + t_2 -$ $t'_2 + a_2 - v_2$
t'_4
$t'_5 :=$ $t_5 + a_3 - v_3$
t'_6
z_0
$\bar{z}_1$
$\bar{z}_2$
$\bar{z}_3$
$\bar{z}_4$
$\bar{z}_5$
$\bar{z}_6 \in Y$
paths $\bar{\sigma}$
$\bar{t}_1 :=$ $t'_1 \pm$ $\delta'/4$
$\bar{t}_2 :=$ $t'_1 - \bar{t}_1 +$ $t'_2 \pm \gamma/2$
$\bar{t}_3 :=$ $t'_1 - \bar{t}_1 +$ $t'_2 - \bar{t}_2 +$ $t'_3 \pm \gamma/2$
$\bar{t}_4 :=$ $t'_1 -$ $\bar{t}_1 + t'_2 -$ $\bar{t}_2 + t'_3 -$ $\bar{t}_3 + t'_4 \pm$ $\delta'/4$
$\bar{t}_5 :=$ $t'_1 - \bar{t}_1 +$ $\cdots +$ $t'_4 - \bar{t}_4 +$ $t'_5 \pm \gamma/2$
$\bar{t}_6 :=$ $t'_1 - \bar{t}_1 +$ $\cdots +$ $t'_5 - \bar{t}_5 +$ t'_6

Figure 5.2: Illustration of paths leading to the set Y. In the original path σ on the top, there are marked the times when the events F_1, F_2, and F_3 get originally scheduled. This path is in the first phase altered to σ' that reaches the target state z (its values v_1,v_2, and v_3 are depicted between σ and σ' by their distance to 0 on the right). In the second phase, a set of paths that reach Y is constructed by allowing imprecision in the waiting times – the transition times are randomly chosen inside the hatched areas. Notice that the random choice within the large intervals does not influence the values in $\bar{z}_6$. The values are only influenced by the choice in the smaller intervals; the size of the smaller intervals is $\gamma/2$, i.e. depends on the size of the d-dimensional hypercube Y. Hence, to get a probability bound linear with respect to $\kappa(Y)$, at most d smaller intervals can be used. Transitions with fixed delay are omitted from the illustration (except for the last transition).

a_i for any $e_i \in F_i$. We call a configuration z a *target* configuration if $z \sim z_p$ and all events $e \in (B \cup C)$ have the same value in z and z_p. We treat a target configuration as a d-dimensional vector describing the distinct values for the sets $F_1, \ldots, F_d$. A δ'-neighbourhood of a target configuration z is the set of configurations $\{z + C \mid C \in (-\delta', \delta')^d\}$. Observe that the δ'-neighbourhood is a d-dimensional space. We set Z to be the $(\delta'/4)$-neighbourhood of z_p (the reason for dividing δ' by 4 is technical and will become clear in the course of this proof). Let κ_d denote the standard Lebesgue measure on the d-dimensional affine space and set $\kappa(Y) := \kappa_d(Y)/\kappa_d(Z)$ for any measurable $Y \subseteq Z$.

In order to prove the probability bound for any measurable $Y \subseteq Z$, it suffices to prove it for the generators of Z, i.e. for d-dimensional hypercubes centred around some state in Z. Let us fix an arbitrary $z \in Z$ and $\gamma < \delta'/4$. We set Y to be the γ-neighbourhood of z. In the rest of the proof we will show how to reach the set Y from the initial state (s_0, ξ_0, t_0) in p steps with high enough probability – linear in $\kappa(Y)$.

We show it by altering the original δ'-wide path $\sigma = (s_0, \xi_0, t_0) \cdots (s_p, \xi_p, t_p)$. In the first phase, we reach the fixed z instead of the configuration z_p. We find waiting times $t'_1, \ldots, t'_p$ that induce a path $\sigma' = (s_0, \xi_0, t_0)\ (s_1, \xi'_1, t'_1)\ \ldots (s_p, \xi'_p, t'_p)$ with trace T such that $(s_p, \xi'_p, t'_p) = z$. In the second phase, we define using σ' a set of paths to Y. We allow for intervals $I_1, \ldots, I_p$ such that for any choice $\bar{t}_1 \in I_1, \ldots, \bar{t}_p \in I_p$ we get a path $\bar{\sigma} = (s_0, \xi_0, t_0)\ (s_1, \bar{\xi}_1, \bar{t}_1)\ \ldots (s_p, \bar{\xi}_p, \bar{t}_p)$ such that $(s_p, \bar{\xi}_p, \bar{t}_p) \in Y$. From the size of the intervals for variable-delay events and from the bound on densities $c_{\mathfrak{D}}$ we get the overall bound on probabilities. Let us start with the first step.

Let $v_1, \ldots, v_d$ be the distinct values of the target configuration z. Recall that $|v_i - a_i| < \delta'/4$ for each i. Let $r(1), \ldots, r(d)$ be the indices such that all events in F_i are originally scheduled in the step $r(i)$. Notice that each $E_{r(i)}$ is a singleton of a variable-delay event. As illustrated in Figure 5.2, we set for each $1 \leq i \leq p$

$$t'_i = \begin{cases} \ell_e - \xi_{i-1}(e) & \text{if } e \in E_i \text{ is fixed-delay,} \\ t_i + \sum_{k=1}^{i-1}(t_k - t'_k) + a_j - v_j & \text{if } i = r(j) \text{ for } 1 \leq j \leq d, \\ t_i + \sum_{k=1}^{i-1}(t_k - t'_k) & \text{otherwise.} \end{cases}$$

Intuitively, we adjust the variable-delays in the steps preceding the original scheduling of sets $F_1, \ldots, F_d$ whereas the remaining variable-delay steps are kept in sync with the original path σ. The absolute time of any transition in σ' (i.e. the position of a line depicting a configuration in Figure 5.2) is not shifted by more than $\delta'/4$ since $|v_i - a_i| < \delta'/4$ for any i. Thus, the difference of any two absolute times is not changed by more than $\delta'/2$. This difference bounds the difference of $|\xi_i(e) - \xi'_i(e)|$ for any i and $e \in \mathscr{E}$. Hence, σ' is $(\delta'/2)$-wide because σ is δ'-wide. Furthermore, σ' goes through the same regions as σ and performs the same sequence of events scheduling. Building on that, the desired property $z'_p = z$ is easy to see.

Next we allow imprecision in the waiting times of σ' so that we get a set of paths of measure linear in $\kappa_d(Y) = \gamma^d$. In each step we compensate for the imprecision of the previous step. Formally, let T_i denote $t'_i + \sum_{k=1}^{i-1}(t'_k - \bar{t}_k)$. For each $1 \leq i \leq p$ we constraint

$$\bar{t}_i \in \begin{cases} [T_i, T_i] & \text{if } E_i \text{ are fixed-delay events,} \\ (T_i - \frac{\gamma}{2},\ T_i + \frac{\gamma}{2}) & \text{if } i = r(j) \text{ for } 1 \leq j \leq d, \\ (T_i - \frac{\delta'}{4},\ T_i + \frac{\delta'}{4}) & \text{otherwise.} \end{cases}$$

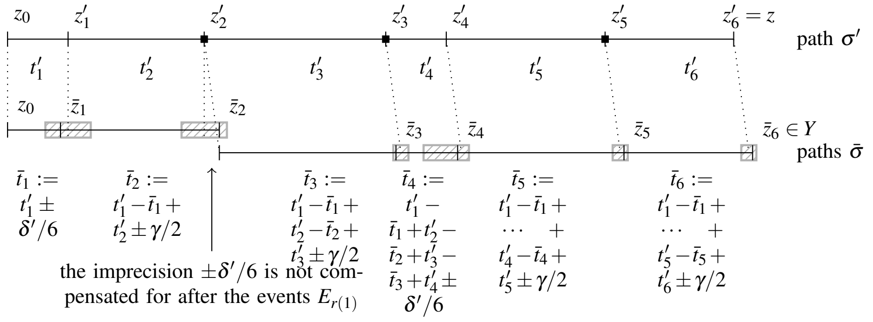
z_0
z'_1
z'_2
z'_3
z'_4
z'_5
$z'_6 = z$
path σ'
t'_1
t'_2
t'_3
t'_4
t'_5
t'_6
z_0
$\bar{z}_1$
$\bar{z}_2$
$\bar{z}_3$
$\bar{z}_4$
$\bar{z}_5$
$\bar{z}_6 \in Y$
paths $\bar{\sigma}$
$\bar{t}_1 := t'_1 \pm \delta'/6$
$\bar{t}_2 := t'_1 - \bar{t}_1 + t'_2 \pm \gamma/2$
$\bar{t}_3 := t'_1 - \bar{t}_1 + t'_2 - \bar{t}_2 + t'_3 \pm \gamma/2$
$\bar{t}_4 := t'_1 - \bar{t}_1 + t'_2 - \bar{t}_2 + t'_3 - \bar{t}_3 + t'_4 \pm \delta'/6$
$\bar{t}_5 := t'_1 - \bar{t}_1 + \cdots + t'_4 - \bar{t}_4 + t'_5 \pm \gamma/2$
$\bar{t}_6 := t'_1 - \bar{t}_1 + \cdots + t'_5 - \bar{t}_5 + t'_6 \pm \gamma/2$
the imprecision $\pm\delta'/6$ is not compensated for after the events $E_{r(1)}$

Figure 5.3: Illustration of construction of $\bar{\sigma}$ for the empty set C and the last step variable.

The difference to σ' of any two absolute times is not changed by more than $\delta'/2$ because the imprecision of any step is bounded by $\delta'/4$. Because σ' is $(\delta'/2)$-wide, any path $\bar{\sigma}$ goes through the same regions as σ'. The difference of the value of events in any F_i in the state $\bar{z}_p$ from the state z is at most $\gamma/2$ because it is only influenced by the imprecision of the step preceding its original scheduling. Hence, $\bar{z}_p \in Y$.

By v we denote the number of variable-delay singletons among $E_1, \dots, E_p$. From the definition of P, it is easy to prove by that

$$P^p(z_0, Y) \quad \geq \quad p_{\min}^p \cdot (c_{\mathfrak{D}} \cdot \gamma)^d \cdot (c_{\mathfrak{D}} \cdot \delta'/2)^{v-d} \quad \geq \quad (p_{\min} \cdot c_{\mathfrak{D}}/2)^p \cdot \gamma^d \cdot \delta'^{p-d}$$

Since $\kappa_d(Y) = (2 \cdot \gamma)^d$ and $\kappa_d(Z) = (2 \cdot \delta'/4)^d$, we have $\kappa(Y) = \kappa_d(Y)/\kappa_d(Z) = (4\gamma/\delta')^d$. We get $P^p(z_0, Y) \geq \kappa(Y) \cdot (\delta' \cdot p_{\min} \cdot c_{\mathfrak{D}}/8)^p$ and setting $\eta = (\delta' \cdot p_{\min} \cdot c_{\mathfrak{D}}/8)^n$ proves the lemma for the case of last step being fixed.

The last step is variable The rest of the proof proceeds in a similar fashion as previously, we reuse the same notions and the same notation. We only redefine the differences: the neighbourhood and the way the paths are altered.

We call $(s, \xi, t) \sim z_p$ a *target* configuration if there is $y \in \mathbb{R}$ such that for all events $e \in C$ we have $\xi(e) - \xi_p(e) = y$ and for all events $e \in B$ we have $\xi(e) = \xi_p(e)$. We treat a target configuration as a g-dimensional vector where $g = d + 1$ if C is non-empty, and $g = d$, otherwise. This vector describes the distinct values for the sets $F_1, \dots, F_d$ and the value y, if necessary. Again, a δ'-neighbourhood of a target configuration z is the set of configuration $\{z + C \mid C \in (-\delta', \delta')^g\}$. We set Z to be the $(\delta'/4)$-neighbourhood of z_p and set $\kappa(Y) := \kappa_g(Y)/\kappa_g(Z)$ for any measurable $Y \subseteq Z$. We fix Y to be a γ-neighbourhood of a fixed $z \in Z$.

The path σ' is obtained from σ in the same way as before. We need to allow imprecision in the waiting times of σ' so that we get a set of paths of measure linear in γ^g.

- For the case $g = d + 1$ it is straightforward as we make the last step also with imprecision $\pm\gamma/2$. We constraint

$$\bar{t}_i \in \begin{cases} [T_i, T_i] & \text{if } E_i \text{ are fixed-delay events,} \\ (T_i - \frac{\gamma}{2},\ T_i + \frac{\gamma}{2}) & \text{if } i = r(j) \text{ for } 1 \leq j \leq d \text{ or } i = m, \\ (T_i - \frac{\delta'}{4},\ T_i + \frac{\delta'}{4}) & \text{otherwise} \end{cases}$$

  where m equals p if E_p contains a variable-delay event and m is the step where E_p were originally scheduled if E_p are fixed-delay events. The difference of the value of events in any F_i in the state $\bar{z}_p$ from the state z is at most γ because it is influenced by the imprecision of the step preceding its original scheduling and also by the imprecision of the last step. Events in C have the difference of the value at most $\gamma/2$ because of the last step. Hence, $\bar{z}_p \in Y$. Again, we get that $P^p(z_0, Y) \ \geq\ \kappa(Y) \cdot (\delta' \cdot p_{\min} \cdot c_{\mathfrak{D}}/8)^p$ and setting $\eta = (\delta' \cdot p_{\min} \cdot c_{\mathfrak{D}}/8)^n$ proves the lemma for the case of the last step being variable and for $g = d + 1$.

- For the case $g = d$ it is somewhat tricky since only at most d choices of waiting times can have their precision dependent on γ. In each step we compensate for the imprecision of the previous step. Only the imprecision of the step preceding the first scheduling E_1 is not compensated for. Otherwise, it would influence the value of events E_1 in $\bar{z}_p$. Let T_i^a denote $t'_i + \sum_{k=a}^{i-1}(t'_k - \bar{t}_k)$. As illustrated in Figure 5.3, we constraint

$$\bar{t}_i \in \begin{cases} [T_i^1, T_i^1] & \text{if } E_i \text{ are fixed-delay events,} \\ (T_i^1 - \frac{\delta'}{6},\ T_i^1 + \frac{\delta'}{6}) & \text{if } i \leq r(1), \\ (T_i^{r(1)+1} - \frac{\gamma}{2},\ T_i^{r(1)+1} + \frac{\gamma}{2}) & \text{if } i = r(j) \text{ for } 2 \leq j \leq d \text{ or } i = m, \\ (T_i^{r(1)+1} - \frac{\delta'}{6},\ T_i^{r(1)+1} + \frac{\delta'}{6}) & \text{otherwise.} \end{cases}$$

  The difference to σ' of any two absolute times is not changed by more than $3 \cdot \delta'/6 = \delta'/2$ because the imprecision of any step is bounded by $\delta'/6$. Because σ' is $(\delta'/2)$-wide, any path $\bar{\sigma}$ goes through the same regions as σ'. The difference of the value of events E_1 in the state $\bar{z}_p$ from the state z is at most $\gamma/2$ because it is only influenced by the imprecision of the last step. The difference of any other event e is at most $2 \cdot \gamma/2$ because it is influenced by the imprecision of the step preceding the original scheduling of e, as well. Hence, $\bar{z}_p \in Y$.

Now, we get that $P^p(z_0, Y) \geq \kappa(Y) \cdot (\delta' \cdot p_{\min} \cdot c_{\mathfrak{D}}/12)^p$ and setting $\eta = (\delta' \cdot p_{\min} \cdot c_{\mathfrak{D}}/12)^p$ proves the lemma for the remaining case of the last step being variable and for $g = d$. □

5.2 GSMP with Timed Automata Objectives

After dealing with GSMP with fixed delay events, we move on to another way to incorporate strict real-time features into GSMP – the timed automata objectives. We show that any DTA observer can be mimicked within the GSMP using fixed-delay events. Furthermore, the resulting GSMP is single-ticking yielding the stability.

Theorem 5.2.1. *Let $\mathscr{M}$ be a single-ticking GSMP, $\mathscr{A}$ be its DTA observer, and q be a location of $\mathscr{A}$. The random variables $\mathbf{d}_q^{\mathscr{A}}$ and $\mathbf{c}_q^{\mathscr{A}}$ are well-defined for almost every run. Furthermore, a single-ticking GSMP $\mathscr{M} \times \mathscr{A}$ can be constructed in time polynomial in $|\mathscr{M}|$ and exponential in $|\mathscr{A}|$ such that $\mathbf{d}_q^{\mathscr{A}}$ and $\mathbf{c}_q^{\mathscr{A}}$ can be expressed by the vectors $\mathbf{d}^{\mathscr{M}\times\mathscr{A}}$ and $\mathbf{c}^{\mathscr{M}\times\mathscr{A}}$. Namely, for any BSCC B of the region graph of $\mathscr{M} \times \mathscr{A}$ with the vectors of frequencies d and c there are sets of states S_d and S_q of $\mathscr{M} \times \mathscr{A}$ such that*

$$\mathbf{d}_q^{\mathscr{A}} = \frac{\sum_{s \in S_d \cap S_q} d_s}{\sum_{s \in S_d} d_s} \quad \textit{and} \quad \mathbf{c}_q^{\mathscr{A}} = \sum_{s \in S_q} c_s \quad \textit{with probability that } \mathscr{M} \times \mathscr{A} \textit{ reaches } B.$$

Notice that the reduction holds even if we start with a GSMP already containing fixed-delay events – if it is single-ticking. The rest of this section forms the proof of Theorem 5.2.1. Let us fix a single-ticking GSMP $\mathscr{M} = (S, \mathscr{E}, \mathbf{E}, \mathrm{Succ}, \alpha_0)$, its DTA observer $\mathscr{A} = (Q, \Sigma, \mathscr{X}, \longrightarrow, q_0)$, and $q \in Q$.

The product GSMP needs to encode the behaviour of $\mathscr{A}$. First, for each clock x of $\mathscr{A}$ and for each $1 \leq i \leq B_{\max}$, we enhance the set of events with a fixed-delay event $e_{x,i}$ where $e_{x,i}$ occurs when the clock x reaches the integral value i. Second, we enrich the state space with the regions of $\mathscr{A}$; we update this component using the newly added fixed-delay events. In this construction, we do not need to care about the "thin" regions with some clock having integral value. More precisely, we add into the state space the set $\mathfrak{R}$ of regions of $\mathscr{A}$ where no clock has integral value. Furthermore, for a region $r \notin \mathfrak{R}$, we denote by $\bar{r}$ the first region from $\mathfrak{R}$ reached from r by flow of time. For a region $r \in \mathfrak{R}$, $\bar{r}$ denotes r itself. Third, as $\mathscr{M} \times \mathscr{A}$ makes artificial steps when updating the current region of $\mathscr{A}$, we also store in the state space whether the last step was a step of $\mathscr{M}$ so that we can express $\mathbf{d}_q^{\mathscr{A}}$ using $\mathbf{d}^{\mathscr{M}\times\mathscr{A}}$. Let us define the product precisely.

Definition 5.2.2. *We set $\mathscr{M} \times \mathscr{A} = (S \times \mathfrak{R} \times \{\mathsf{y}, \mathsf{n}\}, \mathscr{E} \cup \mathscr{E}', \mathbf{E}', \mathrm{Succ}', \alpha_0')$ where*

- *$\mathscr{E}' = \{e_{x,1}, \ldots, e_{x,B_{\max}} \mid x \in \mathscr{X}\}$ where each $e_{x,i}$ has fixed delay 1;*

- $\mathbf{E}'((s,r,b)) = \mathbf{E}(s) \cup \{e_{x,i} \mid x \in \mathscr{X}$, *$x$ satisfies $i-1 < x < i$ in the region r*$\}$;
- $\mathrm{Succ}'((s,r,b),E)$ *is determined as follows. If $E \cap \mathscr{E}' \neq \emptyset$, let r' denote the region reached from r be elapsing the amount of time such that exactly the clocks from $E \cap \mathscr{E}'$ reach the next integral value. Otherwise, let $r' = r$. Then*
  - *if $E \cap \mathscr{E} = \emptyset$,* $\mathrm{Succ}'((s,r,b),E)$ *assigns* 1 *to the state* $(s,\overline{r'},\mathsf{n})$,
  - *if $E \cap \mathscr{E} \neq \emptyset$,* $\mathrm{Succ}'((s,r,b),E)$ *assigns* $\mathrm{Succ}(s,E\cap\mathscr{E})(s')$ *to each state* $(s',\overline{r''},\mathsf{y})$ *for $s' \in S$ where r'' is the region that $\mathscr{A}$ reaches after reading s' in r';*
- $\alpha_0'(s,r_0,1) = \alpha_0(s)$ *for $s \in S$ where r_0 is the region containing* $(q_0,\mathbf{0})$.

Lastly, $S_d = \{(s,r,b) \mid b = \mathsf{y}\}$ *and* $S_q = \{(s,r,b) \mid r$ *has location* $q\}$ *for any* $q \in Q$.

First, we show that $\mathscr{M} \times \mathscr{A}$ is single-ticking, which yields together with Theorem 5.1.2 the first part of Theorem 5.2.1.

Lemma 5.2.3. *The GSMP $\mathscr{M} \times \mathscr{A}$ is single-ticking.*

Proof. Let r be a region of $\mathscr{M} \times \mathscr{A}$. We show that for each set E of ticking events in r, the set $E \cap \mathscr{E}$ is non-empty and ticking in $\mathscr{M}$. This suffices as if there were two sets E_1, E_2 of ticking events, they would be by the maximality in the definition disjoint and thus, the sets $E_1 \cap \mathscr{E}$ and $E_2 \cap \mathscr{E}$ would be disjoint and ticking in $\mathscr{M}$; $\mathscr{M}$ would not be single-ticking. Let us fix a set E of ticking events in r. There is a cycle $r_1, \ldots, r_n$ in the region graph and sets of fixed-delay events $E_1, \ldots, E_n$ that schedule each other.

We first show that for each E_i we have $E_i \cap \mathscr{E} \neq \emptyset$. There must be some j such that $E_j \cap \mathscr{E} \neq \emptyset$. Namely, if all the sets contained only events of $\mathscr{E}'$ modelling clocks of $\mathscr{A}$, we would not obtain a cycle because the clocks of $\mathscr{A}$ can get restarted only by events of $\mathscr{E}$. Observe that each event of $E_j \cap \mathscr{E}$ can be again scheduled only by an occurrence of an event from $\mathscr{E}$. As we have a cycle, all the sets E_i must thus satisfy $E_i \cap \mathscr{E} \neq \emptyset$.

Now we show that the set of events $E \cap \mathscr{E}$ is ticking in the region $\pi(r)$ where the function π maps each region of $\mathscr{M} \times \mathscr{A}$ to its corresponding region of $\mathscr{M}$. Indeed, there is a cycle of regions $k_1, \ldots, k_m$ obtained from $\pi(r_1), \ldots, \pi(r_n)$ by removing repeating occurrences of each region and sets of events $F_1, \ldots, F_m$ obtained from $E_1 \cap \mathscr{E}, \ldots, E_k \cap \mathscr{E}$ by removing the corresponding sets. Each F_i is a maximal set of events with the same fractional value in k_i, otherwise the original set of events E_j is not maximal in r_j. Similarly, each event scheduled to occur in $\pi(r)$, occurs on the cycle $k_1, \ldots, k_m$, otherwise it does not occur on $r_1 \ldots, r_n$. From the definition of $\mathscr{M} \times \mathscr{A}$, it is easy to see that the events $F_1 \ldots, F_m$ schedule each other in the sense of the definition of ticking. In fact, an occurrence of events stemming from $\mathscr{A}$ does not have any influence on the $\mathscr{M}$-part of the product. □

The rest of Theorem 5.2.1 follows again from Theorem 5.1.2 and from observing how $\mathscr{M}$ with $\mathscr{A}$ and $\mathscr{M} \times \mathscr{A}$ correspond on individual runs:

Lemma 5.2.4. *There is a function ρ mapping runs of $\mathscr{M}$ to runs of $\mathscr{M} \times \mathscr{A}$ that preserves the measure and for each run σ it satisfies,*

$$\mathbf{d}_q^{\mathscr{A}}(\sigma) = \frac{\sum_{s \in S_1 \cap S_q} \mathbf{d}_s^{\mathscr{M}\times\mathscr{A}}(\rho(\sigma))}{\sum_{s \in S_1} \mathbf{d}_s^{\mathscr{M}\times\mathscr{A}}(\rho(\sigma))}, \tag{5.2}$$

$$\mathbf{c}_q^{\mathscr{A}}(\sigma) = \sum_{s \in S_q} \mathbf{c}_s^{\mathscr{M}\times\mathscr{A}}(\rho(\sigma)). \tag{5.3}$$

Proof. Intuitively, ρ only adds the deterministic behaviour of $\mathscr{A}$ as encoded into $\mathscr{M} \times \mathscr{A}$. Formally, let $\sigma = (s_0, \xi_0, t_0)(s_1, \xi_1, t_1)\cdots$ be a run of $\mathscr{M}$ which is observed by $\mathscr{A}$ in the run $(q, \nu)s_0(q_0, \nu_0)t_1(q_1', \nu_1')s_1(q_1, \nu_1)\cdots$. The mapping ρ injects in the run intermediate steps whenever $\mathscr{A}$ enters another region. Let us describe it for a fixed $i \in \mathbb{N}_0$. We denote by $r_{i,1}, \cdots, r_{i,k(i)}$ the regions from $\mathfrak{R}$ traversed when waiting for time t_{i+1} in (q_i, ν_i). Further, for $1 \le \ell \le k(i)$, let $t_{i,\ell}$ denote the time spent in $r_{i,\ell}$, $(q_{i,\ell}, \nu_{i,\ell})$ be the configuration of $\mathscr{A}$ when entering $r_{i,\ell}$, and $\xi_{i,\ell}$ be the configuration of $\mathscr{M} \times \mathscr{A}$ obtained from ξ_i and the fractional values from $\nu_{i,\ell}$. We set $\rho(\sigma) = z_0\, \sigma_0\, z_1\, \sigma_1 \cdots$ where

- $z_0 = ((s_0, r_{0,1}, \mathsf{y}), \xi_{0,1}, t_0)$ and $z_i = ((s_i, r_{i,1}, \mathsf{y}), \xi_{i,1}, t_{i-1,k(i-1)})$ for $i > 0$; and
- $\sigma_i = ((s_i, r_{i,2}, \mathsf{n}), \xi_{i,2}, t_{i,1}) \cdots ((s_i, r_{i,k(i)}, \mathsf{n}), \xi_{i,k(i)}, t_{i,k(i)-1})$.

The preservation of measure is straightforward from the determinism both in the definition of $\mathscr{M} \times \mathscr{A}$ as well as the definition of ρ. As regards $\mathbf{d}_q^{\mathscr{A}}$, observe that only the states of the configurations $z_0, z_1, \ldots$ belong to S_d. Hence, to each transition in σ, there is in $\rho(\sigma)$ exactly one transition into a state from S_d. Furthermore, between two visits to y-states, the number of steps (to n-states) is bounded by $|\mathscr{X}| \cdot B_{\max}$; the frequency $\sum_{s \in S_d} \mathbf{d}_s^{\mathscr{M}\times\mathscr{A}}(\rho(\sigma))$ is greater than 0. Therefore, the ratio of visits to S_q out of the visits to S_d equals the ratio of visits to the location q. As regards $\mathbf{c}_q^{\mathscr{A}}$, notice that the transitions to n-states do not change whether the current state belongs to S_q. Therefore, this may get changed only by the transitions to y-states, corresponding to the transitions in $\mathscr{A}$ triggered by the transitions in $\mathscr{M}$ (thanks to the definition of ρ). Therefore, the ratio of time spent in S_q equals the ratio of time spent in to the location q. □

In the following section, we translate the result to deterministic and stochastic Petri nets.

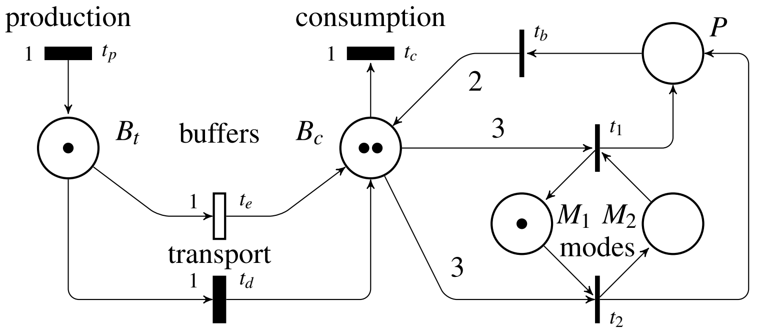

Figure 5.4: DSPN without the time-average limit. Deterministic transitions are denoted by thick bars, exponential transitions are denoted by empty bars, and immediate transitions are denoted by thin bars. The timed frequency of markings with a token in M_1 is not well-defined.

5.3 Deterministic and stochastic Petri nets

As outlined in Section 3.1.1 and formally proven in [Haa10], a formalism of stochastic Petri nets is closely related to GSMP. In this section, we focus on the simplest class of stochastic Petri nets that contain a feature similar to fixed-delay events, namely the deterministic and stochastic Petri nets (DSPN). We use the notation and terminology introduced in Section 3.1.1.

We provide structural conditions upon which a DSPN is stable. To motivate such conditions we first give an example of an unstable DSPN. Notice that the unstable GSMP from Figure 4.2 uses also the uniform distribution. We show that it is not necessary for the instability.

Theorem 5.3.1. *There is a bounded DSPN with two deterministic transitions such that with positive probability its timed frequencies do not exist.*

Proof. In Figure 5.4, we depict a DSPN similar to the GSMP from Figure 4.2. It is a model of a simple producer-consumer system operated by one exponential and three deterministic transitions using buffers B_t and B_c. In addition, on the right, there is a controller with immediate transitions switching a token between modes M_1 and M_2.

Let us explain the behaviour of the model. The transition t_p produces a token exactly every time unit and places it to the transport buffer B_t. As both the exponential transition t_e and the deterministic transition t_d are enabled in parallel, the token is transported into the consumption buffer B_c in *at most* 1 time unit. The consumption of a token by the transition t_c takes again exactly 1 time unit. The consumption buffer B_c can hold at most two tokens. When a third token appears in the buffer, all three tokens are removed by switching the modes. The subsequent immediate transition returns two tokens back, re-enabling t_c. Observe that the modes are switched whenever the current transport takes less time than it has ever taken. The lower the current minimal transport time is, the more time it takes to set a new minimum. Hence, the stays in a single mode get longer and longer.

By similar arguments as in the proof of Theorem 4.2.1, the timed frequency of markings with a token in M_1 is not well-defined for a set of runs with positive measure. Hence, the timed frequencies of the system do not exist. □

Now we turn our attention to structural conditions that guarantee stability. Let us fix a bounded DSPN $\mathscr{N} = (P, T, T', I, H, O, F, \mathbf{p}, m_0)$ with the finite set of markings M. By $M' \subseteq M$ we denote the set of *immediate* markings, i.e. markings with an enabled immediate transition. Furthermore, by $T_d \subseteq T \setminus T'$ we denote the set of deterministic transitions.

Definition 5.3.2. *We say that a transition t* initiates *a transition u if there is a marking $m_0 \in M$ where by firing t (and a sequence of immediate transitions), a sequence of markings $m_1, \ldots, m_k$ is traversed with positive probability such that*

- *$m_i \in M'$ for $1 \leq i < k$ and $m_k \in M \setminus M'$,*
- *u is enabled in m_k and either $u = t$ or u is not enabled in m_i for an $0 \leq i < k$.*

For example, t_d initiates t_c in Figure 5.4, as from the marking $m_0 = (B_t = 1, B_c = 2, M_1 = 1)$ depicted in the figure we visit by firing t_d and immediately t_2 and t_b the markings $m_1 = (B_c = 3, M_1 = 1)$, $m_2 = (P = 1, M_2 = 1)$, and $m_3 = (B_c = 2, M_2 = 1)$ such that t_c is enabled in m_3 but it was not enabled in m_2.

Definition 5.3.3. *We say that a DSPN is* almost-monotone *if there is a strict total order $\prec$ on T_d such that for any transition $u \in T_d$ that initiates a transition $t \in T_d$ we have that either u is minimal w.r.t. $\prec$ or $u \prec t$.*

Note that the DSPN of Figure 5.4 is not almost-monotone as both t_p and t_c initiate themselves. Furthermore, observe that for a given DSPN, is can be easily algorithmically checked whether it is almost-monotone.

Theorem 5.3.4. *For each DSPN $\mathscr{N} = (P, T, T', I, H, O, F, \mathbf{p}, m_0)$ that is almost-monotone and bounded, the timed frequencies almost surely exist.*

Proof. For the proof, we use the construction of a GSMP $\mathscr{M} = (S, \mathscr{E}, \mathbf{E}, \mathrm{Succ}, \alpha_0)$ from [Haa10, Theorem 4.6] that is proven to simulate $\mathscr{N}$ in the sense of Definition 3.1.3. Then we show that $\mathscr{M}$ is single-ticking. Hence all its timed frequencies are almost surely well-defined. Thanks to the definition of simulation, also $\mathscr{N}$ has its timed frequencies almost-surely well-defined.

Let us restate the construction of $\mathscr{M}$.[3] Let $\overline{T} = T \setminus T'$ and $\overline{M} = M \setminus M'$ denote the timed transitions and timed markings of $\mathscr{N}$. We set

- $S = \{(m,u) \mid m \in \overline{M}, u : \overline{T} \to \{\times, e, f\}, u(t) = \times \text{ for each } t \text{ not enabled in } m\}$;
- $\mathscr{E} = \bigcup_{t \in \overline{T}} \{e_t, f_t\}$ contains two events for each timed transition of $\mathscr{N}$;
- $\mathbf{E}((m,u)) = \{e_t \mid t \in \overline{T}, u(t) = e\} \cup \{f_t \mid t \in \overline{T}, u(t) = f\}$;
- $\mathrm{Succ}((m,u),E)$ assigns to (m',u') the probability
$$\sum_{m_1,\ldots,m_k} \mathbf{p}(m_1 - m, E) \cdot \mathbf{p}(m_2 - m_1, \mathbf{E}(m_1) \cap T') \cdot \cdots \\ \cdot \mathbf{p}(m_k - m_{k-1}, \mathbf{E}(m_{k-1}) \cap T'),$$
where $E(m)$ denotes the set of transitions enabled in m and the sum is over all sequences of immediate markings ending with the timed marking $m_k = m'$ that are consistent with u and u': for any transition $t \in \overline{T}$ such that $u(t), u'(t) \in \{e, f\}$ and $u(t) \neq u'(t)$, t was *not* enabled is some immediate marking m_i of the sequence. For each transition $t' \in \overline{T}$ that is newly enabled in m', i.e. $u(t) = \times$ and $u'(t) \neq \times$, we additionally require that $u'(t) = e_t$, otherwise such state obtains zero probability by Succ. Notice that $u(t)$ denotes which of the events $\{e_t, f_t\}$ currently represents t.
- $\alpha_0((m,u)) = 1$ if $m = m_0$ and $u(t) = e_t$ for each transition $t \in \overline{T}$ enabled in m, and $\alpha_0((m,u)) = 0$, otherwise.

The GSMP closely follows the structure of $\mathscr{N}$ but it skips the immediate markings where zero time is spent and takes transitions directly to the first timed marking. Whenever a timed transition t is disabled and enabled again in the intermediate sequence of immediate markings, it is "restarted" in $\mathscr{N}$. The definition of GSMP does not allow restarting an event, we simulate it by switching off one of the events of t and scheduling the other.

Now we show that $\mathscr{M}$ is single-ticking. For a fixed region r in the region graph of $\mathscr{M}$ we show that there is at most one set of ticking events. Let E and E' be two disjoint set of ticking events. According to the definition, there are cycles $r_1, \ldots, r_n$ and $r'_1, \ldots, r'_{n'}$ in the region graph and corresponding sets of fixed-delay events $E_1, \ldots, E_n$ and $E'_1, \ldots, E'_{n'}$. Observe that from the definition of ticking sets, it is possible to pick from these sets (sub)cycles of fixed-delay events $e_1, \ldots, e_m$ and $e'_1, \ldots, e'_{m'}$ such that when each e_i and e'_i occurs in the respective region, e_{i+1} and

[3] Notice that our definition of SPN is simpler than in [Haa10] as it does not allow the distribution of a transition to depend on the marking where it becomes enabled. For the sake of readability, we therefore simplify the construction from [Haa10] of the simulating GSMP $\mathscr{M}$.

e'_{i+1} start to be scheduled, respectively. Let $m_1, \dots, m_m$ and $m'_1, \dots, m'_{m'}$ be cycles of markings corresponding to regions from whose sets we picked the events and $t_1, \dots, t_m$ and $t'_1, \dots, t'_{m'}$ be the corresponding deterministic transitions. Observe that each t_i initiates t_{i+1} and each t'_i initiates t'_{i+1}. Hence, both the cycles must actually contain only one element, i.e. $m = 2$, $m' = 2$, and $t_1 = t_2 = t'_1 = t'_2$; otherwise it is not possible to define a strict total order $\prec$ witnessing that $\mathscr{N}$ is almost-monotone. Therefore, the event e corresponding to t_1 belongs to both E and E' contradicting that E and E' are disjoint. □

Note that for the ease of exposition we did not define the region relation on the configurations of DSPN. Hence, the sufficient condition for a DSPN to be stable that we could provide is a bit weaker than the sufficient condition that we proved for GSMP. However, it is still satisfied by an interesting class of DSPN that allows multiple concurrently enabled deterministic transitions. In the next section, we address the approximation of **d** and **c**.

5.4 Approximations of frequency measures on stable GSMP

In the previous sections we proved that (1) for any $s \in S$ in single-ticking GSMP, $\mathbf{d}_s$ and $\mathbf{c}_s$ are almost surely well-defined and for almost all runs they attain only finitely many values $d_{s,1} \dots, d_{s,k}$ and $c_{s,1}, \dots, c_{s,k}$, respectively; that (2) the same holds for the frequencies $\mathbf{d}_q^{\mathscr{A}}$ and $\mathbf{c}_q^{\mathscr{A}}$ for any observer TA $\mathscr{A}$ and any its location q by a reduction to a single-ticking GSMP; and that (3) timed frequencies in each almost-monotone DSPN can also be analysed by reduction to timed frequencies in a single-ticking GSMP.[4] In this section we show that it is possible to approximate $d_{s,i}$'s and $c_{s,i}$'s and the probabilities that $\mathbf{d}_s$ and $\mathbf{c}_s$ attain these values, respectively.

Theorem 5.4.1. *In a single-ticking GSMP, let $d_{s,1}, \dots, d_{s,k}$ and $c_{s,1}, \dots, c_{s,k}$ be the discrete and timed frequencies of a state $s \in S$, respectively, corresponding to BSCCs of the region graph. For all $1 \le i \le k$, the numbers $d_{s,i}$ and $c_{s,i}$ as well as the probabilities $\mathscr{P}(\mathbf{d}_s = d_{s,i})$ and $\mathscr{P}(\mathbf{c}_s = c_{s,i})$ can be approximated up to any $\varepsilon > 0$.*

The rest of this section deals with the proof of Theorem 5.4.1. Let $X_1, \dots, X_k$ denote the sets of configurations in individual BSCCs and $d_{s,i}$ and $c_{s,i}$ correspond

[4] In order to analyse the discrete frequencies, a more complicated simulating GSMP needs to be used. We restricted to the timed frequencies as it corresponds to the steady-state analysis, the core of the DSPN literature.

to X_i. Since we reach a BSCC almost surely, we have

$$\begin{aligned}\mathscr{P}(\mathbf{d}_s = d_{s,i}) &= \sum_{j=1}^{k} \mathscr{P}(\mathbf{d}_s = d_{s,i} \mid \mathrm{Reach}(X_j)) \cdot \mathscr{P}(\mathrm{Reach}(X_j)) \\ &= \sum_{j=1}^{k} \mathbf{1}[d_{s,j} = d_{s,i}] \cdot \mathscr{P}(\mathrm{Reach}(X_j))\end{aligned}$$

where the second equality follows from the fact that almost all runs in the j-th BSCC yield the discrete frequency $d_{s,j}$. Therefore, $\mathscr{P}(\mathbf{d}_s = d_{s,i})$ and $d_{s,i}$ can be approximated as follows using the methods of [RR04].

Claim 5.4.2. *Let X be a set of all configurations in a BSCC $\mathscr{B}$, $X_s \subseteq X$ the set of configurations with state s, and d_s the frequency corresponding to $\mathscr{B}$. There are computable constants $n_1, n_2 \in \mathbb{N}$ and $p_1, p_2 > 0$ such that for every $i \in \mathbb{N}$ and $z_X \in X$ we have*

$$\begin{aligned}|\mathscr{P}(\mathrm{Reach}(X)) - P^i(z_0, X)| &\leq (1-p_1)^{\lfloor i/n_1 \rfloor} \\ |d_s - P^i(z_X, X_s)| &\leq (1-p_2)^{\lfloor i/n_2 \rfloor}\end{aligned}$$

Proof. Let Y denote the union of regions from which the BSCC $\mathscr{B}$ is reachable. By Lemmata 5.1.6 and 5.1.7 we have $p, q > 0$ and $m \in \mathbb{N}$ and $k < |V|$ such that from any $z \in Y$ we reach X in $m+k$ steps with probability at least $p \cdot q$. We get the first part by setting $n_1 = m + k$ and $p_1 = p \cdot q$. Indeed, if the process stays in Y after n_1 steps, it has the same chance to reach X again, if the process reaches X, it never leaves it, and if the process reaches $\Gamma \setminus (X \cup Y)$, it has no chance to reach X any more.

By Lemma 5.1.10, Γ is (n, ε, κ)-small. By Theorem 8 of [RR04] we thus obtain that for all $x \in \Gamma$ and all $i \in \mathbb{N}$,

$$\sup_{A \in \mathscr{F}} |P^i(x, A) - \pi(A)| \quad \leq \quad (1-\varepsilon)^{\lfloor i/n \rfloor}$$

which yields the second part by setting $A = \{(s', \xi, t) \in \Gamma \mid s' = s\}$ and observing $A \in \mathscr{F}$ and $d_s = \pi(A)$. □

Further, we want to approximate $c_{s,i} = E_\pi[\tau]/E_\pi[W]$, where π is the invariant measure on X_i. In other words, we need to approximate $\int_{X_i} \tau(x)\pi(dx)$ and $\int_{X_i} W(x)\pi(dx)$. An n-th approximation w_n of $E_\pi[W]$ can be gained by discretizing the part of the state space $\{(s', \xi, t) \in \Gamma \mid \forall e \in \mathbf{E}(s') : \xi(e) \leq n\}$ into $1/n$-large hypercubes, where the invariant measure π is approximated using P^n. This approximation converges to $E_\pi[W]$ since W is continuous and $E_\pi[W]$ is finite:

Claim 5.4.3. *On each region, W is continuous, and $E_\pi[W]$ is finite.*

Proof. Let (s',ξ,t) be a configuration, C and D the set of variable-delay and fixed-delay events scheduled in s', respectively. If $D \neq \emptyset$, let $T = \min_{d\in D}(\ell_d - \xi(d))$ be the time the first fixed-delay event can occur; if $D = \emptyset$, let $T = \infty$. The probability that the transition from (s',ξ,t) occurs within time t' is

$$F(t') = \begin{cases} 1 - \prod_{c\in C} \int_{t'}^{\infty} f_{c|\xi(c)}(x)\,dx & \text{for } 0 < t' < T, \\ 1 & \text{for } t \geq T \end{cases}$$

as non-occurrences of variable-delay events are mutually independent. Observe that $F(t')$ is piece-wise differentiable on the interval $(0,T)$, we denote by $f(t')$ its piece-wise derivative. The expected waiting time in (s',ξ,t) is

$$W((s',\xi,t)) = \begin{cases} \int_0^T x\cdot f(x)\,dx + T\cdot(1-F(T)) & \text{for } T < \infty, \\ \int_0^\infty x\cdot f(x)\,dx & \text{for } T = \infty. \end{cases} \tag{5.4}$$

Recall that for each variable-delay event e, the density f_e is continuous and bounded. Therefore, all $f_{e|t'}$ are also continuous, hence F and f are also continuous with respect to ξ and with respect to t' on $(0,T)$. Thus W is continuous for T both finite and infinite. Moreover, for finite T, W is bounded by T which is for any (s',ξ,t) smaller than $\max_{d\in\mathscr{E}_f}\ell_d$. Hence, $E_\pi[W]$ is finite. For $T=\infty$, $E_\pi[W]$ is finite due to the assumption that each f_e has finite expected value. □

This concludes the proof of Theorem 5.4.1 as τ only differs from W in being zero for $s' \neq s$; thus, $E_\pi[\tau]$ can be approximated analogously.

Chapter 6

Qualitative Games over GSMP with Timed Automata Objectives

After dealing with stability in the long-run behaviour of GSMP, we study the stability in the game setting. First, we define a two-player game extension of GSMP, called *generalized semi-Markov games*. In this game extension, we analyse the hard real-time bounds in the form of a DTA observer (and do not allow fixed-delay events).

The core of the chapter is the analysis of the games with the fundamental reachability specification in the observer. We restrict to the *qualitative* case where one of the players has a strategy to reach the target *almost surely*. In such a situation, the observed stability implies that the player can also reach the target almost surely using a strategy with a simple *finite* structure. This is rather surprising given the game has intrinsically uncountable space of configurations. Another interesting insight is that the strategy can be represented using a deterministic timed automaton. We also provide an algorithm to decide whether an almost-sure winning strategy exists and to construct its representing DTA if it exists.

Lastly, we extend the results to the long-run behaviour. Namely, we show that the results for reachability can be employed to the Büchi specification in the observer. Again, if there is a strategy to visit the target infinitely often, there is also such a strategy with a finite representation using a DTA. This chapter is based on [BKK+10b].

6.1 Generalized semi-Markov games

Let us start with the definition of the game formalism.

Definition 6.1.1 (GSMG)**.** *A* generalized semi-Markov game (GSMG) *is a tuple* $\mathcal{G} = (S_\Box, S_\Diamond, M, \mathcal{E}, \mathbf{E}, \mathrm{Succ}, Act, \mathbf{E}, \alpha_0)$ *where*

- $S_\Box$ *and* $S_\Diamond$ *are finite sets of* control states *of player* □ *and* ◊ *where* S *denotes* $S_\Box \cup S_\Diamond$*,* M *is a finite set of* modes*, and* $\alpha_0 \in \mathbf{D}(S)$ *is an* initial *distribution;*

- *$Act \subseteq \mathbf{D}(M)$ is a finite set of* actions *and* $\mathbf{E} : S \to 2^{Act}$ *assigns to each control state s a non-empty set of actions enabled in s;*

- *$\mathscr{E}$ is a finite set of events where to every $e \in \mathscr{E}$ we associate its* lower bound $\ell_e \in \mathbb{N}_0$, upper bound $u_e \in \mathbb{N} \cup \{\infty\}$, *and a* density function $f_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ *which is positive on (ℓ_e, u_e) and satisfies $\int_{\ell_e}^{u_e} f_e(x)\,dx = 1$;*

- $\mathbf{E} : M \to 2^{\mathscr{E}}$ *assigns to each mode m a non-empty set of events scheduled to occur in m, and* Succ : $M \times \mathscr{E} \to S$ *assigns a* successor *state s to each mode m and each event $e \in \mathbf{E}(m)$ occurring in m.*

As a technical assumption, we further require that there is $r \in \mathbb{R}_{>0}$ bounding all conditional expected waiting. Formally, for any event e with $u_e = \infty$ we have $\sup_{b>\ell_e} E[e \mid b] \leq r$ where $E[e \mid b]$ denotes $\int_0^\infty x \cdot f_{e|b}(x)\,dx$. Recall that $f_{e|b}$ here denotes the shifted density function as defined at the beginning of Chapter 2 and already applied in Section 2.1.3.

First, we explain the behaviour of GSMG on an intuitive level. For each event e we keep track how much time $\xi(e)$ has already elapsed since it become scheduled. The game starts in some control state s_0 randomly chosen according to α_0 and we have $\xi_0(e) = 0$ for each $e \in \mathscr{E}$. One step of the game is as follows. Let s_i be a control state and ξ_i be a vector of elapsed time. First, the player controlling s_i chooses a distribution $\mathbf{a} \in \mathbf{D}(\mathbf{E}(s_i))$ and a mode m_i is randomly chosen with probability $\sum_{a\in\mathbf{E}(s_i)} \mathbf{a}(a) \cdot a(m_i)$. Upon entering m_i, some events stop being scheduled or start being scheduled. We denote the vector altered this way by $\xi_i[m_i]$ and set $\xi_i[m_i](e) = 0$ either if e is *old*, i.e. $e \notin \mathbf{E}(m_i)$, or if e is *new*, i.e. $i = 0$ or $e \in \mathbf{E}(m_i) \setminus \mathbf{E}(m_{i-1})$. We set $\xi_i[m_i](e) = \xi_i(e)$ if e is *inherited*, i.e. $i > 0$ and $e \in \mathbf{E}(m_{i-1}) \cap \mathbf{E}(m_i)$. Second, for each event $e \in \mathbf{E}(m_i)$ a random delay t_e is generated according to the density $f_{e|b}$ for $b = \xi_i[m_i](e)$. Let e_i be the event with the minimal delay t_i.[1] After spending t_i time units in m_i, the control state $s_{i+1} = \mathrm{Succ}(m_i, e_i)$ is entered with the vector of elapsed time $\xi_{i+1} = (\xi_i[m_i] \oplus_{m_i} t_i)[e_i := 0]$. Let us define the notation. The value $(\xi \oplus_m t)$ (e) equals $\xi(e) + t$ if $e \in \mathbf{E}(m)$, and $\xi(e)$, otherwise. Furthermore, the value $\xi[e := 0](e')$ equals 0 if $e = e'$, and $\xi(e')$, otherwise. These steps of the game repeat forever forming an infinite *play*.

Formally, we define the semantics of a GSMG as a discrete-time stochastic game over uncountable set of configurations $\Gamma = S \times \mathbb{R}_{\geq 0}^{\mathscr{E}} \times \mathbb{R}_{\geq 0}$ where the second component is the vector of elapsed time and the third component is the time spent in the previous configuration. The set of configurations Γ is endowed with a product σ-field $\mathscr{F}$ of a discrete σ-field over S and Borel σ-fields for each real component. The game starts in a configuration $(s_0, \xi_0, 0)$ with probability $\alpha_0(s_0)$. For each configuration (s, ξ, t), action $a \in \mathbf{E}(s)$, and a measurable set of configurations Y, the

[1] Several events have the same minimal delay with probability 0. This case can be thus ignored.

probability that the next configuration belongs to Y is expressed by the *transition law* $P_{\mathscr{G}}$ defined as follows:

$$P_{\mathscr{G}}((s,\xi,t),a;Y) = \sum_{\substack{m\in M,\\ e\in \mathbf{E}(m)}} a(m) \int_0^\infty \mathrm{Win}(m,e,\xi;t') \cdot \Big[(\mathrm{Succ}(m,e),(\xi[m] \oplus_m t')[e:=0],t') \in Y\Big]\, dt'$$

where $[c]$ denotes the indicator function of the condition c and $\mathrm{Win}(m,e,\xi;t')$ is the density of event e occurring as the first event t' time units after m is entered with vector ξ. Formally,

$$\mathrm{Win}(m,e,\xi;t') = f_{e|\xi[m](e)}(t) \cdot \prod_{e'\in\mathbf{E}(m),e'\neq e} \int_t^\infty f_{e'|\xi[m](e')}(x)dx.$$

A *history* of $\mathscr{G}$ is a finite sequence $\mathfrak{h} = (s_0,\xi_0,t_0)\cdots(s_n,\xi_n,t_n)$ of configurations. The history does not store the actions taken, modes visited, and events triggered. As the sets Act, M, and $\mathscr{E}$ are finite, we assume without loss of generality that this information is encoded in the next control state.

A *strategy* of player $\odot$, where $\odot \in \{\Box,\Diamond\}$, is a measurable[2] function which to every history $(s_0,\xi_0,t_0)\cdots(s_n,\xi_n,t_n)$ with $s_n \in S_\odot$ assigns a probability distribution over the set $A(s_n)$ of actions that are enabled in s_n. The sets of all strategies of player $\Box$ and player $\Diamond$ are denoted by $\Sigma_{\mathscr{G}}$ and $\Pi_{\mathscr{G}}$, respectively. A *play* is an infinite sequence of configurations $\omega = (s_0,\xi_0,t_0)\cdots$. The set of all plays, denoted by $Play$, is endowed with a product σ-field $\mathfrak{P}$ (obtained by the standard cylinder construction). A pair of strategies $\sigma \in \Sigma_{\mathscr{G}}$ and $\pi \in \Pi_{\mathscr{G}}$ together with the transition law $P_{\mathscr{G}}$ and the initial distribution uniquely determine a probability measure $\mathscr{P}_{\mathscr{G}}^{\sigma,\pi}$ over the measurable space $(Play,\mathfrak{P})$. For a more formal treatment, see for example [MP70].

In this chapter, we use DTA for two different purposes. Firstly, we use DTA to encode strategies in GSMG. Secondly, DTA are used as a generic specification language for properties of timed systems in the sense of Section 2.2.3. Using such specifications, we then define *winning conditions* for players $\Box$ and $\Diamond$.

DTA strategies Using DTA we can express a class of strategies with finite representation. The DTA “observes” the history, and the decisions taken by the corresponding strategy depend only on the region of the resulting configuration (q,ν), which makes the encoding finite. Every history $\mathfrak{h} = (s_0,\xi_0,t_0)\cdots(s_n,\xi_n,t_n)$ of $\mathscr{G}$ defines a a (finite) timed word $W(\mathfrak{h}) = s_0 t_0 \cdots s_n t_n$. We define DTA strategies as follows.

[2] It is measurable with respect to the σ-field over the set of all histories obtained as a disjoint union $\bigcup_{i\in\mathbb{N}_0} \mathscr{H}_i$ where each $\mathscr{H}_i$ is a σ-field over histories of length i obtained as an i-fold product of $\mathscr{F}$.

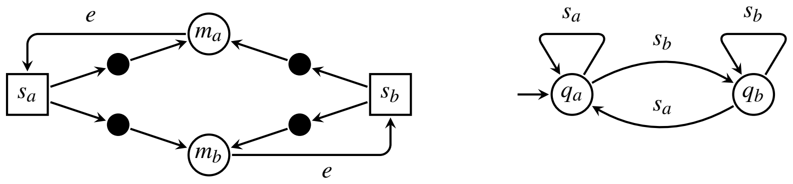

Figure 6.1: Player □ can make the frequency not well-defined. In the game on the left, the density f_e is uniform on $(0,1)$, and the initial distribution α_0 assigns 1 to s_a. In the DTA observer on the right, we measure the discrete or timed frequency of the location q_a.

Definition 6.1.2. *A* DTA strategy *is a strategy τ such that $\tau(\mathfrak{h})$ is rational for every history $\mathfrak{h}$ and there is a DTA $\mathscr{A}$ with alphabet S satisfying the following: Let $\mathfrak{h}$ and $\mathfrak{h}'$ be histories. We have $\tau(\mathfrak{h}) = \tau(\mathfrak{h}')$ if $(q,\nu) \sim (q',\nu')$ where (q,ν) and (q',ν') are the configurations entered by $\mathscr{A}$ after reading $W(\mathfrak{h})$ and $W(\mathfrak{h}')$, respectively.*

Winning conditions For a given specification X, we say that □ wins if the play belongs to X and ◊ wins if the play does not belong to X. As our main concern in the thesis is the stochastic stability, we focus on basic specifications where the stability manifests on the class of strategies sufficient for winning. Namely the reachability and Büchi specification over a set of target locations $T \subseteq Q$ of a DTA observer $\mathscr{A}$.

Remark 6.1.3. *Let us briefly comment the relation of these games to the performance measures studied in previous chapters. A winning condition could be directly based on the performance measures as defined for GSMP. For a given location q, we could say that □ wins if $\mathbf{d}_q$ (or $\mathbf{c}_q$) satisfies $\bowtie$ v with probability $\bar{\bowtie}$ p for $v, p \in [0,1]$ and $\bowtie, \bar{\bowtie} \in \{<, \leq, \geq, >\}$. Notice that such a winning condition does not make sense in the game setting. In Figure 6.1, there is a game and a DTA observer such that the discrete and timed frequency of location q_a depends solely on the strategy of player □. By taking appropriate choices, player □ can make with probability one the frequencies not well-defined. Observe that this phenomenon is not related to continuous-time stochasticity. Standardly, in the games literature, this issue is solved by replacing $\lim_{n\to\infty}$ with $\liminf_{n\to\infty}$ in the definition of the frequency (i.e., in Definition 2.2.9). Such games are then called* mean-payoff games. *Due to its complexity in this setting of continuous-time games, such a winning objective is beyond the scope of the thesis and we suggest it for future work.*

Determinacy As observed in [MS98], the determinacy result for Blackwell games [Mar98] implies determinacy of a large class of stochastic games. This abstract

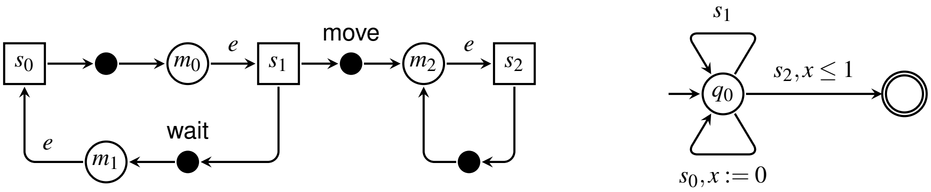

Figure 6.2: Player □ does not have an optimal strategy. In the game on the left, the density f_e is uniform on $(0,1)$, and the initial distribution α_0 assigns 1 to s_0. In the observer on the right, all of the "missing" edges (needed to satisfy the requirement that the guards are total) lead to a "garbage" location which is not depicted.

class includes the games studied in this chapter, and thus we obtain:

Proposition 6.1.4. *Any GSMG $\mathscr{G}$ observed by a DTA $\mathscr{A}$ with a reachability specification and Büchi specification over T is determined, i.e.*

$$\begin{aligned}
\sup_{\sigma\in\Sigma_{\mathscr{G}}} \inf_{\pi\in\Pi_{\mathscr{G}}} \mathscr{P}^{\sigma,\pi}_{\mathscr{G}}(\mathrm{Reach}_{\mathscr{A}}(T)) &= \inf_{\pi\in\Pi_{\mathscr{G}}} \sup_{\sigma\in\Sigma_{\mathscr{G}}} \mathscr{P}^{\sigma,\pi}_{\mathscr{G}}(\mathrm{Reach}_{\mathscr{A}}(T)),\\
\sup_{\sigma\in\Sigma_{\mathscr{G}}} \inf_{\pi\in\Pi_{\mathscr{G}}} \mathscr{P}^{\sigma,\pi}_{\mathscr{G}}(\mathrm{B\ddot{u}chi}_{\mathscr{A}}(T)) &= \inf_{\pi\in\Pi_{\mathscr{G}}} \sup_{\sigma\in\Sigma_{\mathscr{G}}} \mathscr{P}^{\sigma,\pi}_{\mathscr{G}}(\mathrm{B\ddot{u}chi}_{\mathscr{A}}(T)).
\end{aligned}$$

The above equalities define the value *of $\mathscr{G}$ with respect to the reachability and Büchi specifications, denoted by $val(\mathrm{Reach}_{\mathscr{A}}(T))$ or $val(\mathrm{B\ddot{u}chi}_{\mathscr{A}}(T))$, respectively.*

In this chapter we restrict to the qualitative case and assume that the value is 1.

6.2 Reachability specifications

Let us fix a game $\mathscr{G} = (S_\Box, S_\Diamond, M, \mathscr{E}, \mathbf{E}, \mathrm{Succ}, Act, A, \alpha_0)$, its DTA observer $\mathscr{A} = (Q, S, \mathscr{X}, \longrightarrow, q_{init})$, and a reachability specification over $T \subseteq Q$ such that the value $val(\mathrm{Reach}_{\mathscr{A}}(T)) = 1$. The determinacy implies that for any $\varepsilon > 0$, player □ has a strategy that guarantees winning with probability $1-\varepsilon$. However, we show that player □ does not necessarily have any *almost-sure winning* strategy that would guarantee winning with probability 1. Formally, $\sigma^* \in \Sigma_{\mathscr{G}}$ is almost-sure winning if $\mathscr{P}^{\sigma^*,\pi}_{\mathscr{G}}(\mathrm{Reach}_{\mathscr{A}}(T)) = 1$ for any $\pi \in \Pi_{\mathscr{G}}$.

Proposition 6.2.1. *There is a reachability game $\mathscr{G}$ with $S_\Diamond = \emptyset$, a DTA observer $\mathscr{A}$, and a set of target locations T such that the value of the game is* 1 *but there is no almost-sure winning strategy of player □.*

Proof. A simple example is given in Figure 6.2. In that game, $val = 1$ because for every $\varepsilon > 0$, player □ can wait in m_0 until e is fired after a delay smaller than ε

(this eventually happens with probability 1), and then move to m_2. The probability that e is assigned a delay at most $1-\varepsilon$ in m_1 is $1-\varepsilon$, and hence the DTA accepts a play with probability $1-\varepsilon$. However, player $\Box$ has no optimal strategy. Indeed, whenever move is taken, $\Box$ looses with non-zero probability. $\Box$

Therefore we consider the existence and effective constructability of almost-sure winning strategies for player $\Box$. We show that:

Theorem 6.2.2. *If there is (some) strategy of player $\Box$ that is almost-sure winning with respect to* $\mathrm{Reach}_{\mathscr{A}}(T)$*, then there is also a DTA almost-sure winning strategy. Furthermore, there exists an algorithm that in exponential time (1) decides whether there is a DTA almost-sure winning strategy and (2) computes it if it exists.*

A proof of Theorem 6.2.2 is not immediate and requires several steps. First, in Section 6.2.1 we construct a *product game* $\mathscr{G}\times\mathscr{A}$ of $\mathscr{G}$ and $\mathscr{A}$ and show that $\mathscr{G}\times\mathscr{A}$ can be examined instead of $\mathscr{G}$ and $\mathscr{A}$. Intuitively, the product game $\mathscr{G}\times\mathscr{A}$ is constructed by simulating the execution of $\mathscr{A}$ on-the-fly in $\mathscr{G}$. Second, we analyse the strategies in $\mathscr{G}\times\mathscr{A}$. Observe that an almost-sure winning strategy cannot reach with positive probability a "bad" region from that it has zero probability of coming to the target. In other words, from any history in a region that is reached with positive probability it must have non-zero probability of winning. We call every strategy that satisfies this condition a *candidate strategy*. In Section 6.2.2, we show that there is a DTA candidate strategy. However, non-zero probability of winning is not sufficient. In Section 6.2.3, we show that any such strategy is in fact almost-sure winning. Third, in Section 6.2.4 we present the promised algorithm which computes a DTA almost-sure winning strategy if it exists.

6.2.1 The product game

The set of configurations of $\mathscr{G}\times\mathscr{A}$ is obtained as a product of configurations of $\mathscr{G}$ and $\mathscr{A}$, i.e. $S\times(\mathbb{R}_{\geq 0})^{\mathscr{E}}\times\mathbb{R}_{\geq 0}\times Q\times(\mathbb{R}_{\geq 0})^{\mathscr{X}}$. The game starts in a configuration $(s,\mathbf{0},0,q,\mathbf{0})$ with probability $\alpha_0(s)$ where q satisfies $(q_{init},\mathbf{0})\xrightarrow{s}(q,\mathbf{0})$. For each configuration (s,ξ,t,q,ν), action $a\in\mathbf{E}(s)$, and a measurable set of configurations Y, the transition law is defined as follows.

$$P_{\mathscr{G}\times\mathscr{A}}((s,\xi,t,q,\nu),a;Y)=\sum_{\substack{m\in M,\\ e\in\mathbf{E}(m)}} a(m)\int_0^\infty \mathrm{Win}(m,e,\xi;t')$$
$$\cdot\big[(s',\ (\xi[m]\oplus_m t')[e:=0],\ t',\ q',\ \nu')\in Y\big]\,dt'$$

where $s'=\mathrm{Succ}(m,e)$ and $(q,\nu)\xrightarrow{t's'}(q',\nu')$. Intuitively, in each step: (1) the player on turn chooses some action a, (2) a mode m is randomly chosen according to a,

(3) a random time t' is spent waiting in mode m for some event e (4) causing $\mathscr{A}$ to read the time stamp t', and (5) after that a new control state $s' = \mathrm{Succ}(m,e)$ is reached (6) causing $\mathscr{A}$ to read the letter s'.

As $\mathscr{G} \times \mathscr{A}$ is again a stochastic game over the uncountable space of configurations, histories (finite sequences of configurations), plays (infinite sequence of configurations), sets of strategies $\Sigma_{\mathscr{G}\times\mathscr{A}}$ and $\Pi_{\mathscr{G}\times\mathscr{A}}$ and the induced probability measure $\mathscr{P}^{\sigma,\pi}_{\mathscr{G}\times\mathscr{A}}$ are defined analogously to $\mathscr{G}$. Furthermore, for a history $\mathfrak{h} = z_0 \cdots z_n$ we denote by $\mathscr{P}^{\sigma,\pi}_{\mathscr{G}\times\mathscr{A}}[\mathfrak{h}]$ the probability measure of the game that starts in the configuration z_n where the strategies behave as if $\mathfrak{h}$ has been played so far. Formally, for a measurable set A, $\mathscr{P}^{\sigma,\pi}_{\mathscr{G}\times\mathscr{A}}[\mathfrak{h}](A) := \mathscr{P}^{\sigma[\mathfrak{h}],\pi[\mathfrak{h}]}_{(\mathscr{G}\times\mathscr{A})[z_n]}(A)$ where $(\mathscr{G}\times\mathscr{A})[z_n]$ is the game $\mathscr{G}\times\mathscr{A}$ that starts in configuration z_n with probability 1, $\sigma[\mathfrak{h}](\mathfrak{h}') = \sigma(z_0 \cdots z_{n-1}\,\mathfrak{h}')$ for any history $\mathfrak{h}'$, and analogously for $\pi[\mathfrak{h}]$.

The goal of the player $\Box$ is to reach a configuration (s,ξ,t,q,ν) such that $q \in T$ and the goal of the player $\Diamond$ is to avoid it. The set of plays reaching such a configuration is denoted by $\mathrm{Reach}(T)$. A strategy $\sigma \in \Sigma_{\mathscr{G}\times\mathscr{A}}$ is almost-sure winning if $\mathscr{P}^{\sigma,\pi}_{\mathscr{G}\times\mathscr{A}}(\mathrm{Reach}(T)) = 1$ for any $\pi \in \Pi_{\mathscr{G}\times\mathscr{A}}$. In the rest of this subsection, we first relate almost-sure winning strategies in the original game and in the product game. Then, we define a notion of *region-based* strategies in $\mathscr{G}\times\mathscr{A}$ that correspond to DTA strategies in $\mathscr{G}$ and that are studied in the rest of the chapter.

Equivalence of $\mathscr{G}$ and $\mathscr{G}\times\mathscr{A}$ We show that almost-sure winning strategies in $\mathscr{G}$ and $\mathscr{G}\times\mathscr{A}$ coincide. To this end we need to map the strategies from $\mathscr{G}$ to $\mathscr{G}\times\mathscr{A}$ and back. Each strategy τ in $\mathscr{G}$ induces a strategy $\tau^{\uparrow}$ in $\mathscr{G}\times\mathscr{A}$ that simply ignores the additional two components of the state space. I.e., for any history $\mathfrak{h} = (s_0,\xi_0,t_0,q_0,\nu_0)\cdots(s_n,\xi_n,t_n,q_n,\nu_n)$ in $\mathscr{G}\times\mathscr{A}$ we set $\tau^{\uparrow}(\mathfrak{h}) := \tau(\mathfrak{h}')$ where $\mathfrak{h}' = (s_0,\xi_0,t_0)\cdots(s_n,\xi_n,t_n)$. Similarly, each strategy τ in $\mathscr{G}\times\mathscr{A}$ induces a strategy $\tau^{\downarrow}$ in $\mathscr{G}$ that for each history takes the decision that τ takes for the history enhanced with the deterministic behaviour of the DTA. Formally, for any history $\mathfrak{h} = (s_0,\xi_0,t_0)\cdots(s_n,\xi_n,t_n)$ in $\mathscr{G}$ we set $\tau^{\downarrow}(\mathfrak{h}) := \tau(\rho(\mathfrak{h}))$ where $\rho(\mathfrak{h})$ is a history $(s_0,\xi_0,t_0,q_0,\nu_0)\cdots(s_n,\xi_n,t_n,q_n,\nu_n)$ such that $(q_{init},\mathbf{0}) \xrightarrow{s_0} (q_0,\nu_0)$ and furthermore $(q_{i-1},\nu_{i-1}) \xrightarrow{t_i s_i} (q_i,\nu_i)$ for each $0 < i \leq n$.

Proposition 6.2.3. *The almost-sure winning strategies in $\mathscr{G}$ and $\mathscr{G}\times\mathscr{A}$ coincide:*

- *If $\sigma \in \Sigma_{\mathscr{G}}$ is almost-sure winning, then $\sigma^{\uparrow}$ is almost-sure winning in $\mathscr{G}\times\mathscr{A}$.*
- *If $\sigma \in \Sigma_{\mathscr{G}\times\mathscr{A}}$ is almost-sure winning then $\sigma^{\downarrow}$ is almost-sure winning in $\mathscr{G}$.*

Proof. For the first point, it suffices to show that for any $\sigma \in \Sigma_{\mathscr{G}}$ and $\pi \in \Pi_{\mathscr{G}\times\mathscr{A}}$,

$$\mathscr{P}^{\sigma,\pi^{\downarrow}}_{\mathscr{G}}(\mathrm{Reach}_{\mathscr{A}}(T)) \;=\; \mathscr{P}^{\sigma^{\uparrow},\pi}_{\mathscr{G}\times\mathscr{A}}(\mathrm{Reach}(T)).$$

Indeed, $\sigma^{\uparrow}$ is almost-sure winning if $\mathscr{P}^{\sigma^{\uparrow},\pi}_{\mathscr{G}\times\mathscr{A}}(\mathrm{Reach}(T)) = 1$ for any $\pi \in \Pi_{\mathscr{G}\times\mathscr{A}}$.

We extend the mapping ρ to plays. For a play $\omega = (s_0,\xi_0,t_0)\cdots$ of $\mathscr{G}$ we set $\rho(\omega) = (s_0,\xi_0,t_0,q_0,\nu_0)\cdots$ where $(q_{init},\mathbf{0}) \xrightarrow{s_0} (q_0,\nu_0)$ and $(q_{i-1},\nu_{i-1}) \xrightarrow{t_is_i} (q_i,\nu_i)$ for each $i > 0$. We obtain the equality by observing

- $\mathrm{Reach}(T) = \rho(\mathrm{Reach}_{\mathscr{A}}(T))$, which directly follows from the definitions;
- the mapping ρ preserves measure. This follows from the facts that
  1. $\sigma(\mathfrak{h}) = \sigma^{\uparrow}(\rho(\mathfrak{h}))$ and $\pi^{\downarrow}(\mathfrak{h}) = \pi(\rho(\mathfrak{h}))$ for any history $\mathfrak{h}$;
  2. $P_{\mathscr{G}}((s,\xi,t),a;Y) = P_{\mathscr{G}\times\mathscr{A}}((s,\xi,t,q,\nu),a;Y_{q,\nu})$ for any configuration (s,ξ,t,q,ν) of $\mathscr{G}\times\mathscr{A}$, action a, and measurable set of configurations Y of $\mathscr{G}$ where $Y_{q,\nu} = \{(s',\xi',t',q',\nu') \mid (s',\xi',t') \in Y, (q,\nu) \xrightarrow{t's'} (q',\nu')\}$.

For the second point, the proof goes analogously by showing that for any strategy $\sigma \in \Sigma_{\mathscr{G}\times\mathscr{A}}$, and $\pi \in \Pi_{\mathscr{G}}$, we have $\mathscr{P}^{\sigma^{\downarrow},\pi}_{\mathscr{G}}(\mathrm{Reach}_{\mathscr{A}}(T)) = \mathscr{P}^{\sigma,\pi^{\uparrow}}_{\mathscr{G}\times\mathscr{A}}(\mathrm{Reach}(T))$. □

Thanks to Proposition 6.2.3, we can focus on almost-sure winning strategies in $\mathscr{G}\times\mathscr{A}$. The region relation helps us to find a counterpart to DTA strategies in $\mathscr{G}\times\mathscr{A}$.

Region relation Analogously to the definition of the region relation $\sim$ on the configurations of a GSMP and on the configurations of a TA, we define $\sim$ on the configurations of $\mathscr{G}\times\mathscr{A}$. We put $(s,\xi,t,q,\nu) \sim (s',\xi',t',q',\nu')$ if $(s,\xi,t) \sim (s',\xi',t')$ and $(q,\nu) \sim (q',\nu')$. Again, the equivalence classes of $\sim$ are called *regions*. Observe that all configurations of a region r have the same control state, denoted s_r, and the same location, denoted q_r. We say that an action a is *enabled in* r if $a \in \mathbf{E}(s_r)$. Furthermore, r is a *target region* if $q_r \in T$. The sets of all regions and target regions are denoted by $\mathscr{R}$ and $\mathscr{R}_T$, respectively. For $X \subseteq \mathscr{R}$, we denote by $\mathrm{ReachReg}(X)$ the set of plays that reach any configuration in any region in X.

Counterpart to DTA strategies in $\mathscr{G}\times\mathscr{A}$ A strategy $\sigma \in \Sigma_{\mathscr{G}\times\mathscr{A}}$ is *region-based* if $\sigma(\mathfrak{h})$ is rational for any $\mathfrak{h}$ and $\sigma(\mathfrak{h}z) = \sigma(\mathfrak{h}'z')$ for any histories such that $z \sim z'$.

Proposition 6.2.4. *If $\sigma \in \Sigma_{\mathscr{G}\times\mathscr{A}}$ is a region-based strategy, $\sigma^{\downarrow}$ is a DTA strategy.*

Proof. We show that there is a DTA $\mathscr{A}_{\mathscr{G}\times\mathscr{A}}$ such that

1. its regions are in one-to-one correspondence with the regions of $\mathscr{G}\times\mathscr{A}$;
2. it mimics the behaviour of $\mathscr{G}\times\mathscr{A}$ when reading the timed word of the play. Precisely, let $\mathfrak{h}$ and $\mathfrak{h}'$ be histories ending in the same region of $\mathscr{G}\times\mathscr{A}$. After

reading $Ap(\mathfrak{h})$ and $Ap(\mathfrak{h}')$, the automaton $\mathscr{A}_{\mathscr{G}\times\mathscr{A}}$ is defined in such a way that it ends up in configurations (q,ν) and (q',ν'), respectively, that are in the same region of $\mathscr{A}_{\mathscr{G}\times\mathscr{A}}$.

These two points guarantee that $\sigma^{\downarrow}$ is a DTA strategy. Let us define $\mathscr{A}_{\mathscr{G}\times\mathscr{A}}$. Intuitively, it has a special clock for every clock of $\mathscr{A}$ and every event of $\mathscr{E}$, and uses its locations to store also the current control state of the game. Formally, $A_{\mathscr{G}\times\mathscr{A}} = ((S\times Q)\cup\{q_{init}\}, S, \mathscr{X}\cup\mathscr{E}, \hookrightarrow_{init}\cup\hookrightarrow_{play}, q_{init}, \emptyset)$. For a control state s we denote by m_s and e_s the unique mode and event such that $\mathrm{Succ}(m_s,e_s)=s$ (recall that we encode the mode and the event in the successor control state). The set of events *not* scheduled in m is denoted as $N(m)$. We set:

$$\begin{aligned}\hookrightarrow_{init} &= \{\,(q_{init},s,g,X,(s,q')) \mid s\in S,(q_0,s,g,X,q')\in\longrightarrow\}\\ \hookrightarrow_{play} &= \{\,((s,q),s',g,\{e_{s'}\}\cup N(m_{s'})\cup X,(s',q'))\\ &\qquad\qquad\qquad \mid s,s'\in S,q\in Q,(q,s',g,X,q')\in\longrightarrow\}\end{aligned}$$

Note that when entering location (s,q), clocks $N(m_s)$ corresponding to not enabled events are always restarted, i.e. the regions are in one-to-one correspondence to the regions of $\mathscr{G}\times\mathscr{A}$. The second point is easy to show by induction on the length of the history $\mathfrak{h}$. □

In the subsequent proofs we search for a region-based almost-sure winning strategy σ in $\mathscr{G}\times\mathscr{A}$ since $\sigma^{\downarrow}$ is the desired DTA almost-sure winning strategy in $\mathscr{G}$. For the rest of Section 6.2, we deal only with $\mathscr{G}\times\mathscr{A}$. We thus write Σ, Π, $\mathscr{P}^{\sigma,\pi}$, and $\mathscr{P}^{\sigma,\pi}[\mathfrak{h}]$ instead of $\Sigma_{\mathscr{G}\times\mathscr{A}}$, $\Pi_{\mathscr{G}\times\mathscr{A}}$, $\mathscr{P}^{\sigma,\pi}_{\mathscr{G}\times\mathscr{A}}$, and $\mathscr{P}^{\sigma,\pi}_{\mathscr{G}\times\mathscr{A}}[\mathfrak{h}]$, respectively.

6.2.2 The existence of region-based candidate strategies in $\mathscr{G}\times\mathscr{A}$

The next step is to show that there is a region-based candidate strategy. Later in Section 6.2.3 we show that such a strategy is in fact almost sure winning.

Definition 6.2.5. *We call σ a* candidate strategy *if* $\inf_{\pi\in\Pi}\mathscr{P}^{\sigma,\pi}[\mathfrak{h}](\mathrm{Reach}(T))>0$ *for each history $\mathfrak{h}$ ending in a configuration from $\mathscr{S}(\sigma)$ where*

$$\mathscr{S}(\sigma)=\{z\mid\exists\pi\in\Pi,r\in\mathscr{R}:\mathscr{P}^{\sigma,\pi}(\mathrm{ReachReg}(\{r\}))>0,z\in r\}.$$

Proposition 6.2.6. *If there is an almost-sure winning strategy $\sigma\in\Sigma$, then there is a region-based candidate strategy $\sigma^*\in\Sigma$.*

The rest of the subsection forms the proof of Proposition 6.2.6. We fix an almost-sure winning strategy σ and we build the strategy σ^* in iterations. At the beginning, we set $X=\mathscr{R}_T$ and set σ^* to be an arbitrary region-based strategy. In each iteration, we take a non-empty set of regions Y such that $X \cap Y = \emptyset$ and $\bigcup Y \subseteq \mathscr{S}(\sigma)$ and change the decision of σ^* in Y so that

$$\inf_{\pi\in\Pi} \mathscr{P}^{\sigma^*,\pi}[\mathfrak{h}](\mathrm{ReachReg}(X)) > 0 \quad \text{for any } \mathfrak{h} \text{ ending in any region } r \in Y. \tag{6.1}$$

Then we set $X := X \cup Y$. At the end of each iteration, we have by a simple integration that $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma^*,\pi}[\mathfrak{h}](\mathrm{ReachReg}(\mathscr{R}_T)) > 0$ for any history $\mathfrak{h}$ ending in any region from X. We repeat the process until $\bigcup X = \mathscr{S}(\sigma)$. At last we show that $\mathscr{S}(\sigma^*) \subseteq \mathscr{S}(\sigma)$ which concludes the proof that σ^* is a candidate strategy.

Let us now focus on a single iteration. Let X be a set of regions with $\mathscr{R}_T \subseteq X$. We fix an arbitrary region $\bar{r} \subseteq \mathscr{S}(\sigma)$ such that $\bar{r} \notin X$. Intuitively, the set Y contains the regions necessary to reach X from $\bar{r}$ and will be formally defined later in the proof. Let $\bar{\mathfrak{h}}$ be any history ending in $\bar{r}$, for which σ is almost-sure winning. Such a history exists since there is $\pi \in \Pi$ such that $\bar{r}$ is visited with positive probability and since σ is almost-sure winning. Hence, $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma,\pi}[\bar{\mathfrak{h}}](\mathrm{ReachReg}(X)) = 1$ because $\mathscr{R}_T \subseteq X$. For any i, we denote by R_i^X the set of plays that reach X in the first i steps. The first non-trivial observation is that player $\Diamond$ cannot block reaching X arbitrarily long.

Claim 6.2.7. *For any* $\bar{\mathfrak{h}}$ *and set of regions* X *with* $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma,\pi}[\bar{\mathfrak{h}}](\mathrm{ReachReg}(X)) > 0$*, there is* $n \in \mathbb{N}$ *such that* $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma,\pi}[\bar{\mathfrak{h}}](R_n^X) > 0$.

Proof. To prepare an important argument for the actual proof, we first show for any $i \in \mathbb{N}$ that if $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma,\pi}[\bar{\mathfrak{h}}](R_i^X) = 0$ then there is a strategy π such that $\mathscr{P}^{\sigma,\pi}[\bar{\mathfrak{h}}](R_i^X) = 0$. Let us fix $i \in \mathbb{N}$. We show by induction on $0 \leq j \leq i$ that for any $\mathfrak{h}$ extending $\bar{\mathfrak{h}}$ by $(i-j)$ more steps we have

$$\inf_{\pi\in\Pi} \mathscr{P}^{\sigma,\pi}[\mathfrak{h}](R_j^X) = 0 \quad \longrightarrow \quad \exists\, \pi : \mathscr{P}^{\sigma,\pi}[\mathfrak{h}](R_j^X) = 0. \tag{6.2}$$

First, let $j = 0$. If a history satisfies the left hand side of (6.2), it ends in a region not in X, and it also satisfies the right hand side for any $\pi \in \Pi$. Further, let all histories for j satisfy (6.2) and let $\mathfrak{h}$ be a history that extends $\bar{\mathfrak{h}}$ by $i-(j+1)$ steps and $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma,\pi}[\mathfrak{h}](R_j^X) = 0$. Let z be the last configuration of $\mathfrak{h}$. First, let z belong to $\Box$ and a be any action such that σ assigns positive probability to a in $\mathfrak{h}$. There must be a set of configurations A with $P_{\mathscr{G}\times\mathscr{A}}(z,a;A) = 1$ such that all histories of the form $\mathfrak{h}z'$ for $z' \in A$ satisfy the left hand side of (6.2). Indeed if there was a set B with $P_{\mathscr{G}\times\mathscr{A}}(z,a;B) > 0$ such that $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma,\pi}[\mathfrak{h}z'](R_j^X) > 0$ for every $z' \in B$, the left hand side of (6.2) cannot hold for $\mathfrak{h}$. By the induction hypothesis, we have optimal strategies for almost all successors of $\mathfrak{h}$, we easily combine them together and get an optimal strategy for $\mathfrak{h}$. Second, let z belong to $\Diamond$. There must be similarly an action a such that there is a set of configurations A with $P_{\mathscr{G}\times\mathscr{A}}(z,a;A) = 1$ and all histories of the form $\mathfrak{h}z'$ for $z' \in A$ satisfy the left hand side of (6.2). Likewise, we have by the induction hypothesis optimal strategies for almost all successors of $\mathfrak{h}$, we easily combine them together in a strategy that chooses a in $\mathfrak{h}$ and get an optimal strategy for $\mathfrak{h}$.

Now we turn our attention to the claim itself. Let us assume the opposite that for all $i \in \mathbb{N}$ we have $\inf_{\pi \in \Pi} \mathscr{P}^{\sigma,\pi}[\bar{\mathfrak{h}}](R_i^X) = 0$. Thanks to the proof above, for each $i \in \mathbb{N}$ we denote by π_i a strategy such that $\mathscr{P}^{\sigma,\pi_i}[\bar{\mathfrak{h}}](R_i^X) = 0$. Using these strategies, we build a strategy π with $\mathscr{P}^{\sigma,\pi}[\bar{\mathfrak{h}}](\mathrm{ReachReg}(X)) = 0$ contradicting the assumption of the claim. We define π step by step and keep for each $i \in \mathbb{N}_0$ the inductive property that there is a set H_i of histories of the form $\bar{\mathfrak{h}}\mathfrak{h}$ with $|\mathfrak{h}| = i$ such that

1. $\mathscr{P}^{\sigma,\pi}[\bar{\mathfrak{h}}](W_i \setminus \mathrm{ReachReg}(X)) = 1$ where each play from W_i starts with the last $i+1$ configurations of some $\mathfrak{h} \in H_i$;

2. for each $\mathfrak{h} \in H_i$ we have that $\mathscr{P}^{\sigma,\pi_j}[\mathfrak{h}](R_{j-i}^X) = 0$ for infinitely many $j > i$.

We denote the infinite set of strategies for $\mathfrak{h} \in H_i$ in the second point by $\Pi_{\mathfrak{h}}$. Note that the probability in the first condition depends only on first i steps for that π will be defined. The second condition intuitively means that histories in H_i are "promising" for infinitely many strategies π_j w.r.t. their goal $\mathscr{P}^{\sigma,\pi_j}[\bar{\mathfrak{h}}](R_j^X) = 0$.

As regards the base, let $H_0 = \{\bar{\mathfrak{h}}\}$ and $\Pi_{\bar{\mathfrak{h}}} = \{\pi_j \mid j \in \mathbb{N}\}$. The set H_0 satisfies clearly both conditions. Let us have a set H_i, we define π for all histories from H_i and construct H_{i+1}. Let us fix any $\mathfrak{h} \in H_i$ with z being its last configuration.

- First let z belong to $\Diamond$. Because the set of actions is finite, there must be an action a that is assigned positive probability by infinitely many strategies $\Pi' \subseteq \Pi_{\mathfrak{h}}$. We define π to take action a in $\mathfrak{h}$ with probability 1. There must be a set of configurations A such that $P(A \mid z, a) = 1$ and for any $z' \in A$ and $\pi_j \in \Pi'$ it holds $\mathscr{P}^{\sigma,\pi_j}[\mathfrak{h}z'](R_{j-(i+1)}^X) = 0$. Otherwise we get a contradiction with $\pi_j \in H_{\mathfrak{h}}$ because π_j would reach via a with positive probability configurations where it cannot win.

- Second let z belong to $\Box$. By similar arguments as before, there must be a set of configuration A such that $P(A \mid z, a) = 1$ for any action a chosen by σ in $\mathfrak{h}$ with positive probability; and for any $z' \in A$ and $\pi_j \in \Pi'$ it holds $\mathscr{P}^{\sigma,\pi_j}[\mathfrak{h}z'](R_{j-(i+1)}^X) = 0$.

We set $H_{i+1}^{\mathfrak{h}} = \{\mathfrak{h}z' \mid z' \in A\}$ and $\Pi_{\mathfrak{h}z'} = \Pi'$ for any $z' \in A$. Finally, we set $H_{i+1} = \bigcup_{\mathfrak{h} \in H_i} H_{i+1}^{\mathfrak{h}}$. From the observations above, this set satisfies both conditions 1. and 2. When we repeat the inductive steps ad infinitum, we get the contradiction that the probability to reach X using the strategy π is 0. $\Box$

Next, we set which actions need to be taken in which regions to reach X from $\bar{\mathfrak{h}}$ in n steps. For each alternating sequence of regions and actions $\alpha = r_0a_0 \cdots a_{i-1}r_i$ with $i \leq n$, let R_α denote the set of plays $\mathit{Pattern}(r_0a_0 \cdots a_{i-1}r_i) \cap R_n^X$ where the plays $\mathit{Pattern}(r_0a_0 \cdots a_{i-1}r_i)$ visit the respective regions in the first i step where the respective actions are taken. We say that a set A of such sequences is a *guarantee* of X from $\bar{\mathfrak{h}}$ if (1) every sequence contains a region from X, (2) A is followed from $\bar{\mathfrak{h}}$ no matter how $\Diamond$ plays, i.e. $\mathrm{p}(A) > 0$ where $\mathrm{p}(A) = \inf_{\pi \in \Pi} \mathscr{P}^{\sigma,\pi}[\bar{\mathfrak{h}}](\bigcup_{\alpha \in A} R_\alpha)$, and (3) A can be implemented by $\Box$, i.e. for any $j < i$, all sequences in A that agree on the first j regions either agree on the j-th action or the j-th region belongs to $\Diamond$.

Claim 6.2.8. *For any $\bar{\mathfrak{h}}$, set of regions X, and $n \in \mathbb{N}$ with $\inf_{\pi \in \Pi} \mathscr{P}^{\sigma,\pi}[\bar{\mathfrak{h}}](R_n^X) > 0$, there is a guarantee of X from $\bar{\mathfrak{h}}$.*

Proof. We build the guarantee of X from $\mathfrak{h}$ inductively. For $i \in \mathbb{N}$, we say that a set of sequences A is an *i-step guarantee* if (1) every sequence in A has i regions, and A satisfies conditions (2) - (3) of a guarantee, and (4) A is minimal such set, i.e. $\mathrm{p}(A \setminus \{\alpha\}) = 0$ for any $\alpha \in A$, By definition, $\{r\}$ is an 1-step guarantee. We show that for any i-step guarantee A, there is an $(i+1)$-step guarantee A'. Observe that this is sufficient as by this process we obtain an n-step guarantee and any n-step guarantee of X is thanks to (1), (2), and (4) also a guarantee.

Let $A = \{\alpha_1, \cdots, \alpha_\ell\}$ be the i-step guarantee. One by one replacing the sequences of A, we build sets $(A_k)_{k \leq \ell}$ such that the set A_ℓ is the sought A'. We set $A_0 = A$; let $1 \leq k \leq \ell$ and let $\alpha_k = r_0a_0 \cdots a_{i-1}r_i$.

- First, let us assume that r_i belongs to player $\Box$. As there are only finitely many actions and finitely many regions, there must be some action a_i and some region r_{i+1} such that $\mathrm{p}((A_{k-1} \setminus \{\alpha_k\}) \cup \alpha'_k) > 0$ where $\alpha'_k = \alpha_k\, a_i r_{i+1}$. For such α'_k we set $A_k = (A_{k-1} \setminus \{\alpha_k\}) \cup \alpha'_k$ satisfying condition (2). Due to (4) of A, there are no other sequence in A with the same regions as in α_k, hence (3) holds for A_k as well. As regards condition (4), observe that we still have $\mathrm{p}(A_k \setminus \{\alpha\}) = 0$ for any $\alpha \in A_k$: if we remove α'_k, it holds as before, if we remove $\alpha \neq \alpha'_k$, it holds similarly as $R_{\alpha'_k} \subseteq R_{\alpha_k}$.

- Second, let us assume that r_i belongs to player $\Diamond$. We replace α_k with a set of sequences - one for each action of $\Diamond$ enabled in r_i. For each action a enabled in r_i, we denote by $\alpha_{k,a}$ a sequence $\alpha_k\, a\, r_a$ where r_a is some region such that $R_{\alpha_{k,a}}$ has positive measure for any π which chooses a in a subset of R_{α_k} of positive measure. We set $A_k = (A_{k-1} \setminus \alpha_k) \cup \{\alpha_{k,a} \mid a \text{ enabled in } r_i\}$. Hence, the condition (3) holds for A_k. Observe that also (2) holds, i.e. $\mathrm{p}(A_k) > 0$. Indeed, for any strategy π that does not reach the target via α_k, the situation is the same as for A_{k-1}. Any strategy that reaches the target with positive probability via α_k must also reach it with positive probability via some $\alpha_{k,a}$.

Furthermore, we show that (4) holds, i.e. $\mathrm{p}(A_k \setminus \{\alpha\}) = 0$ for any $\alpha \in A_k$. If any $\alpha \in A_{k-1}$ gets removed, we obtain the result as for A_{k-1} as all the sequences in A_k extend the sequences in A_{k-1}. If any $\alpha_{k,a}$ gets removed, there is a sequence of strategies that in the limit do not reach the target via any $A_{k-1} \setminus \{\alpha_k\}$. They may only reach the target via α_k. We alter these strategies so that they choose action a after traversing α_k, thus not reaching the target via $A_k \setminus \{\alpha_{k,a}\}$.

Observe that $A' = A_\ell$ is an $(i+1)$-step guarantee as it satisfies (1) as well. We repeat this process until we get an n-step guarantee A, hence a guarantee. □

With a guarantee A, we finally set $Y = \{r \mid \exists i < n, r_0a_0 \cdots a_{n-1}r_n \in A : r = r_i\}$ to be the set of regions in A. The guarantee A defines on Y a simple strategy that reaches X from $\bar{\mathfrak{h}}$ in n steps with positive probability. However, this strategy (a) does not have to be region-based since it may choose different actions in one region when it is reached via different sequences of regions and (b) does not guarantee reaching X with positive probability from other histories in Y.

To address (a), the guarantee A defines the region-based σ^* on Y as follows. For a region $r \in Y$ of player $\Box$ and a history $\mathfrak{h}$ ending in r, we set $\sigma^*(\mathfrak{h})(a) = 1$ for an action a that appears after a latest occurrence of r in A, i.e. $a = a_i$ and $r = r_i$ for some $r_0a_0 \cdots a_{n-1}r_n \in A$ and $i = \max\{j < n \mid r'_0a'_0 \cdots a'_{n-1}r'_n \in A, r'_j = r\}$.

To address (b), we prove (6.1), i.e. $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma^*,\pi}[\mathfrak{h}](\mathrm{ReachReg}(X)) > 0$, by induction on the distance of $\mathfrak{h}$ to X. For a region $r \in Y$ with control state s we define the distance by

$$\mathrm{dist}(r) = \begin{cases} 0 & \text{if } r \in X, \\ 1 + \min_{r \overset{a}{\rightsquigarrow} r'} \mathrm{dist}(r') & \text{if } r \text{ belongs to } \Box \text{ and } \sigma^* \text{ chooses } a \text{ in } r, \\ 1 + \max_{a\in\mathbf{E}(s)} \min_{r \overset{a}{\rightsquigarrow} r'} \mathrm{dist}(r') & \text{if } r \text{ belongs to } \Diamond. \end{cases}$$

As regards the base, the probability to reach X is 1 from any configuration in X. As regards the induction step, let us have for any region r with distance $\leq i$ that $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma^*,\pi}[\mathfrak{h}](\mathrm{ReachReg}(X)) > 0$ for any $\mathfrak{h}$ ending in r. Let us fix any r with distance $i+1$ and $\mathfrak{h} \in r$. From the definition of distance, by simple integration, and from the fact that the region relation is a congruence w.r.t. one step positive reachability (see the following lemma), we get $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma^*,\pi}[\mathfrak{h}](\mathrm{ReachReg}(X)) > 0$ concluding the proof for one iteration. Note that the following lemma is a game counterpart of Lemma 4.1.2.

Lemma 6.2.9. *Let r_1, r_2 be regions and $z, z' \in r_1$. For any action a enabled in r_1,*

$$P_{\mathscr{G}\times\mathscr{A}}(z, a; r_2) > 0 \quad \textit{iff} \quad P_{\mathscr{G}\times\mathscr{A}}(z', a; r_2) > 0.$$

Furthermore, we write $r_1 \overset{a}{\rightsquigarrow} r_2$ if $P_{\mathscr{G}\times\mathscr{A}}(z, a; r_2) > 0$ for any (hence every) $z \in r_1$.

Proof. For $z = (s, \xi, t, q, \nu)$ with $P_{\mathscr{G}\times\mathscr{A}}(z, a; r_2) > 0$, we show $P_{\mathscr{G}\times\mathscr{A}}(z', a; r_2) > 0$. The other direction follows from symmetry. There must be a mode m, an event $e \in \mathbf{E}(m)$, and some maximal non-empty interval of time I such that

1. when e occurs in mode m after any time $t' \in I$, the game moves to r_2;
2. further, $a(m) > 0$ and $\int_{t'\in I} \mathrm{Win}(m, e, \xi; t')dt' > 0$ since $P_{\mathscr{G}\times\mathscr{A}}(z, a; r_2) > 0$.

Let us discuss closer the bounds of such maximal interval I. Let F be the set $\{\langle \xi[m](e)\rangle \mid e \in \mathbf{E}(m)\} \cup \{\langle \nu(x)\rangle \mid x \in \mathscr{X}\}$ of fractional parts of the elapsed time of the events and clocks. Due to the definition of the region relation, either $\inf I = 0$ or $\inf I = a - g$ for some $a \in \mathbb{N}$ and $g \in A$. The event e can be thus triggered either immediately or after the elapsed time of the event or clock corresponding to g reaches its a-th following integral value. Likewise, either $\sup I = \infty$, or $\sup I = b - h$ for some $b \in \{a, a+1\}$ and $h \in A$; e is triggered either after arbitrary waiting or before the event or clock corresponding to h reaches its b-th following integral value.

Let F' be defined analogously using $z' = (s', \xi', t', q', \nu')$ and $g', h' \in F'$ correspond to the same events or clocks as g, h if defined. Let I' be a non-empty interval with $\inf I' = 0$ if $\inf I = 0$ and $\inf I' = a - g'$, otherwise; and $\sup I' = \infty$ if $\sup I = \infty$ and $\sup I' = b - h'$, otherwise. Due to $z \sim z'$, we have that

1. the game moves to r_2 when e occurs in mode m after any time $t' \in I'$ when starting from z'. This follows from the definition of the region relation that the values from $\xi[m] \cup \nu$ and $\xi'[m] \cup \nu'$ agree on integral values and their fractional values have the same order.
2. $\int_{t'\in I'} \mathrm{Win}(m, e, \xi'; t')dt' > 0$. Indeed, $\mathrm{Win}(m, e, \xi'; t')$ is positive on the whole I' because no events' values may lie in between g' and h', i.e. any event $e' \in \mathbf{E}(m)$ is supported on the whole $(\inf I' + \xi'[m](e'), \sup I' + \xi'[m](e')$. □

After defining σ^* on the whole $\mathscr{S}(\sigma)$, let us prove that $\mathscr{S}(\sigma^*) \subseteq \mathscr{S}(\sigma)$. Observe that the strategy σ^* chooses in any region $r \subseteq \mathscr{S}(\sigma)$ an action from the set $A(r) = \{a \in Act \mid \exists \pi \in \Pi : \mathscr{P}^{\sigma,\pi}(\mathrm{Use}(a\,\mathrm{in}\,r)) > 0\}$ where $\mathrm{Use}(a\,\mathrm{in}\,r)$ is a set of plays that visit r and action a is taken in r. Indeed, only such actions appear in a guarantee using which σ^* is defined.

Claim 6.2.10. *For any σ^* restricted to $A(r)$ in any $r \subseteq \mathscr{S}(\sigma)$, $\mathscr{S}(\sigma^*) \subseteq \mathscr{S}(\sigma)$.*

Proof. For a contradiction, let us assume that there is a region $r \not\subseteq \mathscr{S}(\sigma)$ and $\pi \in \Pi$ such that $\mathscr{P}^{\sigma^*,\pi}(\mathrm{ReachReg}(\{r\})) > 0$. Then there is a sequence of regions $r_0 \cdots r_n$ such that $r_n = r$ and $\mathscr{P}^{\sigma^*,\pi}(Pattern(r_0 \cdots r_n)) > 0$ where $Pattern(r_0 \cdots r_n)$ denotes the plays that visit the respective regions in the first n steps. Indeed, let n be the smallest index such that the set of plays R that reach r in the n-th step satisfies

$\mathscr{P}^{\sigma^*,\pi}(R) > 0$. We partition R according to the sequence of regions visited in the first n steps and obtain finitely many equivalence classes. The sequence of regions $r_0 \cdots r_n$ corresponds to one of the equivalence classes that has positive measure. Let v be the first region in this sequence not contained in $\mathscr{S}(\sigma)$ and let u be the region preceding v.

If u belongs to player $\Diamond$, the strategy π chooses in some histories in u some action a that leads with positive probability to the region v. Let π' be a strategy for which $\mathscr{P}^{\sigma,\pi'}(\text{ReachReg}(\{u\})) > 0$ and that chooses a in the whole region u. By Lemma 6.2.9, we have that $\mathscr{P}^{\sigma,\pi'}(\text{ReachReg}(\{v\})) > 0$ contradicting $v \not\subseteq \mathscr{S}(\sigma)$.

If u belongs to player $\Box$, let $a \in A(u)$ be the action chosen by σ^* in u. For some $\pi \in \Pi$, there must be a set of plays with positive measure w.r.t. $\mathscr{P}^{\sigma,\pi}$ where σ gives positive weight to a. Thus, by Lemma 6.2.9, also v is reached with positive probability, again contradicting $v \not\subseteq \mathscr{S}(\sigma)$. □

The observation that $\mathscr{S}(\sigma^*) \subseteq \mathscr{S}(\sigma)$ concludes the proof of Proposition 6.2.6 because we have shown that σ^* wins with positive probability for any $\mathfrak{h}$ ending in $\mathscr{S}(\sigma)$.

6.2.3 Any region-based candidate strategy is almost-sure winning.

Yet, positive probability of winning of strategy σ^* guaranteed by Proposition 6.2.6 is not sufficient, in the following we need to show that σ^* wins almost surely.

Remark 6.2.11. *If we consider the restricted case of 1-player games with bounded intervals and exponentially distributed unbounded events, we can already easily prove that σ^* is almost-sure winning using [ACD92] as follows. Fixing σ^* resolves all non-determinism and yields a system of the type considered by [ACD92]. Since we are guaranteed the positive probability of reaching the target, we may apply Lemma 3 of [ACD92]. However, in the setting of two-player games, we cannot use this argument directly and some (non-trivial) changes are required.*

To finish the proof of the main theorem for two-player games, we show that every region-based candidate strategy wins with probability 1. This technique is similar to the one used in Section 5.1.1. Note that the probabilities to reach the target, guaranteed to be positive by Proposition 6.2.6, can be arbitrarily small. Assume that these probabilities rapidly decrease as the play goes on; the target is in such case reached with probability < 1. We need to rule out this case and show that the probabilities to reach the target are bounded from below by a positive constant.

In order to bound from below the probabilities of reaching the target, we again use the technique of Alur et al. [ACD92]. We restrict ourselves to δ-*separated* and *bounded* parts of regions where the lower bound exists, see Proposition 6.2.13 below. Because these parts are reached infinitely often with probability one, as shown in Proposition 6.2.18, this restriction is without loss of generality. Note that the definition of δ-separation is analogous to Definition 5.1.5, only extended to the clocks of the observer as well.

Definition 6.2.12. *Let $\delta > 0$. A configuration (s,ξ,t,q,ν) is called* δ-separated *if for any $a,b \in \{0\} \cup \{\xi(e) \mid e \in \mathscr{E}\} \cup \{\nu(x) \mid x \in \mathscr{X}, \nu(x) \leq B\}$ we have either $|\langle a\rangle - \langle b\rangle| \geq \delta$ or $\langle a\rangle = \langle b\rangle$. Furthermore, it is called* bounded by b *if $\xi(e) < b$ for each $e \in \mathscr{E}$.*

Reaching the target from a δ-separated configuration

On the following two pages, we prove there is $n \in \mathbb{N}$ such that the probability of plays that reach the target within n steps is bounded from below. Recall that these plays are denoted by R_n^T.

Proposition 6.2.13. *Let σ^* be a region-based candidate and $\delta > 0$. There is $n \in \mathbb{N}$ and $\varepsilon > 0$ such that $\mathscr{P}^{\sigma^*,\pi}[\mathfrak{h}](R_n^T) > \varepsilon$ for any $\pi \in \Pi$ and any $\mathfrak{h}$ such that its last configuration z belongs to $\mathscr{S}(\sigma^*)$ and is δ-separated and bounded by C.*

The rest of this subsection is devoted to the proof. First, we formally relate δ-separation to a *transition* of bounded size (in Lemma 6.2.15) and to its lower bound on probability (in Lemma 6.2.16).

Definition 6.2.14. *For a configuration z, action a enabled in z, and a region r, a set of configurations $X \subseteq r$ is an a-*transition *from z if there is a mode m with $a(m) > 0$, event e scheduled in m, and interval of time $I = (u, u+\delta)$ such that $u < C$ and*

$$X = \{z' \mid z' \text{ can be entered from } z \text{ by waiting time } t' \text{ in } m \text{ for event } e, t' \in I\}.$$

Furthermore, δ is called the size *of the transition.*

As the target is reached in multiple steps, we need to keep some separation after each transition from a δ-separated configuration.

Lemma 6.2.15. *Let $b > 0$, $\delta > 0$, r_1 and r_2 be regions and a be an action with $r_1 \overset{a}{\rightsquigarrow} r_2$. For any δ-separated $z \in r_1$ bounded by b, there is an a-transition $X \subseteq r_2$ from z of size $\delta/3$ with all $z' \in X$ being $\delta/3$-separated and bounded by $b+C$.*

Proof. Let $z \in r_1$ be a δ-separated configuration bounded by b. By similar arguments as in Lemma 6.2.9, it is easy to see that there is an a-transition $X' \subseteq r$ of size δ with all $z \in X'$ being bounded by $b+C$. Indeed, waiting longer than C leads only to a region that is also reachable by waiting for any time $t \in (C-1, C)$. By taking the middle third of the interval corresponding to X' we get a transition $X \subset X'$ with all $z \in X$ being $(\delta/3)$-separated. □

Lemma 6.2.16. *Let $b > 0$ and $\delta > 0$. There is $c > 0$ such that for any a-transition X of size δ from a configuration z bounded by b, we have $P_{\mathscr{G}\times\mathscr{A}}(z,a;X) > c$.*

Proof. Let $\kappa = \min\{a(m) \mid a \in Act, m \in M, a(m) > 0\}$ be the minimal discrete transition probability in the game. Observe that for every m and $e \in \mathbf{E}(m)$, the function $\mathrm{Win}(m,e,\xi;t')$ is positive and continuous with respect to ξ and t' for ξ and t' such that $(\xi[m] \oplus_m t')(e') \leq u_{e'}$ for every $e' \in \mathbf{E}(m)$. Hence, if we restrict to the compact set of parameters with $0 \leq \xi(e') \leq b$ and $0 \leq t' \leq C + \delta$, this function attains minimum $y_{m,e} > 0$ for every $m \in M$ and $e \in \mathbf{E}(m)$. Let $y = \min\{y_{m,e} \mid m \in M, e \in \mathbf{E}(m)\}$. From the definition of the transition law, we get $P_{\mathscr{G}\times\mathscr{A}}(z,a;X) > c := \kappa \cdot \delta \cdot y$. □

Let us proceed with the proof of Proposition 6.2.13. For any $\mathfrak{h}$ ending in $\mathscr{S}(\sigma^*)$, we have $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma^*,\pi}[\mathfrak{h}](\mathrm{ReachReg}(\mathscr{R}_T)) > 0$ since σ^* is a candidate strategy. From the fact that σ^* is region-based and from Lemma 6.2.9 we have the same for reaching the target in up to $|\mathscr{R}|$ steps, i.e. $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma^*,\pi}[\mathfrak{h}](R^T_{|\mathscr{R}|}) > 0$. For a region $r \subseteq \mathscr{S}(\sigma^*)$ we thus denote by the *distance* of r the minimal $n \leq |\mathscr{R}|$ such that for any $\mathfrak{h}$ ending in r, we have $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma^*,\pi}[\mathfrak{h}](R^T_n) > 0$. The proposition follows directly from an inductive claim.

Claim 6.2.17. *For each $n < |\mathscr{R}|$ there is $\varepsilon_n > 0$ such that for any region r with distance n we have $\mathscr{P}^{\sigma^*,\pi}[\mathfrak{h}](R^T_n) > \varepsilon_n$ where $\mathfrak{h}$ is any history ending in r and its last configuration is $(\delta/3^{|\mathscr{R}|-n})$-separated and bounded by $C \cdot (|\mathscr{R}| - n)$.*

Proof. As regards the induction base, we set $\varepsilon_0 = 1$ and the claim follows. As regards the induction step, let $n < |\mathscr{R}|$, r be a region with distance n and $\mathfrak{h}$ be a history ending in a $(\delta/3^{|\mathscr{R}|-n})$-separated configuration $z \in r$ bounded by $C \cdot (|\mathscr{R}| - n)$. Observe that if r is at region of player □, there must be an action a and a region r' with distance $n-1$ such that σ^* takes a and $r \overset{a}{\rightsquigarrow} r'$. Due to Lemma 6.2.15, there is an a-transition $X \subseteq r'$ of size $\delta' = \delta/3^{|\mathscr{R}|-n+1}$ with all $z' \in X$ being δ'-separated and bounded by $C \cdot (|\mathscr{R}| - n + 1)$. For any $z' \in X$, we have from the induction hypothesis $\mathscr{P}^{\sigma^*,\pi}[\mathfrak{h}z'](R^T_{n-1}) > \varepsilon_{n-1}$. From Lemma 6.2.16, there is a bound c such that the transition is taken with probability at least c. Hence, setting $\varepsilon_n := \varepsilon_{n-1} \cdot c$, we have $\mathscr{P}^{\sigma^*,\pi}[\mathfrak{h}](R^T_n) > \varepsilon_n$.

If r is a region of player ◇, for each action a there must be a region r' with distance lower than n and with $r \overset{a}{\rightsquigarrow} r'$. For each such action we proceed analogously to the previous case using Lemma 6.2.15 and Lemma 6.2.16: there is an a-transition X of size δ' with probability c, each $z' \in X$ is δ'-separated and bounded by $C \cdot (|\mathscr{R}| - n + 1)$. Let ε^a_{n-1} be the probability bound from the induction hypothesis for action a, the claim holds for $\varepsilon_n := c \cdot \min\{\, \varepsilon^a_{n-1} \mid a \text{ enabled in } z\}$. □

Reaching a δ-separated configuration

What happens if in these n steps we do not reach $\mathscr{R}_T$ and leave the δ-separated parts of the regions? The proof that a candidate strategy is almost-sure winning is concluded by observing that regardless of the decisions of the players a δ-separated bounded history is reached almost surely. This way, we repeat infinitely many times a chance (with probability bounded from below) to reach the target.

Proposition 6.2.18. *There is $\delta > 0$ such that for any strategies $\sigma \in \Sigma$ and $\pi \in \Pi$, a δ-separated configuration bounded by C is reached almost surely from any $\mathfrak{h}$.*

We fix δ to $1/(6\cdot(|\mathscr{E}|+|\mathscr{X}|)+2)$. In the rest of this subsection, the proof is split into two phases. In the first phase in time t, with probability at least p a bounded configuration is reached that is *separable*. It means that any scheduled event occurs with probability bounded from below at such time that it *becomes (or stays) δ-separated*. In the second phase, for time t' the game moves with bounded probability among separable configurations; after that a δ-separated configuration is guaranteed to be reached. Each step of the second phase takes at least time δ with probability at least q. Thus, a δ-separated configuration is reached from any configuration in time $t+t'$ with probability at least $p \cdot q^{t'/\delta} > 0$. Hence, it is eventually reached with probability 1.

We say that a configuration $z = (s, \xi, t, q, \nu)$ is *separable* if there is a mapping f, called *separation plan*, that to every event $e \in \mathscr{E}$ with $\xi(e) \in (\ell_e, u_e)$ assigns an interval $(s_e, s_e + \delta) \subseteq [0,1]$ with $\langle s_e \rangle \geq \langle \xi(e) \rangle$ such that for any a, b from the set

$$\begin{aligned}\{0\} \cup \{\xi_n(e) \mid e \in \mathscr{E}\} \cup \{\nu_n(x) \mid x \in \mathscr{X}, \nu_n(x) \leq B\}\\ \cup \{s_e, s_e+\delta \mid e \in \mathscr{E}, \xi(e) \in (\ell_e, u_e)\}\end{aligned}$$

we have either $\langle a \rangle = \langle b \rangle$ or $|\langle a \rangle - \langle b \rangle| \geq \delta$. Further, we say that an event e' (or clock x') is *δ-separated in a configuration* $z = (s, \xi, t, q, \nu)$ if for $a = \xi(e')$ (or $a = \nu(x')$) and for any $b \in \{0\} \cup \{\xi_n(e) \mid e \in \mathscr{E}\} \cup \{\nu_n(x) \mid x \in \mathscr{X}, \nu_n(x) \leq B\}$ we have either $\langle a \rangle = \langle b \rangle$ or $|\langle a \rangle - \langle b \rangle| \geq \delta$. The first phase is formalized as follows.

Claim 6.2.19. *There is $t > 0$ and $\varepsilon > 0$ such that $\mathscr{P}^{\sigma,\pi}[\mathfrak{h}](R'_n) > \varepsilon$ for any $\sigma \in \Sigma$, $\pi \in \Pi$, and history $\mathfrak{h}$ where R'_n are the plays where the first configuration visited after time t is separable and bounded by $2r+C$.*

Proof. We split the first phase in two subphases. In the first subphase we reach a configuration bounded by $2r+C-1$, in the second subphase, we reach a separable configuration bounded by $2r+C$.

As regards the first subphase, recall that r bounds the conditional expected waiting times for any event e with $u_e = \infty$ by $\sup_{b>\ell_e} \int_0^\infty x \cdot f_{e|b}(x)\,dx \leq r$. Let us fix an event e with $u_e = \infty$. From this bound and from the Markov inequality, the probability that e occurs or stops being scheduled within time $2r$ is greater than $1 - r/2r = 1/2$, regardless of the decisions of σ and π. Hence, the probability that starting from $\mathfrak{h}$, all such events occur or stop being scheduled within time $2r$ is bounded by $1/2^{|\mathscr{E}|}$. With this probability, up to this time limit, a configuration (s, ξ, t, q, ν) is visited where (1) any event e with $u_e = \infty$ satisfies $\xi(e) < 2r$ and (2) any event e with $u_e < \infty$ satisfies $\xi(e) < u_e$. Such a configuration is bounded by $2r + C - 1$ since any finite u_e satisfies $u_e \leq C - 1$.

For the second subphase, let (s, ξ, t, q, ν) be any configuration bounded by $2r + C + 1$. Let E be the set of events e with $\ell_e < \xi(e) < u_e < \infty$. Since the configuration is bounded, there is a positive probability p that each event from E either occurs or stops being scheduled within $1/2$ time unit. The resulting configuration (s', ξ', t', q', ν') is bounded by $2r + C$ and all events e with $\ell_e < \xi'(e) < u_e < \infty$ satisfy $\langle \xi'(e) \rangle < 1/2$.

Hence, we set $t = 2r + 1/2$ and we define the mapping f for showing that such a configuration is separable as follows. The interval $[1/2, 1]$ can be partitioned into $3(|\mathscr{E}| + |\mathscr{X}|) + 1$ equally sized intervals of length δ. Some of them are "occupied" by events with infinite upper bound or by clocks, some other will be "occupied" by f. Each occupied interval requires its left and right neighbour to be empty. Each event thus takes up at most 3 intervals and 0 takes up the last interval, there are enough intervals to define f. □

The second phase takes time $t' = \max\{B, C\}$. As each transition in the second phase takes time at least δ, there are in total at most $m = t'/\delta$ transitions, each of bounded probability. Precisely:

Claim 6.2.20. *For any separable configuration z with separation plan f and any a enabled in z, there is an a-transition X from z induced by mode m, event e, and interval I such that*

(1) e is the event to come according to f, i.e. e has greatest $f(e)$ if $\mathrm{dom}(f) \neq \emptyset$,

(2) all events and clocks δ-separated in z, the event e, and all clocks reset by the transition are δ-separated in every $z' \in X$ (as e occurs according to f),

(3) the time the transition takes satisfies $|I| = \delta$ and $I \subseteq (\delta, C)$, and

(4) all $z' \in X$ are separable.

Proof. Let us fix z with separation plan f. Let e be the minimal event in f, i.e. $\inf f(e) < \inf f(e')$ for any other event e' defined by f. We can assume that f is maximal in the sense that adding another interval I into f results in a separation plan only if $\sup I < \inf f(e)$ (if not, we can replace I for $f(e)$ and obtain a greater separation plan). Furthermore, let us fix an action a and an arbitrary mode m with $a(m) > 0$. We fix an event e and interval I as follows.

- If $\mathbf{E}(m) \cap \mathrm{dom}(f) \neq \emptyset$, we set e to be the event from this set with the greatest interval in f and $I := (1 - \sup f(e), 1 - \inf f(e))$.
- Otherwise, let $w \in \mathbb{R}_{\geq 0}$ be the minimal waiting time such that there is an event $e \in \mathbf{E}(m)$ with $\xi(e) + w = \ell_e$. Further, let $a \in [0,1]$ be minimal such that for any time t from the interval $I = (w+a, w+a+\delta)$ we have:
  - $\delta < \langle \xi(e') + t \rangle < 1 - \delta$ for any $e' \in \mathbf{E}(m)$;
  - $\delta < \langle \nu(x) + t \rangle < 1 - \delta$ for any $x \in \mathscr{X}$.

  Such a time a exists since for any placement of $|\mathscr{E}| + |\mathscr{X}|$ "points" on the $[0,1]$ line segment there is at least one continuous interval of length $3\delta = 1/(|\mathscr{E}| + |\mathscr{X}|)$ without any such point inside.

We need to argue that the set of configurations X induced by z, m, e, and I is an a-transition from z. In both cases above, this set of configurations is a subset of one region (due to the definition of separation plan in the first case; due to the defining conditions on I in the second case). Furthermore, observe that e can indeed occur after waiting for any $t \in I$ since no other event runs out of its upper bound sooner.

Point (1) directly follows from the definition of e. As regards point (2), similarly, due to the definition of separation plan and due to the defining conditions on I, the event e and all clocks reset by this transition are δ-separated in every configuration from X.

As regards point (3), the size of I is clearly δ by definition. In the first case $I \subseteq (\delta, 1-\delta) \subseteq (\delta, C)$ thanks to the definition of separation plan. In the second case, all points in I are grater than δ because $a \geq \delta$. Indeed, $\langle \xi(e) + w \rangle = 0$ and $\langle \xi(e) + w + a \rangle \geq \delta$. Furthermore, all points in I are smaller than C since $w \leq C - 1$ and since $a < 1 - \delta$ by similar arguments as above.

As regards point (4), let us fix $z' = (s', \xi', t', q', nu') \in X$ corresponding to the waiting time $t' \in I$. We define f' using f as follows. Let $e' \in \mathbf{E}(m)$ such that $e' \neq e$ and $\xi(e') \in (\ell_{e'}, u_{e'})$. If e is an event previously in the separation plan, i.e. $e' \in \mathrm{dom}(f)$, we set $f'(e') = \{\langle a + t \rangle \mid a \in f(e')\}$.We sort the events $e' \notin \mathrm{dom}(f)$ that exceed *newly* their lower bound in the descending order by $\xi'(e') - \ell_{e'}$ and one by one set $f'(e')$ to the interval $(s_{e'}, s_{e'} + \delta)$ for the greatest $s_{e'} \in [0,1]$ such that f' is still a separation plan. In other words, each new event e' occupies the greatest available free "slot" of length 3δ that is greater than $\langle \xi(e') \rangle$. Such a $s_{e'}$ exists because of the maximality of f. Indeed, if the only empty slot is lower than $\langle \xi(e') \rangle$

this slot could have been taken as a greater slot for the event e in the separation plan f contradicting its maximality. $\square$

After time $t' \geq C$ each event occurs at least once since each event occurs within first time unit after exceeding its lower bound (due to the definition and (1) of Claim 6.2.20). All clocks that get reset in the second phase are δ-separated, the remaining clocks have value above $t' \geq B$. Hence, due to (2) of Claim 6.2.20, the configuration reached after the second phase is δ-separated. This concludes the proof of Proposition 6.2.18.

6.2.4 The algorithm

In this section, we show that the existence of a DTA almost-sure winning strategy is decidable in exponential time, and we also show how to compute such a strategy if it exists.

Due to Propositions 6.2.3 and 6.2.4, this problem can be equivalently considered in the setting of the product game $\mathscr{G} \times \mathscr{A}$ and its region-based strategies. First, we show that this problem can be further reduced to the problem of computing wining strategies in a *finite* stochastic reachability game $[\mathscr{G} \times \mathscr{A}]$ induced by the product game $\mathscr{G} \times \mathscr{A}$. Second, we provide an algorithm solving the problem using this reduction. Let us define the game $[\mathscr{G} \times \mathscr{A}]$:

- The set of vertices of $[\mathscr{G} \times \mathscr{A}]$ is $V = \mathscr{R} \cup \{\, (r,a) \mid r \in \mathscr{R}, a \text{ is enabled in } r \,\}$.
- For each control state s, the game starts with probability $\alpha_0(s)$ in the vertex $v_s \in \mathscr{R}$ such that $v_s = \{(s, \mathbf{0}, q, \mathbf{0})\}$ where q is the location visited after $\mathscr{A}$ reads s in its configuration $(q_0, \mathbf{0})$.
- The game moves from vertex to vertex, forming an infinite play $v_0 v_1 \cdots$. In each vertex $v_i \in \mathscr{R}$, the successor vertex v_{i+1} is chosen by one of the players, whereas in each vertex of the form $v_i = (r,a)$ the successor vertex v_{i+1} is chosen randomly.
  - Player $\odot \in \{\square, \lozenge\}$ controls a vertex $r \in \mathscr{R}$ if its control state s satisfies $s \in S_\odot$. The player $\odot$ chooses in such a vertex r one of the stochastic vertices (r,a) where a is enabled in r.
  - From each stochastic vertex of the form (r,a), there are transitions to all vertices $r' \in \mathscr{R}$, such that $r \overset{a}{\rightsquigarrow} r'$. The probability distribution on the set of outgoing transitions from each stochastic vertex is uniform.

Player $\square$ tries to reach the set $\mathscr{R}_T$ of target regions (which is the same as in the product game) and player $\lozenge$ tries to avoid it. We say that a strategy σ is positional if it satisfies $\sigma(v_1 \cdots v_n) = \sigma(v'_1 \cdots v'_m)$ for any histories with $v_n = v'_m$. Further, we say that a strategy σ of player $\Box$ is almost-sure winning if for any strategy π of player $\Diamond$, the set $\mathscr{R}_T$ is reached almost surely when playing according to σ and π.

The following proposition states the correctness of the reduction. Slightly abusing the notation, we consider region-based strategies to be strategies in both $\mathscr{G} \times \mathscr{A}$ and $[\mathscr{G} \times \mathscr{A}]$. Indeed, a region-based strategy for the product game $\mathscr{G} \times \mathscr{A}$ induces a unique positional strategy for the game $[\mathscr{G} \times \mathscr{A}]$, and vice versa.

Proposition 6.2.21. *Let $\mathscr{G}$ be a game and $\mathscr{A}$ be a DTA. A region-based strategy $\sigma \in \Sigma$ is almost-sure winning in $[\mathscr{G} \times \mathscr{A}]$ iff it is almost-sure winning in $\mathscr{G} \times \mathscr{A}$.*

Proof. Intuitively, it might seem surprising that we set all the probability distributions in $[\mathscr{G} \times \mathscr{A}]$ to be uniform. Since we are interested only in *qualitative* reachability and due to Lemma 6.2.9, we show that this is sufficient for our purposes.

"⇒" Let us fix a region-based strategy $\sigma \in \Sigma$ almost-sure winning in $[\mathscr{G} \times \mathscr{A}]$. For this fixed σ, let $\mathscr{S}_{\mathscr{G} \times \mathscr{A}}$ and $\mathscr{S}_{[\mathscr{G} \times \mathscr{A}]}$ be the regions that are reached with positive probability in $\mathscr{G} \times \mathscr{A}$ and $[\mathscr{G} \times \mathscr{A}]$ for some π, respectively. We show that (a) σ reaches the target in $\mathscr{G} \times \mathscr{A}$ with positive probability from any $\mathfrak{h}$ ending in $\mathscr{S}_{[\mathscr{G} \times \mathscr{A}]}$ and for any $\pi \in \Pi$ and that (b) $\mathscr{S}_{\mathscr{G} \times \mathscr{A}} \subseteq \mathscr{S}_{[\mathscr{G} \times \mathscr{A}]}$. These two points imply that σ is a candidate strategy and hence, by Propositions 6.2.13 and 6.2.18, σ is also an almost-sure winning strategy.

(a) Let $\mathfrak{h}$ be a history ending in $\mathscr{S}_{\mathscr{G} \times \mathscr{A}}$ in a region r. Let us fix an *arbitrary* strategy $\pi \in \Pi$. Let π' be a strategy that (1) satisfies $\pi'(z_0 \cdots z_n) = \pi(z'_0 \cdots z'_n)$ if $z_0 \sim z'_0, \ldots, z_n \sim z'_n$ and (2) after traversing any sequence of regions $r_0 \cdots r_n$ it chooses an action that is chosen with positive probability by π in a set of plays of non-zero measure that traverse the regions $r_0 \cdots r_n$. Such a strategy can be easily built inductively similarly to Claim 6.2.8. The strategy π' is also a (non-positional) strategy in $[\mathscr{G} \times \mathscr{A}]$. Since σ is winning in $[\mathscr{G} \times \mathscr{A}]$, it guarantees reaching the target $\mathscr{R}_T$ from r with positive probability in at most $2 \cdot |\mathscr{R}|$ steps (that correspond to $|\mathscr{R}|$ steps in $\mathscr{G} \times \mathscr{A}$) against any strategy π'. Hence, also in the product game, we have $\mathscr{P}^{\sigma,\pi'}_{\mathscr{G} \times \mathscr{A}}[\mathfrak{h}](R^T_{|\mathscr{R}|}) > 0$ where $R^T_{|\mathscr{R}|}$ are the plays that reach the target within $|\mathscr{R}|$ steps. From the definition of π', we also have $\mathscr{P}^{\sigma,\pi'}_{\mathscr{G} \times \mathscr{A}}[\mathfrak{h}](R'_{|\mathscr{R}|}) > 0$ where $R'_{|\mathscr{R}|}$ are the plays that reach the target within $|\mathscr{R}|$ steps and where the strategies π' and π take the same actions. Hence, $\mathscr{P}^{\sigma,\pi}_{\mathscr{G} \times \mathscr{A}}[\mathfrak{h}](R'_{|\mathscr{R}|}) > 0$ and thus $\mathscr{P}^{\sigma,\pi}_{\mathscr{G} \times \mathscr{A}}[\mathfrak{h}](\text{ReachReg}(T)) > 0$.

(b) For the sake of contradiction, let $\pi \in \Pi$ and $r_0 \cdots r_n$ be a sequence of regions such that $\mathscr{P}^{\sigma,\pi}_{\mathscr{G} \times \mathscr{A}}(\mathit{Pattern}(r_0 \cdots r_n)) > 0$ and r_n is the first region in this sequence not in $\mathscr{S}_{[\mathscr{G} \times \mathscr{A}]}$. If r_{n-1} belongs to $\Box$, r_n is reached also in $[\mathscr{G} \times \mathscr{A}]$ as

$\Box$ plays the same in both games and the construction preserves reachability in one step. If r_{n-1} belongs to $\Diamond$, there must be (from Lemma 6.2.9) an action a such that $r_{n-1} \overset{a}{\rightsquigarrow} r_n$. There is a strategy π using which the region r_{n-1} is reached in $[\mathscr{G} \times \mathscr{A}]$ and that chooses a in r_{n-1}. The strategy π also reaches r_n. In both cases we obtain a contradiction $r_{k+1} \in \mathscr{S}_{[\mathscr{G}\times\mathscr{A}]}$.

"$\Leftarrow$" We show the "if" part by contraposition. Assume that a positional strategy σ is not almost-sure winning in $[\mathscr{G} \times \mathscr{A}]$. Observe that there is π and a region bad with $\mathscr{P}^{\sigma,\pi}_{[\mathscr{G}\times\mathscr{A}]}(\text{ReachReg}(\{bad\})) > 0$ such that from bad the probability to reach $\mathscr{R}_T$ is zero. For this pair of region-based strategies σ and π, the same holds also for $\mathscr{G} \times \mathscr{A}$. Namely, $\mathscr{P}^{\sigma,\pi}_{\mathscr{G}\times\mathscr{A}}(\text{ReachReg}(\{bad\})) > 0$ and from the construction of $[\mathscr{G} \times \mathscr{A}]$ and from Lemma 6.2.9, $\mathscr{P}^{\sigma,\pi}_{\mathscr{G}\times\mathscr{A}}[\mathfrak{h}](\text{ReachReg}(\mathscr{R}_T)) = 0$ for any history $\mathfrak{h}$ ending in bad. Hence, the region-based strategy σ is not almost-sure winning. $\Box$

The reduction to the finite game now straightforwardly yields Algorithm 1 that solves the problem introduced by Theorem 6.2.2. The algorithm uses the following symbolic representation of regions. Similarly to [AD94], a region visited with positive probability can be represented by a triple (s,q,Ξ) where Ξ is called an *area* and contains

- for every element $x \in \mathscr{E} \cup \mathscr{X}$, one constraint from the set
$$\{x = 0\} \cup \{c-1 < x < c \mid 1 \leq c \leq D\} \cup \{x > D\};$$
- for all $x,y \in \mathscr{E} \cup \mathscr{X}$ with constraints $c-1 < x < c$ and $d-1 < y < d$ in (1) for some $c,d \in \mathbb{N}$, one constraint from the set $\{\langle x \rangle < \langle y \rangle, \langle x \rangle = \langle y \rangle, \langle y \rangle < \langle x \rangle\}$.

Observe that the set of all these triples is finite and can be easily constructed. Further, the algorithm uses the following operations over areas. Let Ξ be an area.

Reset Let $X \subseteq (\mathscr{E} \cup \mathscr{X})$ be a reset set. By $\Xi[X := 0]$ we denote an area obtained from Ξ by removing all constraints with any $x \in X$ and adding $x = 0$ for any $x \in X$, $\langle x \rangle = \langle y \rangle$ for any $x,y \in X$, and $\langle x \rangle < \langle y \rangle$ for any $x \in X$ and $y \in (\mathscr{E} \cup \mathscr{X}) \setminus X$.

Guard satisfaction For a clock constraint g, we write $\Xi \models g$ if the set of configurations satisfying the constraints Ξ is the same as the set of configurations satisfying the constraints $\Xi \cup g$. By case distinction, this is guaranteed if

1. $g = g_1 \wedge g_2$ and $\Xi \models g_1$ and $\Xi \models g_2$;
2. $g = x \leq b$ or $g = x < b$ and Ξ contains $x = 0$ or $c-1 < x < c$ for $c \leq b$;
3. $g = x \geq b$ or $g = x > b$ and Ξ contains $x > D$ or $c-1 < x < c$ for $c-1 \geq b$.

Time successor Let m be a mode. The time-successor(Ξ,m) is *undefined* if Ξ contains a constraint $x > D$ for all x from $\mathbf{E}(m) \cup \mathscr{X}$; or Ξ contains a constraint $u_e - 1 < e < u_e$ and no constraint of the form $\langle e\rangle < \langle x\rangle$ for some $e \in \mathbf{E}(m)$. Otherwise, time-successor(Ξ,m) is an area obtained from Ξ as follows.

- If Ξ contains no constraints of the form $c-1 < x < c$, then constraints of the form $x = 0$ are replaced by $0 < x < 1$ for all $x \in \mathbf{E}(m) \cup \mathscr{X}$.
- Otherwise, for every $x \in (\mathbf{E}(m) \cup \mathscr{X})$ such that Ξ contains a constraint of the form $c-1 < x < c$ but no constraint $\langle x\rangle < \langle y\rangle$ (i.e., that has maximal fractional part), we replace in Ξ each $\langle y\rangle < \langle x\rangle$ by $\langle x\rangle < \langle y\rangle$ and we replace $c-1 < x < c$ by $c < x < c+1$ if $c < D$ or by $x > D$, otherwise.

After properly defining all parts of Algorithm 1, observe that its correctness follows from Proposition 6.2.21 and from the following proposition.

Proposition 6.2.22. *Procedure* Construct *terminates and constructs* $[\mathscr{G} \times \mathscr{A}]$.

Proof. The algorithm terminates because for any area Ξ there is a n such that time-successor$^n(\Xi)$ is undefined. Indeed, in each iteration of time-successor either c is incremented in the constraint $c-1 < x < c$ for at least one element x or such a constraint is replaced by $x > D$.

As regards constructing $[\mathscr{G} \times \mathscr{A}]$, the only non-trivial part is that $(r,a) \rightsquigarrow r'$ if and only if $r \overset{a}{\rightsquigarrow} r'$. Observe that the procedure Construct satisfies that if $(r,a) \rightsquigarrow r'$, then for some $z \in r$ the set r' is an a-transition from z of positive size (see Section 6.2.3). Hence $P_{\mathscr{G}\times\mathscr{A}}(z,a;r') > 0$ and thus $r \overset{a}{\rightsquigarrow} r'$. If we have $r \overset{a}{\rightsquigarrow} r'$, then r' is an a-transition of positive size from some $z \in r$. The interval corresponding to the a-transition is abstractly captured by the "flow of time" computed by time-successor$^n(\Xi)$ for some $n \in \mathbb{N}$ and hence, the transition $(r,a) \rightsquigarrow r'$ is added within the procedure Construct. □

We conclude the section by the complexity analysis of Algorithm 1. Since there are exponentially many regions (w.r.t. the number of clocks and events), the size of $[\mathscr{G} \times \mathscr{A}]$ is exponential in the size of $\mathscr{G}$ and $\mathscr{A}$. Note that two-player stochastic games with qualitative reachability objectives are easily solvable in polynomial time [AHK98]. Due to Propositions 6.2.3 and 6.2.4 there is a (DTA) almost-sure winning strategy in $\mathscr{G}$ with $\mathscr{A}$ iff there is a region-based almost-sure winning strategy in $\mathscr{G} \times \mathscr{A}$; and due to Proposition 6.2.21, there is a region-based almost-sure winning strategy in $\mathscr{G} \times \mathscr{A}$ iff there is an almost-sure winning strategy in $[\mathscr{G} \times \mathscr{A}]$. Furthermore, the (positional) almost-sure winning strategy in $[\mathscr{G} \times \mathscr{A}]$ can be transformed into a DTA almost-sure winning strategy in $\mathscr{G}$ with $\mathscr{A}$. Since all transformations of the strategies are trivially effective, we conclude the proof of Theorem 6.2.2 by the following proposition.

Algorithm 1: Decide whether □ has a DTA almost-sure winning strategy

```
input  : $\mathcal{G} = (S_\Box, S_\Diamond, M, \mathcal{E}, \mathbf{E}, \mathrm{Succ}, Act, A, \alpha_0)$, $\mathcal{A} = (Q, \Sigma, \mathcal{X}, \longrightarrow, q_0, T)$
output: YES + an almost-sure winning DTA strategy if it exists; NO if not

$[\mathcal{G} \times \mathcal{A}] \leftarrow$ Construct($\mathcal{G}$, $\mathcal{A}$)
$W \leftarrow$ compute the set of almost-sure winning vertices in $[\mathcal{G} \times \mathcal{A}]$ by [AHK98]
$\sigma \leftarrow$ compute the optimal positional strategy in $[\mathcal{G} \times \mathcal{A}]$ by [AHK98]
if every initial vertex in $[\mathcal{G} \times \mathcal{A}]$ belongs to $W$ then
    return YES + the DTA strategy $\sigma^{\downarrow}$ induced by the region-based $\sigma$
else
    return NO
```

Procedure Construct($\mathcal{G}$,$\mathcal{A}$)

```
input  : $\mathcal{G} = (S_\Box, S_\Diamond, M, \mathcal{E}, \mathbf{E}, \mathrm{Succ}, Act, A, \alpha_0)$, $\mathcal{A} = (Q, \Sigma, \mathcal{X}, \longrightarrow, q_0, T)$
output: game $[\mathcal{G} \times \mathcal{A}] = (V, (V_\Box, V_\Diamond, V_\bigcirc), \rightsquigarrow, P, init)$

$V_\odot \leftarrow S_\odot \times Q \times N$ for both $\odot \in \{\Box, \Diamond\}$ where $N$ is the set of areas
$V_\bigcirc \leftarrow \{(v,a) \mid v \in V_\Box \cup V_\Diamond, a \in Act, a$ is enabled in $v\}$
$V \leftarrow V_\Box \cup V_\Diamond \cup V_\bigcirc$
$\rightsquigarrow \leftarrow \emptyset$
for all $s \in S$ and for $q$ with $(q_0, s, g, X, q) \in \longrightarrow$, and $g \models \mathbf{0}$ do
    $init((s, q, \{x = 0 \mid x \in \mathcal{E} \cup \mathcal{X}\})) \leftarrow \alpha_0(s)$
for all $(v,a) \in V_\bigcirc$ with $v = (s, q, \Xi)$ do
    add to $\rightsquigarrow$ a pair $(v, (v,a))$
    for all modes $m \in M$ with $a(m) > 0$ do
        $\Xi_m \leftarrow \Xi[X := 0]$ where $X = \mathcal{E} \setminus \mathbf{E}(m)$
        while $\Xi_m$ has a time successor in $m$ do
            $\Xi_m \leftarrow$ time-successor$(\Xi_m, m)$
            for all events $e$ with $\Xi_m \models \ell_e \leq e \leq u_e$ do
                $s' \leftarrow \mathrm{Succ}(m, e)$
                for all edges $(q, s', g, X, q')$ in the automaton $\mathcal{A}$ do
                    if $\Xi_m \models g$ then
                        $\Xi' \leftarrow \Xi_m[X \cup \{e\} := 0]$
                        add to $\rightsquigarrow$ a pair $((v,a), (s', q', \Xi'))$
    $P((v,a)) \leftarrow$ uniform distribution
return $(V, (V_\Box, V_\Diamond, V_\bigcirc), \rightsquigarrow, P)$
```

Proposition 6.2.23. *Let $\mathscr{G}$ be a GSMG and $\mathscr{A}$ be a DTA. In time exponential in $|\mathscr{G}|$ and $|\mathscr{A}|$, Algorithm 1 decides whether player $\Box$ has a (DTA) almost-sure winning strategy and computes it if it exists.*

6.3 Büchi specifications

In this section, we show that the results from qualitative reachability specifications can be employed to solving qualitative Büchi specifications. Let us fix a Büchi specification $\text{Büchi}_{\mathscr{A}}(T)$ over a set of locations T and restrict to the value 1. Similarly to Proposition 6.2.1, Figure 6.2 also shows that player $\Box$ is not guaranteed to have an optimal strategy for the Büchi specification.

Theorem 6.3.1. *If there is (some) strategy of player $\Box$ that is almost-sure winning with respect to* $\text{Büchi}_{\mathscr{A}}(T)$*, then there is also a DTA almost-sure winning strategy. Furthermore, there exists an algorithm that in exponential time (1) decides whether there is a DTA almost-sure winning strategy and (2) computes it if it exists.*

Compared to the previous section, the proof is rather straightforward by using the results for reachability. We split the statement into two claims.

Claim 6.3.2. *If there is (some) strategy of player $\Box$ that is almost-sure winning with respect to* $\text{Büchi}_{\mathscr{A}}(T)$*, then there is also a DTA almost-sure winning strategy.*

Proof. First, observe that the product construction can be directly applied to the Büchi specification. Let $\text{Büchi}(T)$ denote the set of plays in the product game that visit T infinitely often. Let us fix a strategy σ almost-sure winning in $\mathscr{G} \times \mathscr{A}$ which we have thanks to Propositions 6.2.3 (easily adapted to the Büchi specification).

Observe that there must be a set of configurations from the target regions $Z \subseteq \bigcup_{r \in \mathscr{R}_T} r$ such that

1. the set Z is reached with probability one, i.e. $\inf_{\pi \in \Pi} \mathscr{P}^{\sigma,\pi}_{\mathscr{G}\times\mathscr{A}}(\text{Reach}(Z)) = 1$;
2. from Z, the set Z is reached again with probability one. For the following proof it suffice to claim that for every $z \in Z$, there is a history $\mathfrak{h}z$ such that $\inf_{\pi \in \Pi} \mathscr{P}^{\sigma,\pi}_{\mathscr{G}\times\mathscr{A}}[\mathfrak{h}z](\text{Reach}'(Z)) = 1$ where $\text{Reach}'(Z)$ is the set of plays that reach a configuration in Z after at least one step.

Firstly, let R be the smallest set of regions containing Z, i.e. for each $r \in R$ there is some $z \in Z$ such that $z \in r$. From the first point and from Propositions 6.2.6, 6.2.13, and 6.2.18, there is a region-based strategy σ' such that

$$\inf_{\pi \in \Pi} \mathscr{P}^{\sigma',\pi}_{\mathscr{G}\times\mathscr{A}}(\text{ReachReg}(R)) = 1. \tag{6.3}$$

Secondly, let us fix any region $r \in R$ and any $z \in Z \cap r$. Furthermore, let $\mathfrak{h}z$ be the history satisfying the second point above. Let us alter the game a bit: we add into the game another non-target copy $\bar{r}$ of the region r. The altered game, denoted by $(\mathscr{G} \times \mathscr{A})_r$, starts in the configuration $\bar{z} \in \bar{r}$ corresponding to z. From the fact that $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma,\pi}_{\mathscr{G}\times\mathscr{A}}[\mathfrak{h}z](\mathrm{Reach}'(Z)) = 1$, we get $\inf_{\pi\in\Pi} \mathscr{P}^{\sigma',\pi}_{(\mathscr{G}\times\mathscr{A})_r}(\mathrm{Reach}(Z)) = 1$ where $\sigma'(\mathfrak{h}') = \sigma(\mathfrak{h}\mathfrak{h}')$. Hence, we again obtain by Propositions 6.2.6, 6.2.13, and 6.2.18 a region-based strategy σ_r that guarantees reaching R in $(\mathscr{G} \times \mathscr{A})_r$ almost surely. Furthermore, observe that by Lemma 6.2.9, σ_r also guarantees

$$\inf_{\pi\in\Pi} \mathscr{P}^{\sigma_r,\pi}_{\mathscr{G}\times\mathscr{A}}[\mathfrak{h}'](\mathrm{ReachReg}'(R)) = 1 \qquad \text{for any } \mathfrak{h}' \text{ ending in } r \tag{6.4}$$

where $\mathrm{ReachReg}'(R)$ are the set of plays that reach a region in R after at least one step. We define the region-based strategy σ^* using the strategy σ' and the strategies $(\sigma_r)_{r\in R}$. The strategy behaves as σ' until R is reached; then it behaves as σ_r if the last visited region from R is r. As region-based strategies take actions only based on the *current* region, we need to encode into the regions the information necessary for σ^* to decide.

To this end, let us create another DTA observer $\mathscr{A}'$ with set of locations $Q \times \mathscr{R} \times (\{0\} \cup R)$ that in the first component behaves as $\mathscr{A}$; in the second component it stores the current region of $\mathscr{G} \times \mathscr{A}$ which it simulates on the fly; and in the third component it stores the last visited region from R (initially set to 0). The construction to update the second component is straightforward and fully explained in the proof of Proposition 6.2.4, hence we do not repeat it here. Once the second component stores the current region, the construction to update the third component is straightforward as well. We base the target locations $T' = \{(q, r, r') \mid q \in T\}$ on the first component.

We define the region-based strategy σ^* for $\mathscr{G} \times \mathscr{A}'$. For a region r of $\mathscr{G} \times \mathscr{A}'$, let $[r]$ denote the corresponding region of $\mathscr{G} \times \mathscr{A}$ by projecting out the second and the third component of $\mathscr{A}'$. In every region r of $\mathscr{G} \times \mathscr{A}'$ with location of the form $(q, r', 0)$, we set $\sigma^*(r) = \sigma'([r])$. In every region r with location of the form (q, r', r''), we set $\sigma^*(r) = \sigma_{r''}([r])$. Thanks to (6.3) and (6.4) we show that

$$\begin{aligned} &\inf_{\pi\in\Pi} \mathscr{P}^{\sigma^*,\pi}_{\mathscr{G}\times\mathscr{A}'}(\mathrm{ReachReg}(R)) = 1, \\ &\inf_{\pi\in\Pi} \mathscr{P}^{\sigma^*,\pi}_{\mathscr{G}\times\mathscr{A}'}[\mathfrak{h}'](\mathrm{ReachReg}'(R)) = 1 \qquad \text{for any } \mathfrak{h}' \text{ ending in } R. \end{aligned}$$

Assume the opposite that there is a strategy π in the game $\mathscr{G} \times \mathscr{A}'$ (and history $\mathfrak{h}'$) contradicting one of the equalities. As the additional information in the locations of $\mathscr{A}'$ is deterministic and based on the history, we easily obtain a strategy π' in $\mathscr{G} \times \mathscr{A}$ that contradicts one of the equalities (6.3) and (6.4).

The two equalities above imply that σ^* is almost sure winning in $\mathscr{G} \times \mathscr{A}'$ with respect to Büchi(T'). Thanks to Proposition 6.2.4 (easily adapted to the Büchi specification), we get a DTA strategy that is almost sure winning w.r.t. Büchi$_{\mathscr{A}'}(T')$. As for each play ω we have $\omega \in$ Büchi$_{\mathscr{A}'}(T')$ iff $\omega \in$ Büchi$_{\mathscr{A}}(T)$, this strategy is also almost-sure winning w.r.t. Büchi$_{\mathscr{A}}(T)$. □

The algorithm to solve the problem is the same as Algorithm 1, only the polynomial algorithm to solve qualitative stochastic reachability games on lines 2 and 3 is replaced by the polynomial algorithm to solve qualitative stochastic Büchi games [JKH02]. Hence, the complexity is the same as in Theorem 6.2.2. It remains to show that this algorithm is correct.

Claim 6.3.3. *With respect to visiting T infinitely often, a region-based strategy $\sigma \in \Sigma_{\mathscr{G}\times\mathscr{A}}$ is almost-sure winning in $[\mathscr{G} \times \mathscr{A}]$ iff it is almost-sure winning in $\mathscr{G} \times \mathscr{A}$.*

Proof. We again reuse the results from reachability specifications.

"⇒" Let us fix a region-based strategy $\sigma \in \Sigma_{\mathscr{G}\times\mathscr{A}}$ almost-sure winning in $[\mathscr{G} \times \mathscr{A}]$. Then there is a set of regions $R \subseteq \mathscr{R}_T$ such that

1. σ is almost-sure winning with respect to reaching R and
2. from any history ending in R it guarantees revisiting R.

Similarly to the previous proof, let $(\mathscr{G} \times \mathscr{A})_z$ denote a product game where we duplicate the region of z, make it non-target, and start in configuration of the duplicate region corresponding to z. From Proposition 6.2.21, we get that σ is almost-sure winning in $\mathscr{G} \times \mathscr{A}$ with respect to visiting R infinitely often because

1. σ is also almost-sure winning in $\mathscr{G} \times \mathscr{A}$ with respect to reaching R.
2. Let z be any configuration from a region from R. The strategy σ is also almost-sure winning in $[(\mathscr{G} \times \mathscr{A})_z]$ with respect to reaching R. Hence, σ is also almost-sure winning in $(\mathscr{G} \times \mathscr{A})_z$ with respect to reaching R.

"⇐" Assume that a positional strategy σ is not almost-sure winning in $[\mathscr{G} \times \mathscr{A}]$. Observe that there is π and a region $bad \in T$ with $\mathscr{P}^{\sigma,\pi}_{[\mathscr{G}\times\mathscr{A}]}(\text{ReachReg}(\{bad\})) > 0$ such that from bad the probability to revisit T is smaller than 1. For this pair of region-based strategies σ and π, the same holds also for $\mathscr{G} \times \mathscr{A}$. Namely, the probability $\mathscr{P}^{\sigma,\pi}_{\mathscr{G}\times\mathscr{A}}(\text{ReachReg}(\{bad\})) > 0$ and from the construction of $[\mathscr{G} \times \mathscr{A}]$ and from Lemma 6.2.9, $\mathscr{P}^{\sigma,\pi}_{\mathscr{G}\times\mathscr{A}}[\mathfrak{h}](\text{ReachReg}'(T)) = 0$ for any history $\mathfrak{h}$ ending in bad. Hence, the region-based strategy σ is not almost-sure winning. □

This claim concludes the proof of Theorem 6.3.1 and also the whole Chapter 6.

Chapter 7

Conclusions

In the thesis we have studied stochastic stability of discrete-event systems enriched with hard real-time bounds. Most of the text has addressed the formalism of generalized semi-Markov processes with fixed-delay events.

We have discovered unstable behaviour previously unnoticed, thus contradicting several previous results. Namely, we have found that the previous verification algorithm of GSMP with fixed-delay events against qualitative Büchi specifications is incorrect. Furthermore, by showing that GSMP with fixed-delay events (and various related formalism) do not need to *have* their steady-state distribution, we have shown that various previous algorithms for *approximating* the steady-state distribution are incorrect on these models.

Then, we have described a stable subclass of single-ticking GSMP with fixed-delay events. By a technically demanding proof, we have shown that every single-ticking GSMP has its frequency measures almost-surely well-defined. We have also proven stability of GSMP observed by DTA by reducing them to single-ticking GSMP. Furthermore, we have proven stability of almost-monotone DSPN again by reducing them to single-ticking GSMP. Finally, we have shown that our stability results imply effective approximability of the frequency measures.

Finally, we have defined a novel game extension of GSMP, namely the generalized semi-Markov games. It is a two-player turn-based complete-information game where each player controls its set of control states and chooses one of finitely many actions. The players control the game only after each occurrence of an event. The stochastic nature of the model when waiting for events is the same as in GSMP. We have addressed the qualitative analysis of GSMG observed by DTA with reachability and Büchi specifications. As our goal has been better understanding of the structure of the game, we have studied the structure of almost-sure winning strategies. By a delicate proof, we have shown that strategies of finite structure – that can be captured by a DTA – suffice for almost-sure winning of $\Box$. A part of the proof techniques are similar to the techniques used in the previous chapter. However, various additional insights were necessary for the proofs in the game setting.

Overall, we believe that the thesis has filled a gap in the literature by providing a fundamental material on DES with fixed-delay events. Fixed-delay events are an important modelling concept for probabilistic verification as well as performance evaluation. Therefore, better understanding the structure of the models and the boundaries of what can be analysed is crucial.

7.1 Future and ongoing work

The subclass of single-ticking GSMP is not an exact characterization of stable GSMP. One promising area for future work is exploring the models beyond this subclass. There are many questions to ask:

1. Is it possible to provide an exact characterization of models where all the frequencies are almost-surely well-defined?
2. If yes, is the region graph a sufficient structure to characterize the qualitative behaviour of all stable models? This question is important from the algorithmic point of view as finding regions (or elements of some other finite partition of the space of configurations) that are revisited with probability one is a crucial step in the quantitative analysis of the models.
3. Given a model and one of its states such that the frequencies of this state are almost-surely well-defined, is it possible to approximate the distribution of the frequencies?

Based on our ongoing work, we conjecture that the questions 1. and 3. have a negative answer and that these problems are undecidable.

Another wide area for future research are the two-player games over discrete-event systems.

- Inspired by the recent application of two player games to compositional verification of Interactive Markov Chains [BHK+12; HKK13], we propose to adapt this approach to the non-Markovian setting. The external player $\Diamond$ in their setting also controls timing of some (external) events and models this way the unknown environment of the component to verify. As the non-Markovian behaviour brings many complications in the analysis, it may be advisable to start with a restricted class of models, such as only with exponential and fixed-delay events.
- Another possibility is to extend the analysis of GSMG to the *quantitative* case. First of all, is the quantitative analysis decidable? Are there any efficient algorithms for interesting types of specifications? Are there ε-optimal strategies with some reasonably small representation?

All in all, the are various promising directions for extending the research presented in the thesis.

Bibliography

[AB06] R. Alur and M. Bernadsky. “Bounded Model Checking for GSMP Models of Stochastic Real-Time Systems”. In: *Proceedings of 9th International Workshop Hybrid Systems: Computation and Control (HSCC)*. Vol. 3927. LNCS. Springer, 2006, pp. 19–33.

[ACB84] M. Ajmone Marsan, G. Conte, and G. Balbo. “A class of generalized stochastic Petri nets for the performance evaluation of multiprocessor systems”. In: *ACM Transactions on Computer Systems (TOCS)* 2.2 (1984), pp. 93–122.

[ACD91] R. Alur, C. Courcoubetis, and D. Dill. “Model-Checking for Probabilistic Real-Time Systems”. In: *Proceedings of ICALP’91*. Vol. 510. LNCS. Springer, 1991, pp. 115–136.

[ACD92] R. Alur, C. Courcoubetis, and D. Dill. “Verifying Automata Specifications of Probabilistic Real-Time Systems”. In: *Real-Time: Theory in Practice*. Vol. 600. LNCS. Springer, 1992, pp. 28–44.

[AD94] R. Alur and D. Dill. “A Theory of Timed Automata”. In: *Theoretical Computer Science* 126.2 (1994), pp. 183–235.

[AHK98] L. de Alfaro, T. Henzinger, and O. Kupferman. *Concurrent Reachability Games*. Tech. rep. UCB/ERL M98/33. EECS Department, University of California, Berkeley, 1998.

[Alf97] L. de Alfaro. “Formal verification of probabilistic systems”. PhD thesis. Stanford University, 1997.

[Alf98] L. de Alfaro. “How to Specify and Verify the Long-Run Average Behavior of Probabilistic Systems”. In: *Proceedings of LICS’98*. IEEE, 1998, pp. 454–465.

[AM97] S. Amari and R. Misra. “Closed-Form Expressions for Distribution of Sum of Exponential Random Variables”. In: *IEEE Transactions on Reliability* 46 (1997), pp. 519–522.

[ASS+00] A. Aziz, K. Sanwal, V. Singhal, and R. Brayton. “Model-checking continuous-time Markov chains”. In: *ACM Transactions on Computational Logic* 1.1 (2000), pp. 162–170.

[BA07] M. Bernadsky and R. Alur. “Symbolic Analysis for GSMP Models with One Stateful Clock”. In: *Proceedings of 10th International Workshop Hybrid Systems: Computation and Control (HSCC)*. Vol. 4416. LNCS. Springer, 2007, pp. 90–103.

[BBB+07] C. Baier, N. Bertrand, P. Bouyer, T. Brihaye, and M. Größer. “Probabilistic and topological semantics for timed automata”. In: *Proceedings of the 27th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*. Springer, 2007, pp. 179–191.

[BBB+08a] C. Baier, N. Bertrand, P. Bouyer, T. Brihaye, and M. Grosser. “Almost-sure model checking of infinite paths in one-clock timed automata”. In: *Proceedings of the 23th Annual IEEE Symposium on Logic in Computer Science (LICS)*. IEEE. 2008, pp. 217–226.

[BBB+08b] N. Bertrand, P. Bouyer, T. Brihaye, and N. Markey. “Quantitative model-checking of one-clock timed automata under probabilistic semantics”. In: *Proceedings of the 5th International Conference on Quantitative Evaluation of Systems (QEST)*. IEEE. 2008, pp. 55–64.

[BBJ+12] P. Bouyer, T. Brihaye, M. Jurdzinski, and Q. Menet. “Almost-sure model-checking of reactive timed automata”. In: *Proceedings of the 9th International Conference on Quantitative Evaluation of Systems (QEST)*. IEEE. 2012, pp. 138–147.

[BC84] A. Bobbio and A. Cumani. “Discrete state stochastic systems with phase type distributed transition times”. In: *Proceedings of the AMSE International Conference on Modelling and Simulation*. 1984, pp. 173–192.

[BCR+10] G. Bucci, L. Carnevali, L. Ridi, and E. Vicario. “Oris: a tool for modeling, verification and evaluation of real-time systems”. In: *International journal on software tools for technology transfer* 12.5 (2010), pp. 391–403.

[BdH+06] H. Bohnenkamp, P. R. d’Argenio, H. Hermanns, and J.-P. Katoen. “MoDeST: A compositional modeling formalism for hard and softly timed systems”. In: *Software Engineering, IEEE Transactions on* 32.10 (2006), pp. 812–830.

[BF09] P. Bouyer and V. Forejt. “Reachability in stochastic timed games”. In: *Automata, Languages and Programming* (2009), pp. 103–114.

[BFK+09] T. Brázdil, V. Forejt, J. Krčál, J. Křetínský, and A. Kučera. “Continuous-Time Stochastic Games with Time-Bounded Reachability”. In: *Proceedings of the 29th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*. LIPIcs. Schloss Dagstuhl, 2009, pp. 61–72.

[BFK+13] T. Brázdil, V. Forejt, J. Krčál, J. Křetínský, and A. Kučera. “Continuous-time stochastic games with time-bounded reachability”. In: *Information and Computation* 224 (2013), pp. 46–70.

[BG02] M. Bravetti and R. Gorrieri. “The theory of interactive generalized semi-Markov processes”. In: *Theoretical Computer Science* 282.1 (2002), pp. 5–32.

[BH03] C. Baier and B. Haverkort. “Model-checking algorithms for continuous-time Markov chains”. In: *IEEE Transactions on software engineering* 29.6 (2003), pp. 524–541.

[BHH+05] C. Baier, B. Haverkort, H. Hermanns, and J. Katoen. “Model checking meets performance evaluation”. In: *ACM SIGMETRICS Performance Evaluation Review* 32.4 (2005), pp. 10–15.

[BHH+11] P. Buchholz, E. M. Hahn, H. Hermanns, and L. Zhang. “Model Checking Algorithms for CTMDPs”. In: *Proceedings of 23rd International Conference on Computer Aided Verification (CAV)*. Springer, 2011, pp. 225–242.

[BHK+05] C. Baier, H. Hermanns, J. Katoen, and B. Haverkort. “Efficient computation of time-bounded reachability probabilities in uniform continuous-time Markov decision processes”. In: *Theoretical Computer Science* 345.1 (2005), pp. 2–26.

[BHK+12] T. Brázdil, H. Hermanns, J. Krčál, J. Křetínský, and V. Řehák. “Verification of Open Interactive Markov Chains”. In: *Proceedings of 32th International Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*. LIPIcs. Schloss Dagstuhl, 2012, pp. 474–485.

[BCH+07] C. Baier, L. Cloth, B. Haverkort, M. Kuntz, and M. Siegle. “Model checking Markov chains with actions and state labels”. In: *IEEE Transactions on Software Engineering* (2007), pp. 209–224.

[BKK+10a] T. Brázdil, J. Krčál, J. Křetínský, A. Kučera, and V. Řehák. “Stochastic Real-Time Games with Qualitative Timed Automata Objectives”. In: *Proceedings of CONCUR 2010*. Vol. 6269. LNCS. Springer, 2010, pp. 207–221.

[BKK+10b] T. Brázdil, J. Krčál, J. Křetínský, A. Kučera, and V. Řehák. “Stochastic real-time games with qualitative timed automata objectives”. In: *Proceedings of 21st International Conference on Concurrency Theory (CONCUR)* (2010), pp. 207–221.

[BKK+11a] T. Brázdil, J. Krčál, J. Křetínský, A. Kučera, and V. Řehák. “Measuring Performance of Continuous-Time Stochastic Processes using Timed Automata”. In: *Proceedings of 14th International Conference on Hybrid Systems: Computation and Control (HSCC’11).* ACM Press, 2011, pp. 33–42.

[BKK+11b] T. Brázdil, J. Krčál, J. Křetínský, and V. Řehák. “Fixed-delay events in generalized semi-Markov processes revisited”. In: *Proceedings of 22nd International Conference on Concurrency Theory (CONCUR).* Springer, 2011, pp. 140–155.

[BKK+13] T. Brázdil, Ľ. Korenčiak, J. Krčál, J. Křetínský, and V. Řehák. “On time-average limits in deterministic and stochastic Petri nets”. In: *Proceedings of the ACM/SPEC International conference on performance engineering (ICPE).* Poster paper. ACM. 2013, pp. 421–422.

[BL13] L. Bortolussi and R. Lanciani. “Model Checking Markov Population Models by Central Limit Approximation”. In: *Proceedings of the 10th International Conference on Quantitative Evaluation of Systems (QEST).* Vol. 8054. LNCS. Springer, 2013, pp. 123–138.

[BPS+05] G. Bucci, R. Piovosi, L. Sassoli, and E. Vicario. “Introducing probability within state class analysis of dense-time-dependent systems”. In: *Proceedings of the 2nd International Conference on Quantitative Evaluation of Systems (QEST).* IEEE. 2005, pp. 13–22.

[BPS+98] A. Bobbio, A. Puliafito, M. Scarpa, and M. Telek. “Webspn: A web-accessible Petri net tool”. In: *Proceedings of the conference on Web-based Modeling & Simulation.* 1998.

[BS11] P. Buchholz and I. Schulz. “Numerical analysis of continuous time Markov decision processes over finite horizons”. In: *Computers & OR* 38.3 (2011), pp. 651–659.

[BS81] A. Barbour and R. Schassberger. “Insensitive average residence times in generalized semi-Markov processes”. In: *Advances in Applied Probability* (1981), pp. 720–735.

[CGL94] G. Ciardo, R. German, and C. Lindemann. “A characterization of the stochastic process underlying a stochastic Petri net”. In: *IEEE Transactions on Software Engineering* 20.7 (1994), pp. 506–515.

[CGV09] L. Carnevali, L. Grassi, and E. Vicario. “State-Density Functions over DBM Domains in the Analysis of Non-Markovian Models”. In: *IEEE Transactions on Software Engineering* 35.2 (2009), pp. 178–194.

[CKT93] H. Choi, V. G. Kulkarni, and K. S. Trivedi. “Transient analysis of deterministic and stochastic Petri nets”. In: *Proceedings of the 14th International Conference on Application and Theory of Petri Nets (ICATPN)*. Springer, 1993, pp. 166–185.

[CKT94] H. Choi, V. G. Kulkarni, and K. S. Trivedi. “Markov regenerative stochastic Petri nets”. In: *Performance Evaluation* 20.1 (1994), pp. 337–357.

[CL08] C. Cassandras and S. Lafortune. *Introduction to discrete event systems*. Springer, 2008.

[Cox55a] D. R. Cox. “A use of complex probabilities in the theory of stochastic processes”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 51 (02 Apr. 1955), pp. 313–319.

[Cox55b] D. R. Cox. “The analysis of non-Markovian stochastic processes by the inclusion of supplementary variables”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 51 (03 July 1955), pp. 433–441.

[CT92] A. Coyle and P. Taylor. “Bounds on the sensitivity of generalised semi-Markov processes with a single generally distributed lifetime”. In: *Mathematics of operations research* 17.1 (1992), pp. 132–147.

[DHS07] S. Donatelli, S. Haddad, and J. Sproston. “CSL^{TA}: an Expressive Logic for Continuous-Time Markov Chains”. In: *Proceedings of the 4th International Conference on Quantitative Evaluation of Systems (QEST)*. IEEE. 2007, pp. 31–40.

[Dil90] D. L. Dill. “Timing assumptions and verification of finite-state concurrent systems”. In: *Proceedings of the International Workshop on Automatic verification methods for finite state systems*. Springer. 1990, pp. 197–212.

[DKB97] P. R. D’Argenio, J.-P. Katoen, and H. Brinksma. “A stochastic automata model and its algebraic approach”. In: *Proceedings of the 5th International Workshop on Process Algebras and Performance Modeling*. Centre for Telematics and Information Technology University of Twente, 1997.

[DLL+11] A. David, K. G. Larsen, A. Legay, M. Mikučionis, and Z. Wang. "Time for statistical model checking of real-time systems". In: *Proceedings of 23rd International Conference on Computer Aided Verification (CAV)*. Springer. 2011, pp. 349–355.

[Dos79] B. T. Doshi. "Generalized semi-Markov decision processes". In: *Journal of Applied Probability* (1979), pp. 618–630.

[GGH+12] C. C. Guet, A. Gupta, T. A. Henzinger, M. Mateescu, and A. Sezgin. "Delayed continuous-time Markov chains for genetic regulatory circuits". In: *Proceedings of 24th International Conference on Computer Aided Verification (CAV)*. Springer. 2012, pp. 294–309.

[GL94] R. German and C. Lindemann. "Analysis of stochastic Petri nets by the method of supplementary variables". In: *Performance Evaluation* 20.1-3 (1994), pp. 317–335.

[Gly83] P. W. Glynn. "On the role of generalized semi-Markov processes in simulation output analysis". In: *Proceedings of the 15th conference on Winter simulation-Volume 1*. IEEE Press. 1983, pp. 39–44.

[Gly89] P. Glynn. "A GSMP formalism for discrete event systems". In: *Proceedings of the IEEE* 77 (1989), pp. 14–23.

[GY94] P. Glasserman and D. D. Yao. *Monotone structure in discrete-event systems*. John Wiley & Sons, Inc., 1994.

[Haa10] P. Haas. *Stochastic Petri Nets: Modelling, Stability, Simulation*. Springer Series in Operations Research and Financial Engineering. Springer, 2010.

[Her02] H. Hermanns. *Interactive Markov Chains: The Quest for Quantified Quality*. Vol. 2428. Lecture Notes in Computer Science. Springer, 2002.

[HG01] S. G. Henderson and P. W. Glynn. "Regenerative steady-state simulation of discrete-event systems". In: *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 11.4 (2001), pp. 313–345.

[HG02] P. Haas and P. Glynn. "On Simulation Output Analysis for Generalized Semi-Markov Processes". In: *Performance Evaluation Review* 30 (2002). Special issue on the 4th Workshop on Mathematical Performance Modeling and Analysis (MAMA 2002), pp. 34–37.

[HH13] H. Hafeti and H. Hermanns. “Improving Time Bounded Computations in Interactive Markov Chain”. In: *Proceedings of the 5th IPM International Conference on Fundamentals of Software Engineering (FSEN)*. Springer, 2013, pp. 250–266.

[Hil96] J. Hillston. *A compositional approach to performance modelling*. Cambridge University Press New York, NY, USA, 1996.

[HKK13] H. Hermanns, J. Krčál, and J. Křetínský. “Compositional Verification and Optimization of Interactive Markov Chains”. In: *Proceedings of 24th International Conference on Concurrency Theory (CONCUR)*. 2013, pp. 364–379.

[HLG+09] M. Heiner, S. Lehrack, D. Gilbert, and W. Marwan. “Extended stochastic Petri nets for model-based design of wetlab experiments”. In: *Transactions on Computational Systems Biology XI*. Springer, 2009, pp. 138–163.

[HMC97] S. Haddad, P. Moreaux, and G. Chiola. “Efficient handling of phase-type distributions in generalized stochastic Petri nets”. In: *Proceedings of the 18th International Conference on Application and Theory of Petri Nets (ICATPN)*. Springer, 1997, pp. 175–194.

[HMM05] S. Haddad, L. Mokdad, and P. Moreaux. “Performance evaluation of non Markovian stochastic discrete event systems-a new approach”. In: *Proceedings of the 7th IFAC Workshop on Discrete Event Systems 2004 (WODES’04)*. Elsevier Science Limited. 2005, p. 243.

[HMM06] S. Haddad, L. Mokdad, and P. Moreaux. “A new approach to the evaluation of non Markovian stochastic Petri nets”. In: *Proceedings of the 27th International Conference on Applications and Theory of Petri Nets and Other Models of Concurrency (ICATPN)*. Springer, 2006, pp. 221–240.

[Hor02] G. Horton. “A new paradigm for the numerical simulation of stochastic Petri nets with general firing times”. In: *Proceedings of the European Simulation Symposium*. 2002, pp. 129–136.

[HPR+12] A. Horváth, M. Paolieri, L. Ridi, and E. Vicario. “Transient analysis of non-Markovian models using stochastic state classes”. In: *Performance Evaluation* 69.7 (2012), pp. 315–335.

[HRV10a] A. Horváth, L. Ridi, and E. Vicario. “Approximating distributions and transient probabilities of Markov chains by Bernstein expolynomial functions”. In: *Proceedings of the International Conference on the Numerical Solution of Markov Chains*. 2010, pp. 52–55.

[HRV10b] A. Horváth, L. Ridi, and E. Vicario. “Transient analysis of generalised semi-Markov processes using transient stochastic state classes”. In: *Proceedings of the 7th International Conference on Quantitative Evaluation of Systems (QEST)*. IEEE. 2010, pp. 231–240.

[HS87] P. Haas and G. Shedler. “Regenerative Generalized Semi-Markov Processes”. In: *Stochastic Models* 3.3 (1987), pp. 409–438.

[HTT00] C. Hirel, B. Tuffin, and K. S. Trivedi. “Spnp: Stochastic Petri nets. version 6.0”. In: *Proceedings of the 13th International Conference on Computer Performance Evaluation, Modelling Techniques and Tools*. Springer, 2000, pp. 354–357.

[CHK+09] T. Chen, T. Han, J. Katoen, and A. Mereacre. “Quantitative model checking of continuous-time Markov chains against timed automata specifications”. In: *Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science (LICS)*. IEEE. 2009, pp. 309–318.

[JC01] R. Jones and G. Ciardo. “On phased delay stochastic Petri nets: Definition and an application”. In: *Proceedings of the 9th International Workshop on Petri Nets and Performance Models*. IEEE. 2001, pp. 165–174.

[JCL+09] S. K. Jha, E. M. Clarke, C. J. Langmead, A. Legay, A. Platzer, and P. Zuliani. “A bayesian approach to model checking biological systems”. In: *Proceedings of the 7th International Conference on Computational Methods in Systems Biology*. Springer. 2009, pp. 218–234.

[Jen53] A. Jensen. “Markoff chains as an aid in the study of Markoff processes”. In: *Scandinavian Actuarial Journal* 1953.1 (1953), pp. 87–91.

[JKH02] M. Jurdziński, O. Kupferman, and T. A. Henzinger. “Trading probability for fairness”. In: *Proceedings of the 16th International Workshop on Computer Science Logic (CSL)*. Springer. 2002, pp. 292–305.

[JS89] H. Jochens and G. Shedler. *Modelling and simulation of stochastic systems with SPSIM*. Tech. rep. RJ 6825. IBM Almaden Research Center, San Jose, CA, 1989.

[KH09] C. Krull and G. Horton. “Proxel-based simulation: Theory and applications”. In: *Proceedings of the 6th St. Petersburg Workshop on Simulation*. Citeseer. 2009, pp. 709–714.

[KNP11] M. Kwiatkowska, G. Norman, and D. Parker. “PRISM 4.0: Verification of probabilistic real-time systems”. In: *Proceedings of 23rd International Conference on Computer Aided Verification (CAV)*. Springer. 2011, pp. 585–591.

[KNS+00] M. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. “Verifying quantitative properties of continuous probabilistic timed automata”. In: *Proceedings of 11th International Conference on Concurrency Theory (CONCUR)* (2000), pp. 123–137.

[Kre13] J. Kretinsky. “Verification and Optimization of Time-un/bounded Properties of Stochastic Systems”. To appear. PhD thesis. TU Munich, 2013.

[Kul95] V. G. Kulkarni. *Modeling and analysis of stochastic systems*. Chapman & Hall, 1995.

[KZH+09] J.-P. Katoen, I. S. Zapreev, E. M. Hahn, H. Hermanns, and D. N. Jansen. “The Ins and Outs of The Probabilistic Model Checker MRMC”. In: *Proceedings of the 6th International Conference on Quantitative Evaluation of Systems (QEST)*. `www.mrmc-tool.org`. IEEE Computer Society, 2009, pp. 167–176.

[Laz05] S. Lazarova-Molnar. “The proxel-based method: Formalisation, analysis and applications”. PhD thesis. Otto-von-Guericke-Universität Magdeburg, Universitätsbibliothek, 2005.

[LDB10] A. Legay, B. Delahaye, and S. Bensalem. “Statistical model checking: An overview”. In: *Proceedings of the 1st International Conference on Runtime Verification*. Springer. 2010, pp. 122–135.

[LHK01] G. López, H. Hermanns, and J. Katoen. “Beyond memoryless distributions: Model checking semi-Markov chains”. In: *Proceedings of the Joint International Workshop on Process Algebra and Probabilistic Methods, Performance Modelling and Verification*. Springer, 2001, pp. 57–70.

[Lin93] C. Lindemann. “An improved numerical algorithm for calculating steady-state solutions of deterministic and stochastic Petri net models”. In: *Performance Evaluation* 18.1 (1993), pp. 79–95.

[LRT99] C. Lindemann, A. Reuys, and A. Thummler. “The DSPNexpress 2.000 performance and dependability modeling environment”. In: *Fault-Tolerant Computing, 1999. Digest of Papers. Twenty-Ninth Annual International Symposium on.* IEEE. 1999, pp. 228–231.

[LS96] C. Lindemann and G. Shedler. “Numerical analysis of deterministic and stochastic Petri nets with concurrent deterministic transitions”. In: *Performance Evaluation* 27 (1996), pp. 565–582.

[LT99] C. Lindemann and A. Thümmler. “Transient analysis of deterministic and stochastic Petri nets with concurrent deterministic transitions”. In: *Performance Evaluation* 36 (1999), pp. 35–54.

[Mar98] D. Martin. “The Determinacy of Blackwell Games”. In: *Journal of Symbolic Logic* 63.4 (1998), pp. 1565–1581.

[Mat62] K. Matthes. “Zur theorie der bedienungsprozesse”. In: *Transactions of the Third Prague Conference on Information Theory, Statistical Decision Functions, Random Processes.* 1962, pp. 513–528.

[MC87] M. Marsan and G. Chiola. “On Petri nets with deterministic and exponentially distributed firing times”. In: *Advances in Petri Nets 1987* (1987), pp. 132–145.

[Mol85] M. K. Molloy. “Discrete time stochastic Petri nets”. In: *IEEE Transactions on Software Engineering* 4 (1985), pp. 417–423.

[MP70] A. Maitra and T. Parthasarathy. “On stochastic games”. In: *Journal of Optimization Theory and Applications* 5.4 (1970), pp. 289–300.

[MS98] A. Maitra and W. Sudderth. “Finitely Additive Stochastic Games with Borel Measurable Payoffs”. In: *International Journal of Game Theory* 27 (1998), pp. 257–267.

[MT09] S. Meyn and R. Tweedie. *Markov Chains and Stochastic Stability.* Cambridge University Press, 2009.

[MU49] N. Metropolis and S. Ulam. “The Monte Carlo method”. In: *Journal of the American statistical association* 44.247 (1949), pp. 335–341.

[Neu10] M. R. Neuhäußer. “Model checking nondeterministic and randomly timed systems”. PhD thesis. RWTH Aachen University and University of Twente, 2010.

[Neu81] M. F. Neuts. *Matrix-geometric solutions in stochastic models: an algorithmic approach.* Courier Dover Publications, 1981.

[NSK09] M. Neuhäußer, M. Stoelinga, and J. Katoen. “Delayed nondeterminism in continuous-time Markov decision processes”. In: *Proceedings of the 12th International Conference on Foundations of Software Science and Computational Structures (FoSSaCS)* (2009), pp. 364–379.

[Ros06] J. Rosenthal. *A first look at rigorous probability theory*. World Scientific Publishing, 2006.

[RR04] G. Roberts and J. Rosenthal. “General state space Markov chains and MCMC algorithms”. In: *Probability Surveys* 1 (2004), pp. 20–71.

[RS11] M. N. Rabe and S. Schewe. “Finite optimal control for time-bounded reachability in CTMDPs and continuous-time Markov games”. In: *Acta Informatica* 48.5-6 (2011), pp. 291–315.

[RS13] M. N. Rabe and S. Schewe. “Optimal time-abstract schedulers for CTMDPs and continuous-time Markov games”. In: *Theoretical Computer Science* 467 (2013), pp. 53–67.

[SB98] M. Scarpa and A. Bobbio. “Kronecker representation of stochastic Petri nets with discrete PH distributions”. In: *Proceedings of the IEEE International Symposium on Computer Performance and Dependability (IPDS)*. IEEE. 1998, pp. 52–62.

[SDP03] M. Scarpa, S. Distefano, and A. Puliafito. “A parallel approach for the solution of non-Markovian Petri nets”. In: *Proceedings of the 10th European PVM/MPI Users’ Group Meeting on Recent Advances in Parallel Virtual Machine and Message Passing Interface*. Springer, 2003, pp. 196–203.

[SV07] L. Sassoli and E. Vicario. “Close form derivation of state-density functions over DBM domains in the analysis of non-Markovian models”. In: *Proceedings of the 4th International Conference on Quantitative Evaluation of Systems (QEST)*. IEEE. 2007, pp. 59–68.

[SVA04] K. Sen, M. Viswanathan, and G. Agha. “Statistical model checking of black-box probabilistic systems”. In: *Proceedings of 16th International Conference on Computer Aided Verification (CAV)*. Springer. 2004, pp. 202–215.

[Tay89] P. Taylor. “Insensitivity in processes with zero speeds”. In: *Advances in Applied Probability* (1989), pp. 612–628.

[Whi80a] W. Whitt. “Continuity of Generalized Semi-Markov Processes”. In: *Mathematics of Operations Research* 5.4 (1980), pp. 494–501.

[Whi80b] W. Whitt. “Continuity of generalized semi-Markov processes”. In: *Mathematics of Operations Research* 5.4 (1980), pp. 494–501.

[YKN+06] H. L. Younes, M. Kwiatkowska, G. Norman, and D. Parker. “Numerical vs. statistical probabilistic model checking”. In: *International Journal on Software Tools for Technology Transfer* 8.3 (2006), pp. 216–228.

[You05] H. Younes. “Ymer: A statistical model checker”. In: *Proceedings of 17th International Conference on Computer Aided Verification (CAV)*. Springer. 2005, pp. 429–433.

[YS02] H. Younes and R. Simmons. “Probabilistic verification of discrete event systems using acceptance sampling”. In: *Proceedings of 14th International Conference on Computer Aided Verification (CAV)*. Springer. 2002, pp. 23–39.

[YS04] H. L. Younes and R. G. Simmons. “Solving generalized semi-Markov decision processes using continuous phase-type distributions”. In: *Proceedings of the 19th National Conference on Artificial Intelligence*. Vol. 4. 2004, p. 742.

[ZFG+00] A. Zimmermann, J. Freiheit, R. German, and G. Hommel. “Petri net modelling and performability evaluation with TimeNET 3.0”. In: *Proceedings of the 13th International Conference on Computer Performance Evaluation, Modelling Techniques and Tools*. Springer, 2000, pp. 188–202.

[ZFH01] A. Zimmermann, J. Freiheit, and G. Hommel. “Discrete Time Stochastic Petri Nets for the Modeling and Evaluation of Real-Time Systems.” In: *Proceedings of the 15th IEEE International Symposium on Parallel & Distributed Processing*. 2001, p. 100.

[ZN10] L. Zhang and M. Neuhäußer. “Model checking interactive Markov chains”. In: *Proceedings of the 16th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*. Springer, 2010, pp. 53–68.