

When are Stochastic Transition Systems Tameable?

Nathalie Bertrand¹, Patricia Bouyer^{*2}, Thomas Brihaye³, and Pierre Carlier^{*2,3}

- 1 Inria Rennes, France
- 2 LSV, CNRS & ENS Cachan, France
- 3 Université de Mons, Belgium

Abstract

A decade ago, Abdulla, Ben Henda and Mayr introduced the elegant concept of decisiveness for denumerable Markov chains [1]. Roughly speaking, decisiveness allows one to lift most good properties from finite Markov chains to denumerable ones, and therefore to adapt existing verification algorithms to infinite-state models. Decisive Markov chains however do not encompass stochastic real-time systems, and general stochastic transition systems (STSs for short) are needed. In this article, we provide a framework to perform both the qualitative and the quantitative analysis of STSs. First, we define various notions of decisiveness (inherited from [1]), notions of fairness and of attractors for STSs, and make explicit the relationships between them. Then, we define a notion of abstraction, together with natural concepts of soundness and completeness, and we give general transfer properties, which will be central to several verification algorithms on STSs. We further design a generic construction which will be useful for the analysis of ω -regular properties, when a finite attractor exists, either in the system (if it is denumerable), or in a sound denumerable abstraction of the system. We next provide algorithms for qualitative model-checking, and generic approximation procedures for quantitative model-checking. Finally, we instantiate our framework with stochastic timed automata (STA), generalized semi-Markov processes (GSMPs) and stochastic time Petri nets (STPNs), three models combining dense-time and probabilities. This allows us to derive decidability and approximability results for the verification of these models. Some of these results were known from the literature, but our generic approach permits to view them in a unified framework, and to obtain them with less effort. We also derive interesting new approximability results for STA, GSMPs and STPNs.

Digital Object Identifier 10.4230/LIPIcs...

* Supported by ERC project EQUALIS



licensed under Creative Commons License CC-BY
Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Contents

1	Introduction	4
2	Preliminaries	6
2.1	Stochastic transition systems	6
2.2	A σ -algebra for measuring sets of infinite paths	8
2.3	STs as transformers of probability measures	9
2.4	Basic properties of paths in STs	9
2.5	Labelled STs and their properties	10
3	Nice properties of STs	12
3.1	Several decisiveness notions	13
3.2	Attractors	14
3.3	Fairness	15
3.4	Relationships between the various properties	16
4	Abstractions between STs	17
4.1	Abstraction	17
4.2	Transfer of properties through abstractions	19
4.2.1	The case of sound abstractions	19
4.2.2	Trickier transfers of properties	20
4.3	Conditions for completeness and soundness	22
5	Using attractors for analyzing STs	23
5.1	The case of DMCs with a finite attractor	23
5.2	General STs via an abstraction	25
6	Qualitative analysis	28
6.1	Basic properties under decisiveness hypotheses	28
6.2	Basic properties through abstractions	29
6.3	ω -regular properties in DMCs with a finite attractor	29
6.4	ω -regular properties of general STs via abstraction and finite attractor	29
7	Approximate quantitative analysis	30
7.1	Quantitative reachability under decisiveness hypotheses	30
7.2	Quantitative repeated reachability under decisiveness hypotheses	31
7.3	ω -regular properties in DMC with a finite attractor	32
7.4	ω -regular properties of general STs via abstraction and finite attractor	32
7.5	Time-bounded verification of stochastic real-time systems	32
8	Applications	33
8.1	Stochastic timed automata	33
8.1.1	Definition	34
8.1.2	The thick graph abstraction	36
8.1.3	Reactive STA	37
8.1.4	Single-clock STA	39
8.2	Generalized semi-Markov processes	41
8.2.1	Definition	41
8.2.2	The refined region graph abstraction	42

8.2.3 Analyzing GSMPs	44
8.3 Stochastic time Petri nets	45
9 A guided tour of STSs	46
10 Conclusion	47
A Technical results of Section 2	53
A.1 Additional technical results for Subsection 2.2	53
A.2 Missing proofs of Subsections 2.3 and 2.5	56
B Technical results of Section 3	58
B.1 Proof of Lemma 14	58
B.2 Proof of Proposition 23	60
C Technical results of Section 4	62
C.1 Additional technical results for Subsection 4.1	62
C.2 Missing proofs in Subsection 4.2	65
D Technical results of Section 5	71
E Technical results of Section 6	73
E.1 Additional technical results for Subsection 6.1	73
E.2 Qualitative analysis of simple properties	74
F Technical results of Section 7	75
G Technical results of Section 8	76



1 Introduction

Given its success for finite-state systems, the model checking approach to verification has been extended to various models based on automata, and including features such as time, probability and infinite data structures. These models allow one to represent software systems more faithfully, by representing timing constraints, randomization, and *e.g.* unbounded call stacks. At the same time, they often offer the possibility to consider *quantitative* verification questions, such as whether the best execution time meets a requirement, or whether the system is reliable with high probability. Quantitative verification is notably hard for infinite-state systems, and often requires the development of techniques dedicated to each class of models.

A decade ago, Abdulla, Ben Henda and Mayr introduced the concept of decisiveness for denumerable Markov chains [1]. Formally, a Markov chain is decisive w.r.t. a set of states F if runs almost-surely reach F or a state from which F can no longer be reached. The concept of decisiveness thus forbids some weird behaviours in denumerable Markov chains, and allows one to lift most good properties from finite Markov chains to denumerable ones, and therefore to adapt existing verification algorithms to infinite-state models. In particular, assuming decisiveness enables the quantitative model checking of (repeated) reachability properties, by providing an approximation scheme, which is guaranteed to terminate for any given precision for decisive Markov chains. Decisiveness also elegantly subsumes other concepts such as the existence of finite attractors, or coarseness [1].

Decisive Markov chains however are not general enough to represent stochastic real-time systems. Indeed, to faithfully model time in real-time systems, it is adequate to use dense time [4], that is, timestamps of events are taken from a dense domain (like the set of rational or of real numbers). This source of infinity for the state-space of the system is particularly difficult to handle: the state-space is non-denumerable (even continuous), the branching in the transition system is also non-denumerable, *etc.* For those reasons, stochastic real-time systems do not fit in the framework of decisive Markov chains of [1].

Also, standard analysis techniques for non-stochastic real-time systems (when they exist) cannot be easily adapted to stochastic real-time systems. Traditionally, these techniques rely on the design of appropriate finite abstractions, which preserve good properties of the original model. A prominent example of such an abstraction is that of the region automaton for timed automata [6]. However, these abstractions do not preserve all quantitative properties and, in the context of stochastic systems they may be too coarse already for the evaluation of the probability of properties as simple as reachability properties.

A general framework to analyse a large class of stochastic real-time systems, or more generally continuous stochastic systems, is thus lacking. In this article, we face this issue and provide a framework to perform the *stochastic transition systems* (STSs for short). To do so, we generalize the main concepts of [1] (such as decisiveness, attractors), and standard notions for Markov chains (like fairness). STSs are purely stochastic (i.e. without non-determinism) Markov processes [37, 38], that is, Markov chains with a continuous state-space. Note that, while this journal version builds on the conference paper [15], we choose here to phrase our results for time-homogeneous and Markovian models. As mentioned in [37], the Markovian assumption is not a severe restriction since many apparently non Markovian processes can be recast to Markovian models by changing the state space. In our opinion, this choice furthermore enables the design of a richer and more elegant theory (compared to [15]).

Our first contribution is to define various notions of decisiveness (inherited from [1]),

notions of fairness and of attractors in the general context of STSs. To complete the semantical picture, we make explicit the relationships between these notions, in the general case of STSs, and also when restricting to denumerable Markov chains. Decisiveness or the existence of attractors will be later exploited to analyze properties for STSs.

As mentioned earlier, the analysis of real-time systems often requires the development of abstractions. As a second contribution, we define a notion of abstraction, which makes sense for STSs. Concepts of soundness and completeness are naturally defined for those abstractions, and general transfer properties are given, which will be central to several verification algorithms on STSs. The special case of denumerable abstractions is discussed, since it allows one to transfer more properties from the abstract system to the concrete one.

We then focus on denumerable Markov chains with a finite attractor, or more generally STSs admitting a sound abstraction satisfying this property, and an ω -regular property represented by a deterministic Muller automaton. Our third contribution consists in building a graph for the attractor, which contains enough information to analyze the probability that the STS satisfies the property. This is completely new compared to the original results of [1] and our conference paper [15]. It is inspired by a procedure of [2] for probabilistic lossy channel systems, a special class of denumerable Markov chains with a finite attractor.

Our fourth contribution concerns the qualitative model checking problem for various properties. In particular, we extend the results of [1] and show that, under some decisiveness assumptions, the almost-sure model checking of (repeated) reachability properties reduces to a simpler problem, namely to a reachability problem with probability 0. We advocate that this reduction simplifies the problem: in countable models, the 0-reachability amounts to the non existence of a path, in the underlying non-probabilistic system; beyond countable models, checking that a reachability property is satisfied with probability 0 amounts to exhibiting a somehow regular set of executions with positive measure. Beyond (repeated) reachability properties, we apply our above-mentioned approach via the graph of an attractor for the qualitative analysis of ω -regular properties.

Our fifth contribution is the design of generic approximation procedures for the quantitative model-checking problem, inspired by the path enumeration algorithm of Purushothoman Iyer and Narashima [33]. Under some decisiveness assumptions, we prove that these approximation schemes are guaranteed to terminate. Assuming the STSs can be represented finitely and enjoy some smooth effectiveness properties, one derives approximation algorithms: one can approximate, up to a desired (arbitrary) precision, the probability of (repeated) reachability properties. Note that without these effectiveness properties, one cannot hope for algorithms, and this motivates our above formulation of “procedures”. Further, once again via the use of the graph of an attractor, we design an approximation algorithm for ω -regular properties; this algorithm reduces the quantitative analysis of an ω -regular property to the quantitative verification of a reachability property in the concrete model. Up to our knowledge, this approach is completely new, and provides an interesting framework for quantitative verification of stochastic systems.

Our last contribution consists in instantiating our framework with high-level stochastic models, stochastic timed automata (STA), generalized semi-Markov processes (GSMP) and stochastic time Petri Nets (STPN), which are three classes of models combining dense-time and probabilities. This allows us to derive decidability and approximability results for their verification. Some of these results were known from the literature, *e.g.* the ones from [17], but our generic approach permits to view them in a unified framework, and to obtain them with less effort. We also derive interesting new approximability results for STA and GSMPs. In particular, the approximability results implied by this paper for STA are far

more general than those obtained using an *ad-hoc* approach in [16]. In the case of STPNs, we also interestingly embed the framework of [32, 39] into our setting, which allows to show that we can relax some assumptions while preserving approximability results.

The paper concludes with an overview of our main results, organized as a guided tour of the STSs: it summarizes the relationships between all notions, and provides the reader recipes to analyze STSs.

In the interest of readability, most technical proofs are postponed to the appendix, with clear pointers. The emphasis is put on our new approach to the analysis of ω -regular properties, which remains in the core of the paper (Section 5).

Other related works. Apart from direct related works that we have already mentioned (like [1]), let us review related work from the literature. First, in [22], approximants are given, which rely on refinements of a partitioning of the state-space of an STS (via conditional expectations). However, there is no stopping criteria if we want to turn these approximants to a proper approximation scheme. And the approach is also very different.

Then, for specific classes of stochastic systems, approximation algorithms exist, which do however focus more on expressing mathematical properties of (integral) equations that one should solve, not really on convergence of the schemes. Sometimes, strong conditions are put on the system, so that convergence is obvious. This is for instance the case of [5, 13, 39].

The literature on stochastic hybrid systems is very rich, and since there is little hope to have some decidability results, approximation methods are very much developed. We give here some examples of works that have been done, but this is obviously not exhaustive. In [26], an over-approximation method based on a discrete abstraction is proposed for stochastic hybrid automata, but no converging approximation scheme is provided. In [43], an approximation (with some guarantee on the error made) is provided, which can be used for time-bounded verification of safety properties. Some other papers focus on discrete-time, allowing the use of constraint-solving methods, see e.g. [27].

Continuous stochastic systems as mentioned above are hard to analyze: first, it is difficult (and sometimes even impossible) to compute the exact value of the probability of some property (as simple as a reachability property) in such a system; and, for such complex systems, there is no generic proofs of convergence for approximation schemes. The key contribution of the current paper is to identify conditions to have correct decision procedures and approximation schemes, and to provide full proofs of convergence and correctness.

2 Preliminaries

In this section, we define the general model of stochastic transition systems, which are Markov chains with a continuous state-space. This model corresponds to labelled Markov processes of [37] with a single action (hence removing non-determinism). We then define several probability measures, on infinite paths, but also on the state-space, which give different point of views over the behaviour of the systems. We continue by defining regular measurable events, and end up with defining deterministic Muller automata, and technical material for handling properties specified by these automata. In the interest of space, for basics on probability and measure theory, we refer the reader to [37].

2.1 Stochastic transition systems

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

the σ -algebra and simply write this set as $\text{Dist}(S)$.

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


Note that it is sufficient to define $\kappa(s, \cdot)$ (for every $s \in S$) over a subset which generates the σ -algebra Σ .

Observe that if S is a denumerable set and $\Sigma = 2^S$, then \mathcal{T} is a denumerable Markov chain (DMC for short). If S is finite, the kernel κ then coincides with the standard probability matrix of the Markov chain. We now give two examples of STSs.

► **Example 2** (Denumerable Markov chain). *The first example is the DMC depicted in Figure 1. We consider here $\mathcal{T}_{\text{RW}} = (S_{\text{RW}}, \Sigma_{\text{RW}}, \kappa_{\text{RW}})$ where*

- $S_{\text{RW}} = \mathbb{N}$,
- $\Sigma_{\text{RW}} = 2^{S_{\text{RW}}}$,
- for each $i \geq 1$, $\kappa_{\text{RW}}(i, \{i + 1\}) = p$ and $\kappa_{\text{RW}}(i, \{i - 1\}) = 1 - p$ with $p \in]0, 1[$, and
- $\kappa_{\text{RW}}(0, \{1\}) = 1$.

This represents a random walk – hence the index RW – over the natural numbers.



■ **Figure 1** Random walk over \mathbb{N} .

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

► **Example 3** (Continuous-time Markov chain). *We now give a continuous variant of the previous random walk which models a simple queueing system. Precisely, we consider a queueing system with a single queue, a parameter λ for arrivals and ν for serving times. Each state $i \in \mathbb{N}$ is equipped with a non-negative real number that corresponds to the time that has elapsed since the beginning. Formally, we consider $\mathcal{T}_{\text{QS}} = (S_{\text{QS}}, \Sigma_{\text{QS}}, \kappa_{\text{QS}})$ with $S_{\text{QS}} = \mathbb{N} \times \mathbb{R}_+$. We equip S_{QS} with the σ -algebra generated by $2^{\mathbb{N}} \times \mathcal{B}(\mathbb{R}_+)$ where $\mathcal{B}(\mathbb{R}_+)$ is the Borel σ -algebra on \mathbb{R}_+ . Then intuitively, κ_{QS} describes how the length of the queue evolves with time. Formally, for each $t, d \in \mathbb{R}_+$,*

$$\kappa_{\text{QS}}((0, t), \{1\} \times [0, t + d]) = \kappa_{\text{QS}}((0, t), \{1\} \times [t, t + d]) = \int_0^d \lambda e^{-\lambda x} dx$$

and for every $i \geq 1$,

$$\begin{aligned} \kappa_{\text{QS}}((i, t), \{i + 1\} \times [0, t + d]) &= \kappa_{\text{QS}}((i, t), \{i + 1\} \times [t, t + d]) = \int_0^d \lambda e^{-(\lambda + \nu)x} dx \\ \kappa_{\text{QS}}((i, t), \{i - 1\} \times [0, t + d]) &= \kappa_{\text{QS}}((i, t), \{i - 1\} \times [t, t + d]) = \int_0^d \nu e^{-(\lambda + \nu)x} dx \end{aligned}$$

There will be more examples of STSs with a continuous set of states in Section 8.

In the sequel, we fix an STS $\mathcal{T} = (S, \Sigma, \kappa)$. We will give two semantical views on the behaviour of \mathcal{T} : the first one is operational, in the sense that \mathcal{T} generates executions, with a measure over these executions; the second one observes how the state-space evolves over time. The first point-of-view is the standard semantics of probabilistic systems and is widely used in the model-checking community, where the temporal aspects are important. From a state, a probabilistic transition is performed according to a fixed distribution, and the system resumes from one of the successor states. Among others, *e.g.* [9], uses this semantics for the standard model of continuous-time Markov chains. The second point-of-view bloomed more recently. It proposes to view probabilistic systems as transformers of probability distributions. Compared to the previous point-of-view, here one is interested not in states the system can be in, but rather in how the probability mass evolves along steps. This semantics was motivated by the ability to express different properties than the previous one [12]. It has been considered for both discrete-time Markov chains [35, 3] and continuous-time models, *e.g.* those induced by stochastic Petri nets [32]. These two point-of-views are two sides of the same coin, and we will use both in the following, though we are ultimately interested in properties related to the operational semantics.

2.2 A σ -algebra for measuring sets of infinite paths

The objective is now to interpret \mathcal{T} in an operational manner, and to define a probability measure over the set of infinite paths of \mathcal{T} . We follow the lines of [24]. A *finite (resp. infinite) path* of \mathcal{T} is a finite (resp. infinite) sequence of states. We write $\text{Paths}(\mathcal{T})$ for the set of infinite paths of \mathcal{T} . In order to get a probability measure over $\text{Paths}(\mathcal{T})$, we need to equip this set with a σ -algebra. We therefore define for each finite sequence of measurable sets $(A_i)_{0 \leq i \leq n} \in \Sigma^{n+1}$ the following set of infinite paths:

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

This set is called a *cylinder*. We then equip $\text{Paths}(\mathcal{T})$ with the σ -algebra generated by the cylinders. We denote it by $\mathcal{F}_{\mathcal{T}}$.

Let μ be a(n initial) probability measure over Σ , that is, $\mu \in \text{Dist}(S)$. We define a probability measure $\text{Prob}_{\mu}^{\mathcal{T}}$ as follows. First we inductively define a probability measure over the cylinders. For every finite sequence of measurable subsets $(A_i)_{0 \leq i \leq n} \in \Sigma^{n+1}$, we set:

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

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

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

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

► **Lemma 4.** $\text{Prob}_{\mu}^{\mathcal{T}}$ defines a probability measure over $(\text{Paths}(\mathcal{T}), \mathcal{F}_{\mathcal{T}})$.

The proof of Lemma 4 is classical and we omit it here. The interested reader may *e.g.* refer to the proof of [17, Proposition 3.2], which can easily be adapted to our context.

2.3 STSs as transformers of probability measures

One can also interpret the dynamics of \mathcal{T} as a transformer of probability measures over (S, Σ) . For μ a probability measure over Σ , its transformation through \mathcal{T} can be defined as the measure $\Omega_{\mathcal{T}}(\mu)$ defined for every $A \in \Sigma$ by:

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

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

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

► **Lemma 5.** *Let $\mu \in \text{Dist}(S)$ be an initial distribution and let $(A_i)_{0 \leq i \leq n}$ be a sequence of measurable sets. Write $\nu_0 = \mu_{A_0}$, the conditional probability of μ given A_0 , and for every $1 \leq j \leq n-1$, write $\nu_j = (\Omega_{\mathcal{T}}(\nu_{j-1}))_{A_j}$. Then, for every $0 \leq j \leq n$:*

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

The proof of this result is postponed to the technical appendix (page 56).

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

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

This alternative view of stochastic processes as transformers of probability measures is heavily used by Paolieri *et al.* in their time-bounded analysis of stochastic Petri nets [39]: the evolution of the probability distributions is tracked through stochastic state classes. We advocate that the two views (behavioural and probability transformers) need to be used at the same time. One of the first results we establish (see Lemma 19 later) is a witness of their interplay. Also, the measure transformer view will prove quite useful when it comes to abstraction in Section 4. Beyond that, it has been observed by Kwon *et al.* that the classes of properties one can express in both views are incomparable [35, 34], and depending on the application, one or the other can be more appropriate.

2.4 Basic properties of paths in STSs

To define properties on the STS \mathcal{T} , we use LTL-like notations, that will be interpreted as measurable subsets of $\text{Paths}(\mathcal{T})$. Let $\mathcal{L}_{S, \Sigma}$ be the set of formulas defined by the following grammar:

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

where $B \in \Sigma$, $\bowtie \in \{\geq, \leq, =\}$ is a comparison operator and $k \in \mathbb{N}$ is a natural number. Given $\rho = (s_n)_{n \geq 0}$ we write $\rho_{\geq i} = (s_n)_{n \geq i} \in \text{Paths}(\mathcal{T})$ for each $i \geq 0$. Then the satisfaction

relation of paths formulas is given as follows:

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

We write $\text{Ev}_{\mathcal{T}}(\varphi)$ for the set of infinite paths ρ in \mathcal{T} such that $\rho \models \varphi$. It is standard to show that the event $\text{Ev}_{\mathcal{T}}(\varphi)$ is a measurable subset of $(\text{Paths}(\mathcal{T}), \mathcal{F}_{\mathcal{T}})$ (see *e.g.* [44]). In particular, for every initial probability measure μ , $\text{Prob}_{\mu}^{\mathcal{T}}(\text{Ev}_{\mathcal{T}}(\varphi))$ is well-defined. In the sequel, for simplicity, we often write $\text{Prob}_{\mu}^{\mathcal{T}}(\varphi)$ instead of $\text{Prob}_{\mu}^{\mathcal{T}}(\text{Ev}_{\mathcal{T}}(\varphi))$.

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

► **Example 6.** We illustrate some properties on STSs \mathcal{T}_{RW} and \mathcal{T}_{QS} of Examples 2 and 3. Consider first \mathcal{T}_{RW} and $B = \{0\}$. The event $\text{Ev}_{\mathcal{T}_{\text{RW}}}(\mathbf{F} B)$ corresponds to the set of runs that eventually reach state 0, *i.e.* $\text{Ev}_{\mathcal{T}_{\text{RW}}}(\mathbf{F} B) = \{\rho = (s_n)_{n \geq 0} \in \text{Paths}(\mathcal{T}_{\text{RW}}) \mid \exists n \geq 0, s_n = 0\}$, while the event $\text{Ev}_{\mathcal{T}_{\text{RW}}}(\mathbf{G} \mathbf{F} B)$ corresponds to the set of runs that infinitely often visit 0, *i.e.* $\text{Ev}_{\mathcal{T}_{\text{RW}}}(\mathbf{G} \mathbf{F} B) = \{\rho = (s_n)_{n \geq 0} \in \text{Paths}(\mathcal{T}_{\text{RW}}) \mid \forall n \geq 0, \exists k \geq n, s_k = 0\}$. Now in \mathcal{T}_{QS} , we could consider the set of states $B' = \{0\} \times [0, T]$ with $T > 0$. Then the event $\text{Ev}_{\mathcal{T}_{\text{QS}}}(\mathbf{F} B')$ corresponds to the set of runs that eventually reach 0 within T time units, *i.e.* $\text{Ev}_{\mathcal{T}_{\text{QS}}}(\mathbf{F} B') = \{\rho = ((s_n, t_n))_{n \geq 0} \in \text{Paths}(\mathcal{T}_{\text{QS}}) \mid \exists n \geq 0, (s_n = 0 \wedge t_n \leq T)\}$.

2.5 Labelled STSs and their properties

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

► **Definition 7.** A *labelled STS* (LSTS for short) is a tuple $\mathcal{T} = (S, \Sigma, \kappa, \text{AP}, \mathcal{L})$, where (S, Σ, κ) is an STS, AP is a finite set of atomic propositions, and $\mathcal{L} : S \rightarrow 2^{\text{AP}}$ is a measurable labelling function.

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

A property over AP is a subset P of $(2^{\text{AP}})^{\omega}$. An infinite path $\rho = s_0 s_1 \dots$ of \mathcal{T} satisfies the property P whenever $\mathcal{L}(s_0) \mathcal{L}(s_1) \mathcal{L}(s_2) \dots \in P$, written $\rho \models P$. ω -regularity is a standard notion in computer science to characterise simple sets of infinite behaviours, and typical ω -regular properties are Büchi and Muller properties. In order to express such properties, we introduce a new notation for the set of atomic propositions that are true infinitely often along a sequence of labels: for $\varpi = u_0 u_1 u_2 \dots \in (2^{\text{AP}})^{\omega}$, we define $\text{Inf}(\varpi) = \{a \in \text{AP} \mid |\{j \in \mathbb{N} \mid a \in u_j\}| = \infty\}$. We extend this notation to paths in a natural way: if $\rho = s_0 s_1 s_2 \dots \in S^{\omega}$, writing $\varpi = \mathcal{L}(s_0) \mathcal{L}(s_1) \mathcal{L}(s_2) \dots$, we define (with a slight abuse of notation) $\text{Inf}(\rho) = \text{Inf}(\varpi)$. A *Büchi property* P over AP can be specified by a subset of atomic propositions $F \subseteq \text{AP}$ as $P = \{\varpi \in (2^{\text{AP}})^{\omega} \mid \text{Inf}(\varpi) \cap F \neq \emptyset\}$. A *Muller property* over AP is a property P such that there exists $\mathcal{F} \subseteq 2^{\text{AP}}$ with $P = \{\varpi \in (2^{\text{AP}})^{\omega} \mid \text{Inf}(\varpi) \in \mathcal{F}\}$.

► **Remark 8.** It should be noted that the set of infinite paths satisfying Büchi or Muller properties can be expressed using events as in Section 2.4. Indeed, for $F \subseteq \text{AP}$ we write $2_F^{\text{AP}} = \{u \in 2^{\text{AP}} \mid u \cap F \neq \emptyset\}$ and given $a \in \text{AP}$, $2_a^{\text{AP}} = \{u \in 2^{\text{AP}} \mid a \in u\}$. Then,

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

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

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

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

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

► **Definition 9.** A *deterministic Muller automaton* (DMA) over AP is a tuple $\mathcal{M} = (Q, q_0, E, \mathcal{F})$ where:

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

and such that

- \mathcal{M} is deterministic: for all pair of edges (q, u, q_1) and (q, u, q_2) in E , $q_1 = q_2$;
- \mathcal{M} is complete: for every $q \in Q$, for every $u \in 2^{\text{AP}}$, there exists $(q, u, q') \in E$.

A deterministic Muller automaton \mathcal{M} naturally gives rise to a property $P_{\mathcal{M}}$ defined by the language (over 2^{AP}) accepted by \mathcal{M} . Its semantics over infinite paths of \mathcal{T} is derived from that of property $P_{\mathcal{M}}$: if $\rho \in \text{Paths}(\mathcal{T})$, we write $\rho \models \mathcal{M}$ whenever $\rho \models P_{\mathcal{M}}$. Expanding Remark 8, one derives the standard fact that the set $\mathcal{T}[\mathcal{M}] \stackrel{\text{def}}{=} \{\rho \in \text{Paths}(\mathcal{T}) \mid \rho \models \mathcal{M}\}$ is measurable, and we write $\text{Prob}_{\mu}^{\mathcal{T}}(\mathcal{M})$ for $\text{Prob}_{\mu}^{\mathcal{T}}(\mathcal{T}[\mathcal{M}])$.

► **Remark 10.** It is well known (see [45] and [30, Chapter 3]) that for any LTL formula φ (the syntax given in the previous subsection, where we replace sets B by inverse images by \mathcal{L} of atomic propositions from AP), there is a deterministic Muller automaton \mathcal{M}_{φ} that characterises φ , that is: for every run ρ , $\rho \models \varphi$ iff $\rho \models \mathcal{M}_{\varphi}$.

Product STS

To measure the probability of properties specified by a DMA $\mathcal{M} = (Q, q_0, E, \mathcal{F})$, it is standard to build a new STS consisting of the product of \mathcal{T} with \mathcal{M} [10]. To this aim, we consider the discrete σ -algebra 2^Q on the finite set of locations Q of \mathcal{M} . The product $S \times Q$ can then be equipped with the product σ -algebra $\Sigma \times 2^Q$ defined as the smallest σ -algebra containing all rectangles, that is, all sets of the form $A_1 \times A_2$ with $A_1 \in \Sigma$ and $A_2 \in 2^Q$. Let us notice that given X an element of the product σ -algebra $\Sigma \times 2^Q$, for all $q \in Q$, the set $\pi_q(X) \stackrel{\text{def}}{=} \{x \in S \mid \exists Q' \subseteq Q \text{ s.t. } q \in Q' \text{ and } (x, Q') \in X\}$ belongs to Σ . We therefore write (abusively) $X = \bigcup_{q \in Q} \pi_q(X) \times \{q\}$. Then, the product σ -algebra $\Sigma \times 2^Q$ coincides with Σ' , the set of all subsets of $S \times Q$ of the form $\bigcup_{q \in Q} C_q \times \{q\}$, where $C_q \in \Sigma$ for every $q \in Q$ (see the proof in the appendix, page 57). Note that in the sequel, we will sometimes write (C_q, q) instead of $C_q \times \{q\}$.


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

► **Definition 11.** Given $\mathcal{T} = (S, \Sigma, \kappa, \text{AP}, \mathcal{L})$ an LSTS and $\mathcal{M} = (Q, q_0, E, \mathcal{F})$ a DMA over AP, we define the product of \mathcal{T} with \mathcal{M} as the LSTS $\mathcal{T} \times \mathcal{M} = (S', \Sigma', \kappa', \text{AP}', \mathcal{L}')$ such that:

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

- $\kappa'((s, q), (A, q')) = \begin{cases} \kappa(s, A) & \text{if } (q, \mathcal{L}(s), q') \in E, \text{ and} \\ 0 & \text{otherwise;}^1 \end{cases}$
- $\text{AP}' = Q$;
- $\mathcal{L}'(s, q) = q$.

► **Example 12.** We consider again \mathcal{T}_{RW} the random walk over \mathbb{N} of Example 2. We assume that it is equipped with the simple set of atomic propositions $\text{AP} = \{a\}$ and we assume that each state of the STS is labelled with a . Let \mathcal{M} be the DMA depicted on the left-hand side of Figure 2. The winning condition is given by $\mathcal{F} = \{\{q_1, q_2\}\}$. The product $\mathcal{T}_{\text{RW}} \times \mathcal{M}$ is then depicted on the right-hand side of Figure 2. It should be noted that we there assume that the system starts at $(0, q_0)$ however, there should be similar chains starting in (i, q_0) for each $i \geq 1$. Note also that we did not specify the labels on the states: according to the definition each state is labelled with its current position in \mathcal{M} .



■ **Figure 2** A Muller automaton \mathcal{M} and the product $\mathcal{T}_{\text{RW}} \times \mathcal{M}$.

We define on $\mathcal{T} \times \mathcal{M}$ a Muller condition which is inherited from the one of \mathcal{M} via the new labelling function \mathcal{L}' : a run ρ satisfies the Muller condition \mathcal{F}' whenever $\mathcal{L}'(\rho)$ satisfies the Muller condition \mathcal{F} . We thus later simply use \mathcal{F} instead of \mathcal{F}' .

We now explain how initial distributions for \mathcal{T} are lifted to the product $\mathcal{T} \times \mathcal{M}$. The idea is simple: the \mathcal{T} -component is inherited from \mathcal{T} , and the \mathcal{M} -component is a Dirac² distribution over q_0 , the initial state of \mathcal{M} . In other words, when an initial distribution $\mu \in \text{Dist}(S)$ is fixed for \mathcal{T} , the initial distribution of $\mathcal{T} \times \mathcal{M}$ will be $\mu \times \delta_{q_0}$. We show that this allows to properly express the probability of a property described by a DMA, with the following correspondence (whose proof is given in the appendix, page 57).

► **Proposition 13.** Let $\mu \in \text{Dist}(S)$ be an initial distribution for \mathcal{T} , and $\mathcal{M} = (Q, q_0, E, \mathcal{F})$ be a DMA. Then:

$$\text{Prob}_{\mu}^{\mathcal{T}}(\mathcal{T}[\mathcal{M}]) = \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathcal{M}}(\{\rho \in \text{Paths}(\mathcal{T} \times \mathcal{M}) \mid \rho \models \mathcal{F}\}) .$$

3 Nice properties of STSs

In [1], Abdulla *et al.* introduced the elegant concept of decisive Markov chain. Intuitively, decisiveness allows one to lift the good properties of finite Markov chains to infinite (but denumerable) Markov chains. We explain here how to extend and refine this concept and some related concepts to general STSs, and we establish relationships between these properties.

¹ Note that the above definition of κ' extends naturally to all elements of the σ -algebra Σ' : for each pair (q, u) with $q \in Q$ and $u \in 2^{\text{AP}}$, there is a unique $q' \in Q$ such that $(q, u, q') \in E$. Fix $(s, q) \in S \times Q$, write q' for the unique location such that $(q, \mathcal{L}(s), q') \in E$. Then for each $A = \bigcup_{q \in Q} C_q \times \{q\}$, $\kappa'((s, q), A) = \kappa'((s, q), (C_{q'}, q')) = \kappa(s, C_{q'})$.

² We recall that given a state s , the *Dirac distribution over s* is defined by $\delta_s(A) = 1$ if $s \in A$ and $\delta_s(A) = 0$ otherwise, for every $A \in \Sigma$.

3.1 Several decisiveness notions

Decisiveness has been defined in [1] as a desirable property of denumerable Markov chains, since it implies that they behave essentially like finite Markov chains.

For $B \in \Sigma$ a measurable set, we define its *avoid-set* $\widetilde{B} = \{s \in S \mid \text{Prob}_{\delta_s}^T(\mathbf{F}B) = 0\}$. It corresponds to the set of states from which the system will always avoid the set B with probability 1 (or equivalently, reach B with probability 0). The set \widetilde{B} enjoys the following properties, that obviously hold also in the context of denumerable Markov chains, but require proofs in our general context of STSs (those proofs are postponed to the technical appendix, page 58).

► **Lemma 14.** *Given $B \in \Sigma$, it holds that:*

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

Let us comment on the third and fourth properties stated in this lemma. The third item indicates that if we start from outside \widetilde{B} , then we will always have a positive probability to hit B . The fourth property says that \widetilde{B} is some kind of sink: once we hit \widetilde{B} , we cannot escape it. The other properties are rather straightforward to understand (even though proving the first property requires some technical developments).

We are now ready to define different decisiveness concepts. Two stem from [1] (though no initial distribution was fixed) while the third one was identified in [15] as a useful alternative.

► **Definition 15.** Let μ be an initial probability distribution ($\mu \in \text{Dist}(S)$). Then:

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

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

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

Note that if \mathcal{T} is finite, then it is decisive, strongly decisive and persistently decisive. Let us now illustrate the subtleties of the various decisiveness notions on examples.

► **Example 16.** *Let us consider again the STS \mathcal{T}_{RW} of Example 2, representing a discrete-time random walk, and assume $p > 1/2$. As a classical result, the random walk is then diverging, see e.g. [20]. Since the chain is strongly connected, for each $B \subseteq \mathbb{N}$, $\widetilde{B} = \emptyset$. Let us first assume that the initial distribution is $\mu = \delta_0$, the Dirac distribution over state 0. Then it can be shown that for each set of states B , $\text{Prob}_{\mu}^{\mathcal{T}_{\text{RW}}}(\mathbf{F} B) = 1$ and thus, \mathcal{T}_{RW} is $\text{Dec}(\mu, B)$.*

Assuming now that the initial distribution is $\mu' = \delta_1$, if $B = \{0\}$, then $\text{Prob}_{\mu'}^{\mathcal{T}_{\text{RW}}}(\mathbf{F} \{0\}) < 1$. But since $\widetilde{B} = \emptyset$, we derive that \mathcal{T}_{RW} is not $\text{Dec}(\mu', B)$.

Consider now for each $i \geq 0$, $B_i = \{i\}$. Since $p > 1/2$, classical results on random walks imply that for each i , $\text{Prob}_{\mu}^{\mathcal{T}_{\text{RW}}}(\mathbf{G} \mathbf{F} B_i) = 0$. And since $\widetilde{B}_i = \emptyset$, we obtain that \mathcal{T}_{RW} is not $\text{StrDec}(\mu, B_i)$.

Consider now the STS \mathcal{T}_{QS} of Example 3. Assume that $\lambda > \nu$ and that $\mu = \delta_{(0,0)}$ and fix some $T > 0$. We consider $B_1 = \{1\} \times [0, T]$. Then one can compute $\widetilde{B} = \mathbb{N} \times]T, \infty[$. Note that here, as time almost-surely always progresses, $\text{Prob}_{\mu}^{\mathcal{T}_{\text{QS}}}(\mathbf{F} \widetilde{B}) = 1$. It thus follows that \mathcal{T}_{QS} is $\text{Dec}(\mu, B)$ and $\text{StrDec}(\mu, B)$.

3.2 Attractors

The notion of finite attractor has been used in several contexts like probabilistic lossy channel systems (see e.g. [2, 42]) and abstracted in [1] in the context of denumerable Markov chains. A finite attractor is a finite set of states which is reached almost-surely from every state of the system. We lift this definition to our context, obviously relaxing the finiteness assumption, since it is very unlikely that systems with a continuous state-space will have finite attractors. Since the whole set of states is a trivial attractor, this general definition will prove useful once we are able to define attractors with some finiteness property, which will be done through *abstractions* in Section 4.

► **Definition 17.** Let $\mu \in \text{Dist}(S)$ be an initial distribution. $B \in \Sigma$ is a μ -attractor for \mathcal{T} if $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} B) = 1$. Further, B is an *attractor* for \mathcal{T} if it is a μ -attractor for every $\mu \in \text{Dist}(S)$.

► **Example 18.** *Consider again the random walk \mathcal{T}_{RW} of Example 2 and assume again that $p > 1/2$. For $B = \{5\}$, it can be shown that B is a μ -attractor for $\mu = \delta_0$. However, for any distribution $\mu' \in \text{Dist}(\mathbb{N}_{\geq 6})$ over natural numbers larger than or equal to 6, $\text{Prob}_{\mu'}^{\mathcal{T}_{\text{RW}}}(\mathbf{F} B) < 1$ and thus B is not a μ' -attractor.*

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

The existence of an attractor is a rather strong property for an STS. It will actually imply that attractors are almost-surely visited infinitely often. While this result was already known for DMCs (see [1, Lemmas 3.2 and 3.4]), the general case of STSs requires a specific treatment. We choose to present the proof of this result in the core of the paper, since it illustrates how to use at the same time the two views on STSs.

► **Lemma 19.** *If B is an attractor for \mathcal{T} then for every initial distribution $\mu \in \text{Dist}(S)$,*

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

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

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

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

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

From Lemma 5, if we write $\nu_0 = \mu$ and $\nu_j = \Omega_\mathcal{T}(\nu_{j-1})$ for each $1 \leq j \leq n-1$, we get that for each $m \geq 1$,

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

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

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

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

3.3 Fairness

Fairness is a standard notion in probabilistic systems [40, 41, 11], saying that something which is allowed infinitely often should happen infinitely often almost-surely. This can for instance be instantiated in denumerable Markov chains as follows: if a state s is visited infinitely often, and the probability to move from s to s' is positive, then, almost-surely, infinitely often the state s' is visited. Not all Markov chains are fair, but finitely-branching Markov chains are. Note that other notions of fairness are discussed in [11].

The above notion of fairness cannot be lifted directly to continuous state-space STSs (since for two states s and s' , the probability to move from s to s' is likely to be 0). A more careful definition of this notion must be provided for general STSs. For $B \in \Sigma$, we define

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

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

► **Definition 20.** Let $\mu \in \text{Dist}(S)$ be some initial distribution, and $B \in \Sigma$. The STS \mathcal{T} is *fair w.r.t. B from μ* , written \mathcal{T} is *fair*(μ, B), if for every $B' \in \text{PreProb}^\mathcal{T}(B)$, $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B') > 0$ implies


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

As for decisiveness, we extend this definition to sets $\mathcal{B} \subseteq \Sigma$, and use similar notations when we relax the fixed initial measure μ . Finally, we say that \mathcal{T} is *strongly fair* whenever it is fair w.r.t. B from μ for every $B \in \Sigma$ and every $\mu \in \text{Dist}(S)$.

► **Example 21.** Consider again the random walk of Example 2. In \mathcal{T}_{RW} , there is a positive lower bound on each single move in the chain: take $\varepsilon = \min(p, 1 - p) > 0$, and realize then that for every $B \subseteq S$, for every $B' \in \text{PreProb}^{\mathcal{T}_{\text{RW}}}(B)$ and for every $s \in B'$, $\kappa_1(s, B) \geq \varepsilon$. Then using a proof similar to [11, Theo. 2], we deduce that \mathcal{T}_{RW} is fair w.r.t. B from any initial distribution. Hence, \mathcal{T}_{RW} is strongly fair.

► **Example 22 (Counter-example).** Consider now the DMC $\mathcal{T}_{\text{unfair}}$ depicted in Figure 3. From each state a_n ($n \geq 1$), the probability to move to b is $\frac{1}{3^n}$ whereas the probability to move to a_{n+1} is $1 - \frac{1}{3^n}$. From b , the probability to move to a_1 is 1.

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



■ **Figure 3** A denumerable Markov chain \mathcal{T}_3 that is not strongly fair.

3.4 Relationships between the various properties

In this section, we compare all the notions, and give the precise links between all of them. We first analyze the general case, and reinforce the results in the case of DMCs.

► **Proposition 23.** Let $\mathcal{B} \subseteq \Sigma$ and $\mu \in \text{Dist}(S)$. The following implications hold:

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

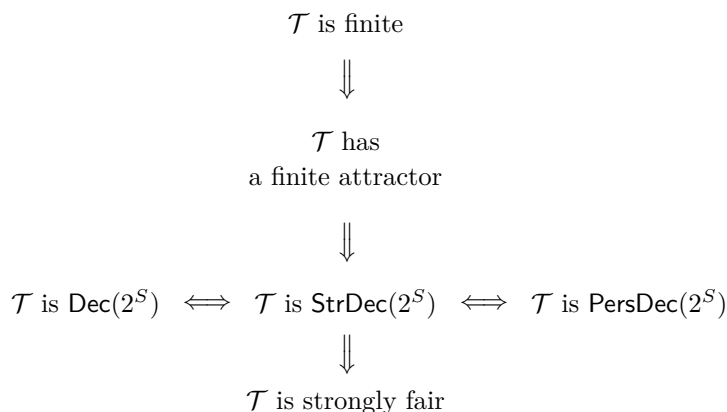
The three missing implications in the above proposition do actually not hold, as proved by the following example. We also illustrate the fact that $\text{Dec}(\mu, \mathcal{B})$ and $\text{fair}(\mu, \mathcal{B})$ are incomparable.

► **Example 24 (Counter-example).** Consider again the random walk \mathcal{T}_{RW} of Example 2. We have shown in Example 21 that \mathcal{T}_{RW} is strongly fair, whatever the choice of p . Now let us assume that $p > 1/2$ and let us consider the initial distribution $\mu = \delta_0$, the Dirac distribution over 0. Then from Example 16, \mathcal{T}_{RW} is decisive from μ w.r.t. any set of states. Again in this example, we have observed that it is not strongly decisive w.r.t. any set of the form $B = \{i\}$ with $i \geq 0$. This shows that \mathcal{T}_{RW} is neither $\text{Dec}(\mu, \mathcal{B}) \implies \text{StrDec}(\mu, \mathcal{B})$, nor $\text{fair}(\mu, \mathcal{B}) \implies \text{StrDec}(\mu, \mathcal{B})$, nor $\text{fair}(\mathcal{B}) \implies \text{StrDec}(\mathcal{B})$. And since \mathcal{T}_{RW} is not decisive from δ_1 w.r.t. $\{0\}$, this also proves that $\text{fair}(\mu, \mathcal{B})$ does not imply $\text{Dec}(\mu, \mathcal{B})$.

In order to illustrate that $\text{Dec}(\mu, \mathcal{B})$ does not imply $\text{fair}(\mu, \mathcal{B})$ in general, we consider the DMC chain $\mathcal{T}_{\text{unfair}}$ of Example 22. We consider $B = \{b\}$ and $\mu = \delta_b$. It is easily observed that $\mathcal{T}_{\text{unfair}}$ is $\text{Dec}(\mu, B)$ as we start in b with probability 1, but we have shown that $\mathcal{T}_{\text{unfair}}$ is not fair(μ, B).

If \mathcal{T} is a DMC, *i.e.* if S is at most denumerable and $\Sigma = 2^S$, we can complete the picture using furthermore the result of [1, Lemma 3.4], which says that any DMC with a finite attractor is decisive w.r.t. any set of states.

► **Proposition 25.** Assume \mathcal{T} is a DMC. The following implications hold:



4 Abstractions between STSs

While decisiveness is well-defined for general STSs, proving that a given STS \mathcal{T} is decisive might be technical in general. A standard approach in model-checking to avoid such difficulties is to abstract the system into a simpler one, that can be analyzed and provides guarantees on the concrete system. We thus propose a notion of abstraction, which will help proving properties of general STSs. Also, through abstractions, we will be able to characterize meaningful attractors.

Abstractions only preserve the positivity of probabilities, which is rather weak. However, by strengthening the notion, we will see that they will be very useful, not only for the qualitative analysis of STSs, but also for their quantitative analysis.

4.1 Abstraction

Let $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1)$ and $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2)$ be two STSs. Let $\alpha : (S_1, \Sigma_1) \rightarrow (S_2, \Sigma_2)$ be a measurable function. A set $B \in \Sigma_1$ is said α -closed whenever $B = \alpha^{-1}(\alpha(B))$: for every $s, s' \in S_1$, if $s \in B$ and $\alpha(s) = \alpha(s')$, then $s' \in B$. Following [28], we define the pushforward of α as $\alpha_{\#} : \text{Dist}(S_1) \rightarrow \text{Dist}(S_2)$ by $\alpha_{\#}(\mu)(M_2) = \mu(\alpha^{-1}(M_2))$ for every $\mu \in \text{Dist}(S_1)$ and for every $M_2 \in \Sigma_2$. The role of the pushforward $\alpha_{\#}$ is to transfer the measures from (S_1, Σ_1) to (S_2, Σ_2) . In the following we will say that two probability distributions μ and ν over some probability space (S, Σ) are *qualitatively equivalent* if for each $A \in \Sigma$, $\mu(A) = 0 \iff \nu(A) = 0$.

► **Definition 26.** \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 if

$$\forall \mu \in \text{Dist}(S_1), \alpha_{\#}(\Omega_{\mathcal{T}_1}(\mu)) \text{ is equivalent to } \Omega_{\mathcal{T}_2}(\alpha_{\#}(\mu)) .$$

Later, one may speak of abstraction instead of α -abstraction if α is clear in the context.

From the definitions of $\Omega_{\mathcal{T}}$, $\alpha_{\#}$ and equivalent measures, the notion of α -abstraction equivalently requires that for every $\mu \in \text{Dist}(S_1)$ and every $A \in \Sigma_2$,


$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(S_1, \alpha^{-1}(A))) > 0 \iff \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(S_2, A)) > 0 .$$

Intuitively, the two STSs have the same “qualitative” single steps.

Let us now provide two examples of α -abstraction.

► **Example 27.** Consider the two STSs $\mathcal{T}_{RW} = (S_{RW}, \Sigma_{RW}, \kappa_{RW})$ –the discrete random walk on \mathbb{N} from Example 2– and $\mathcal{T}_{QS} = (S_{QS}, \Sigma_{QS}, \kappa_{QS})$ –its continuous variant from Example 3. Recall that $S_{RW} = \mathbb{N}$ and $S_{QS} = \mathbb{N} \times \mathbb{R}_+$. Letting $\alpha : S_{QS} \rightarrow S_{RW}$ be the function defined by $\alpha((n, t)) = n$, one can easily be convinced that \mathcal{T}_{RW} is an α -abstraction of \mathcal{T}_{QS} .

► **Example 28.** Consider the random walk on \mathbb{N} from Example 2 $\mathcal{T}_{RW} = (S_{RW}, \Sigma_{RW}, \kappa_{RW})$, and the finite Markov chain $\mathcal{T}_f = (S_f, \Sigma_f, \kappa_f)$ depicted on Fig. 4. We define a function $\alpha : S_{RW} \rightarrow S_f$ by $\alpha(n) = s_n$ if $n \in \{0, 1\}$, and $\alpha(n) = 2$ otherwise. Clearly enough \mathcal{T}_f is an α -abstraction of \mathcal{T}_{RW} .



■ **Figure 4** A finite Markov chain.

The notion of α -abstraction naturally extends to labelled STSs. A labelled STS $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2, \text{AP}_2, \mathcal{L}_2)$ is an α -abstraction of $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1, \text{AP}_1, \mathcal{L}_1)$ whenever:

- $(S_2, \Sigma_2, \kappa_2)$ is an α -abstraction of $(S_1, \Sigma_1, \kappa_1)$;
- $\text{AP}_1 = \text{AP}_2$;
- for every $s_1, s'_1 \in S_1$, $\alpha(s_1) = \alpha(s'_1) \Rightarrow \mathcal{L}_1(s_1) = \mathcal{L}_1(s'_1)$;
- for every $s \in S_1$, $\mathcal{L}_1(s) = a \Rightarrow \mathcal{L}_2(\alpha(s)) = a$.

The two last conditions imply that for each $a \in 2^{\text{AP}}$, $\mathcal{L}_1^{-1}(\{a\})$ is α -closed. Moreover, for each $a \in 2^{\text{AP}}$, $\alpha^{-1}(\mathcal{L}_2^{-1}(\{a\})) = \mathcal{L}_1^{-1}(\{a\})$.

By definition, abstractions preserve the positivity of the probability of single-step moves. More generally, one easily shows (see the appendix page 62 for details) that the positivity of reachability properties or more generally of properties with bounded witnesses is also preserved through α -abstractions: assuming \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 , for every $\mu \in \text{Dist}(S_1)$, for every $A, B \in \Sigma_2$:

$$\text{Prob}_\mu^{\mathcal{T}_1}(\text{Ev}_{\mathcal{T}_1}(\alpha^{-1}(A) \mathbf{U} \alpha^{-1}(B))) > 0 \iff \text{Prob}_{\alpha\#(\mu)}^{\mathcal{T}_2}(\text{Ev}_{\mathcal{T}_2}(A \mathbf{U} B)) > 0 \quad (1)$$

However this does not apply to liveness properties, such as $\text{Ev}_{\mathcal{T}}(\mathbf{G} \mathbf{F} A)$ with $A \in \Sigma_2$, or to other qualitative questions like almost-sure reachability (reach B with probability 1). To ensure that these more involved properties are preserved via abstraction, we will strengthen the assumptions on the abstraction and on the STSs.

Soundness and completeness of abstractions

We now define soundness and completeness of abstractions, that allow one to lift properties of an abstraction \mathcal{T}_2 to the concrete STS \mathcal{T}_1 . For the rest of this section, we fix an STS \mathcal{T}_1 and let \mathcal{T}_2 be an α -abstraction of \mathcal{T}_1 .

► **Definition 29.** Let $\mu \in \text{Dist}(S_1)$. The α -abstraction \mathcal{T}_2 is μ -sound whenever for every $B \in \Sigma_2$:

$$\text{Prob}_{\alpha\#(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) = 1 \implies \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1 .$$

\mathcal{T}_2 is a *sound* α -abstraction of \mathcal{T}_1 if it is μ -sound for every $\mu \in \text{Dist}(S_1)$.

Assuming soundness, it will thus be sufficient to prove that a reachability property holds almost surely in the abstraction to derive that the corresponding reachability property also holds almost surely in the concrete STS.

► **Definition 30.** Let $\mu \in \text{Dist}(S_1)$. The α -abstraction \mathcal{T}_2 is μ -complete whenever for every $B \in \Sigma_2$,

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1 \implies \text{Prob}_{\alpha\#(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) = 1$$

\mathcal{T}_2 is a complete α -abstraction of \mathcal{T}_1 if it is μ -complete for every $\mu \in \text{Dist}(S_1)$.

Assuming completeness, if a reachability property holds with probability smaller than 1 in the abstraction, then the corresponding property will also have probability smaller than 1 in the concrete model.

Altogether, sound and complete abstractions will guarantee that, up to α , the same reachability properties are satisfied almost-surely in \mathcal{T}_1 and in \mathcal{T}_2 .

► **Example 31.** We continue here Example 27 with \mathcal{T}_{RW} (with parameter $0 < p < 1$) and \mathcal{T}_{QS} (with parameters $\lambda > 0$ and $\nu > 0$). One can show that \mathcal{T}_{RW} is a sound and complete abstraction of \mathcal{T}_{QS} whenever $p > 1/2 \iff \lambda > \nu$.

► **Example 32.** We continue here Example 28 with \mathcal{T}_{RW} (with parameter $0 < p < 1$) and \mathcal{T}_f (with parameter $0 < q < 1$). One can show that \mathcal{T}_f is always a complete abstraction of \mathcal{T}_{RW} ; Moreover, \mathcal{T}_f is also sound if and only if $p \leq 1/2$.

4.2 Transfer of properties through abstractions

In this section, we explain how and under which conditions one can transfer interesting decisiveness, attractor and fairness properties of STSs through abstractions.

4.2.1 The case of sound abstractions

► **Proposition 33.** If \mathcal{T}_2 is a μ -sound α -abstraction of \mathcal{T}_1 , then for every $B \in \Sigma_2$:

$$\mathcal{T}_2 \text{ is Dec}(\alpha\#(\mu), B) \implies \mathcal{T}_1 \text{ is Dec}(\mu, \alpha^{-1}(B)) .$$

Using equivalences between the various properties stated in Proposition 23, we can extend the above result to other decisiveness properties: assuming \mathcal{T}_2 is a sound α -abstraction of \mathcal{T}_1 , for every $B \in \Sigma_2$, if \mathcal{T}_2 is $\text{Dec}(B)$ (or equiv. $\text{StrDec}(B)$, $\text{PersDec}(B)$) then \mathcal{T}_1 is $\text{Dec}(\alpha^{-1}(B))$ (or equiv. $\text{StrDec}(\alpha^{-1}(B))$, $\text{PersDec}(\alpha^{-1}(B))$).

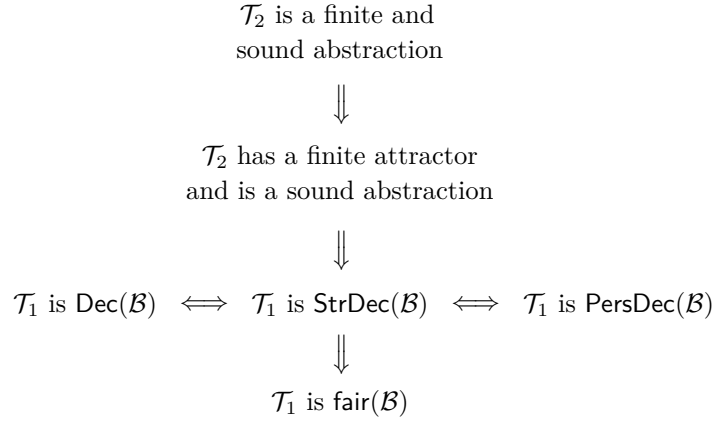
The definitions of attractor and of sound α -abstraction yield a similar result:

► **Proposition 34.** If \mathcal{T}_2 is a sound α -abstraction of \mathcal{T}_1 and if $A \in \Sigma_2$ is an attractor for \mathcal{T}_2 , then $\alpha^{-1}(A)$ is an attractor for \mathcal{T}_1 .

Denumerable (and in particular finite) abstractions play an important role, hence we summarize all interesting and useful implications and equivalences for DMCs, which are direct consequences of Propositions 25 and 33.

► **Proposition 35.** Assume \mathcal{T}_2 is a DMC, which is an α -abstraction of \mathcal{T}_1 . Let $\mathcal{B} = \{\alpha^{-1}(B) \mid B \in \Sigma_2\}$ be the set of α -closed sets of Σ_1 . The following implications and equivalences hold true:





4.2.2 Trickier transfers of properties

We established that decisiveness properties could be transferred through sound abstractions. However, proving soundness of an abstraction is not easy in general, and one way to do it is by proving some decisiveness properties. It is therefore relevant to explore alternatives to prove decisiveness properties. We provide here two such alternatives.

First, we assume a denumerable abstraction, and lower bounds on probabilities of reachability properties.

► **Proposition 36.** Let \mathcal{T}_2 be a DMC such that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 .

1. Assume that there is a finite set $A_2 = \{s_1, \dots, s_n\} \subseteq S_2$ such that A_2 is an attractor for \mathcal{T}_2 and $A_1 = \bigcup_{i=1}^n \alpha^{-1}(s_i) = \alpha^{-1}(A_2)$ is an attractor for \mathcal{T}_1 .
2. Assume moreover that for every $1 \leq i \leq n$, for every α -closed set B in Σ_1 , there exist $p > 0$ and $k \in \mathbb{N}$ such that:
 - for every $\mu \in \text{Dist}(\alpha^{-1}(s_i))$, $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F}_{\leq k} B) \geq p$, or
 - for every $\mu \in \text{Dist}(\alpha^{-1}(s_i))$, $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} B) = 0$.

Then \mathcal{T}_1 is decisive w.r.t. every α -closed set.

While the first condition on transfer of attractors is easily readable, let us discuss the second one. It intuitively says that, whenever some (α -closed) set B can be reached with positive probability from some distribution μ with support $\alpha^{-1}(s_i)$, where s_i is an element of the finite attractor of \mathcal{T}_2 , then it should be reachable with a probability lower-bounded by some p , p being independent of the choice of μ ; furthermore an upper bound k on the number of steps for reaching B is technically required in the proof, but we do not know whether it is needed for the result to hold.

We write (†) for the hypotheses over \mathcal{T}_1 in this proposition. The idea behind this result is that, with probability 1, the attractor of \mathcal{T}_1 will be visited infinitely often, and, if at each visit of the attractor, there is a positive probability to reach some (α -closed) set B , since that probability is by assumption bounded from below, then B will indeed be visited infinitely often with probability 1. This will allow to show the dichotomy between reachability of B and reachability of \tilde{B} , which is required for proving the decisiveness property. The full proof is given in the appendix, page 65, but we give here a sketch. Note that this kind of proofs appears quite often in the literature (see *e.g.* [1, Lemma 3.4], but we have to do it carefully here, since the framework is rather general).

Sketch of proof. Fix $B \subseteq S_2$ and $\mu \in \text{Dist}(S_1)$. Towards a contradiction, assume that \mathcal{T}_1 is not μ -decisive w.r.t. B : this means that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \alpha^{-1}(\widetilde{B}^c)) > 0$.

Since A_1 is an attractor, we deduce from Lemma 19 that

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \alpha^{-1}(\widetilde{B}^c) \wedge \mathbf{G} \mathbf{F} A_1) > 0 .$$

We write $A'_2 \subseteq A_2$ for the non-empty³ set of states s of A_2 such that

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \alpha^{-1}(\widetilde{B}^c) \wedge \mathbf{G} \mathbf{F} \alpha^{-1}(s)) > 0$$

Then obviously $A'_2 \subseteq B^c \cap (\widetilde{B}^c)^c$.

Since $A'_2 \subseteq (\widetilde{B}^c)^c$, from Lemma 14 (third item), the hypothesis (\dagger) and the finiteness of A'_2 , we get that there is $p > 0$ and $k \in \mathbb{N}$ such that for every $s \in A'_2$ and every $\mu \in \text{Dist}(\alpha^{-1}(s))$,

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F}_{\leq k} B) \geq p .$$

Writing A'_1 for $\alpha^{-1}(A'_2)$, we can show that

$$\begin{aligned} 0 &< \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \alpha^{-1}(\widetilde{B}^c) \wedge \mathbf{G} \mathbf{F} A'_1) \\ &\leq \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \mathbf{F} A'_1) \\ &\leq \lim_{n \rightarrow \infty} (1 - p)^n = 0 \end{aligned}$$

which is the required contradiction. \blacktriangleleft

Second, we strengthen the hypothesis on the abstraction, assuming it is finite, but we relax the condition on \mathcal{T}_1 , requiring only a fairness property.

► **Proposition 37.** Let \mathcal{T}_2 be a finite Markov chain such that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Fix $\mu \in \text{Dist}(S_1)$, and assume that \mathcal{T}_1 is μ -fair w.r.t. every α -closed set. Then \mathcal{T}_1 is μ -decisive w.r.t. every α -closed set.

Sketch of proof. We give here the main steps of the proof, the details are postponed to the appendix (page 69).

A key element of the proof relies on the fact that, since \mathcal{T}_2 is a finite MC, it can be viewed as a graph and we can talk of the *bottom strongly connected components (BSCC)* of \mathcal{T}_2 . The first step of the proof aims at showing that, roughly speaking, the union of all BSCCs of \mathcal{T}_2 is a μ -attractor for \mathcal{T}_1 . More precisely, if $\mathcal{C} = \{s \in S_2 \mid \exists C \in \text{BSCC}(\mathcal{T}_2), s \in C\}$, we prove that $\text{Prob}_\mu^{\mathcal{T}_1}(\alpha^{-1}(\mathcal{C})) = 1$. This is shown thanks to the following arguments:

- for each $s \in S_2$, $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(s)) > 0$ implies that $s \in \mathcal{C}$ – this uses the μ -fairness assumption of \mathcal{T}_1 w.r.t. α -closed sets, and the core property of BSCCs (we cannot escape from them);
- using Bayes formula, one can decompose the set of paths according to the states which are visited infinitely often (which corresponds to a decomposition according to the BSCC the path ultimately visit).

Once we have shown that $\alpha^{-1}(\mathcal{C})$ is a μ -attractor for \mathcal{T}_1 , it suffices to observe that for each $B \subseteq S_2$ and each BSCC C of \mathcal{T}_2 , either $B \cap C \neq \emptyset$, or $C \subseteq \widetilde{B}$. Transferring those observations to \mathcal{T}_1 and using Bayes formula to decompose $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B}))$ according to which BSCC is reached, it is easy to check that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B})) = 1$. \blacktriangleleft

³ Since $A_1 = \bigcup_{s \in A_2} \alpha^{-1}(s)$.

4.3 Conditions for completeness and soundness

In our applications (Section 8), completeness will be for free. Indeed, a simple condition (finiteness) implies completeness as stated in the next lemma.

► **Lemma 38.** *If \mathcal{T}_2 is a finite Markov chain and an α -abstraction of \mathcal{T}_1 , then \mathcal{T}_2 is complete.*

Proof. Pick $s_0 \in S_2$, and $\mu \in \text{Dist}(\alpha^{-1}(\{s_0\}))$ (in particular, $\alpha_{\#}(\mu) = \delta_{s_0}$, the Dirac measure over $\{s_0\}$). Assume that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1$ but $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) < 1$.

Since \mathcal{T}_2 is a finite Markov chain, there are $s_1, \dots, s_n \in S_2$ such that

$$\text{Prob}_{\delta_{s_0}}^{\mathcal{T}_2}(\text{Cyl}(s_0, s_1, \dots, s_n)) > 0$$

and for each $\rho = (s_i)_{i \geq 0} \in \text{Cyl}(s_0, \dots, s_n)$ and for each $i \geq 0$, $s_i \notin B$.

For each $0 \leq i \leq n$, we write $A_i = \alpha^{-1}(\{s_i\})$. Then, following Equation (1) (page 18), we get that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(A_0, A_1, \dots, A_n)) > 0$. However, $\text{Cyl}(A_0, A_1, \dots, A_n) \cap \text{Ev}_{\mathcal{T}}(\mathbf{F} \alpha^{-1}(B)) = \emptyset$, yielding a contradiction. ◀

Note that the above lemma does not hold for denumerable abstractions. To illustrate this, any two random walks over \mathbb{N} are abstractions of each other, and it is well-known that almost-sure reachability depends on the probability values.

In general, completeness can be guaranteed by some decisiveness condition on the abstract system. Note that, since finite Markov chains are always decisive, the next lemma actually subsumes the latter one, that we however found interesting to have as such.

► **Lemma 39.** *Let $\mu \in \text{Dist}(S_1)$. Assume that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 and that \mathcal{T}_2 is $\text{Dec}(\alpha_{\#}(\mu))$. Then, \mathcal{T}_2 is a μ -complete α -abstraction.*

Proof. Fix $B \in \mathcal{B}$ and assume that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1$ but $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) < 1$.

Since \mathcal{T}_2 is $\text{Dec}(\alpha_{\#}(\mu))$, we infer from Lemma 14 (fifth item) that $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}((\neg B) \mathbf{U} \widetilde{B}) > 0$, and applying Equation (1) (page 18), we get that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\alpha^{-1}(\neg B) \mathbf{U} \alpha^{-1}(\widetilde{B})) > 0$. This contradicts the hypothesis that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1$. ◀

Proving soundness is more delicate. We nevertheless show that a decisiveness condition on the concrete system will ensure soundness.

► **Proposition 40.** *Let \mathcal{T}_2 be an α -abstraction of \mathcal{T}_1 . Assume \mathcal{T}_1 is decisive w.r.t. every α -closed set. Then \mathcal{T}_2 is a sound α -abstraction of \mathcal{T}_1 .*

Proof. Towards a contradiction assume that $B \in \Sigma_2$ is such that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) < 1$. Since \mathcal{T}_1 is decisive w.r.t. $\alpha^{-1}(B)$ from μ , it holds from Lemma 14 (fifth item) that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\neg \alpha^{-1}(B) \mathbf{U} \widetilde{\alpha^{-1}(B)}) > 0$. Applying Equation (1) again (page 18), we get that $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}((\neg B) \mathbf{U} \widetilde{B}) > 0$, which contradicts the assumption that $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) = 1$. ◀

For \mathcal{T}_2 an α -abstraction of \mathcal{T}_1 , notice that completeness is ensured by a decisiveness assumption on \mathcal{T}_2 , whereas soundness requires \mathcal{T}_1 being decisive w.r.t. every α -closed set. While these conditions look very similar, the condition for soundness is actually harder to check since the abstract STS \mathcal{T}_2 is expected to be simpler than the original concrete STS \mathcal{T}_1 .

5 Using attractors for analyzing STSs

In this section we emphasize a generic approach to the analysis of STSs w.r.t. properties given by a DMA, when the STS satisfies some attractor-based property. This approach is inspired by the works of [2, 14] on lossy channel systems, but is new (as far as we know) in the general context of DMCs, and *a fortiori* of STSs.

As we will see in the next sections, this will yield procedures (which can be turned to effective algorithms for some classes of systems) for the qualitative as well as for the approximate quantitative analysis of STSs. Many of the proofs are inlined here, since they convey interesting ideas.

5.1 The case of DMCs with a finite attractor

We fix a finite set of atomic propositions AP, and we let $\mathcal{T} = (S, \Sigma, \kappa, \text{AP}, \mathcal{L})$ be a labelled DMC. We also let $\mathcal{M} = (Q, q_0, E, \mathcal{F})$ be a DMA. The product $\mathcal{T} \times \mathcal{M}$ has been defined in Section 2.5. First, attractors transfer from \mathcal{T} to the product $\mathcal{T} \times \mathcal{M}$, as stated below, and proven in the technical appendix (page 71).

► **Lemma 41.** *Assume that A is an attractor for \mathcal{T} . Then $A \times Q$ is an attractor for $\mathcal{T} \times \mathcal{M}$. Furthermore, if A is finite, then so is $A \times Q$.*

For the rest of this subsection, we assume that \mathcal{T} has a finite attractor. Applying Lemma 41, the product $\mathcal{T} \times \mathcal{M}$ admits a finite attractor that we denote B . We write $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$ (or simply $\text{Graph}(B)$ when \mathcal{T} and \mathcal{M} are clear from the context) for the finite graph whose vertices are states of B , and in which there is an edge from (s_1, q_1) to (s_2, q_2) if there is a path from (s_1, q_1) to (s_2, q_2) in $\mathcal{T} \times \mathcal{M}$. The *bottom strongly connected components* (BSCCs)⁴ of the graph $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$ play a central role in the model checking of ω -regular properties of \mathcal{T} . Let us first discuss the relationships between the BSCCs and attractors for $\mathcal{T} \times \mathcal{M}$.

► **Lemma 42.** *The following properties are satisfied:*

- *The set $\{(s, q) \in C \mid C \text{ BSCC of } \text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)\}$ is an attractor of $\mathcal{T} \times \mathcal{M}$.*
- *If C and C' are two distinct BSCCs of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$, for every $\mu \in \text{Dist}(S \times Q)$, $\text{Prob}_{\mu}^{\mathcal{T} \times \mathcal{M}}(\mathbf{F} C \wedge \mathbf{F} C') = 0$.*
- *If C is a BSCC of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$, for every $\mu \in \text{Dist}(C)$, $\text{Prob}_{\mu}^{\mathcal{T} \times \mathcal{M}}(\mathbf{G} \mathbf{F} C) = 1$.*

Proof. The first property is obvious. The second property is a consequence of the fact that there is no path between two states of two different BSCCs. This second property implies that for each BSCC $C' \neq C$ of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$ and for each $\mu \in \text{Dist}(C)$, $\text{Prob}_{\mu}^{\mathcal{T} \times \mathcal{M}}(\mathbf{F} C') = 0$. From the first property and Lemma 19, we know that for each $\mu \in \text{Dist}(S \times Q)$, $\text{Prob}_{\mu}^{\mathcal{T} \times \mathcal{M}}(\mathbf{G} \mathbf{F} \bigvee_{C' \in \text{BSCC}(\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B))} C') = 1$. This holds true in particular for each $\mu \in \text{Dist}(C)$ and thus, from the previous observation for such initial distributions, we get that $\text{Prob}_{\mu}^{\mathcal{T} \times \mathcal{M}}(\mathbf{G} \mathbf{F} C) = 1$ for each $\mu \in \text{Dist}(C)$. ◀

From Lemma 42, the BSCCs of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$ form an attractor, and once the system enters a BSCC C , only that BSCC will be visited again, and this will happen infinitely often with probability 1. In particular, the satisfaction of the Muller condition in $\mathcal{T} \times \mathcal{M}$, inherited from \mathcal{F} , can be characterized by the BSCCs satisfying the Muller condition \mathcal{F} (in a sense that we will make precise).

⁴ Those are strongly connected components which cannot be left.

► **Definition 43** (Good BSCC). A BSCC C of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$ is *good for \mathcal{F}* , written $C \in \text{Good}_{\mathcal{T} \times \mathcal{M}}^B(\mathcal{F})$, if there exists $F \in \mathcal{F}$ such that

- (a) for every state $(s, q) \in S \times Q$, if there exists $(r, p) \in C$ with a path from (r, p) to (s, q) in $\mathcal{T} \times \mathcal{M}$, then $q \in F$; and
- (b) for every $q \in F$ there exists $s \in S$, there exists a state $(r, p) \in C$ with a path from (r, p) to (s, q) in $\mathcal{T} \times \mathcal{M}$.

Let C be an arbitrary BSCC of $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$. We define the set $F_C = \{q \in Q \mid \exists s \in S, \exists (r, p) \in C \text{ s.t. there is a path from } (r, p) \text{ to } (s, q)\}$ as the set of states of the Muller automaton that can be reached from C . Within a BSCC, all reachable states will actually be visited infinitely often almost-surely. More precisely, we state the following result:

► **Lemma 44.** *For every $(s, q) \in C$, $\text{Prob}_{\delta_{(s, q)}^{\mathcal{T} \times \mathcal{M}}}(\text{Inf} = F_C) = 1$.*⁵

Proof. Let $(s, q) \in C$, and $\rho = (s, q)(s_1, q_1)(s_2, q_2) \dots$ a path in $\mathcal{T} \times \mathcal{M}$ starting at (s, q) . By definition of F_C , all q_i 's are in F_C , hence $\text{Prob}_{\delta_{(s, q)}^{\mathcal{T} \times \mathcal{M}}}(\text{Inf} \subseteq F_C) = 1$.

We now argue why all elements of F_C are actually almost-surely visited infinitely often. Fix $p \in F_C$ and (r, p) that is reachable from C . All two states of C are reachable one from each other; thus, from every state of C , (r, p) is reachable through a finite path. Hence there is some $\iota > 0$ and $k \in \mathbb{N}$ such that for every state $(s', q') \in C$,

$$\text{Prob}_{\delta_{(s', q')}^{\mathcal{T} \times \mathcal{M}}}(\mathbf{F}_{\leq k}(r, p)) \geq \iota .$$

Applying a reasoning similar to the proof of Proposition 36, we get that $\text{Prob}_{\delta_{(s, q)}^{\mathcal{T} \times \mathcal{M}}}(\mathbf{G} \mathbf{F}(r, p) \mid \mathbf{G} \mathbf{F} C) = 1$. Indeed, $\text{Prob}_{\delta_{(s, q)}^{\mathcal{T} \times \mathcal{M}}}(\mathbf{F} \mathbf{G} \neg(r, p) \wedge \mathbf{G} \mathbf{F} C) \leq \lim_{n \rightarrow \infty} (1 - \iota)^n = 0$. Thanks to the third item of Lemma 42, we obtain that

$$\text{Prob}_{\delta_{(s, q)}^{\mathcal{T} \times \mathcal{M}}}(\mathbf{G} \mathbf{F}(r, p)) = 1 .$$

We conclude that $\text{Prob}_{\delta_{(s, q)}^{\mathcal{T} \times \mathcal{M}}}(\text{Inf} \supseteq F_C) = 1$, which completes the proof. ◀

As a consequence:

► **Corollary 45.** *For every initial distribution $\mu \in \text{Dist}(S)$ for \mathcal{T} and for every $q \in Q$, $\text{Prob}_{\mu \times \delta_{q_0}^{\mathcal{T} \times \mathcal{M}}}(\mathbf{F} C) > 0$ implies $\text{Prob}_{\mu \times \delta_{q_0}^{\mathcal{T} \times \mathcal{M}}}(\text{Inf} = F_C \mid \mathbf{F} C) = 1$.*

We can now completely characterize the probability of satisfying an ω -regular property.

► **Theorem 46.** *Let \mathcal{T} be a labelled DMC with a finite attractor, and $\mathcal{M} = (Q, q_0, E, \mathcal{F})$ be a DMA. Then, for every initial distribution $\mu \in \text{Dist}(S)$ for \mathcal{T} :*

$$\text{Prob}_{\mu \times \delta_{q_0}^{\mathcal{T} \times \mathcal{M}}}(\text{Inf} \in \mathcal{F}) = \sum_{C \in \text{Good}_{\mathcal{T} \times \mathcal{M}}^B(\mathcal{F})} \text{Prob}_{\mu \times \delta_{q_0}^{\mathcal{T} \times \mathcal{M}}}(\mathbf{F} C)$$

where B is an attractor for $\mathcal{T} \times \mathcal{M}$.

⁵ We recall that $\text{Inf} = F_C$ characterizes the set of runs ρ' in $\mathcal{T} \times \mathcal{M}$ such that $\text{Inf}(\mathcal{L}'(\rho')) = F_C$ (\mathcal{L}' is the labelling function of $\mathcal{T} \times \mathcal{M}$ such that $\mathcal{L}'(s, q) = q$).

Proof. Fix $\mu \in \text{Dist}(S)$ and $q \in Q$. As stated in Lemma 42, the BSCCs of $\text{Graph}(B)$ form an attractor for $\mathcal{T} \times \mathcal{M}$, and two BSCCs are probabilistically disjoint. Using Bayes formula with a disjunction over the BSCCs, we can write:

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathcal{M}}(\text{Inf} \in \mathcal{F}) = \sum_{\substack{C \text{ BSCC of } \text{Graph}(B) \\ C \text{ } \mu \times \delta_{q_0}\text{-reachable}}} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathcal{M}}(\mathbf{F} C) \cdot \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathcal{M}}(\text{Inf} \in \mathcal{F} \mid \mathbf{F} C)$$

where we say that C is $\mu \times \delta_{q_0}$ -reachable whenever $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathcal{M}}(\mathbf{F} C) > 0$. Hence we deduce that:

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathcal{M}}(\text{Inf} \in \mathcal{F}) = \sum_{C \text{ BSCC of } \text{Graph}(B)} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathcal{M}}(\mathbf{F} C) \cdot \mathbf{1}_{\mathcal{F}}(F_C)$$

thanks to Corollary 45, where $\mathbf{1}_{\mathcal{F}}$ is the characteristic function of \mathcal{F} (that is, $\mathbf{1}_{\mathcal{F}}(F) = 1$ if $F \in \mathcal{F}$, and $\mathbf{1}_{\mathcal{F}}(F) = 0$ otherwise). This concludes the proof of the theorem. ◀

5.2 General STSs via an abstraction

While the previous approach is adapted to DMCs, it does not apply directly to general STSs: indeed, it is unlikely that general STSs have *finite* attractors, and finiteness of the attractor is fundamental for the correctness of the approach. The idea will then be to rely on an abstraction that admits a finite attractor, and to transfer properties through that abstraction.

Let $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1, \text{AP}, \mathcal{L}_1)$ and $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2, \text{AP}, \mathcal{L}_2)$ be two LSTSs such that \mathcal{T}_2 is a DMC, which is an α -abstraction of \mathcal{T}_1 . Under certain conditions, we show how to express the probability of satisfying an ω -regular property represented by a DMA $\mathcal{M} = (Q, q_0, E, \mathcal{F})$ in \mathcal{T}_1 using the abstraction \mathcal{T}_2 . We consider both the product $\mathcal{T}_1 \times \mathcal{M}$ and the product $\mathcal{T}_2 \times \mathcal{M}$.

First we justify why, within a slight abuse of terminology, $\mathcal{T}_2 \times \mathcal{M}$ can be viewed as an α -abstraction of $\mathcal{T}_1 \times \mathcal{M}$. We also exhibit a sufficient condition under which it is sound.

► **Lemma 47.** *Let $\alpha_{\mathcal{M}} : S_1 \times Q \rightarrow S_2 \times Q$ be the lifting of α such that $\alpha_{\mathcal{M}}(s, q) = (\alpha(s), q)$. If \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 , then $\mathcal{T}_2 \times \mathcal{M}$ is an $\alpha_{\mathcal{M}}$ -abstraction of $\mathcal{T}_1 \times \mathcal{M}$. Furthermore, if $\mathcal{T}_1 \times \mathcal{M}$ is $\text{Dec}(\mathcal{B})$ where $\mathcal{B} = \{\alpha_{\mathcal{M}}^{-1}(B) \mid B \in \Sigma'_2\}$, then $\mathcal{T}_2 \times \mathcal{M}$ is a sound $\alpha_{\mathcal{M}}$ -abstraction of $\mathcal{T}_1 \times \mathcal{M}$.*

While the proof of the first part of the lemma is technical hence postponed to the appendix (page 72), the second part of the lemma is a consequence of Proposition 40.

► **Remark 48.** In the sequel, our applications will be smooth enough to meet the hypothesis: $\mathcal{T}_1 \times \mathcal{M}$ is decisive w.r.t. $\alpha_{\mathcal{M}}$ -closed sets. However we still have several open questions. The first one is the following: does soundness between \mathcal{T}_2 and \mathcal{T}_1 imply soundness between $\mathcal{T}_2 \times \mathcal{M}$ and $\mathcal{T}_1 \times \mathcal{M}$? While this seems quite natural, it is surprisingly tricky. Although we did not manage to find a counter-example for this general question, we found one for a fixed initial distribution. It is described in Example 96 in the appendix (page 72) and highlights some difficulties we encounter when aiming at transferring analysis from the abstraction to the concrete model.

This justifies the fact that we assumed decisiveness. As we already know, if \mathcal{T}_2 is a sound α -abstraction of \mathcal{T}_1 and \mathcal{T}_2 is decisive w.r.t. any set of states, then \mathcal{T}_1 is decisive w.r.t. any α -closed sets. Then the second natural question is the following: does decisiveness w.r.t. α -closed sets for \mathcal{T}_1 imply decisiveness w.r.t. $\alpha_{\mathcal{M}}$ -closed sets for $\mathcal{T}_1 \times \mathcal{M}$? Again, we do not have a general counter-example, but we have one for a fixed initial distribution. This is described in Example 96 in the appendix (page 72).

From now on, whenever $\mathcal{T}_1 \times \mathcal{M}$ is decisive w.r.t. $\alpha_{\mathcal{M}}$ -closed sets and thus the previous result is applicable, we will abusively write α for $\alpha_{\mathcal{M}}$.

We focus now on the case where \mathcal{T}_2 has a finite attractor.⁶ Applying Lemma 41, $\mathcal{T}_2 \times \mathcal{M}$ has also a finite attractor, which we denote B_2 . We reuse notations of the previous subsection, in particular the graph of the attractor $\text{Graph}_{\mathcal{T}_2 \times \mathcal{M}}(B_2)$, and the set F_C of recurring states when C is a BSCC of that graph.

The following lemma is a counterpart to Lemma 42 for \mathcal{T}_1 . Under the hypothesis that $\mathcal{T}_1 \times \mathcal{M}$ is decisive w.r.t. α -closed sets, even though $\mathcal{T}_1 \times \mathcal{M}$ does not have a finite attractor, it has an attractor with an interesting structure inherited from $\mathcal{T}_2 \times \mathcal{M}$. In the sequel, we write $\mathcal{B} = \{\alpha^{-1}(B) \mid B \in \Sigma'_2\}$.

► **Lemma 49.** *Assume \mathcal{T}_2 has a finite attractor, and assume that $\mathcal{T}_2 \times \mathcal{M}$ is a sound α -abstraction of $\mathcal{T}_1 \times \mathcal{M}$. Write B_2 for an attractor of $\mathcal{T}_2 \times \mathcal{M}$. The following properties are satisfied:*

- *The set $\alpha^{-1}(\{(s, q) \in C \mid C \text{ BSCC of } \text{Graph}_{\mathcal{T}_2 \times \mathcal{M}}(B_2)\})$ is an attractor of $\mathcal{T}_1 \times \mathcal{M}$.*
- *If C and C' are two distinct BSCCs of $\text{Graph}_{\mathcal{T}_2 \times \mathcal{M}}(B_2)$, for every $\mu \in \text{Dist}(S_1 \times Q)$, $\text{Prob}_{\mu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \alpha^{-1}(C) \wedge \mathbf{F} \alpha^{-1}(C')) = 0$.*
- *If C is a BSCC of $\text{Graph}_{\mathcal{T}_2 \times \mathcal{M}}(B_2)$, for every $\mu \in \text{Dist}(\alpha^{-1}(C))$, $\text{Prob}_{\mu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{G} \mathbf{F} \alpha^{-1}(C)) = 1$.*

Proof. Since $\mathcal{T}_2 \times \mathcal{M}$ is a sound α -abstraction of $\mathcal{T}_1 \times \mathcal{M}$, the first property is derived from Proposition 34 and Lemma 42. The second property is a consequence of Lemma 42, and of the fact that $\mathcal{T}_2 \times \mathcal{M}$ is an α -abstraction of $\mathcal{T}_1 \times \mathcal{M}$. Finally, the third property is, as in the proof of Lemma 42, a consequence of the second point and of Lemma 19. ◀

We then prove a counterpart to Lemma 44 for \mathcal{T}_1 , which shows that a BSCC is characterized by the set F_C of states that are visited infinitely often from C .

► **Lemma 50.** *Assume \mathcal{T}_2 has a finite attractor, and assume that $\mathcal{T}_2 \times \mathcal{M}$ is a sound α -abstraction of $\mathcal{T}_1 \times \mathcal{M}$. Let C be a BSCC of $\text{Graph}_{\mathcal{T}_2 \times \mathcal{M}}(B_2)$, and $\mu \in \text{Dist}(\alpha^{-1}(C))$. Then:*

$$\text{Prob}_{\mu}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Inf} = F_C) = 1.$$

Proof. As already argued in the proof of Lemma 44, for every $p \in F_C$, for every state $s_2 \in C$, $\text{Prob}_{\delta_{s_2}}^{\mathcal{T}_2 \times \mathcal{M}}(\mathbf{F} p) = 1$ (we abusively write p for the measurable set $S_2 \times \{p\}$). Since $\mathcal{T}_2 \times \mathcal{M}$ is a sound α -abstraction of $\mathcal{T}_1 \times \mathcal{M}$, we derive for every $\nu \in \text{Dist}(\alpha^{-1}(C))$ that $\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} p) = 1$ (as before we abusively write p for $S_1 \times \{p\} = \alpha^{-1}(S_2 \times \{p\})$). We can then show that for each $\nu \in \text{Dist}(\alpha^{-1}(C))$ and for each $p \in F_C$,

$$\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{G} \mathbf{F} p) = 1.$$

Indeed, towards a contradiction, assume that there is a distribution $\nu \in \text{Dist}(\alpha^{-1}(C))$ such that $\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{G} \mathbf{F} p) < 1$, i.e. $\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \mathbf{G} \neg p) > 0$. From the third point of Lemma 49, we get that $\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{G} \mathbf{F} \alpha^{-1}(C) \wedge \mathbf{F} \mathbf{G} \neg p) > 0$. Now, observe that

$$\text{Ev}_{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{G} \mathbf{F} \alpha^{-1}(C) \wedge \mathbf{F} \mathbf{G} \neg p) \subseteq \text{Ev}_{\mathcal{T}_1 \times \mathcal{M}}\left(\bigcup_{n \in \mathbb{N}} (\mathbf{F}_{=n} \alpha^{-1}(C) \wedge \mathbf{G}_{\geq n} \neg p)\right).$$

⁶ As \mathcal{T}_2 has a finite attractor, it is decisive and thus \mathcal{T}_2 is a complete α -abstraction of \mathcal{T}_1 by Lemma 39.

It follows that there is $n \in \mathbb{N}$ such that $\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F}_{=n} \alpha^{-1}(C) \wedge \mathbf{G}_{\geq n} \neg p) > 0$. From Lemma 5, we get that there is $\nu' \in \text{Dist}(S'_1)$ (with $S'_1 = S_1 \times Q$) such that

$$\begin{aligned}
 & \text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F}_{=n} \alpha^{-1}(C) \wedge \mathbf{G}_{\geq n} \neg p) \\
 &= \lim_{m \rightarrow \infty} \text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Cyl}(\overbrace{S'_1, \dots, S'_1}^{n \text{ times}}, \alpha^{-1}(C) \wedge \neg p, \overbrace{\neg p, \dots, \neg p}^{m \text{ times}})) \\
 &\leq \lim_{m \rightarrow \infty} \text{Prob}_{\nu'}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Cyl}(\alpha^{-1}(C) \wedge \neg p, \overbrace{\neg p, \dots, \neg p}^{m \text{ times}})) \quad \text{from Lemma 5} \\
 &= \lim_{m \rightarrow \infty} \nu'(\alpha^{-1}(C)) \cdot \text{Prob}_{\nu'_{\alpha^{-1}(C)}}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Cyl}(\neg p, \overbrace{\neg p, \dots, \neg p}^{m \text{ times}})) \\
 &= \nu'(\alpha^{-1}(C)) \cdot \text{Prob}_{\nu'_{\alpha^{-1}(C)}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{G} \neg p).
 \end{aligned}$$

From the assumption, we thus get that $\text{Prob}_{\nu'_{\alpha^{-1}(C)}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{G} \neg p) > 0$ where $\nu'_{\alpha^{-1}(C)} \in \text{Dist}(\alpha^{-1}(C))$ which is the required contradiction. Hence, for each $\nu \in \text{Dist}(\alpha^{-1}(C))$ and for each $p \in F_C$, $\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{G} \mathbf{F} p) = 1$.

It now suffices to show that, from any $\nu \in \text{Dist}(\alpha^{-1}(C))$, no other state is visited almost-surely infinitely often. Fix $p' \notin F_C$. Then, by definition of F_C , we have that $\text{Prob}_{\alpha_{\#}(\nu)}^{\mathcal{T}_2 \times \mathcal{M}}(\mathbf{F} p') = 0$. Since $\mathcal{T}_2 \times \mathcal{M}$ is an α -abstraction of $\mathcal{T}_1 \times \mathcal{M}$, we deduce that $\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} p') = 0$.

We conclude that $\text{Prob}_{\nu}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Inf} = F_C) = 1$, which is the claim of the lemma. \blacktriangleleft

We are now in a position to decompose the probability to satisfy the Muller condition \mathcal{F} in $\mathcal{T}_1 \times \mathcal{M}$ into the reachability probability of good BSCCs in $\text{Graph}_{\mathcal{T}_2 \times \mathcal{M}}(B_2)$.

► Theorem 51. *Let \mathcal{T}_1 and \mathcal{T}_2 be two LSTSs such that \mathcal{T}_2 is a DMC with a finite attractor, and \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Let $\mathcal{M} = (Q, q_0, E, \mathcal{F})$ be a DMA. Assume moreover that $\mathcal{T}_2 \times \mathcal{M}$ is an α -sound abstraction of $\mathcal{T}_1 \times \mathcal{M}$, and that B_2 is a finite attractor of $\mathcal{T} \times \mathcal{M}$. Then, for every initial distribution μ for \mathcal{T}_1 :*

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Inf} \in \mathcal{F}) = \sum_{C \in \text{Good}_{\mathcal{T}_2 \times \mathcal{M}}^{B_2}(\mathcal{F})} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \alpha^{-1}(C)) .$$

Proof. Applying Lemma 50, for every $\mu \in \text{Dist}(S_1)$, assuming $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \alpha^{-1}(C)) > 0$, then $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Inf} = F_C \mid \mathbf{F} \alpha^{-1}(C)) = 1$. By the two first properties of Lemma 49, we can write the following Bayes formula, with a disjunction over the BSCCs of $\text{Graph}_{\mathcal{T}_2 \times \mathcal{M}}(B_2)$:

$$\begin{aligned}
 \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Inf} \in \mathcal{F}) &= \sum_{\substack{C \text{ BSCC of } \text{Graph}_{\mathcal{T}_2 \times \mathcal{M}}(B_2) \\ C \text{ } \mu \times \delta_{q_0}\text{-reachable}}} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \alpha^{-1}(C)) \cdot \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Inf} \in \mathcal{F} \mid \mathbf{F} \alpha^{-1}(C)) \\
 &= \sum_{C \text{ BSCC of } \text{Graph}_{\mathcal{T}_2 \times \mathcal{M}}(B_2)} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \alpha^{-1}(C)) \cdot \mathbb{1}_{\mathcal{F}}(F_C) \\
 &= \sum_{C \in \text{Good}_{\mathcal{T}_2 \times \mathcal{M}}^{B_2}(\mathcal{F})} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \alpha^{-1}(C))
 \end{aligned}$$

This concludes the proof of the theorem. \blacktriangleleft

6 Qualitative analysis

In this section, we rely on the notions previously introduced and studied to design generic procedures for the qualitative analysis of properties of STSs, under some assumptions that will be made precise. We emphasize that these are procedures rather than algorithms, since algorithms would require some effectiveness conditions on the STSs (numerical conditions, or decidability of some graph properties in the underlying non-stochastic model). Next, we will make explicit necessary conditions to obtain algorithms from the generic procedures. For most natural STSs (and in particular for our applications – see Section 8), these conditions will be satisfied.

For the next two subsections, we fix an STS $\mathcal{T} = (S, \Sigma, \kappa)$.

6.1 Basic properties under decisiveness hypotheses

Our objective here is to describe generic procedures that capture the qualitative (almost-sure and positive) satisfaction of reachability and repeated reachability properties.

Given $B \in \Sigma$ a measurable set, recall that $\widetilde{B} = \{s \in S \mid \text{Prob}_{\delta_s}^{\mathcal{T}}(\mathbf{F} B) = 0\}$ denotes its avoid-set. Some properties of this set, while not crucial for the understanding but required for the proofs, are given in Appendix E.1 (page 73).

Extending the approach of [1], we establish characterizations of the qualitative satisfaction of (repeated) reachability properties in terms of the positive satisfaction of reachability-like properties. We advocate that these are simpler to check on STSs: positive reachability amounts to guessing a “symbolic” path (or cylinder) leading to the target, and to showing that this path has a positive measure. The next proposition is stated in greater details as Proposition 99 in the technical Appendix E (page 74).

► **Proposition 52.** Let $\mu \in \text{Dist}(S)$. Then we have the following implications, yielding various characterizations for the qualitative analysis of STSs (under specified assumptions):

Almost-sure reachability If \mathcal{T} is $\text{Dec}(\mu, B)$, then:

$$\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} B) = 1 \iff \text{Prob}_{\mu}^{\mathcal{T}}(\neg B \mathbf{U} \widetilde{B}) = 0 .$$

Almost-sure repeated reachability If \mathcal{T} is $\text{StrDec}(\mu, B)$, then:

$$\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B) = 1 \iff \text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \widetilde{B}) = 0 .$$

Positive repeated reachability If \mathcal{T} is $\text{Dec}(\mu, \widetilde{B})$ and $\text{PersDec}(\mu, B)$, then:

$$\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{G} \mathbf{F} B) > 0 \iff \text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F} \widetilde{\widetilde{B}}) > 0 .$$

While the two first characterizations are quite intuitive under the corresponding decisiveness assumptions, let us comment on the characterization of positive repeated reachability: the set $\widetilde{\widetilde{B}}$ is the set from which one cannot reach \widetilde{B} (or with probability 0), hence from which we will be able to revisit B again and again. With this interpretation in mind, the characterization is somewhat natural.

This reduces all these problems to checking the (non-)positivity of some reachability, or a slight generalization thereof, property in the STS. Those are the simplest properties one can hope to be decidable in a class of models. Effectiveness hence relies here on the computation of avoid-sets, avoid-sets of avoid-sets, and on the decidability of the positive reachability (or Until) problem.

6.2 Basic properties through abstractions

Via abstractions, one can reduce the qualitative analysis of basic properties (reachability and repeated reachability) from the concrete model to the abstract model. Indeed, one can use the previous results (Propositions 23 and 33 together with Proposition 52), and show:

► **Proposition 53.** Assume \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 , and fix $B \in \Sigma_2$.

- Let $\mu \in \text{Dist}(S_1)$ be an initial distribution for \mathcal{T}_1 . Assume that \mathcal{T}_2 is μ -sound and μ -complete. Then:

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1 \text{ iff } \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) = 1 .$$

- Assume that \mathcal{T}_2 is sound and complete, and that \mathcal{T}_2 is $\text{StrDec}(B)$. Then for every $\mu \in \text{Dist}(S_1)$:

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(B)) = 1 \text{ iff } \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{G} \mathbf{F} B) = 1 .$$

- Assume that \mathcal{T}_2 is sound and complete, and that \mathcal{T}_2 is $\text{PersDec}(B)$ and $\text{Dec}(\widetilde{B})$. Then for every $\mu \in \text{Dist}(S_1)$:

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(B)) = 1 \text{ iff } \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{G} \mathbf{F} B) = 1 .$$

This allows one to perform the qualitative analysis of (repeated) reachability properties in \mathcal{T}_1 on its abstraction \mathcal{T}_2 , which is quite useful since \mathcal{T}_2 will usually be simpler than \mathcal{T}_1 .

6.3 ω -regular properties in DMCs with a finite attractor

Following Section 5.1, under the assumption that the STS has a finite attractor, we have completely characterized the probability of satisfying the property defined by a DMA using the probability of reaching BSCCs of a finite graph (Theorem 46). Using that result, we get the following characterization of the almost-sure satisfaction relation.

► **Corollary 54** (Almost-sure ω -regular property). *Let \mathcal{T} be a labelled DMC with a finite attractor, and $\mathcal{M} = (Q, q_0, E, \mathcal{F})$ be a DMA. Let B be a finite attractor for $\mathcal{T} \times \mathcal{M}$. For every initial distribution $\mu \in \text{Dist}(S)$ for \mathcal{T} :*

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathcal{M}}(\text{Inf} \in \mathcal{F}) = 1 \text{ if and only if}$$

$$\text{every BSCC } C \text{ of } \text{Graph}_{\mathcal{T} \times \mathcal{M}}(B) \text{ such that } \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathcal{M}}(\mathbf{F} C) > 0 \text{ is good for } \mathcal{F}.$$

In order to turn this characterization into a decision procedure, we need to be able to compute the attractor B for $\mathcal{T} \times \mathcal{M}$, and to build the graph $\text{Graph}_{\mathcal{T} \times \mathcal{M}}(B)$; also one needs to be able to compute for every BSCC C the set F_C .

6.4 ω -regular properties of general STSs via abstraction and finite attractor

Following Section 5.2, under several assumptions over an abstraction, we have completely characterized the probability for a concrete system to satisfy a property given as a DMA using a decomposition of a graph defined for the abstract system (Theorem 51). From that result, we deduce the following characterization of the almost-sure satisfaction relation via an abstraction. It turns out that the value resulting from the decomposition is equal to 1 if, and only if, the property is almost-surely satisfied by the abstract system.

► **Corollary 55.** *Let \mathcal{T}_1 and \mathcal{T}_2 be two LSTSs such that \mathcal{T}_2 is a DMC with a finite attractor, and \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Let $\mathcal{M} = (Q, q_0, E, \mathcal{F})$ be a DMA. Assume moreover that $\mathcal{T}_2 \times \mathcal{M}$ is an α -sound abstraction of $\mathcal{T}_1 \times \mathcal{M}$. Then, for every initial distribution μ for \mathcal{T}_1 :*

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Inf} \in \mathcal{F}) = 1 \quad \text{if and only if} \quad \text{Prob}_{\alpha_{\#}(\mu \times \delta_{q_0})}^{\mathcal{T}_2 \times \mathcal{M}}(\text{Inf} \in \mathcal{F}) = 1 .$$

Hence, this reduces the almost-sure model-checking of a property given by \mathcal{M} in \mathcal{T}_1 to the almost-sure model-checking of a reachability property (applying Corollary 54). For the approach to be effective, it is sufficient that the analysis at the level of $\mathcal{T}_2 \times \mathcal{M}$ is effective.

As already quickly mentioned, under the hypotheses of Corollary 55, the abstraction $\mathcal{T}_2 \times \mathcal{M}$ is complete (since it has a finite attractor). Though it is not explicitly used, we could not have such an equivalence without some completeness of the abstraction.

► **Remark 56** (Discussion on the approach of [17]). While the notion of abstraction was not precisely defined in [17] for stochastic timed automata, it was implicitly already there. Also, decidability of the almost-sure satisfaction was ensured thanks to a fairness condition. Using the terminology of the current paper, the framework was the following: \mathcal{T}_1 and \mathcal{T}_2 are two STSs such that \mathcal{T}_2 is a *finite* Markov chain which is an α -abstraction of \mathcal{T}_1 . Then the condition for the abstraction to yield interesting results was that \mathcal{T}_1 should be fair w.r.t. every α -closed sets (the latter condition implying the fairness of $\mathcal{T}_1 \times \mathcal{M}$, for \mathcal{M} a DMA). Thanks to Proposition 37, this implies that $\mathcal{T}_1 \times \mathcal{M}$ is actually decisive w.r.t. α -closed sets. Applying Proposition 40, we get that $\mathcal{T}_2 \times \mathcal{M}$ is sound abstraction of $\mathcal{T}_1 \times \mathcal{M}$. Given that \mathcal{T}_2 is finite, it trivially has a finite attractor. Hence, the conditions of Theorem 51 are satisfied, and the approach of [17] was then a particular case of that theorem, when applied to specific subclasses of stochastic timed automata (further details are provided in Subsection 8.1).

7 Approximate quantitative analysis

Beyond qualitative analysis, we are interested in quantitative analysis of stochastic systems, that is, in computing the probability that an STS satisfies a given property. No generic decidability results can be stated in the very general context of STSs,⁷. We thus focus here on approximate analysis and develop generic approximation procedures which, under reasonable assumptions, allow one to compute within arbitrary precision, the probability of a given property. As for properties, we consider first reachability, then repeated reachability, later ω -regular properties, and finally some timed properties.

For the next two subsections, we fix an STS $\mathcal{T} = (S, \Sigma, \kappa)$, and a distribution $\mu \in \text{Dist}(S)$.

7.1 Quantitative reachability under decisiveness hypotheses

In order to approximate the reachability probability of a set $B \in \Sigma$ in \mathcal{T} , we define the two following sequences, similar to the ones given for decisive Markov chains [1]. For every $n \in \mathbb{N}$:

$$\begin{cases} p_n^{\text{Yes}} &= \text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F}_{\leq n} B); \\ p_n^{\text{No}} &= \text{Prob}_{\mu}^{\mathcal{T}}(\neg B \mathbf{U}_{\leq n} \widetilde{B}). \end{cases}$$

⁷ Note already that there is no uniform effective way to represent STSs, so that we can hardly expect generic (and effective) procedures or algorithms.

Since the sequences of events $(\mathbf{F}_{\leq n} B)_{n \in \mathbb{N}}$ and $(\neg B \mathbf{U}_{\leq n} \widetilde{B})_{n \in \mathbb{N}}$ are non-decreasing and converge respectively to $\mathbf{F} B$ and $\neg B \mathbf{U} \widetilde{B}$, the sequences $(p_n^{\text{Yes}})_n$ and $(p_n^{\text{No}})_n$ are non-decreasing and converge respectively to $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B)$ and $\text{Prob}_\mu^{\mathcal{T}}(\neg B \mathbf{U} \widetilde{B})$. Assuming now that \mathcal{T} is decisive w.r.t. B , the two limits are related, as stated below. The proof of this proposition can be found in Appendix, page 75.

► **Proposition 57** (Approximation scheme for reachability properties). If \mathcal{T} is $\text{Dec}(\mu, B)$, then the two sequences $(p_n^{\text{Yes}})_n$ and $(1 - p_n^{\text{No}})_n$ are adjacent⁸ and converge to $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B)$.

To obtain an ε -approximation for $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B)$, it suffices to evaluate p_n^{Yes} and p_n^{No} for larger and larger values of n , until $1 - p_n^{\text{No}} - p_n^{\text{Yes}} < \varepsilon$, and to return p_n^{Yes} . This scheme is effective as soon as one can compute \widetilde{B} , and the probability (from μ) of cylinders of the forms $\text{Cyl}(\underbrace{S, \dots, S}_{n \text{ times}}, B)$ and $\text{Cyl}(\underbrace{\neg B, \dots, \neg B}_{n \text{ times}}, \widetilde{B})$. In case p_n^{Yes} and p_n^{No} cannot be computed exactly, but can only be approximated up to any desired error bound, this scheme can be refined to obtain a 2ε -approximation for $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B)$.

► **Remark.** The above approximation scheme can be adapted to Until properties of the form $B' \mathbf{U} B$ (for $B, B' \in \Sigma$) in a straightforward way as follows: for every $n \in \mathbb{N}$,

$$\begin{cases} \hat{p}_n^{\text{Yes}} &= \text{Prob}_\mu^{\mathcal{T}}(B' \mathbf{U}_{\leq n} B); \\ \hat{p}_n^{\text{No}} &= \text{Prob}_\mu^{\mathcal{T}}(\neg B \mathbf{U}_{\leq n} (\widetilde{B} \vee \neg B')). \end{cases}$$

Convergence of that scheme here also relies on a decisiveness property w.r.t. B .

7.2 Quantitative repeated reachability under decisiveness hypotheses

We now define two sequences that will yield an approximation scheme for a repeated reachability probability, under stronger assumptions on the model. For every $n \in \mathbb{N}$:

$$\begin{cases} q_n^{\text{Yes}} &= \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F}_{\leq n} \widetilde{B}); \\ q_n^{\text{No}} &= \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F}_{\leq n} \widetilde{B}). \end{cases}$$

Here again, with no assumption on \mathcal{T} , clearly enough, the sequences $(q_n^{\text{Yes}})_n$ and $(q_n^{\text{No}})_n$ are non-decreasing and converge respectively to $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} \widetilde{B})$ and $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} \widetilde{B})$. Assuming now that \mathcal{T} is persistently decisive w.r.t to B and decisive w.r.t. \widetilde{B} , the two sequences are closely related, as stated below. The proof of this result can be found page 76.

► **Proposition 58** (Approximation scheme for repeated reachability). If \mathcal{T} is $\text{PersDec}(\mu, B)$ and $\text{Dec}(\mu, \widetilde{B})$, then the two sequences $(q_n^{\text{Yes}})_n$ and $(1 - q_n^{\text{No}})_n$ are adjacent and converge to $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B)$.

Effectiveness of the scheme relies on the computability of the avoid sets \widetilde{B} and $\widetilde{\widetilde{B}}$, and on the effective computation of the probability of cylinders of the forms $\text{Cyl}(\underbrace{\neg \widetilde{B}, \dots, \neg \widetilde{B}}_{n \text{ times}}, \widetilde{\widetilde{B}})$ and $\text{Cyl}(\underbrace{\widetilde{\widetilde{B}}, \dots, \widetilde{\widetilde{B}}}_{n \text{ times}}, \widetilde{B})$. Similarly as before, in case q_n^{Yes} and q_n^{No} cannot be computed exactly, but can only be approximated up to any desired error bound, this scheme can be refined to obtain a 2ε -approximation for $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B)$.

⁸ Recall that two sequences $(a_n)_{n \in \mathbb{N}}$ and $(b_n)_{n \in \mathbb{N}}$ are said *adjacent* if w.l.o.g. (a_n) is non-decreasing, (b_n) is non-increasing and the sequence $(a_n - b_n)_{n \in \mathbb{N}}$ converges to 0.

7.3 ω -regular properties in DMC with a finite attractor

To go beyond reachability and repeated reachability, we now consider an ω -regular property given by a DMA $\mathcal{M} = (Q, q_0, E, \mathcal{F})$. We assume that $\mathcal{T} = (S, \Sigma, \kappa, \text{AP}, \mathcal{L})$ is a labelled DMC.

In order to approximate the probability that the model satisfies this external specification, we assume that \mathcal{T} has a finite attractor. Following Section 6.3, we consider the finite attractor B of $\mathcal{T} \times \mathcal{M}$, and we apply Theorem 46: for each $\mu \in \text{Dist}(S)$,

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

Thus, the computation of the probability that a given model satisfies a given external specification is reduced to the computation of a reachability probability. Now, given that \mathcal{T} and hence $\mathcal{T} \times \mathcal{M}$ has a finite attractor, $\mathcal{T} \times \mathcal{M}$ is $\text{Dec}(\mu \times \delta_{q_0}, B)$ for any measurable set B , so that we can apply the approximation scheme from Section 7.1 to obtain an approximation of the desired value.

The effectiveness of the approach relies on the effectiveness of the scheme for reachability, but also on the computability of an attractor for \mathcal{T} , and of the set of good BSCCs of the graph of the attractor.

7.4 ω -regular properties of general STSs via abstraction and finite attractor

We assume the same framework as in Section 6.4, that is $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1, \text{AP}, \mathcal{L}_1)$ and $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2, \text{AP}, \mathcal{L}_2)$ are two LSTSs such that:

- \mathcal{T}_2 is a sound α -abstraction of \mathcal{T}_1
- \mathcal{T}_2 is a DMC with a finite attractor B_2 .

We consider again a DMA $\mathcal{M} = (Q, q_0, E, \mathcal{F})$, as well as the products $\mathcal{T}_1 \times \mathcal{M}$ and $\mathcal{T}_2 \times \mathcal{M}$. Writing $\mathcal{B} = \{\alpha_{\mathcal{M}}^{-1}(B) \mid B \in \Sigma_2'\}$, we assume that $\mathcal{T}_1 \times \mathcal{M}$ is $\text{Dec}(\mathcal{B})$. Remember that this implies, from Lemma 47, that $\mathcal{T}_2 \times \mathcal{M}$ is a sound $\alpha_{\mathcal{M}}$ -abstraction of $\mathcal{T}_1 \times \mathcal{M}$.

Fix an initial distribution μ for \mathcal{T}_1 . Thanks to Theorem 51:

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Inf} \in \mathcal{F}) = \sum_{C \in \text{Good}_{\mathcal{T}_2 \times \mathcal{M}}^{B_2}(\mathcal{F})} \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \alpha^{-1}(C)) .$$

Thus, as previously, the computation of the probability that a given model satisfies a given external specification is reduced to the computation of a reachability probability. Since we assumed $\mathcal{T}_1 \times \mathcal{M}$ to be $\text{Dec}(\mathcal{B})$, we can use the approximation scheme from Section 7.1 to approximate the searched value.

Effectiveness of the approach requires effective numerical computations for the distributions, as well as good constructivity properties for various sets, like the BSCCs of the graph of the attractor, and avoid-sets of these, etc.

7.5 Time-bounded verification of stochastic real-time systems

The initial motivation to consider general STSs stems from real-time stochastic systems, that is, systems with both timing constraints and probabilistic choices. While everything which precedes holds for any kind of STSs, we highlight now some specific features of real-time stochastic systems.

► **Definition 59.** A *real-time stochastic transition system* (RT-STTS) is an STS $\mathcal{T} = (\widehat{S}, \widehat{\Sigma}, \kappa)$ such that (i) there is a measurable space (S, Σ) with $\widehat{S} = S \times \mathbb{R}_{\geq 0}$, and $\widehat{\Sigma}$ is the product σ -algebra of Σ and the Borel sets of $\mathbb{R}_{\geq 0}$; and (ii) for every $(s, t) \in S \times \mathbb{R}_{\geq 0}$, $\kappa((s, t), \{(s', t') \in S \times \mathbb{R}_{\geq 0} \mid t' > t\}) = 1$.

The first condition makes explicit the time component of the system (given by $\mathbb{R}_{\geq 0}$; S then contains the spatial information), while the second condition imposes the time to increase almost-surely. By explicitly integrating absolute time into DMCs (where it is increased by one at each new event) or CTMCs (where it is increased by the time elapsed in each state – hence it represents absolute time since the start of the system), they can be interpreted as RT-STTSs. All other examples that we will consider in Section 8 can also be seen as RT-STTSs, after explicit integration of absolute time in the state-space.

Let $\mathcal{T} = (\widehat{S}, \widehat{\Sigma}, \kappa)$ be an RT-STTS. A desirable property of a real-time system is that it should be (almost-surely) *non-Zeno*: a path $\rho = (s_0, t_0)(s_1, t_1) \dots \in \text{Paths}(\mathcal{T})$ is non-Zeno whenever $\lim_{n \rightarrow +\infty} t_n = +\infty$. Under such an hypothesis, we first identify natural attractors of an RT-STTS.

► **Lemma 60.** *Assume that \mathcal{T} is almost-surely non-Zeno. Then, for every $\Delta \in \mathbb{Q}_{\geq 0}$, the set $A_\Delta = \{(s, t) \in S \times \mathbb{R}_{\geq 0} \mid t > \Delta\}$ is an attractor of \mathcal{T} .*

As a consequence, as soon as it is almost-surely non-Zeno, an RT-STTS is (strongly) decisive w.r.t. every bounded measurable subset and every initial distribution. Fix μ an initial distribution assigning 0 to the initial timestamp t_0 . Assume one wants to compute the probability of property $B' \mathbf{U}_I B$ from μ , where I is some bounded interval of $\mathbb{R}_{\geq 0}$ with rational bounds, and $B, B' \in \Sigma$; this is the probability of the following set of paths:

$$\{\rho = (s_0, t_0)(s_1, t_1) \dots \in \text{Paths}(\mathcal{T}) \mid \exists n \in \mathbb{N} \text{ s.t. } t_n \in I, s_n \in B \text{ and } \forall j < n, s_j \in B'\} .$$

Then, for any $\Delta \in \mathbb{Q}_{\geq 0}$ with $\Delta > \sup I$, A_Δ is included in the avoid-set of $B \times I$, and \mathcal{T} is therefore decisive w.r.t. $B \times I$. In particular, the approximation scheme of Subsection 7.1 applies.

Hence, assuming the RT-STTS \mathcal{T} is almost-surely non-Zeno (which needs to be proven “by hand”, or structurally obvious), and under some effectiveness assumption on \mathcal{T} , the quantitative analysis of time-bounded until or reachability properties is doable.

8 Applications

The general approach to the qualitative and quantitative analysis of stochastic systems over a possibly continuous state-space can be instantiated in multiple frameworks. To demonstrate its versatility, we present three types of models to which it applies: stochastic timed automata, generalized semi-Markov processes and stochastic time Petri nets. These models are taken from the literature without further motivations. This section is technical (since the models themselves are complex), and can be skipped by the reader not necessarily interested in these models. However, it is interesting to observe that several results from the literature can be recovered (and extended) via our generic approach.

8.1 Stochastic timed automata

Stochastic timed automata (STA) [17] are stochastic real-time processes derived from timed automata [6] by randomizing both the delays and the edge choices. The semantics of a STA

is naturally given via a STS as defined in this paper, although this had not been formulated this way originally.

Several decidability results have been proven for subclasses of STA, requiring the development of ad-hoc methods [7, 8, 16, 18], and in [17], we proposed the first unifying method capturing all known decidability results for the qualitative model-checking problem: the so-called *thick graph* is a finite graph based on the standard region automaton construction for timed automata [6], which allows one to infer good transfer properties from this finite graph to the original STA when some *fairness* property is satisfied. The current work improves our understanding of [15] and allows us both to unify all decidability and approximability results that were known, and to get new approximability results for the quantitative model-checking problem (of ω -regular properties).

8.1.1 Definition

To define the model properly, we first give some notations. Let X be a finite set of clocks. We write $G(X)$ for the set of guards defined as finite conjunctions of constraints of the form $x \bowtie c$, where $x \in X$, $\bowtie \in \{<, \leq, =, \geq, >\}$ and $c \in \mathbb{N}$. Guards are interpreted over clock valuations $\nu: X \rightarrow \mathbb{R}_{\geq 0}$ in a natural way – we then write $\nu \models g$. Also, for ν a valuation we define $[Y \leftarrow 0](\nu)$ the valuation assigning 0 to every $x \in Y$ and $\nu(x)$ to each other clock, and if $d \in \mathbb{R}_{\geq 0}$, we write $\nu + d$ for the valuation assigning $\nu(x) + d$ to every clock $x \in X$.

► **Definition 61.** A *stochastic timed automaton* (STA) is a tuple

$$\mathcal{A} = (L, \ell_0, X, E, (\mu_\gamma)_{\gamma \in L \times \mathbb{R}_{\geq 0}^X}, (w_e)_{e \in E})$$

where:

- L is a finite set of states (or locations);
- $\ell_0 \in L$ is the initial state;
- X is a finite set of clocks;
- $E \subseteq L \times G(X) \times 2^X \times L$ is a finite set of edges; and
- for every configuration $\gamma \in L \times \mathbb{R}_{\geq 0}^X$, μ_γ is a(n a priori) continuous distribution over possible delays from $\gamma = (\ell, \nu)$, that is, the support of distribution μ_γ is precisely $I(\gamma) \stackrel{\text{def}}{=} \{d \in \mathbb{R}_{\geq 0} \mid \exists e = (\ell, g, Y, \ell') \in E \text{ s.t. } \nu + d \models g\}$;
- and for every $e \in E$, $w_e \in \mathbb{N}_{\geq 0}$ is a positive weight.

Originally, the semantics of an STA $\mathcal{A} = (L, \ell_0, X, E, (\mu_\gamma)_{\gamma \in L \times \mathbb{R}_{\geq 0}^X}, (w_e)_{e \in E})$ was defined as a probability measure on the set of possible runs of the underlying timed automaton (L, ℓ_0, X, E) : a run in such a timed automaton is an alternating sequence of delay transitions and of discrete transitions. A delay transition is of the form $(\ell, \nu) \xrightarrow{d} (\ell, \nu + d)$, where $\gamma \stackrel{\text{def}}{=} (\ell, \nu) \in L \times \mathbb{R}_{\geq 0}^X$ is a configuration and $d \in \mathbb{R}_{\geq 0}$,⁹ and a discrete transition is of the form $(\ell, \nu) \xrightarrow{e} (\ell', \nu')$ where $e = (\ell, g, Y, \ell') \in E$ is such that $\nu \models g$, and $[Y \leftarrow 0](\nu) = \nu'$. When $\nu \models g$, we say that e is enabled at γ .

The probability measure was obtained by sampling delay transitions from a configuration γ following distribution μ_γ , and by sampling discrete transitions using the weights: the probability to take edge e from configuration γ is given by $p_\gamma(e) \stackrel{\text{def}}{=} \frac{w_e}{\sum_{\{w_{e'} \mid e' \text{ enabled at } \gamma\}} w_{e'}}$ if e is enabled at γ , and by $p_\gamma(e) \stackrel{\text{def}}{=} 0$ otherwise.

⁹ Later we will also write $\gamma + d$ for the configuration $(\ell, \nu + d)$.

To have properly-defined measures we need some sanity assumptions on distributions $(\mu_\gamma)_{\gamma \in L \times \mathbb{R}_{\geq 0}^X}$: If we write λ for the Lebesgue measure over \mathbb{R}_+ , it must be the case that for each $\gamma \in L \times \mathbb{R}_+^X$, if $\lambda(I(\gamma)) > 0$ then μ_γ is equivalent to the restriction of λ on $I(\gamma)$; Otherwise, it is the uniform distribution over the points of $I(\gamma)$.

We now give the semantics of an STA $\mathcal{A} = (L, \ell_0, X, E, (\mu_\gamma)_{\gamma \in L \times \mathbb{R}_{\geq 0}^X}, (w_e)_{e \in E})$ as an STS $\mathcal{T}_\mathcal{A} = (S_\mathcal{A}, \Sigma_\mathcal{A}, \kappa_\mathcal{A})$ as follows. The set $S_\mathcal{A}$ is the set of configurations $L \times \mathbb{R}_{\geq 0}^X$, $\Sigma_\mathcal{A}$ is the σ -algebra product between 2^L and the Borel σ -algebra on $\mathbb{R}_{\geq 0}^{|X|}$, and the kernel $\kappa_\mathcal{A}$ is defined by:


$$\kappa_\mathcal{A}(\gamma, B) = \sum_{e=(\ell, g, Y, \ell') \in E} \int_{d \in \mathbb{R}_{\geq 0}} \mathbb{1}_B(\ell', [Y \leftarrow 0](\nu + d)) \cdot p_{\gamma+d}(e) \, d\mu_\gamma(d)$$

where $\mathbb{1}_B$ is the characteristic function of B . It gives the probability to hit set $B \subseteq S_\mathcal{A}$ from configuration γ in one step (composed of a delay transition followed by a discrete transition).

The probability measure on paths derived from $\mathcal{T}_\mathcal{A}$ in Section 2.2 coincides with the original definition of [17].

We fix for the rest of this section an STA $\mathcal{A} = (L, \ell_0, X, E, (\mu_\gamma)_{\gamma \in L \times \mathbb{R}_{\geq 0}^X}, (w_e)_{e \in E})$, and $\mathcal{T}_\mathcal{A} = (S_\mathcal{A}, \Sigma_\mathcal{A}, \kappa_\mathcal{A})$ its corresponding STS.

► **Example 62** (A stochastic timed automaton with an “unfair” convergence behaviour). Consider the STA \mathcal{A} of Figure 5, with: $L = \{\ell_0, \dots, \ell_4\}$, $X = \{x, y\}$ and the set of edges E as described on the figure. We assume that each edge has a weight of 1 and that each location is either equipped with a uniform distribution over possible delays (in ℓ_0, ℓ_2 and ℓ_4) or a Dirac distribution over the unique possible delay (in ℓ_1 and ℓ_3).¹⁰ As said previously, it can be considered as an STS $\mathcal{T}_\mathcal{A}$ where the set of states is given by $L \times \mathbb{R}_+^2$ and the Markov kernel is computed according to the distributions over the edges and the delays.



■ **Figure 5** A two-clock STA \mathcal{A} with an unfair convergence behaviour

We would like to stress that \mathcal{A} suffers from a time-convergence phenomenon. This convergence phenomenon (that we make precise in the following lines), is due to the timing constraints, and is in fact inherent to the underlying timed automaton (without the stochastic aspects). We will see later (Example 64) how it impacts the stochastic behaviour of the STA \mathcal{A} . Let us now discuss the (non stochastic) time-convergence phenomenon. In order to do so, for the rest of the paragraph, we see \mathcal{A} as a timed automaton and forget about the stochastic aspects. Let us imagine that we enter location ℓ_0 with the value of clock x (resp. y) being 0 (resp. $0 < \nu < 1$), we are thus in configuration $(\ell_0, (0, \nu))$. Let us consider the case where we take the right loop. We thus first enter location ℓ_1 and then ℓ_2 , that we reached with configuration $(\ell_2, 1 - \nu, 0)$, after a total delay of $1 - \nu$ time units spent since

¹⁰ When we reach ℓ_1 , the value v_y of clock y is smaller than 1; since the constraint on the edge between ℓ_1 and ℓ_2 is constrained by $y = 1$, there is a single possible delay for taking this edge: wait d such that $v_y + d = 1$.

the last arrival in ℓ_0 . In order to return to ℓ_0 , we have to wait a delay ν' such that $\nu' < 1$ (because of the guard $y < 1$) and $\nu < \nu'$ (because of the guard $x > 1$). We thus return to ℓ_0 with configuration $(\ell_0, (0, \nu'))$, where $\nu < \nu'$. One can check that a similar situation occurs when taking the left loop. Thus when considering an infinite path of \mathcal{A} , if we denote by $(\ell_0, (0, \nu_n))$ its configuration at the n -th passage in ℓ_0 , we can infer that the sequence $(\nu_n)_{n \in \mathbb{N}}$ is increasing (and bounded by 1), and thus converging.

8.1.2 The thick graph abstraction

The thick graph of [17] is an abstraction in our context. To see this, we recall the concept of regions, that have been designed for standard timed automata [6]. We write $M_{\mathcal{A}}$ for the maximal integer appearing in a guard of \mathcal{A} . Let $\nu, \nu' \in \mathbb{R}_{\geq 0}^X$ be two valuations over X . We say that ν and ν' are *region-equivalent* for \mathcal{A} whenever the following conditions hold:

1. for every $x \in X$, either both $\nu(x)$ and $\nu'(x)$ are strictly larger than $M_{\mathcal{A}}$, or the integral parts of $\nu(x)$ and $\nu'(x)$ coincide;
2. for every $x, y \in X$ such that $\nu(x), \nu(y) \leq M_{\mathcal{A}}$, writing $\{\cdot\}$ for the fractional part, $\{\nu(x)\} \leq \{\nu(y)\}$ if and only if $\{\nu'(x)\} \leq \{\nu'(y)\}$.

This region-equivalence has finite-index, and partitions the set of valuations $\mathbb{R}_{\geq 0}^X$ into classes which are called *regions*, and we write $R_{\mathcal{A}}$ for the set of regions. If $\nu \in \mathbb{R}_{\geq 0}^X$, we write $[\nu]_{\mathcal{A}}$ for the region to which ν belongs.

We define the abstraction $\alpha : L \times \mathbb{R}_{\geq 0}^X \rightarrow L \times R_{\mathcal{A}}$ as the projection which associates (ℓ, ν) onto $(\ell, [\nu]_{\mathcal{A}})$. We then define the finite Markov chain $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ as follows:

- its set of states is $L \times R_{\mathcal{A}}$;
- there is an edge from (ℓ, r) to (ℓ', r') whenever there exists some $\nu \in r$ such that $\kappa_{\mathcal{A}}((\ell, \nu), \{\ell'\} \times r') > 0$;¹¹
- from each state $(\ell, r) \in L \times R_{\mathcal{A}}$, we associate the uniform distribution over $\{(\ell', r') \in L \times R_{\mathcal{A}} \mid \text{there is an edge from } (\ell, r) \text{ to } (\ell', r')\}$.

By construction, we get:

► **Lemma 63.** $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ is a finite α -abstraction of $\mathcal{T}_{\mathcal{A}}$.


Let us notice that finiteness of the abstraction implies completeness (Lemma 38).

As witnessed in [17, Appendix D.2], this abstraction may not give much information in general about the probability of linear-time properties in the original STA (see Example 64). However we will see that, in several cases, it helps to obtain decidability and approximability results (among which some are new).


► **Example 64** (A stochastic timed automaton with an “unfair” convergence behaviour (continued)). We know that the thick graph viewed as a finite Markov chain, is an α -abstraction of the original STA, but it can be shown that in general it is not sound. Let us denote $\mathcal{T}_{\mathcal{A}}$ the STS naturally associated with the STA of Example 62 (see Fig. 5). One can show that $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$, the thick graph associated with $\mathcal{T}_{\mathcal{A}}$, is the one provided on Fig. 6, starting from a Dirac distribution $\delta_{(\ell_0, (0, \nu))}$ with $0 < \nu < 1$. The regions are the following ones: $r_0 = \{(x, y) \mid x = 0 \wedge 0 < y < 1\}$, $r_1 = \{(x, y) \mid 0 < x < y < 1\}$, $r_2 = \{(x, y) \mid y = 0 \wedge 0 < x < 1\}$, $r_3 = \{(x, y) \mid 1 < x < y < 2\}$, $r_4 = \{(x, y) \mid y = 0 \wedge 1 < x < 2\}$. We clearly have that $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ is an α -abstraction of $\mathcal{T}_{\mathcal{A}}$.

¹¹Note that it is a local condition which is easy to check.

However, it can be shown that $\mathcal{T}_A^{\text{tg}}$ is not a sound abstraction of \mathcal{T}_A . The time-convergence phenomenon of \mathcal{A} (described in Example 62) implies that each time we return to location ℓ_0 , the probability to take the right loop decreases while the probability to take the left loop increases. More precisely, it has been shown that in the original STA \mathcal{A} , the probability to reach ℓ_2 from ℓ_0 is strictly lower than 1. This has been done formally via a tedious and technical calculation of Taylor series in [17, Section 6.2.2]. This implies that $\mathcal{T}_A^{\text{tg}}$ is not a sound abstraction of \mathcal{T}_A (since the probability to reach (ℓ_2, r_2) from (ℓ_0, r_0) is 1 in $\mathcal{T}_A^{\text{tg}}$). In fact, a sound abstraction of \mathcal{T}_A would rather behave as the **non-homogeneous**¹² finite Markov chain of Fig. 7, where n represents the n -th passage in (ℓ_0, r_0) . This shows in particular that general STA are not fair (and thus not decisive). This is why we focus on two subclasses of STA in the rest of this section.



■ **Figure 6** $\mathcal{T}_A^{\text{tg}}$, the thick graph (viewed as a finite Markov chain) associated with the two-clock STA with an “unfair” convergence behaviour (Fig. 5).



■ **Figure 7** A non-homogeneous finite Markov chain which is in some sense equivalent to the STA with an “unfair” convergence behaviour

8.1.3 Reactive STA

Following [18], the STA \mathcal{A} is *reactive* whenever for every configuration $\gamma = (\ell, \nu) \in S_{\mathcal{A}}$, $I(\gamma) = \mathbb{R}_{\geq 0}$, and for every ℓ , there exists a distribution μ_{ℓ} with support $\mathbb{R}_{\geq 0}$ such that for every $\nu \in \mathbb{R}_{\geq 0}^X$, $\mu_{(\ell, \nu)} = \mu_{\ell}$. Note that we do not make any Markovian hypothesis on time elapsing, and μ_{ℓ} does not need to be exponential.

We take the notations used in the previous subsection for defining the thick-graph abstraction. A region r is *memoryless* whenever for every clock $x \in X$, either $\nu(x) = 0$ for every $\nu \in r$, or $\nu(x) > M_{\mathcal{A}}$ for every $\nu \in r$. We write $R_{\mathcal{A}}^{\text{mem}}$ for the set of memoryless regions.

From [17, Lemma 13], which states that the set of memoryless regions is visited infinitely often almost-surely from every configuration $\gamma \in S_{\mathcal{A}}$,¹³ we get:

¹² Although we only consider homogeneous systems, the non-homogeneous ones can fit our general model of STS by unfolding it. For instance Example 22 can be seen as the unfolding of a finite non-homogeneous Markov chain with two states.

¹³ To give all arguments, it is easy to see that, in one step, one can ensure reaching a memoryless region by delaying at least $M_{\mathcal{A}} + 1$ time units; since there is one single distribution which is applied at every configuration of a given location, the probability to do so is uniformly bounded from below from every configuration.

► **Proposition 65.** The set $\alpha^{-1}(L \times R_{\mathcal{A}}^{\text{mem}})$ is an attractor for $\mathcal{T}_{\mathcal{A}}$.

Using Propositions 36 and 40, we also get that:

► **Proposition 66.** $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ is a sound α -abstraction of $\mathcal{T}_{\mathcal{A}}$.

Proof. It can easily be shown that $L \times R_{\mathcal{A}}^{\text{mem}}$ is a finite attractor of $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$. Thanks to Proposition 65, $\alpha^{-1}(L \times R_{\mathcal{A}}^{\text{mem}})$ is an attractor for $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$. It remains to show the last condition of the hypotheses of Proposition 36. We therefore need to prove that for each $(\ell_m, r_m) \in L \times R_{\mathcal{A}}^{\text{mem}}$, there are $p > 0$ and $k \in \mathbb{N}$ such that for each region $(\ell, r) \in L \times R_{\mathcal{A}}$:

- for each $\mu \in \text{Dist}(\alpha^{-1}(\ell_m, r_m))$, $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F}_{\leq k} \alpha^{-1}(\ell, r)) \geq p$, or
- for each $\mu \in \text{Dist}(\alpha^{-1}(\ell_m, r_m))$, $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F} \alpha^{-1}(\ell, r)) = 0$.

This is a consequence of [17, Lemma F.4] which says that from a memoryless region, the future (and its probability) is independent of the precise current configuration. This in particular implies that for two configurations $\gamma, \gamma' \in \alpha^{-1}(\ell_m, r_m)$, for every α -closed set B , for every integer k , $\text{Prob}_{\delta_{\gamma}}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F}_{=k} B) = \text{Prob}_{\delta_{\gamma'}}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F}_{=k} B)$. By extension, for every $\mu \in \text{Dist}(\alpha^{-1}(\ell_m, r_m))$, $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F}_{=k} B) = \text{Prob}_{\delta_{\gamma}}^{\mathcal{T}_{\mathcal{A}}}(\mathbf{F}_{=k} B)$. This implies the expected bounds, by taking $B = \alpha^{-1}(\ell, r)$. ◀

Similarly to labelled STS, we consider labelled STA, where each location is labelled by atomic propositions. As consequences of Sections 6 and 7, we get the following decidability and approximability results for reactive STA:

► **Corollary 67.** Let \mathcal{A} be a reactive labelled STA, and \mathcal{M} a DMA. Then:

1. we can decide whether \mathcal{A} satisfies almost-surely \mathcal{M} ;
2. for every initial distribution μ which is numerically amenable w.r.t. \mathcal{A}^{14} , we can compute arbitrary approximations of $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{A}}}(\mathcal{M})$.

Proof. This is an application of Theorem 51, Corollary 55 and of Sections 7.1 and 7.4. It should be noted that all the hypotheses are met:

- $\mathcal{T}_{\mathcal{A}}^{\text{tg}} \times \mathcal{M}$ has a finite attractor: since $\mathcal{T}_{\mathcal{A}}^{\text{tg}}$ is a finite MC then so is $\mathcal{T}_{\mathcal{A}}^{\text{tg}} \times \mathcal{M}$ and we get a trivial finite attractor;
- $\mathcal{T}_{\mathcal{A}} \times \mathcal{M}$ is decisive w.r.t. any $\alpha_{\mathcal{M}}$ -closed sets.

This second point is a little more tricky. First one should realise that since $\mathcal{T}_{\mathcal{A}}$ is reactive, then $\mathcal{T}_{\mathcal{A}} \times \mathcal{M}$ is also reactive, since the condition to be reactive concerns only the distributions over the delays on each location of the STA and those distributions are not modified from the product with \mathcal{M} . It should be noted that $\mathcal{T}_{\mathcal{A}}^{\text{tg}} \times \mathcal{M}$ corresponds to the thick region graph abstraction of $\mathcal{T}_{\mathcal{A}} \times \mathcal{M}$ since \mathcal{M} does not influence the behaviour of $\mathcal{T}_{\mathcal{A}}$. Then from Proposition 66, we know that $\mathcal{T}_{\mathcal{A}}^{\text{tg}} \times \mathcal{M}$ is a sound $\alpha_{\mathcal{M}}$ -abstraction of $\mathcal{T}_{\mathcal{A}} \times \mathcal{M}$. Since $\mathcal{T}_{\mathcal{A}}^{\text{tg}} \times \mathcal{M}$ is a finite MC, we get that it is decisive w.r.t. any set of states. We can thus conclude from Proposition 33. ◀

► **Remark 68.** We believe that the proposed approach through abstractions and finite attractors simplifies drastically the proof of decidability of almost-sure model-checking, and in particular avoids the ad-hoc but long and technical proof of [17, Lemma 7.14]. Furthermore, we obtain interesting approximability results, some of them being consequences of [15], but

¹⁴We say that a distribution μ is numerically amenable w.r.t. \mathcal{A} if, given $k \in \mathbb{N}$, given $\varepsilon > 0$ and given a sequence of locations and regions $(\ell_0, r_0), (\ell_1, r_1), \dots, (\ell_k, r_k)$, one can compute a numerical approximation $\text{Prob}_{\mu}^{\mathcal{A}}(\text{Cyl}((\ell_0, r_0), (\ell_1, r_1), \dots, (\ell_k, r_k)))$ up to ε .

the general case of ω -regular properties (in particular LTL properties) being new to this paper.

► **Remark 69.** Corollary 67 can be extended to properties expressed as deterministic and complete Muller *timed* automata (DCMTA), which are standard deterministic and complete¹⁵ timed automata [6] with a Muller accepting condition. Indeed, the product of a reactive STA with such a DCMTA is reactive. Hence, the whole theory that we have developed applies: the STS of the product admits a sound finite abstraction. DCMTA allows one to express rich properties with timing constraints and one can evaluate their likelihood in the STA. It should be noticed that the convergence proof of the approximation scheme of [21] can be recovered as a byproduct, since CTMCs are particular cases of reactive STA. Another way to obtain this result would have been to apply the approach of Subsection 7.5 on time-bounded verification.

8.1.4 Single-clock STA

We will apply a similar reasoning to single-clock STA. We therefore assume that \mathcal{A} is now a single-clock STA. As in [17, Section 7.1], we assume the following conditions:

- (i) for all $\ell \in L$, for all $[a, b] \subseteq \mathbb{R}_+$, the function $\nu \mapsto \mu_{(\ell, \nu)}([a, b])$ is continuous;
- (ii) if $\gamma' = \gamma + t$ for some $t \geq 0$, and if $0 \notin I(\gamma + t', e)$ for each $0 \leq t' \leq t$, then $\mu_\gamma(I(\gamma, e)) \leq \mu_{\gamma'}(I(\gamma', e))$;
- (iii) there is $0 < \lambda_0 < 1$ such that for every state γ with $I(\gamma)$ unbounded, $\mu_\gamma([0, \frac{1}{2}]) \leq \lambda_0$, where for each $\gamma = (\ell, \nu) \in L \times \mathbb{R}_+^X$ and for each $e = (\ell, g, Y, \ell') \in E$, $I(\gamma, e) = \{d \in \mathbb{R}_+ \mid \nu + d \models g\}$. These requirements are technical, but they are rather natural and easily satisfiable. For instance, a timed automaton equipped with uniform (resp. exponential) distributions on bounded (resp. unbounded) intervals satisfy these conditions. If we assume exponential distributions on unbounded intervals, the very last requirement corresponds to the bounded transition rate condition in [24], required to have reasonable and realistic behaviours.

In [17, Section 7.1], there is no clear attractor property. From the details of the proofs we can nevertheless define $A_{\mathcal{A}}^{\max} = \{(\ell, r_0) \mid \ell \in L\} \cup \{(\ell, r) \in L \times R_{\mathcal{A}} \mid \forall (\ell', r') \in L \times R_{\mathcal{A}}, (\ell, r) \rightarrow^* (\ell', r') \text{ in } \mathcal{T}_{\mathcal{A}}^{\text{tg}} \text{ implies } r' = r\}$ where r_0 is the region composed of the single null valuation.

► **Proposition 70.** The set $\alpha^{-1}(A_{\mathcal{A}}^{\max})$ is an attractor for $\mathcal{T}_{\mathcal{A}}$.

Proof. Let $C = \{0\} \cup \{c \mid c \text{ constant appearing in a guard of } \mathcal{A}\} \stackrel{\text{def}}{=} \{c_0 < c_1 < \dots < c_h\}$. The set of regions for \mathcal{A} can be chosen as $\{\{c_i\} \mid 0 \leq i \leq h\} \cup \{c_{i-1}; c_i \mid 1 \leq i \leq h\}$ (see [36]).

Following the proof of [17, Theorem 7.2], the set of infinite paths in \mathcal{A} can be divided into (a) the set of paths that take resetting edges infinitely often, and (b) the set of paths that take resetting edges only finitely often.

We assume that the probability that (a) happens is positive, and we reason now in the σ -algebra which is conditioned by (a). Then under condition (a), $\alpha^{-1}(\{(\ell, r_0) \mid \ell \in L\})$ is reached almost-surely.

We assume that the probability that (b) happens is positive, and we reason now in the σ -algebra which is conditioned by (b). Under condition (b), almost-surely the value of the

¹⁵ In this context, complete means that from every configuration, for every subset of AP, and every $t \in \mathbb{R}_{\geq 0}$, there is an edge labelled by that subset which is enabled after t time units. So this is complete w.r.t time and actions.

clock is non-decreasing along the path, and almost-surely a final region r is reached (that is, ultimately the value of the clock along the path belongs to r forever). We fix such a region r , and we condition again with regard to that “final region” r . We write E_r for the event (b) intersected with “the path ends up in r ”. Let r' be a strict successor region of r , with dimension at least as big as that of r (if r is an open interval, then r' has to be an open interval). There exists $\alpha > 0$ such that for every $\nu \in r$, for every $\ell \in L$, for every $e = (\ell, g, Y, \ell')$ with $r' \subseteq g$, $\text{Prob}_{\delta(\ell, \nu)}^{\mathcal{T}_A}((\ell, \nu) \xrightarrow{e}) \geq \alpha$. Hence, using standard technics, we show that with probability 1, if infinitely often such edges are enabled, infinitely often they will be taken; this contradicts hypothesis E_r . Hence, under condition E_r , with probability 1, one cannot visit infinitely often configurations enabling edges guarded by some strict time-successor r' of r . Once this is assumed, we can then show that almost-surely, only finitely many resetting edges can be enabled. This means that, under condition E_r , almost-surely, ultimately only states of $\alpha^{-1}(\{(\ell, r) \in L \times R_A \mid \forall(\ell', r') \in L \times R_A, (\ell, r) \rightarrow^* (\ell', r') \text{ in } \mathcal{T}_A^{\text{tg}} \text{ implies } r' = r\})$ are visited. Hence, that set is an attractor, under condition (b).

Using some Bayes formula w.r.t. conditions (a) and (b), we conclude that $\alpha^{-1}(A_A^{\text{max}})$ is an attractor; this ends the proof. ◀

As before, we get:

► **Proposition 71.** $\mathcal{T}_A^{\text{tg}}$ is a sound α -abstraction of \mathcal{T}_A .

Proof. We easily get that A_A^{max} is a finite attractor for $\mathcal{T}_A^{\text{tg}}$, whereas $\alpha^{-1}(A_A^{\text{max}})$ is an attractor for \mathcal{T}_A (Proposition 70).

As for reactive STA, it remains to show the last property appearing in the hypotheses of Proposition 36. The required bounds obviously exist for the region r_0 (since only a single valuation belongs to r_0). Furthermore, as argued in the proof of Proposition 70, when condition (b) is assumed, ultimately, the paths almost surely end up in $\alpha^{-1}(\{(\ell, r) \in L \times R_A \mid \forall(\ell', r') \in L \times R_A, (\ell, r) \rightarrow^* (\ell', r') \text{ in } \mathcal{T}_A^{\text{tg}} \text{ implies } r' = r\})$, hence, ultimately, the STA behaves like a **finite** Markov chain. The required bounds can be inferred.

This allows to conclude that $\mathcal{T}_A^{\text{tg}}$ is a sound α -abstraction of \mathcal{T}_A (using Propositions 36 and 40). ◀

As a consequence, we get the following decidability and approximability results for one-clock STA:

► **Corollary 72.** *Let \mathcal{A} be a one-clock labelled STA, and \mathcal{M} a DMA. Then:*

1. *we can decide whether \mathcal{A} satisfies almost-surely \mathcal{M} ;*
2. *for every initial distribution μ which is numerically amenable w.r.t. \mathcal{A} , we can compute arbitrary approximations of $\text{Prob}_{\mu}^{\mathcal{T}_A}(\mathcal{M})$.*

Proof. Similarly to the proof of Corollary 67, this is an application of Theorem 51, Corollary 55 and of Sections 7.1 and 7.4. The facts that:

- $\mathcal{T}_A^{\text{tg}} \times \mathcal{M}$ has a finite attractor, and
- $\mathcal{T}_A \times \mathcal{M}$ is decisive w.r.t. any α -closed sets.

can be deduced by similar arguments. We only observe that if \mathcal{T}_A is a single-clock STA, then so is $\mathcal{T}_A \times \mathcal{M}$ and that hypotheses (i), (ii) and (iii) are preserved through the product with \mathcal{M} as those only concern distributions over the STA which are not altered from the product with \mathcal{M} . ◀

► **Remark 73.** The proof of the existence of an attractor is very similar to the one we used for proving the fairness property in [17, Section 7.1]. However, for free, we get all the approximation results (as previously only few results could be inferred from [15])! It is worth noting that these results encompass the results of [16], where a strong assumption on cycles of the STA was made (but a closed-form for the probability could be computed). We remark here that the graph used in [16] is actually the graph of the attractor, as done in Section 6.3.

► **Remark 74 (Time-bounded analysis).** Let us finish by discussing how and when the approach of Subsection 7.5 for timed properties can be applied to STA. Similarly to CTMCs [9] that they extend, reactive STA are almost-surely non-Zeno. Hence one can apply the approximation scheme of Subsection 7.1 to (time-)bounded until formulas or time-bounded reachability properties.

One can decide whether a single-clock STA is almost-surely non-Zeno [8]. In the positive case, the approximation scheme of Subsection 7.1 can therefore be applied as well.

8.2 Generalized semi-Markov processes

A generalized semi-Markov process [19, 29] is a stochastic process with a finite control, built on a set of events. Each event is equipped with a random variable representing its duration: an event can either be a variable-delay event, whose duration is given by a probability distribution defined by a density function, or be a fixed-delay event, modelled by a Dirac distribution. A transition is characterized by a set of events which expire, and schedules a set of new events. This model is known to generalize CTMCs. In this section, we show how to exploit our techniques to recover and generalize results from the literature on quantitative verification of generalized semi-Markov processes.

8.2.1 Definition

► **Definition 75.** A *generalized semi-Markov process* (GSMP) is a tuple $\mathcal{G} = (Q, \mathcal{E}, \ell, u, f, \mathbf{E}, \text{Succ})$ where

- Q is a finite set of states;
- $\mathcal{E} = \{e_1, \dots, e_p\}$ is a finite set of events;
- $\ell : \mathcal{E} \rightarrow \mathbb{N}_{\geq 0}$ and $u : \mathcal{E} \rightarrow \mathbb{N}_{>0} \cup \{\infty\}$ are bounds such that for every $e \in \mathcal{E}$, $\ell(e) \leq u(e)$;
- $f : \mathcal{E} \rightarrow \text{Dist}([\ell(e); u(e)])$ assigns distributions to every event $e \in \mathcal{E}$;
- $\mathbf{E} : Q \rightarrow 2^{\mathcal{E}}$ assigns to each state q a set of events enabled (or active) in q ;
- $\text{Succ} : Q \times 2^{\mathcal{E}} \rightarrow \text{Dist}(Q)$ is the successor function defined for (q, E) whenever $E \subseteq \mathbf{E}(q)$;

Each event $e \in \mathcal{E}$ has an upper (resp. lower) bound $u_e \stackrel{\text{def}}{=} u(e)$ (resp. $\ell_e \stackrel{\text{def}}{=} \ell(e)$) on its delay: the duration of event e is randomly chosen in the interval $[\ell_e, u_e]$ according to density $f_e \stackrel{\text{def}}{=} f(e)$. In contrast to *fixed-delay* events, e is called a *variable-delay* event, if $\ell_e < u_e$. Events can alternatively be seen as random variables: with a variable-delay event is associated a density function and with a fixed-delay event is associated the corresponding Dirac distribution.

The semantics of a GSMP \mathcal{G} is given as an STS $\mathcal{T}_{\mathcal{G}} = (S_{\mathcal{G}}, \Sigma_{\mathcal{G}}, \kappa_{\mathcal{G}})$. There are two points-of-view to define the semantics of \mathcal{G} , one is through a residual-time semantics using races between events [19] (clocks behave like in timed automata), and the other is to sample the delay of an event once, when it is scheduled [23] (clocks are “countdown”). Though the results of [19] are stated using the first convention, we prefer the second option, since it is

easier to understand the semantics. Note that the duality between the two allows obviously to interpret the results of [19] in our setting.

Let $q \in Q$ be a state; a valuation $\nu \in (\mathbb{R}_{\geq 0}^{\perp})^{\mathcal{E}}$, where $\mathbb{R}_{\geq 0}^{\perp} = \mathbb{R}_{\geq 0} \cup \{\perp\}$, is compatible with q whenever $\nu(e) = \perp$ if $e \notin \mathbf{E}(q)$, and $\nu(e) \in \mathbb{R}_{\geq 0}$ otherwise; in the latter case, $\nu(e)$ is the remaining time for e before expiring. Configurations of \mathcal{G} are then given by:

$$S_{\mathcal{G}} = \{(q, \nu) \in Q \times (\mathbb{R}_{\geq 0}^{\perp})^{\mathcal{E}} \mid \nu \text{ is compatible with } q\}.$$

Let $\gamma = (q, \nu) \in S_{\mathcal{G}}$ be a configuration, and define $E_0(\gamma) = \{e \in \mathbf{E}(q) \mid \forall e' \in \mathbf{E}(q), \nu(e) \leq \nu(e')\}$ and $d(\gamma) = \nu(e)$ for $e \in E_0(\gamma)$. From configuration γ , there is a transition to configuration $\gamma' = (q', \nu')$ on occurrence of the set of events $E_0(\gamma)$ after delay $d(\gamma)$ whenever:

$$\nu'(e) = \begin{cases} \perp & \text{if } e \notin \mathbf{E}(q') \\ \nu(e) - d(\gamma) & \text{if } e \in (\mathbf{E}(q) \cap \mathbf{E}(q')) \setminus E_0(\gamma) \\ t & \text{otherwise, with } \ell_e \leq t \leq u_e. \end{cases}$$

The σ -algebra $\Sigma_{\mathcal{G}}$ is obtained as the product between $2^{\mathcal{E}}$ and the Borel σ -algebra on $(\mathbb{R}_{\geq 0}^{\perp})^{\mathcal{E}}$. Let $\gamma = (q, \nu)$, and $B = \{q'\} \times B'$. Then, assuming $\text{Succ}(q, E_0(\gamma))(q') > 0$, we define the Markov kernel $\kappa_{\mathcal{G}}$ by:

$$\kappa_{\mathcal{G}}(\gamma, B) = \text{Succ}(q, E_0(\gamma))(q') \cdot \int_{(t_1, \dots, t_p) \in B'} \left(\prod_{e \in \mathbf{E}(q')} g_e(t_e) \right) dt_{e_1} \dots dt_{e_p}$$

where $g_e(t) = f_e(t)$ if $e \in \mathbf{E}(q') \setminus (\mathbf{E}(q) \setminus E_0(\gamma))$; $g_e(t) = \delta_{\nu(e) - d(\gamma)}$ if $e \in (\mathbf{E}(q) \cap \mathbf{E}(q')) \setminus E_0(\gamma)$; $g_e(t) = \delta_{\perp}$ if $e \notin \mathbf{E}(q')$. In other words, for a newly activated event e , its timestamp t_e is sampled (independently from the other events) according to density f_e ; for events inherited from the previous state, the delay which has elapsed is applied (hence the Dirac distributions in the definition).

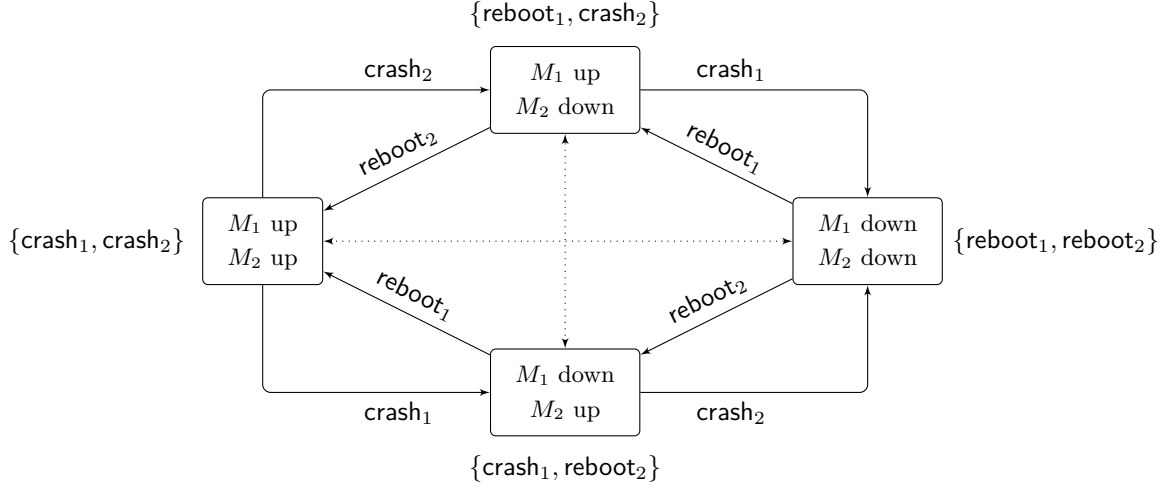
► **Example 76.** Consider a two-machine network (call M_1 and M_2 the two machines), in which crash times (event denoted crash_i for machine M_i) follow an exponential distribution with parameter λ_i ($\lambda_i \in \mathbb{R}_{>0}$) and reboot times (event denoted reboot_i for machine M_i) follow a uniform distribution over interval $[0, U_i]$ for some positive integer U_i . A GSMP model for the network is given on Figure 8.

When machine M_i is up (resp. down), a delay before event crash_i (resp. reboot_i) occurs and is sampled according to distribution $f(\text{crash}_i)$ (resp. $f(\text{reboot}_i)$). After that delay the event is triggered, unless it is preempted by some concurrent event. For instance, in state “ M_1 and M_2 up” (leftmost state), delays for the two events crash_1 and crash_2 are sampled, and the shortest delay decides the next transition to be taken and the next state which is reached. Note that dotted arrows represent events that happen with probability 0 (for instance, the very same delays are sampled for the two actions crash_1 and crash_2); we have omitted the corresponding labels (for readability).

We fix for the rest of this section a GSMP $\mathcal{G} = (Q, \mathcal{E}, \ell, u, f, \mathbf{E}, \text{Succ})$, and $\mathcal{T}_{\mathcal{G}} = (S_{\mathcal{G}}, \Sigma_{\mathcal{G}}, \kappa_{\mathcal{G}})$ its corresponding STS. To avoid too much technicalities, we assume that \mathcal{G} has no fixed-delay events, that is, for every event e , $\ell_e < u_e$. What we will present here would nevertheless extend to so-called *single-ticking* GSMPs [19].

8.2.2 The refined region graph abstraction

Due to the choice of the countdown-clock semantics (“clock values” decrease down to 0), the thick graph defined in subsection 8.1.2 has to be twisted a bit. Furthermore standard



$f(\text{crash}_i)$ is an exponential distribution with parameter λ_i
 $f(\text{reboot}_i)$ is a uniform distribution with support $[0, U_i]$

■ **Figure 8** An example GSMP: a network with two machines

regions will not be fine enough to yield an interesting abstraction. We will therefore refine regions using sets of separated configurations, that we define now.

Let $\varepsilon > 0$. We say that configuration $\gamma = (q, \nu)$ is ε -separated if for every $a, b \in \{0\} \cup \{\nu(e) \mid e \in \mathbf{E}(q)\}$, either $a = b$ or $|a - b| > \varepsilon$. We write C_G^ε for the set of ε -separated configurations.

The following lemma, stated and proven in [19], allows us to find an adequate granularity for a refined region abstraction.

► **Lemma 77** (Lemma 1 of [19]). *There exists $\varepsilon > 0$, $m \in \mathbb{N}$ and $p > 0$ such that for every $\gamma \in S_G$, $\text{Prob}_{\delta_\gamma}(\mathbf{F}_{\leq m} C_G^\varepsilon) \geq p$.*

We select $\varepsilon > 0$ following Lemma 77, and w.l.o.g. we assume ε is of the form $\frac{1}{d}$ with $d \in \mathbb{N}_{>0}$. We let M_G be the maximal constant appearing in constants $\{\ell_e \mid e \in \mathcal{E}\}$ and $\{u_e \mid e \in \mathcal{E} \text{ and } u_e < \infty\}$. Each event $e \in \mathcal{E}$ is virtually assigned a clock variable x_e , and we consider a refinement of the region equivalence for clocks $\{x_e \mid e \in \mathcal{E}\}$ w.r.t. maximal constant M_G and granularity $\frac{1}{d}$ as follows. Two valuations $\nu, \nu' \in (\mathbb{R}_{\geq 0}^1)^\mathcal{E}$ are equivalent whenever the following conditions hold:

1. for every $e \in \mathcal{E}$, either both $\nu(e)$ and $\nu'(e)$ are strictly larger than M_G , or the integral parts of $d \cdot \nu(e)$ and $d \cdot \nu'(e)$ coincide;
2. for every $e_1, e_2 \in \mathcal{E}$, for every $c \in \frac{1}{d} \cdot \mathbb{N} \cap [-M_G; M_G]$, for every $\bowtie \in \{<, \leq, =, \geq, >\}$, $\nu(e_1) - \nu(e_2) \bowtie c$ if and only if $\nu'(e_1) - \nu'(e_2) \bowtie c$.

Note that the above conditions refine the ones given in subsection 8.1.2 using diagonal constraints ([6]), and w.r.t. the granularity $\frac{1}{d}$ as well. We write R_G^ε for the set of equivalence classes, also called regions. We realize that any region $r \in R_G^\varepsilon$ has either only ε -separated configurations, or only non- ε -separated configurations. In particular, C_G^ε is a finite union of such regions.

We then define the abstraction $\alpha : Q \times \mathbb{R}_{\geq 0}^\mathcal{E} \rightarrow Q \times R_G^\varepsilon$ by projection, and the finite Markov chain $\mathcal{T}_G^{\text{rg}, \varepsilon}$ as follows:

- its set of states is $Q \times R_G^\varepsilon$;
- there is an edge from (q, r) to (q', r') whenever there exists $\nu \in r$ such that $\kappa_G((q, \nu), \{q'\} \times r') > 0$;
- from each state $(q, r) \in Q \times R_G^\varepsilon$, we associate the uniform distribution over $\{(q', r') \in Q \times R_G^\varepsilon \mid \text{there is an edge from } (q, r) \text{ to } (q', r')\}$.

Since $\mathcal{T}_G^{\text{rg}, \varepsilon}$ is just a rescaling of a standard region automaton, we immediately get:

► **Lemma 78.** $\mathcal{T}_G^{\text{rg}, \varepsilon}$ is a finite α -abstraction of \mathcal{T}_G .

As previously, we notice that the above abstraction is obviously complete (since it is finite).

8.2.3 Analyzing GSMPs

Let $A_G^\varepsilon = \{(q, r) \in Q \times R_G^\varepsilon \mid \alpha^{-1}(q, r) \subseteq C_G^\varepsilon\}$. As a consequence of Lemma 77 we get:

► **Proposition 79.** The set $\alpha^{-1}(A_G^\varepsilon)$ is a finite attractor for \mathcal{T}_G .

Finally, as for STA and using [19, Lemma 2] (which allows to prove that the conditions of Proposition 36 are actually satisfied), we also get:

► **Proposition 80.** $\mathcal{T}_G^{\varepsilon, \text{rg}}$ is a sound α -abstraction of \mathcal{T}_G .

As consequences, we get the following decidability and approximability results for GSMPs:

► **Corollary 81.** Let \mathcal{G} be a labelled GSMP (with no fixed-delay event), and \mathcal{M} be a DMA. Then:

1. we can decide whether \mathcal{G} satisfies almost-surely \mathcal{M} ;
2. for every initial distribution μ which is numerically amenable w.r.t. \mathcal{G} ,¹⁶ we can compute arbitrary approximations of $\text{Prob}_\mu^{\mathcal{T}_G}(\mathcal{M})$.

Proof. Again, the proof is similar to the ones of Corollaries 67 and 72. We just notice that it is obvious that if \mathcal{T}_G is a GSMP with no fixed-delay events, then so is $\mathcal{T}_G \times \mathcal{M}$. ◀

► **Remark 82.** We believe our approach gives new hints into the approximate quantitative model-checking of GSMPs, for which, up to our knowledge, only few results are known. For instance in [5, 13], the authors approximate the probability of until formulas of the form “the system reaches a target before time T within k discrete events, while staying within a set of safe states” (resp. “the system reaches a target while staying within a set of safe states”) for GSMPs (resp. a restricted class of GSMPs), and study numerical aspects. However one can notice that the contributions of [5, 13] are more the mathematical analysis of integral equations that need to be solved than convergence of approximation schemes.

Our approach permits to have approximation algorithms (with an arbitrary precision) for reachability or until properties for single-ticking GSMPs, or time-bounded reachability or until properties for the class of GSMPs with no cycle of immediate events (immediate events are fixed-delay events with delay 0).¹⁷ The numerical aspects in our computations can be dealt with as in [5, 13].

¹⁶We say that a distribution μ is numerically amenable w.r.t. \mathcal{G} if, given $k \in \mathbb{N}$, given $\varepsilon > 0$ and given a sequence of states and refined regions $(q_0, r_0), (q_1, r_1), \dots, (q_k, r_k)$, one can approximate $\text{Prob}_\mu^{\mathcal{T}_G}(\text{Cyl}((q_0, r_0), (q_1, r_1), \dots, (q_k, r_k)))$ up to ε .

¹⁷We recall the discussion on non-Zeno real-time systems of Section 7.5 (page 32), and realize that such GSMPs are almost-surely non-Zeno (see Appendix 8).

8.3 Stochastic time Petri nets

As a last instantiation of our general framework, we briefly discuss a particular class of general semi-Markov processes, namely GSMPs induced by *stochastic time Petri nets* (STPNs) [?]. Explaining how STPNs fit in our framework is interesting on its own since, on the one hand they admit a simpler abstraction than GSMPs, and on the other hand it allows us to compare to existing work on the transient analysis of STPNs [32, 39]. STPNs form a probabilistic and timed extension of the well-known Petri nets. These models are natural to represent concurrent systems in which events have random durations. An STPN is defined by a finite set of places, and a finite set of transitions, each equipped with a probability distribution over $\mathbb{R}_{\geq 0}$. The transitions in STPNs correspond to events in GSMPs. Now, the set of enabled events (here enabled transitions), is determined by the current marking. As for standard Petri nets, a *marking* maps places to natural numbers, representing the number of tokens in each place of the net. A transition is then *enabled* when the marking contains at least one token per input place. The difference with standard Petri nets lies in the choice of which transition will fire and when. When a transition has just fired, one token is consumed in each input place, one is added to each output place, and for every newly enabled transition, a delay is fired according to its associated probability distribution. The transition with minimum delay fires after this delay elapses, possibly disabling transitions or enabling new ones. Given an initial marking, the semantics of an STPN is thus an uncountable stochastic transition system where the states are tuples consisting of a marking, and a time to fire (the remaining delay) for each enabled transition.

In order to fit in our framework, we consider STPNs with the following two restrictions, as in [32, 39]: on the one hand, the underlying Petri net is bounded, *i.e.* the number of reachable markings is finite; and on the other hand, the stochastic process defined by the STPN is Markov regenerative. Let us explain in details this last assumption. When using arbitrary distributions to model the random duration of the transitions in the STPN, the remaining delays of transitions that stay enabled (when a transition just fired) have to be stored. Indeed, they will be compared with newly sampled durations to determine the next minimum delay, and thus impact the time to elapse as well as the next transition to fire. A *regeneration point* in the stochastic process defined by an STPN is a state for which the remaining duration of each enabled transition is non-zero only for transitions with exponentially distributed durations [32] (in [39], the condition is slightly more general, but more technical to explain, hence we focus on [32]). For such states, the non-zero remaining durations can indeed be forgotten and re-sampled using the same exponential distribution without altering the STPN semantics, thanks to the memoryless property of exponential distributions. An STPN is said to be *Markov regenerative* when almost surely regeneration points are encountered infinitely often with probability 1. Under our assumptions of boundedness of the Petri net, and Markov regeneration, clearly enough the STPN admits a finite attractor, namely the set of its regeneration points.

To apply our results to stochastic processes generated by STPNs, we first identify a finite state abstraction which will happen to be a sound abstraction. We let \mathcal{T}_1 be the stochastic process defined by an STPN, and define a finite-state Markov chain \mathcal{T}_2 , that corresponds to the state-class graph (see *e.g.* [32]) equipped with uniform discrete distributions. More precisely, the states of the abstraction \mathcal{T}_2 consist of state classes, that are the equivalent of regions for time Petri nets (we consider state classes just before sampling newly enabled transitions): state classes gather configurations with same marking and same ordering on remaining delays for transitions. Given that Petri nets are bounded, and their set of transitions is finite, the state space of \mathcal{T}_2 is finite. Some states correspond to regeneration points: those

in which the only transitions with non-zero remaining delay are exponentially distributed; in that case we assume w.l.o.g. that they will be resampled, which allows one to describe such a state in \mathcal{T}_2 using only its marking. Further, writing α for the abstraction, to build \mathcal{T}_2 , there is a transition from state s to state t as soon as there exists a transition in \mathcal{T}_1 between some state of $\alpha^{-1}(s)$ and $\alpha^{-1}(t)$. If $A = \{s_1, \dots, s_n\}$ denotes the set of regeneration points, it is important to realize that under this abstraction α , for every $s_i \in A$, $\alpha^{-1}(s_i) = s_i$. Due to the hypothesis that regeneration points are encountered almost-surely infinitely often, the set A is an attractor both in the concrete stochastic process, and in the finite-state abstraction. Moreover, for every $s_i \in A$ and every α -closed set B of the concrete stochastic process \mathcal{T}_1 such that there is a path from s_i to B , we let $k_{i,B}$ be its length and $p_{i,B}$ be its probability, so that $\text{Prob}_{\delta_{s_i}}^{\mathcal{T}_1}(\mathbf{F}_{\leq k_{i,B}} B) \geq p_{i,B}$. Letting $k = \max_{i,B} k_{i,B}$ and $p = \min_{i,B} p_{i,B}$, we obtain that for every $s_i \in A$ and every α -closed set B , either $\text{Prob}_{\delta_{s_i}}^{\mathcal{T}_1}(\mathbf{F}_{\leq k} B) \geq p$, or $\text{Prob}_{\delta_{s_i}}^{\mathcal{T}_1}(\mathbf{F}_{\leq k} B) = 0$. Therefore, we can apply Proposition 36 (see page 20), to derive that \mathcal{T}_1 is decisive w.r.t. every α -closed set. Further, by Proposition 40, we obtain that \mathcal{T}_2 is a sound α -abstraction of \mathcal{T}_1 .

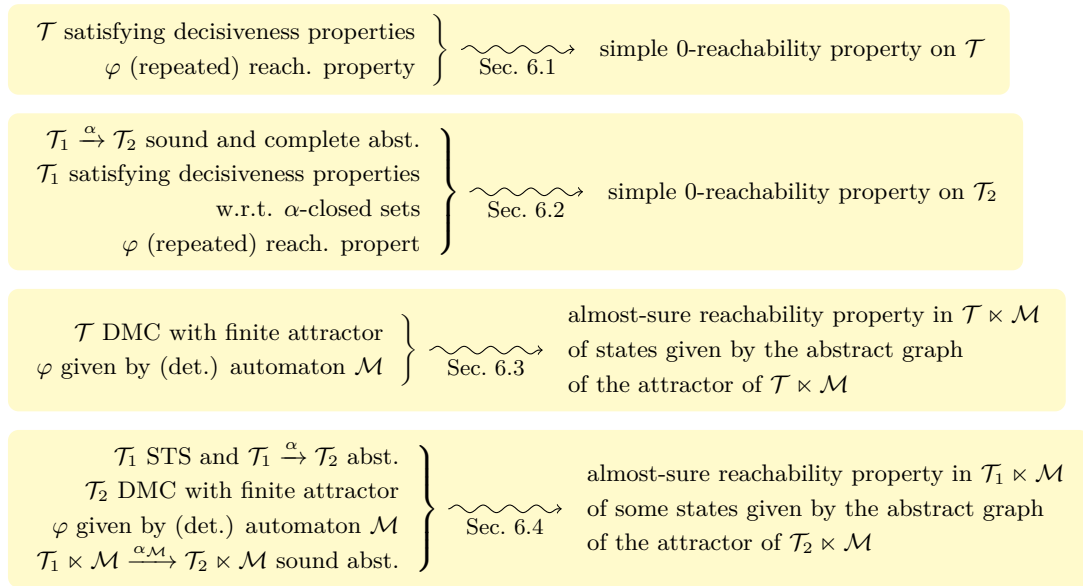
Similarly to the general case of GSMPs, we can thus conclude that the approximate quantitative model checking of STPNs can be done, provided numerical computations are amenable.

In [39], the authors provide a technique to perform the quantitative verification of time-bounded until properties in bounded STPNs that are Markov regenerative, with the assumption of a bound on the number of transitions between regeneration points (we call it strong Markov regenerative). As discussed above, we can relax the hypothesis on the bounded number of transitions between regeneration points, since the mere Markov regeneration property suffices to guarantee the existence of an attractor. A simple criterion to ensure Markov regeneration is to assume that every cycle in the state class graph contains a regeneration point (it is in fact equivalent to the strong Markov regeneration property). Note also that already the (weak) Markov regeneration implies that the stochastic process is almost-surely non-Zeno (since almost-surely states with exponentially distributed events are visited infinitely often), and therefore the approximation schemes for time-bounded reachability or until properties apply. For a fair comparison with the work of [32, 39], note however that their contribution focuses on *efficient* numerical computations of probabilities for time-bounded reachability or until properties. In contrast, we totally ignore the efficiency of probability computations and focus on sufficient conditions that enable our general framework to apply.

9 A guided tour of STSs

We now give an overview of the results presented in this paper. For improving readability, not all precise statements are listed. For instance, we omit the results which assume a fixed initial distribution. Also, few notations are borrowed from the paper, yet the global picture is almost self-contained.

The guide should be read as follows. Given an STS \mathcal{T} and a property φ , Figures 9 and 10 provide the assumptions on \mathcal{T} to be able to perform the qualitative or quantitative analysis of φ on \mathcal{T} . Note that when we consider abstractions $\mathcal{T}_1 \xrightarrow{\alpha} \mathcal{T}_2$, then we assume $\mathcal{T}_1 = \mathcal{T}$. Then, Figures 11, 12 and 13 summarize the relationships between the various notions. They should be used to know how to prove the properties that are expected of the model, either directly or via an abstraction (which needs to be designed).



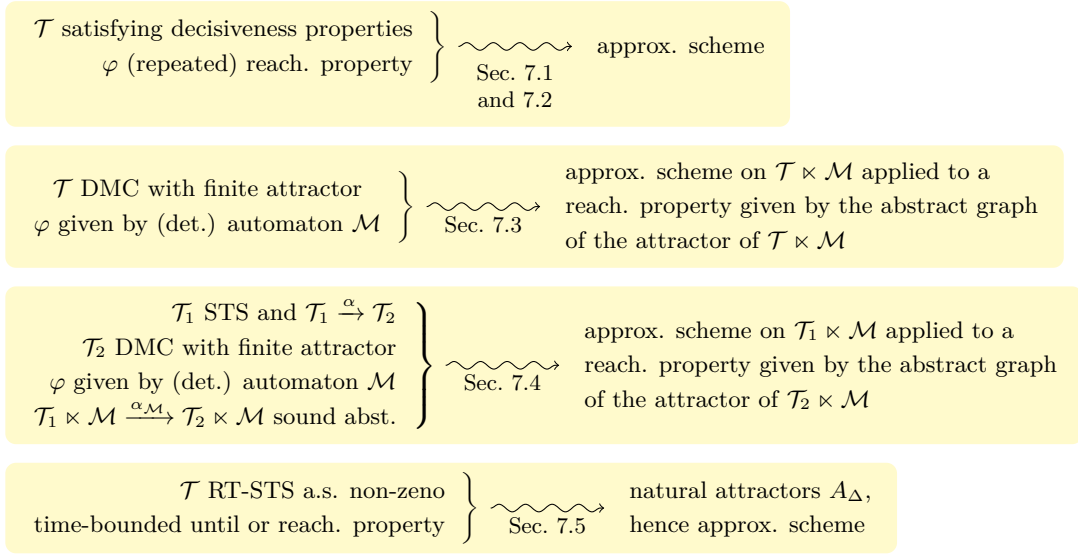
■ **Figure 9** Qualitative analysis: given \mathcal{T} an STS and φ a property, decide whether $\text{Prob}^{\mathcal{T}}(\varphi) = 0$ or $= 1$. \mathcal{T} is replaced with \mathcal{T}_1 in case of an abstraction. The edge \rightsquigarrow reads “amounts to”.

10 Conclusion

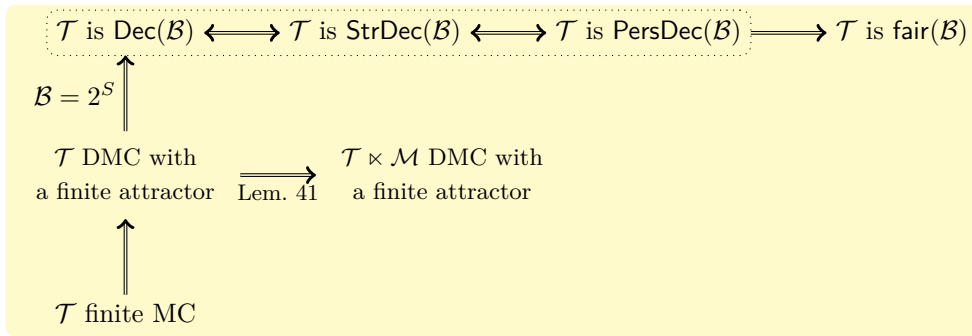
This paper deals with general stochastic transition systems (hence possibly continuous state-space Markov chains). We defined abstract properties of such stochastic processes, which allow one to design general procedures for their qualitative or quantitative analysis. The effectiveness of the approach requires some effectiveness assumption on specific high-level formalisms that are used to describe the stochastic process. We have demonstrated the effectiveness of the approach on three classes of systems: stochastic timed automata, generalized semi-Markov processes and stochastic time Petri nets. In both cases, we recover known results; but our approach yields further approximability results, which, up to our knowledge, are new.

We believe that, more importantly, we provide in this paper a methodology to understand stochastic models from a verification and algorithmics point-of-view. Section 9 gives a high-level description of our results, and of properties that should be satisfied by the stochastic model in order to apply our algorithms. In many cases, we showed that the hypotheses were really necessary to get the expected results, by providing counter-examples when the hypotheses are relaxed.

As future work, we plan to investigate new applications, such as infinite-state systems that occur in parameterized verification. Applying our results to stochastic hybrid systems is very tempting. However, our approach heavily relies on the existence of a sound finite or countable abstraction of the system. In case of stochastic timed automata, we noticed that although a finite time-abstract bisimulation always exists, decisiveness is not guaranteed (see Example 62). It thus seems hard to identify decisive subclasses of stochastic hybrid systems: most often, the underlying hybrid system does not admit a finite or countable time-abstract bisimulation [31]. Nevertheless, as mentioned at the end of Section 7, one can consider revisiting results on the time-bounded analysis of stochastic hybrid systems (see *e.g.* [21, 43]).



■ **Figure 10** Quantitative analysis (Section 7): given \mathcal{T} an STS and φ a property, compute (or approximate) $\text{Prob}^{\mathcal{T}}(\varphi)$. In case of the abstraction, $\mathcal{T}_1 = \mathcal{T}$. The edge \rightsquigarrow reads “amounts to”.



■ **Figure 11** Properties of STS (Section 3.4 and Lemma 41)

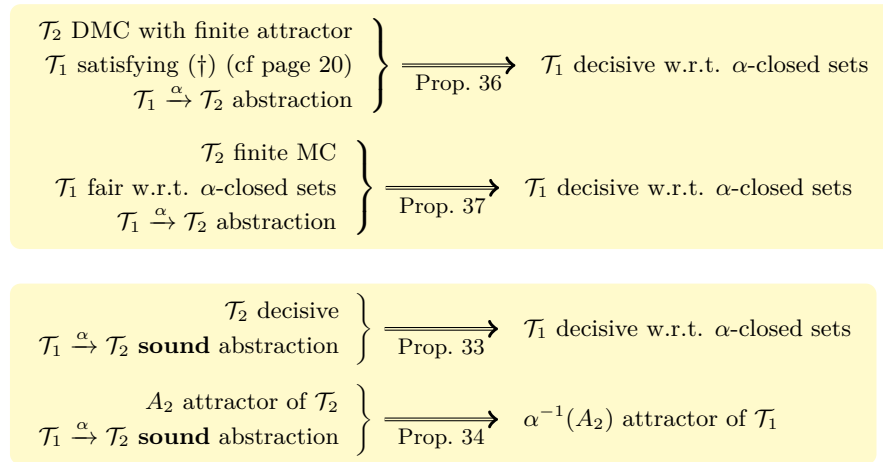
Also, we would like to adopt a similar generic approach for processes with non-determinism like Markov decision processes, or even stochastic two-player games.

Finally, let us mention that our approach could be interpreted in the context of stochastic relations [25], and that, for instance, the pushforward $\alpha_{\#}$ corresponds to the Giry monad applied to the abstraction α . We thank an anonymous reviewer to point us [25], and believe this may inspire further work to better understand abstractions of STSs.


Acknowledgement. We would like to warmly thank the anonymous referees, who provided exceptionally detailed reviews which greatly helped improving the paper.

References


- 1 Parosh A. Abdulla, Noomene Ben Henda, and Richard Mayr. Decisive Markov chains. *Logical Methods in Computer Science*, 3(4), 2007.
- 2 Parosh A. Abdulla, Nathalie Bertrand, Alexander Rabinovich, and Philippe Schnoebelen. Verification of probabilistic systems with faulty communication. *Information and Computation*, 202(2):141–165, 2005.




■ **Figure 12** Transfer of properties via abstractions



Condition for completeness:



Condition for soundness:



■ **Figure 13** Completeness and soundness of abstractions

- 3 Manindra Agrawal, S. Akshay, Blaise Genest, and P. S. Thiagarajan. Approximate verification of the symbolic dynamics of Markov chains. In *Proc. 27th Annual Symposium on Logic in Computer Science (LICS'12)*. IEEE Computer Society, 2012.
- 4 Rajeev Alur. *Techniques for Automatic Verification of Real-Time Systems*. PhD thesis, Stanford University, Stanford, CA, USA, 1991.
- 5 Rajeev Alur and Mikhail Bernadsky. Bounded model checking for GSMP models of stochastic real-time systems. In *Proc. 9th International Workshop on Hybrid Systems: Computation and Control (HSCC'06)*, volume 3927 of *Lecture Notes in Computer Science*, pages 19–33. Springer, 2006.
- 6 Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
- 7 Christel Baier, Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, and Marcus Größer. Probabilistic and topological semantics for timed automata. In *Proc. 27th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'07)*, volume 4855 of *Lecture Notes in Computer Science*, pages 179–191. Springer, 2007.
- 8 Christel Baier, Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, and Marcus Größer. Almost-sure model checking of infinite paths in one-clock timed automata. In *Proc. 23rd*

- Annual Symposium on Logic in Computer Science (LICS'08)*, pages 217–226. IEEE Computer Society Press, 2008.
- 9 Christel Baier, Boudewijn Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Model-checking algorithms for continuous-time Markov chains. *IEEE Transactions on Software Engineering*, 29(7):524–541, 2003.
 - 10 Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
 - 11 Christel Baier and Marta Z. Kwiatkowska. On the verification of qualitative properties of probabilistic processes under fairness constraints. *Information Processing Letters*, 66(2):71–79, 1998.
 - 12 Danièle Beauquier, Alexander Moshe Rabinovich, and Anatol Slissenko. A logic of probability with decidable model-checking. In *Proc. 16th International Workshop on Computer Science Logic (CSL'02)*, volume 2471 of *Lecture Notes in Computer Science*, pages 306–321. Springer, 2002.
 - 13 Mikhail Bernadsky and Rajeev Alur. Symbolic analysis for GSMP models with one stateful clock. In *Proc. 10th International Workshop on Hybrid Systems: Computation and Control (HSCC'07)*, volume 4416 of *Lecture Notes in Computer Science*, pages 90–103. Springer, 2007.
 - 14 Nathalie Bertrand. *Modèles stochastiques pour les pertes de messages dans les protocoles asynchrones et techniques de vérification automatique*. PhD thesis, École Normale Supérieure de Cachan, Cachan, France, 2006.
 - 15 Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, and Pierre Carlier. Analysing decisive stochastic processes. In *Proc. 43rd International Colloquium on Automata, Languages and Programming (ICALP'16)*, volume 55 of *LIPIcs*, pages 101:1–101:14. Leibniz-Zentrum für Informatik, 2016.
 - 16 Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, and Nicolas Markey. Quantitative model-checking of one-clock timed automata under probabilistic semantics. In *Proc. 5th International Conference on Quantitative Evaluation of Systems (QEST'08)*, pages 55–64. IEEE Computer Society Press, 2008.
 - 17 Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, Quentin Menet, Christel Baier, Marcus Größer, and Marcin Jurdziński. Stochastic timed automata. *Logical Methods in Computer Science*, 10(4):1–73, 2014.
 - 18 Patricia Bouyer, Thomas Brihaye, Marcin Jurdziński, and Quentin Menet. Almost-sure model-checking of reactive timed automata. In *Proc. 9th International Conference on Quantitative Evaluation of Systems (QEST'12)*, pages 138–147. IEEE Computer Society Press, 2012.
 - 19 Tomáš Brázdil, Jan Krčál, Jan Křetínský, and Vojtěch Řehák. Fixed-delay events in generalized semi-Markov processes revisited. In *Proc. 22nd International Conference on Concurrency Theory (CONCUR'11)*, volume 6901 of *Lecture Notes in Computer Science*, pages 140–155. Springer, 2011.
 - 20 Pierre Brémaud. *Markov Chains*. Springer, 1999.
 - 21 Taolue Chen, Marco Diciolla, Marta Z. Kwiatkowska, and Alexandru Mereacre. Time-bounded verification of CTMCs against real-time specifications. In *Proc. 9th International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS'11)*, volume 6919 of *Lecture Notes in Computer Science*, pages 26–42. Springer, 2011.
 - 22 Vincent Danos, Josée Desharnais, and Prakash Panangaden. Conditional expectation and the approximation of labelled Markov processes. In *Proc. 14th International Conference on Concurrency Theory (CONCUR'03)*, volume 2761 of *Lecture Notes in Computer Science*, pages 468–482. Springer, 2003.
 - 23 Pedro R. D'Argenio and Joost-Pieter Katoen. A theory of stochastic systems Part I: Stochastic automata. *Information and Computation*, 203(1):1–38, 2005.

- 24 Josée Desharnais and Prakash Panangaden. Continuous stochastic logic characterizes bisimulation of continuous-time Markov processes. *Journal of Logic and Algebraic Programming*, 56:99–115, 2003.
- 25 Ernst-Erich Doberkat. *Stochastic Relations – Foundations for Markov Transition Systems*. Chapman & Hall/CRC, 2007.
- 26 Martin Fränzle, Ernst Moritz Hahn, Holger Hermanns, Nicolás Wolovick, and Lijun Zhang. Measurability and safety verification for stochastic hybrid systems. In *Proc. 14th International Conference on Hybrid Systems: Computation and Control (HSCC’11)*, pages 43–52, 2011.
- 27 Martin Fränzle, Holger Hermanns, and Tino Teige. Stochastic satisfiability modulo theory: A novel technique for the analysis of probabilistic hybrid systems. In *Proc. 11th International Workshop on Hybrid Systems: Computation and Control (HSCC’08)*, pages 172–186, 2008.
- 28 Daniel Gburek, Christel Baier, and Sascha Klüppelholz. Composition of stochastic transition systems. In *Proc. 43rd International Colloquium on Automata, Languages and Programming (ICALP’16)*, volume 55 of *LIPICs*, pages 102:1–102:15. Leibniz-Zentrum für Informatik, 2016.
- 29 Peter W. Glynn. A GSMP formalism for discrete event systems. *Proc. of the IEEE*, 77(1):14–23, 1989.
- 30 Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002.
- 31 Thomas A. Henzinger. The theory of hybrid automata. In *Proceedings, 11th Annual IEEE Symposium on Logic in Computer Science, New Brunswick, New Jersey, USA, July 27-30, 1996*, pages 278–292. IEEE Computer Society, 1996.
- 32 András Horváth, Marco Paolieri, Lorenzo Ridi, and Enrico Vicario. Transient analysis of non-Markovian models using stochastic state classes. *Performance Evaluation*, 69(7-8):315–335, 2012.
- 33 S. Purushothaman Iyer and Murali Narasimha. Probabilistic lossy channel systems. In *Proc. 7th International Joint Conference on Theory and Practice of Software Development (TAPSOFT’97)*, volume 1214 of *Lecture Notes in Computer Science*, pages 667–681. Springer, 1997.
- 34 Vijay Anand Korthikanti, Mahesh Viswanathan, Gul Agha, and YoungMin Kwon. Reasoning about MDPs as transformers of probability distributions. In *Proc. 7th International Conference on Quantitative Evaluation of Systems (QEST’10)*, pages 199–208. IEEE Computer Society Press, 2010.
- 35 YoungMin Kwon and Gul Agha. Linear inequality LTL (iLTL): A model checker for discrete time Markov chains. In *Proc. 6th International Conference on Formal Engineering Methods (ICFEM’04)*, volume 3308 of *Lecture Notes in Computer Science*, pages 194–208. Springer, 2004.
- 36 François Laroussinie, Nicolas Markey, and Philippe Schnoebelen. Model checking timed automata with one or two clocks. In *Proc. 15th International Conference on Concurrency Theory (CONCUR’04)*, volume 3170 of *Lecture Notes in Computer Science*, pages 387–401. Springer, 2004.
- 37 Marco Ajmone Marsan, Gianni Conte, and Gianfranco Balbo. A class of generalized stochastic Petri nets for the performance evaluation of multiprocessor systems. *ACM Transactions on Computer Systems*, 2(2):93–122, 1984.
- 38 Prakash Panangaden. Measure and probability for concurrency theorists. *Theor. Comput. Sci.*, 253(2):287–309, 2001.
- 39 Prakash Panangaden. *Labelled Markov Processes*. Imperial College Press, 2009.

- 40 Marco Paolieri, András Horváth, and Enrico Vicario. Probabilistic model checking of regenerative concurrent systems. *IEEE Transactions on Software Engineering*, 42(2):153–169, 2016.
- 41 Amir Pnueli. On the extremely fair treatment of probabilistic algorithms. In *Proc. 15th Ann. Symp. Theory of Computing (STOC'83)*, pages 278–290. ACM Press, 1983.
- 42 Amir Pnueli and Lenore D. Zuck. Probabilistic verification. *Information and Computation*, 103(1):1–29, 1993.
- 43 Alexander Rabinovich. Quantitative analysis of probabilistic lossy channel systems. *Information and Computation*, 204(5):713–740, 2006.
- 44 Sadegh Soudjani, Rupak Majumdar, and Alessandro Abate. Safety verification of continuous-space pure jump Markov processes. In *Proc. 22nd International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'16)*, volume 9636 of *Lecture Notes in Computer Science*, pages 147–163. Springer, 2016.
- 45 Moshe Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proc. 26th Annual Symposium on Foundations of Computer Science (FOCS'85)*, pages 327–338. IEEE Computer Society Press, 1985.
- 46 Moshe Y. Vardi and Pierre Wolper. Reasoning about infinite computations. *Information and Computation*, 115(1):1–37, 1994.

– Appendix –

Former results already stated in the core of the paper are put in a box. New results are normally stated without box.

A

 Technical results of Section 2

A.1 Additional technical results for Subsection 2.2

We discuss some properties on the probability measures that we defined on paths of an STS. While not essential for the global understanding of the paper, they are useful in some of the coming proofs.

Recall that, if $s \in S$, the *Dirac distribution over s* , denoted δ_s , is defined for every measurable set A , by $\delta_s(A) = 1$ if $s \in A$, and $\delta_s(A) = 0$ otherwise.

Given any initial distribution μ , we can decompose the probability measure $\text{Prob}_\mu^\mathcal{T}$ into the various probability measures $\text{Prob}_{\delta_s}^\mathcal{T}$ for $s \in S$.

► **Lemma 83.** *For every $\varpi \in \mathcal{F}_\mathcal{T}$,*

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

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

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

It follows that for every $\mu \in \text{Dist}(S)$, we can write

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

and thus, by uniqueness of the measure extension, for every $\varpi \in \mathcal{F}_\mathcal{T}$,

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

This concludes the proof of the lemma. ◀

Recall that given two probability distributions μ and ν over some probability space (S, Σ) , μ and ν are *qualitatively equivalent* if for each $A \in \Sigma$, $\mu(A) = 0 \Leftrightarrow \nu(A) = 0$. The next lemma establishes that two qualitatively equivalent initial distributions yield two qualitatively equivalent distributions over paths.

► **Lemma 84.** *Let μ and ν be two probability measures over (S, Σ) . If μ and ν are qualitatively equivalent, then $\text{Prob}_\mu^\mathcal{T}$ and $\text{Prob}_\nu^\mathcal{T}$ are also qualitatively equivalent.*

Proof. We have to show that for each $\pi \in \mathcal{F}_\mathcal{T}$, $\text{Prob}_\mu^\mathcal{T}(\pi) = 0 \Leftrightarrow \text{Prob}_\nu^\mathcal{T}(\pi) = 0$. Since the complement of each cylinder is a finite union of cylinders and since each denumerable unions of cylinders can be written as a denumerable disjoint union of cylinders, it suffices to show

this for each cylinder $\text{Cyl}(A_0, \dots, A_n)$ with $A_0, \dots, A_n \in \Sigma$. We have to show that for each $A_0, \dots, A_n \in \Sigma$,

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

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

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

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

$$\int_{s_0 \in A_0} \kappa(s_0, A_1) \mu(ds_0) = 0. \quad (2)$$

Write $B = \{s_0 \in A_0 \mid \kappa(s_0, A_1) > 0\}$. We can write $B = \kappa(\cdot, A_1)^{-1}([0, 1]) \cap A_0$ which is in Σ from the hypotheses over κ . From (2), we can easily check that $\mu(B) = 0$, which implies that $\nu(B) = 0$ and thus

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

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

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

We inductively define:

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

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

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

We thus get that

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

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

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

(b)

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

where if $i = 0$, $\kappa(s_{i-1}, \cdot)$ will stand for the initial distribution μ . Point (a) is here in order to establish that the sets $\{s_i \in S \mid \text{Prob}_{\kappa(s_i, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{i+1}, \dots, B_{n-1}, A_n)) > 0\} \cap A_i$ are measurable, and point (b) aims at reducing our integrals to sets whose images have positive values. It should be observed that the second point is an immediate consequence of the first point. We thus only need to prove point (a). We do this by induction over i . First, if $i = n - 2$, we show that

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

A.2 Missing proofs of Subsections 2.3 and 2.5

► **Lemma 5.** *Let $\mu \in \text{Dist}(S)$ be an initial distribution and let $(A_i)_{0 \leq i \leq n}$ be a sequence of measurable sets. Write $\nu_0 = \mu_{A_0}$, the conditional probability of μ given A_0 , and for every $1 \leq j \leq n-1$, write $\nu_j = (\Omega_{\mathcal{T}}(\nu_{j-1}))_{A_j}$. Then, for every $0 \leq j \leq n$:*

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

Proof. The proof is by induction on j . Assume that $j = 0$, we have to show:

$$\text{Prob}_{\mu}^{\mathcal{T}}(\text{Cyl}(A_0, A_1, \dots, A_n)) = \mu(A_0) \cdot \text{Prob}_{\Omega_{\mathcal{T}}(\nu_0)}^{\mathcal{T}}(\text{Cyl}(A_1, \dots, A_n)).$$

First:

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

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

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

This concludes the proof for $j = 0$.

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

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

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

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

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

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

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

Proof of the fact that the σ -algebra $\Sigma \times 2^Q$ coincides with Σ' , the set of all subsets of $S \times Q$ of the form $\bigcup_{q \in Q} C_q \times \{q\}$, where $C_q \in \Sigma$ for every $q \in Q$ (stated page 11).

Proof. It suffices to show that

- (i) Σ' contains all rectangles;
- (ii) $\Sigma' \subseteq \Sigma \times 2^Q$; and
- (iii) Σ' is a σ -algebra.

Property (i) follows from the decomposition of any rectangle $A_1 \times A_2$ into elements of Σ' :

$$A_1 \times A_2 = \bigcup_{q \in A_2} A_1 \times \{q\} \cup \bigcup_{q \in A_2^c} \emptyset \times \{q\}.$$

Property (ii) is straightforward since for every $q \in Q$, $C_q \times \{q\} \in \Sigma \times 2^Q$, and thus, the union $\bigcup_{q \in Q} C_q \times \{q\}$ also belongs to the σ -algebra $\Sigma \times 2^Q$.

We finally establish property (iii). First Σ' is non-empty as $\emptyset \in \Sigma'$. Then, for $A = \bigcup_{q \in Q} C_q \times \{q\} \in \Sigma'$, the complement $A^c = \bigcup_{q \in Q} C_q^c \times \{q\}$ still belongs to Σ' since Σ is a σ -algebra and hence for each q , $C_q^c \in \Sigma$. Similarly, we get that Σ' is closed under denumerable unions. \blacktriangleleft

► **Proposition 13.** Let $\mu \in \text{Dist}(S)$ be an initial distribution for \mathcal{T} , and $\mathcal{M} = (Q, q_0, E, \mathcal{F})$ be a DMA. Then:

$$\text{Prob}_{\mu}^{\mathcal{T}}(\mathcal{T}[\mathcal{M}]) = \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T} \times \mathcal{M}}(\{\rho \in \text{Paths}(\mathcal{T} \times \mathcal{M}) \mid \rho \models \mathcal{F}\}) .$$

Proof. We will establish a link between distributions over $\text{Paths}(\mathcal{T})$ and distributions over $\text{Paths}(\mathcal{T} \times \mathcal{M})$. In order to do so, we introduce some notations. Given $A_0, A_1, \dots, A_n \in \Sigma'$ we write for each i , $A_i = \bigcup_{q \in Q} A_{i,q} \times \{q\}$. Also given $u_1, \dots, u_n \in 2^{\text{AP}}$ and $q \in Q$ we inductively define

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

Observe that since \mathcal{M} is deterministic, those states are uniquely defined. We then have the following result.

► **Lemma 85.** For each initial distribution $\mu \in \text{Dist}(S)$ for \mathcal{T} , for each state $q \in Q$ of \mathcal{M} , for each $n \in \mathbb{N}$ and for each $A_0, \dots, A_n \in \Sigma'$, it holds that

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

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

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

which is trivial from the definition of $\mu \times \delta_q$. Now fix $n \geq 0$. Assume that for each $0 \leq i \leq n$, the above property holds true and show that it is still the case for $i = n + 1$. Let $\mu \in \text{Dist}(S)$, $q \in Q$ and $A_0, \dots, A_{n+1} \in \Sigma'$. We have that

$$\begin{aligned} & \text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times \mathcal{M}}(\text{Cyl}(A_0, \dots, A_{n+1})) \\ &= \int_{(s_0, q') \in A_0} \text{Prob}_{\kappa'((s_0, q'), \cdot)}^{\mathcal{T} \times \mathcal{M}}(\text{Cyl}(A_1, \dots, A_{n+1})) d(\mu \times \delta_q)((s_0, q')) \\ &= \int_{s_0 \in A_{0,q}} \text{Prob}_{\kappa'((s_0, q), \cdot)}^{\mathcal{T} \times \mathcal{M}}(\text{Cyl}(A_1, \dots, A_{n+1})) d\mu(s_0) \\ &= \sum_{u_1 \in 2^{\text{AP}}} \int_{s_0 \in A_{0,q} \cap \mathcal{L}^{-1}(u_1)} \text{Prob}_{\kappa'((s_0, q), \cdot)}^{\mathcal{T} \times \mathcal{M}}(\text{Cyl}(A_1, \dots, A_{n+1})) d\mu(s_0) \\ &= \sum_{u_1 \in 2^{\text{AP}}} \int_{s_0 \in A_{0,q} \cap \mathcal{L}^{-1}(u_1)} \text{Prob}_{\kappa(s_0, \cdot) \times \delta_{q_{u_1}}}^{\mathcal{T} \times \mathcal{M}}(\text{Cyl}(A_1, \dots, A_{n+1})) d\mu(s_0) \end{aligned} \quad (3)$$

from unicity of q_{u_1} .

Using the induction hypothesis, we get that

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

Combining with (3), we thus obtain that

$$\begin{aligned} & \text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times \mathcal{M}}(\text{Cyl}(A_0, \dots, A_{n+1})) = \\ & \sum_{u_1, \dots, u_{n+1} \in 2^{\text{AP}}} \text{Prob}_{\kappa(s_0, \cdot)}^{\mathcal{T}}(\text{Cyl}(A_{0,q} \cap \mathcal{L}^{-1}(u_1), \dots, A_{n,q_{u_1} \dots u_n} \cap \mathcal{L}^{-1}(u_{n+1}), A_{n+1,q_{u_1} \dots u_{n+1}})) \end{aligned}$$

which concludes the proof. ◀

The proposition is a direct consequence of the previous lemma. ◀

B Technical results of Section 3

We give here the missing proofs of Section 3.

B.1 Proof of Lemma 14

► **Lemma 14.** *Given $B \in \Sigma$, it holds that:*

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

Proof. We first prove the first point. Remember that given $B \in \Sigma$,

$$\widetilde{B} = \{s \in S \mid \text{Prob}_{\delta_s}^{\mathcal{T}}(\mathbf{F} B) = 0\}.$$

Observe that we can write:

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

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

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

We will use similar arguments as in the proof of Lemma 84. Remember that if $n \geq 1$, it

holds that $\text{Prob}_{\delta_s}^{\mathcal{T}}(\text{Cyl}(\overbrace{S, \dots, S}^{n \text{ times}}, B)) = \text{Prob}_{\kappa(s, \cdot)}^{\mathcal{T}}(\text{Cyl}(\overbrace{S, \dots, S}^{n-1 \text{ times}}, B))$. First, if $n = 0$ then this set corresponds to the set $\{s \in S \mid \delta_s(B) = 0\} = B^c$ which is in Σ . Now if $n = 1$ then

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

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

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

We inductively define:

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

From the hypotheses over κ , it holds that $B_i \in \Sigma$ for each $0 \leq i < n$. In the sequel, s_0 denotes s . As in the proof of Lemma 84, we can show that firstly, $\int_{s_{n-1} \in S} \kappa(s_{n-1}, B) \kappa(s_{n-2}, ds_{n-1}) = \text{Prob}_{\kappa(s_{n-2}, \cdot)}^{\mathcal{T}}(\text{Cyl}(B_{n-1}, B))$ and that for each $1 \leq i \leq n-2$,

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

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

It follows that

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

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

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

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

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

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

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

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

$$\begin{aligned} \text{Ev}_T(\mathbf{F} \widetilde{B} \wedge \mathbf{G} \mathbf{F}(\widetilde{B})^c) &\subseteq \text{Ev}_T\left(\bigvee_{n \geq 0} (\mathbf{F}_{=n} \widetilde{B} \wedge \mathbf{F}_{>n}(\widetilde{B})^c)\right) \\ &= \bigcup_{n \geq 0} \bigcup_{m > 0} \text{Cyl}(\overbrace{S, \dots, S}^{n-1 \text{ times}}, \widetilde{B}, \overbrace{S, \dots, S}^{m \text{ times}}, (\widetilde{B})^c) \end{aligned}$$

it follows that there is $n \in \mathbb{N}$ and $m > 0$ such that

$$\text{Prob}_\mu^T(\text{Cyl}(\overbrace{S, \dots, S}^{n-1 \text{ times}}, \widetilde{B}, \overbrace{S, \dots, S}^{m \text{ times}}, (\widetilde{B})^c)) > 0.$$

From Lemma 5, writing $\nu = \Omega_{\mathcal{T}}^{(n)}(\mu)$, we get that

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

And from the third property proven previously, we deduce that

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

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

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

$$\text{Ev}_T(\mathbf{F} B \vee \mathbf{F} \widetilde{B}) = \text{Ev}_T(\mathbf{F} B \vee (\neg B \cup \widetilde{B})).$$

◀

B.2 Proof of Proposition 23

► **Proposition 23.** Let $\mathcal{B} \subseteq \Sigma$ and $\mu \in \text{Dist}(S)$. The following implications hold:

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

Proof. From the definitions, the following implications obviously hold true. For each $\mathcal{B} \subseteq \Sigma$ and for each $\mu \in \text{Dist}(S)$:

- \mathcal{T} is $\text{StrDec}(\mu, \mathcal{B}) \implies \mathcal{T}$ is $\text{Dec}(\mu, \mathcal{B})$, and
- \mathcal{T} is $\text{PersDec}(\mu, \mathcal{B}) \implies \mathcal{T}$ is $\text{Dec}(\mu, \mathcal{B})$.

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

► **Lemma 86.** *For each $\mathcal{B} \subseteq \Sigma$ and for each $\mu \in \text{Dist}(S)$, it holds that $\text{StrDec}(\mu, \mathcal{B})$ is equivalent to $\text{PersDec}(\mu, \mathcal{B})$.*

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

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

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

Now fix again $B \in \mathcal{B}$ and assume that \mathcal{T} is $\text{StrDec}(\mu, B)$, i.e. $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$. From Lemma 14 (fourth item), we get that $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \vee \mathbf{G} \mathbf{F} \widetilde{B}) = 1$ and it is then straightforward to establish that for each $p \geq 0$, $\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq p} B \vee \mathbf{F}_{\geq p} \widetilde{B}) = 1$. We hence deduce that \mathcal{T} is $\text{PersDec}(\mu, B)$ and thus $\text{PersDec}(\mu, \mathcal{B})$ as it holds true for each $B \in \mathcal{B}$. This concludes the proof. ◀

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

► **Lemma 87.** *For each $\mathcal{B} \subseteq \Sigma$, it holds that all three notions $\text{PersDec}(\mathcal{B})$, $\text{StrDec}(\mathcal{B})$ and $\text{Dec}(\mathcal{B})$ are equivalent.*

Proof. Fix $\mathcal{B} \subseteq \Sigma$. From the above results, it only remains to prove that $\text{Dec}(\mathcal{B}) \implies \text{StrDec}(\mathcal{B})$ or $\text{Dec}(\mathcal{B}) \implies \text{PersDec}(\mathcal{B})$. We prove the last one. We pick $B \in \mathcal{B}$ and assume that \mathcal{T} is $\text{Dec}(B)$, i.e. for each $\mu \in \text{Dist}(S)$, $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$. Pick $\mu \in \text{Dist}(S)$ and $i \geq 0$. We get that

$$\begin{aligned} \text{Prob}_\mu^\mathcal{T}(\mathbf{G}_{\geq i} B^c \wedge \mathbf{G}_{\geq i} (\widetilde{B})^c) &\leq \text{Prob}_{\mu_i}^\mathcal{T}(\mathbf{G}(B^c \cap (\widetilde{B})^c)) \\ &\quad \text{where } \mu_i = \Omega_{\mathcal{T}}^{(i)}(\mu), \text{ from Lemma 5} \\ &\quad \text{and from a similar argument as in the proof of Lemma 19} \\ &\leq 0 \text{ since } \mathcal{T} \text{ is } \text{Dec}(B). \end{aligned}$$

Hence for each $i \geq 0$, $\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq i} B \vee \mathbf{F}_{\geq i} \widetilde{B}) = 1$ and since it holds true for each $\mu \in \text{Dist}(S)$ and each $B \in \mathcal{B}$, we get that \mathcal{T} is $\text{PersDec}(\mathcal{B})$. ◀

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

► **Lemma 88.** *For each $\mathcal{B} \subseteq \Sigma$ and for each $\mu \in \text{Dist}(S)$, it holds that $\text{StrDec}(\mu, \mathcal{B})$ implies $\text{fair}(\mu, \mathcal{B})$, and $\text{StrDec}(\mathcal{B})$ implies $\text{fair}(\mathcal{B})$.*

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

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

$$\text{Prob}_\mu^\mathcal{T}(\mathbf{G F} B' \wedge \mathbf{F} \widetilde{B}) = 0. \quad (4)$$

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

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

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

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

It follows, from Lemma 5 like seen previously, that there is $\nu \in \text{Dist}(S)$ ($\nu = \Omega_\mathcal{T}^{(n)}(\mu)$), such that

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

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

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

Hence, $\nu(\widetilde{B}) > 0$ and we can apply Lemma 14 (second item) to obtain a contradiction. Hence, equation (4) holds. We then write:

$$\begin{aligned} 1 &= \text{Prob}_\mu^\mathcal{T}(\mathbf{G F} B \vee \mathbf{F} \widetilde{B} \mid \mathbf{G F} B') \quad \text{from strong decisiveness} \\ &= \frac{\text{Prob}_\mu^\mathcal{T}((\mathbf{G F} B \vee \mathbf{F} \widetilde{B}) \wedge \mathbf{G F} B')}{\text{Prob}_\mu^\mathcal{T}(\mathbf{G F} B')} \\ &= \frac{\text{Prob}_\mu^\mathcal{T}((\mathbf{G F} B \wedge \mathbf{G F} B') \vee (\mathbf{F} \widetilde{B} \wedge \mathbf{G F} B'))}{\text{Prob}_\mu^\mathcal{T}(\mathbf{G F} B')} \\ &= \frac{\text{Prob}_\mu^\mathcal{T}(\mathbf{G F} B \wedge \mathbf{G F} B')}{\text{Prob}_\mu^\mathcal{T}(\mathbf{G F} B')} \quad \text{from (4)} \\ &= \text{Prob}_\mu^\mathcal{T}(\mathbf{G F} B \mid \mathbf{G F} B') \end{aligned}$$

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

This concludes the proof of the proposition. ◀

C Technical results of Section 4

C.1 Additional technical results for Subsection 4.1

We now establish several technical results, which make explicit how STSs are related through an α -abstraction. The relationship is only qualitative, in the sense that it only relates positive reachability probabilities, but does not relate almost-sure or lower-bounded probabilities.

► **Lemma 89.** *Let $\alpha : (S_1, \Sigma_1) \rightarrow (S_2, \Sigma_2)$ be a measurable function. Then for every $s \in S_2$ and every $\mu \in \text{Dist}(\alpha^{-1}(\{s\}))$, $\alpha_{\#}(\mu) = \delta_s$.*

Proof. Fix $s \in S_2$ and $\mu \in \text{Dist}(\alpha^{-1}(\{s\}))$. For each $A \in \Sigma_2$, we have that $(\alpha_{\#}(\mu))(A) = \mu(\alpha^{-1}(A))$. If $s \in A$, then $\alpha^{-1}(\{s\}) \subseteq \alpha^{-1}(A)$ and thus $\mu(\alpha^{-1}(A)) = 1$. Otherwise, if $s \notin A$, then $\alpha^{-1}(\{s\}) \cap \alpha^{-1}(A) = \emptyset$ and thus $\mu(\alpha^{-1}(A)) = 0$. This directly implies that $\alpha_{\#}(\mu) = \delta_s$. ◀

► **Lemma 90.** *Assume that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Then, for every $i \in \mathbb{N}$, for every $\mu \in \text{Dist}(s_1)$, $\alpha_{\#}(\Omega_{\mathcal{T}_1}^{(i)}(\mu))$ is equivalent to $\Omega_{\mathcal{T}_2}^{(i)}(\alpha_{\#}(\mu))$.*

Proof. We show this by induction on i . Case $i = 1$ is by definition. Fix some $i \geq 1$ and assume that the statement holds true for each $1 \leq j \leq i$. By induction hypothesis, we have that $\alpha_{\#}(\Omega_{\mathcal{T}_1}^{(i)}(\mu))$ is equivalent to $\Omega_{\mathcal{T}_2}^{(i)}(\alpha_{\#}(\mu))$. We want to show that $\alpha_{\#}(\Omega_{\mathcal{T}_1}^{(i+1)}(\mu))$ is equivalent to $\Omega_{\mathcal{T}_2}^{(i+1)}(\alpha_{\#}(\mu))$.

We first notice that $\Omega_{\mathcal{T}_2}(\alpha_{\#}(\Omega_{\mathcal{T}_1}^{(i)}(\mu)))$ is equivalent to $\Omega_{\mathcal{T}_2}^{(i+1)}(\alpha_{\#}(\mu))$. Indeed write $\nu = \alpha_{\#}(\Omega_{\mathcal{T}_1}^{(i)}(\mu))$ and $\nu' = \Omega_{\mathcal{T}_2}^{(i)}(\alpha_{\#}(\mu))$. From the induction hypothesis, we know that ν and ν' . Following a similar argument as in the proof of Lemma 84 and from the definition of $\Omega_{\mathcal{T}_2}$, we can deduce that $\Omega_{\mathcal{T}_2}(\nu)$ is equivalent to $\Omega_{\mathcal{T}_2}(\nu')$. So it remains to show that $\Omega_{\mathcal{T}_2}(\alpha_{\#}(\mu'))$ is equivalent to $\alpha_{\#}(\Omega_{\mathcal{T}_1}(\mu'))$, when $\mu' = \Omega_{\mathcal{T}_1}^{(i)}(\mu)$. This is by definition of an α -abstraction. ◀

In other words, the above lemma states that for each $A \in \Sigma_2$ and for each $i \in \mathbb{N}$,

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F}_{=i} \alpha^{-1}(A)) > 0 \iff \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F}_{=i} A) > 0 .$$

This can even be generalized to cylinders:

► **Lemma 91.** *Assume that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Then for every $\mu \in \text{Dist}(S_1)$, for every $(A_i)_{0 \leq i \leq n} \in \Sigma_2^{n+1}$,*

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_0), \dots, \alpha^{-1}(A_n))) > 0 \iff \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(A_0, \dots, A_n)) > 0 .$$

Proof. We do the proof by induction on n . The case $n = 0$ is obvious from the definition of $\alpha_{\#}$. Now fix $n \geq 1$ and assume that for each $0 \leq k \leq n - 1$, for each $\mu \in \text{Dist}(S_1)$ and for each $(A_i)_{0 \leq i \leq k} \in \Sigma_2^{k+1}$,

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_0), \dots, \alpha^{-1}(A_k))) > 0 \iff \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(A_0, \dots, A_k)) > 0.$$

We show that it is still the case for n . Fix $\mu \in \text{Dist}(S_1)$ and $(A_i)_{i \geq n+1} \in \Sigma_2^{n+2}$. We let $\nu_0 = \mu_{\alpha^{-1}(A_0)}$ and $\nu'_0 = (\alpha_{\#}(\mu))_{A_0}$. Note that we hence assume that $\mu(\alpha^{-1}(A_0)) > 0$. We first realize that $\nu'_0 = \alpha_{\#}(\nu_0)$. Indeed for each $A \in \Sigma_2$,

$$(\alpha_{\#}(\nu_0))(A) = \nu_0(\alpha^{-1}(A)) = \frac{\mu(\alpha^{-1}(A \cap A_0))}{\mu(\alpha^{-1}(A_0))} = \frac{(\alpha_{\#}(\mu))(A \cap A_0)}{(\alpha_{\#}(\mu))(A_0)} = \nu'_0(A).$$

Then, applying Lemma 5, we get:

$$\begin{aligned} & \text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_0), \alpha^{-1}(A_1), \dots, \alpha^{-1}(A_n))) \\ &= \mu(\alpha^{-1}(A_0)) \cdot \text{Prob}_{\Omega_{\mathcal{T}_1}^{\mathcal{T}_1}(\nu_0)}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_1), \dots, \alpha^{-1}(A_n))) \end{aligned}$$

and

$$\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(A_0, A_1, \dots, A_n)) = (\alpha_{\#}(\mu))(A_0) \cdot \text{Prob}_{\Omega_{\mathcal{T}_2}^{\mathcal{T}_1}(\nu'_0)}^{\mathcal{T}_1}(\text{Cyl}(A_1, \dots, A_n)).$$

By definition of an α -abstraction, the measures $\Omega_{\mathcal{T}_2}(\nu'_0)$ and $\alpha_{\#}(\Omega_{\mathcal{T}_1}(\nu_0))$ are equivalent. Hence from Lemma 84,

$$\text{Prob}_{\Omega_{\mathcal{T}_2}(\nu'_0)}^{\mathcal{T}_2}(\text{Cyl}(A_1, \dots, A_n)) > 0 \Leftrightarrow \text{Prob}_{\alpha_{\#}(\Omega_{\mathcal{T}_1}(\nu_0))}^{\mathcal{T}_2}(\text{Cyl}(A_1, \dots, A_n)) > 0.$$

From the hypothesis of induction, we get that

$$\text{Prob}_{\alpha_{\#}(\Omega_{\mathcal{T}_1}(\nu_0))}^{\mathcal{T}_2}(\text{Cyl}(A_1, \dots, A_n)) > 0 \Leftrightarrow \text{Prob}_{\Omega_{\mathcal{T}_1}(\nu_0)}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_1), \dots, \alpha^{-1}(A_n))) > 0.$$

Since $(\alpha_{\#}(\mu))(A_0) = \mu(\alpha^{-1}(A_0))$, we conclude:

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_0), \alpha^{-1}(A_1), \dots, \alpha^{-1}(A_n))) > 0 \Leftrightarrow \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(A_0, A_1, \dots, A_n)) > 0.$$

We still have to consider the case where $\mu(\alpha^{-1}(A_0)) = 0$. In that case, $(\alpha_{\#}(\mu))(A_0) = 0$ and thus

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(A_0), \alpha^{-1}(A_1), \dots, \alpha^{-1}(A_n))) = 0 = \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(A_0, A_1, \dots, A_n))$$

which terminates the proof. \blacktriangleleft

As an immediate consequence, the positivity of properties with bounded witnesses are preserved through α -abstractions:

► **Corollary 92.** *Assume that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Then for every $\mu \in \text{Dist}(S_1)$, for every $A, B \in \Sigma_2$:*

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Ev}_{\mathcal{T}_1}(\alpha^{-1}(A) \text{ U } \alpha^{-1}(B))) > 0 \iff \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Ev}_{\mathcal{T}_2}(A \text{ U } B)) > 0 .$$

Soundness and completeness of abstractions

When the abstract system \mathcal{T}_2 is a DMC, soundness and completeness have a simpler characterization, which will be useful in the proofs.

► **Lemma 93.** *Assume \mathcal{T}_2 is a DMC. Then:*

■ \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 iff for every $s, s' \in S_2$,

$$\kappa_2(s, \{s'\}) > 0 \iff \forall \mu \in \text{Dist}(\alpha^{-1}(\{s\})), \text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(S_1, \alpha^{-1}(\{s'\}))) > 0 .$$

■ \mathcal{T}_2 is sound iff for every $s \in S_2$ and every $B \in \Sigma_2$,

$$\text{Prob}_{\delta_s}^{\mathcal{T}_2}(\mathbf{F} B) = 1 \implies \forall \mu \in \text{Dist}(\alpha^{-1}(\{s\})), \text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1 .$$

■ \mathcal{T}_2 is complete iff for every $s \in S_2$ and every $B \in \Sigma_2$,

$$\forall \mu \in \text{Dist}(\alpha^{-1}(\{s\})), \text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1 \implies \text{Prob}_{\delta_s}^{\mathcal{T}_2}(\mathbf{F} B) = 1 .$$

Proof. We handle the case of soundness. Indeed assume that for each $s \in S_2$ and for each $B \in \Sigma_2$, the condition presented in the statement (second item) holds true. Then fix $\mu \in \text{Dist}(S_1)$, $B \in \Sigma_2$ and assume that $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) = 1$ and show that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 1$. Towards a contradiction, assume that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) < 1$. Then, since \mathcal{T}_2 is a DMC, there is $s \in S_2$ such that $\mu(\alpha^{-1}(s)) > 0$ and

$$\text{Prob}_{\mu_{\alpha^{-1}(s)}}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) < 1.$$

From the hypothesis, it follows that $\text{Prob}_{\delta_s}^{\mathcal{T}_2}(\mathbf{F} B) < 1$. Observe that since $\mu(\alpha^{-1}(s)) > 0$, we have that $(\alpha_{\#}(\mu))(s) > 0$. Hence we get a contradiction by noticing:

$$\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F} B) \leq (\alpha_{\#}(\mu))(s) \cdot \text{Prob}_{\delta_s}^{\mathcal{T}_2}(\mathbf{F} B) < 1.$$

\blacktriangleleft

C.2 Missing proofs in Subsection 4.2

► **Proposition 33.** If \mathcal{T}_2 is a μ -sound α -abstraction of \mathcal{T}_1 , then for every $B \in \Sigma_2$:

$$\mathcal{T}_2 \text{ is Dec}(\alpha_{\#}(\mu), B) \implies \mathcal{T}_1 \text{ is Dec}(\mu, \alpha^{-1}(B)) .$$

In order to prove Proposition 33, we first show the following technical lemma, which relates avoid-sets in \mathcal{T}_1 and in \mathcal{T}_2 .

► **Lemma 94.** Let \mathcal{T}_2 be an α -abstraction of \mathcal{T}_1 . Then, for every $B \in \Sigma_2$: $\widetilde{\alpha^{-1}(B)} = \alpha^{-1}(\widetilde{B})$.

Proof. Fix $B \in \Sigma_2$. We have the series of equivalences:

$$\begin{aligned} s \in \widetilde{\alpha^{-1}(B)} &\iff \text{Prob}_{\delta_s}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) = 0 \\ &\iff \text{Prob}_{\alpha_{\#}(\delta_s)}^{\mathcal{T}_2}(\mathbf{F} B) = 0 \quad (\text{Corollary 92}). \end{aligned}$$

Now from Lemma 89, one can show that $\alpha_{\#}(\delta_s) = \delta_{\alpha(s)}$ by noticing that $\delta_s \in \text{Dist}(\alpha^{-1}(\alpha(s)))$. Hence $s \in \widetilde{\alpha^{-1}(B)}$ iff $\alpha(s) \in \widetilde{B}$ (i.e. $s \in \alpha^{-1}(\widetilde{B})$), which concludes the proof. ◀

We are now ready to prove Proposition 33.

Proof of Proposition 33. Fix $B \in \Sigma_2$ and assume that \mathcal{T}_2 is $\text{Dec}(\alpha_{\#}(\mu), B)$, i.e.

$$\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F} B \vee \mathbf{F} \widetilde{B}_2) = 1 . \quad (5)$$

To show that \mathcal{T}_1 is $\text{Dec}(\mu, \alpha^{-1}(B))$, by Lemma 94, it suffices to prove that

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B}_2)) = 1 .$$

The latter is immediate by (5) since \mathcal{T}_2 is μ -sound. ◀

► **Proposition 36.** Let \mathcal{T}_2 be a DMC such that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 .

1. Assume that there is a finite set $A_2 = \{s_1, \dots, s_n\} \subseteq S_2$ such that A_2 is an attractor for \mathcal{T}_2 and $A_1 = \bigcup_{i=1}^n \alpha^{-1}(s_i) = \alpha^{-1}(A_2)$ is an attractor for \mathcal{T}_1 .
2. Assume moreover that for every $1 \leq i \leq n$, for every α -closed set B in Σ_1 , there exist $p > 0$ and $k \in \mathbb{N}$ such that:
 - for every $\mu \in \text{Dist}(\alpha^{-1}(s_i))$, $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F}_{\leq k} B) \geq p$, or
 - for every $\mu \in \text{Dist}(\alpha^{-1}(s_i))$, $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} B) = 0$.

Then \mathcal{T}_1 is decisive w.r.t. every α -closed set.

Proof. Fix $B \subseteq S_2$ and $\mu \in \text{Dist}(S_1)$. We want to show that \mathcal{T}_1 is μ -decisive w.r.t. $\alpha^{-1}(B)$. We therefore have to show that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B})) = 1$. Towards a contradiction we assume that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{G}(\neg \alpha^{-1}(B)) \wedge \mathbf{G}(\neg \alpha^{-1}(\widetilde{B}))) > 0$, i.e. $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \alpha^{-1}(\widetilde{B}^c)) > 0$. Since $A_1 = \alpha^{-1}(A_2)$ is an attractor of \mathcal{T}_1 , we deduce from Lemma 19 that $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{G} \mathbf{F} \alpha^{-1}(A_2)) = 1$, hence:

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \alpha^{-1}(\widetilde{B}^c) \wedge \mathbf{G} \mathbf{F} \alpha^{-1}(A_2)) > 0 . \quad (6)$$

We let $A'_2 \subseteq A_2$ be the subset of states s of A_2 such that:

$$\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \alpha^{-1}(\widetilde{B}^c) \wedge \mathbf{G} \mathbf{F} \alpha^{-1}(\{s\})) > 0 .$$

Due to equation (6), A'_2 is non-empty, and furthermore every such s belongs to B^c and $(\widetilde{B})^c$. We set $A'_1 = \alpha^{-1}(A'_2)$.

In particular, $A'_1 \subseteq \alpha^{-1}((\widetilde{B})^c)$, hence from Lemma 14 (third item) we get that for every $\nu \in \text{Dist}(A'_1)$, $\text{Prob}_\nu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B)) > 0$. According to hypothesis (†), for every $s \in A'_2$, we can find $p_s > 0$ and $k_s \in \mathbb{N}$ such that for every $\nu_s \in \text{Dist}(\alpha^{-1}(s))$,

$$\text{Prob}_{\nu_s}^{\mathcal{T}_1}(\mathbf{F}_{\leq k_s} \alpha^{-1}(B)) \geq p_s.$$

Then taking $p = \min\{p_s \mid s \in A'_2\} > 0$ and $k = \max\{k_s \mid s \in A'_2\} \in \mathbb{N}$ (since A'_2 is finite), it holds that for every $\nu \in \text{Dist}(A'_1)$,

$$\text{Prob}_\nu^{\mathcal{T}_1}(\mathbf{F}_{\leq k} \alpha^{-1}(B)) \geq p \quad \text{hence} \quad \text{Prob}_\nu^{\mathcal{T}_1}(\mathbf{G}_{\leq k} \alpha^{-1}(B^c)) \leq 1 - p. \quad (7)$$

From (6), we deduce that:

$$0 < \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \alpha^{-1}((\widetilde{B})^c) \wedge \mathbf{G} \mathbf{F} A'_1) \leq \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \mathbf{F} A'_1)$$

Standardly in the literature (see *e.g.* [1, Lemma 3.4]), one infers immediately from (7) that

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{G} \alpha^{-1}(B^c) \wedge \mathbf{G} \mathbf{F} A'_1) \leq \lim_{n \rightarrow \infty} (1 - p)^n = 0$$

However we believe this is not so immediate, especially in our general setting, and we develop a complete proof below. Note that with this result, we exhibit a contradiction, which will conclude the proof.

It remains to show the last inequality. First we introduce some useful notations. Observe that from the definition of A'_1 , it holds that $A'_1 \subseteq \alpha^{-1}(B^c)$. Then for each $j \in \mathbb{N}$, we will write $B_{[j]}^c$ for the finite sequence $\alpha^{-1}(B^c), \dots, \alpha^{-1}(B^c)$ where $\alpha^{-1}(B^c)$ occurs exactly j times, and similarly we will write $(B^c \setminus A'_1)_{[j]}$ for the finite sequence $\alpha^{-1}(B^c) \setminus A'_1, \dots, \alpha^{-1}(B^c) \setminus A'_1$ where $\alpha^{-1}(B^c) \setminus A'_1$ occurs exactly j times. Then observe that

$$\begin{aligned} \text{Ev}_{\mathcal{T}_1}(\mathbf{G} \mathbf{F} A'_1 \wedge \mathbf{G} \alpha^{-1}(B^c)) = \\ \bigcap_{n \in \mathbb{N}} \bigcup_{j_0 \in \mathbb{N}} \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_{[j_0]}^c, A'_1, B_{[j_1]}^c, A'_1, B_{[j_2]}^c, \dots, B_{[j_{n-1}]}^c, A'_1, B_{[j_n]}^c), \end{aligned} \quad (8)$$


where $\mathbb{N}_{\geq k}$ denotes the set of natural numbers larger than or equal to k . We depict such a cylinder and what we can infer on the probabilities on Figure 14. As all behaviours are always in $\alpha^{-1}(B^c)$, the big rectangle represents this set, while the small one represents $A'_1 \subseteq \alpha^{-1}(B^c)$ which we know is reached infinitely often with probability 1. The behaviours are thus decomposed according to each visit in A'_1 followed by at least k moves (while staying in A'_1). The dashed arrows represent the k first steps. Note that within those k steps, A'_1 could be reached but it has no importance. What matters here is the fact that from A'_1 , the probability of the next k steps within $\alpha^{-1}(B^c)$ is upper bounded by $1 - p$. The curled arrows hold for the next visit to A'_1 which we hence know that it will happen with probability 1.

We will prove by induction over n that for each $n \geq 0$ and for each $\nu \in \text{Dist}(S_1)$,

$$\text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j_0 \in \mathbb{N}} \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_{[j_0]}^c, A'_1, B_{[j_1]}^c, A'_1, B_{[j_2]}^c, \dots, B_{[j_{n-1}]}^c, A'_1, B_{[j_n]}^c) \right) \leq (1 - p)^n.$$

Observe that for each $n \geq 0$, it holds that

$$\begin{aligned} \bigcup_{j_0 \in \mathbb{N}} \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_{[j_0]}^c, A'_1, B_{[j_1]}^c, A'_1, B_{[j_2]}^c, \dots, B_{[j_{n-1}]}^c, A'_1, B_{[j_n]}^c) \subseteq \\ \bigcup_{j_0 \in \mathbb{N}} \bigcup_{(j_1, \dots, j_{n-1}) \in \mathbb{N}_{\geq k}^{n-1}} \text{Cyl}(B_{[j_0]}^c, A'_1, B_{[j_1]}^c, A'_1, B_{[j_2]}^c, \dots, B_{[j_{n-1}]}^c, A'_1, B_{[k]}^c). \end{aligned}$$



■ **Figure 14** Scheme for the proof of Proposition 36.

Hence it is enough to demonstrate that for each $n \geq 0$ and for each $\nu \in \text{Dist}(S_1)$,

$$\text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j_0 \in \mathbb{N}} \bigcup_{(j_1, \dots, j_{n-1}) \in \mathbb{N}_{\geq k}^{n-1}} \text{Cyl}(B_{[j_0]}^c, A'_1, B_{[j_1]}^c, A'_1, B_{[j_2]}^c, \dots, B_{[j_{n-1}]}^c, A'_1, B_{[k]}^c) \right) \leq (1-p)^n. \quad (9)$$

First fix $n = 0$ and $\nu \in \text{Dist}(S_1)$. It corresponds to the two first arrows on Figure 14. We will show that for each $m \geq 0$,

$$\text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \text{Cyl}(B_{[j]}^c, A'_1, B_{[k]}^c) \right) \leq 1-p,$$

that is we decompose Figure 14 according to the length of the first curled arrow. We first prove cases $m = 0$ and $m = 1$ in order to illustrate what is happening, and then we will make the general case. If $m = 0$, it then holds that

$$\begin{aligned} \text{Prob}_\nu^{\mathcal{T}_1} (\text{Cyl}(A'_1, B_{[k]}^c)) &= \nu(A'_1) \cdot \text{Prob}_{\nu_{A'_1}}^{\mathcal{T}_1} (\text{Cyl}(A'_1, B_{[k]}^c)) \\ &\leq \text{Prob}_{\nu_{A'_1}}^{\mathcal{T}_1} (\text{Cyl}(B_{[k+1]}^c)) \leq 1-p \end{aligned} \quad (10)$$

where the first inequality holds from the fact that $A'_1 \subseteq \alpha^{-1}(B^c)$, and the second one from (7). Note that we assumed here that $\nu(A'_1) > 0$, but it has no importance since if $\nu(A'_1) = 0$, then the inequality trivially holds. Now if $m = 1$, first observe that

$$\text{Cyl}(A'_1, B_{[k]}^c) \cup \text{Cyl}(B^c, A'_1, B_{[k]}^c) = \text{Cyl}(A'_1, B_{[k]}^c) \cup \text{Cyl}(B^c \setminus A'_1, A'_1, B_{[k]}^c)$$

where in the second member of the equality, the union is disjoint. It follows that, writing $\nu'_0 = \nu_{B^c \setminus A'_1}$ and $\nu_1 = (\Omega_{\mathcal{T}_1}(\nu'_0))_{A'_1}$:

$$\begin{aligned} &\text{Prob}_\nu^{\mathcal{T}_1} (\text{Cyl}(A'_1, B_{[k]}^c) \cup \text{Cyl}(B^c, A'_1, B_{[k]}^c)) \\ &= \text{Prob}_\nu^{\mathcal{T}_1} (\text{Cyl}(A'_1, B_{[k]}^c)) + \text{Prob}_\nu^{\mathcal{T}_1} (\text{Cyl}(B^c \setminus A'_1, A'_1, B_{[k]}^c)) \\ &\leq \nu(A'_1) \cdot (1-p) + \nu(B^c \setminus A'_1) \cdot (\Omega_{\mathcal{T}_1}(\nu'_0))(A'_1) \cdot \text{Prob}_{\nu_1}^{\mathcal{T}_1} (\text{Cyl}(A'_1, B_{[k]}^c)) \text{ from Lemma 5} \\ &\leq \nu(A'_1) \cdot (1-p) + \nu(B^c \setminus A'_1) \cdot (\Omega_{\mathcal{T}_1}(\nu'_0))(A'_1) \cdot (1-p) \leq (1-p). \end{aligned}$$

Note that we again assumed here that $\nu(B^c \setminus A'_1) > 0$ and $(\Omega_{\mathcal{T}_1}(\nu'_0))(A'_1) > 0$, which has again no importance since otherwise, the probability of one of the cylinders would be equal

to 0 and which would thus not interfere on the above inequality. We now prove the general case for $m \geq 2$. Again, we can decompose the union of the cylinders into a disjoint one as follows:

$$\bigcup_{j=0}^m \text{Cyl}(B_j^c, A'_1, B_{[k]}^c) = \bigcup_{j=0}^m \text{Cyl}((B^c \setminus A'_1)_j, A'_1, B_{[k]}^c).$$

We use the following notations: $\nu'_0 = \nu_{B^c \setminus A'_1}$, $\nu_0 = \nu_{A'_1}$, and

- for each $1 \leq i \leq m-1$, $\nu'_i = (\Omega_{\mathcal{T}_1}(\nu'_{i-1}))_{B^c \setminus A'_1}$ and
- for each $1 \leq i \leq m$, $\nu_i = (\Omega_{\mathcal{T}_1}(\nu'_{i-1}))_{A'_1}$.

Note that we assume again that the conditional probability are well-defined, but like in cases $m = 0$ and $m = 1$, we can make this supposition w.l.o.g. Then using Lemma 5 and the observation (7), we get that:

$$\begin{aligned} \text{Prob}_{\nu}^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \text{Cyl}(B_{[j]}^c, A'_1, B_{[k]}^c) \right) &= \sum_{j=0}^m \text{Prob}_{\nu}^{\mathcal{T}_1} (\text{Cyl}(B_{[j]}^c, A'_1, B_{[k]}^c)) \\ &= \nu(A'_1) \cdot \text{Prob}_{\nu_0}^{\mathcal{T}_1} (\text{Cyl}(A'_1, B_{[k]}^c)) \\ &+ \sum_{j=1}^m \left(\nu(B^c \setminus A'_1) \cdot \prod_{i=1}^{j-1} (\Omega_{\mathcal{T}_1}(\nu'_i))(B^c \setminus A'_1) \cdot (\Omega_{\mathcal{T}_1}(\nu'_{j-1}))(A'_1) \cdot \overbrace{\text{Prob}_{\nu'_j}^{\mathcal{T}_1} (\text{Cyl}(A'_1, B_{[k]}^c))}^{\leq 1-p} \right) \\ &\leq (1-p) \cdot \left(\nu(A'_1) + \sum_{j=1}^m \left(\nu(B^c \setminus A'_1) \cdot \prod_{i=1}^{j-1} (\Omega_{\mathcal{T}_1}(\nu'_i))(B^c \setminus A'_1) \cdot (\Omega_{\mathcal{T}_1}(\nu'_{j-1}))(A'_1) \right) \right) \\ &= (1-p) \cdot \text{Prob}_{\nu}^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \text{Cyl}(B_{[j]}^c, A'_1) \right) \leq 1-p \end{aligned}$$

where the last equality comes again from Lemma 5, but in the other sense this time. Finally through the limit over m , we obtain that (9) is true when $n = 0$.

Now fix $n \geq 0$ and assume that for $0 \leq l \leq n$ and for each $\nu \in \text{Dist}(S_1)$, the inequality (9) holds true. We get in particular that for each $\nu \in \text{Dist}(S_1)$,

$$\text{Prob}_{\nu}^{\mathcal{T}_1} \left(\bigcup_{j_0 \in \mathbb{N}} \bigcup_{(j_1, \dots, j_{n-1}) \in \mathbb{N}_{\geq k}^{n-1}} \text{Cyl}(B_{[j_0]}^c, A'_1, B_{[j_1]}^c, A'_1, B_{[j_2]}^c, \dots, B_{[j_{n-1}]}^c, A'_1, B_{[k]}^c) \right) \leq (1-p)^n.$$

We want to show that (9) is still satisfied for $n+1$. Like in case $n = 0$, we will show that for each $m \geq 0$,

$$\text{Prob}_{\nu}^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_{[j]}^c, A'_1, B_{[j_1]}^c, A'_1, B_{[j_2]}^c, \dots, B_{[j_n]}^c, A'_1, B_{[k]}^c) \right) \leq (1-p)^{n+1}.$$

We thus again decompose the scheme of Figure 14 according to the length of the first arrow. In fact the proof is very similar to the case $n = 0$ as once you hit for the second time $\alpha^{-1}(B^c)$ in the scheme (*i.e.* after the first dashed arrow), the induction hypothesis can be applied. What happens before is the exact same behaviour as in the case for $n = 0$. For each $m \geq 0$ this finite union of cylinders can be decomposed into a finite union of disjoint sets as follows:

$$\begin{aligned} \bigcup_{j=0}^m \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_{[j]}^c, A'_1, B_{[j_1]}^c, A'_1, B_{[j_2]}^c, \dots, B_{[j_n]}^c, A'_1, B_{[k]}^c) &= \\ \bigcup_{j=0}^m \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}((B^c \setminus A'_1)_{[j]}, A'_1, B_{[j_1]}^c, A'_1, B_{[j_2]}^c, \dots, B_{[j_n]}^c, A'_1, B_{[k]}^c). \end{aligned}$$

Then using Lemma 5 and this decomposition into a disjoint union, it holds that

$$\begin{aligned} \text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_{[j]}^c, A'_1, B_{[j_1]}^c, A'_1, B_{[j_2]}^c, \dots, B_{[j_n]}^c, A'_1, B_{[k]}^c) \right) = \\ \sum_{j=0}^m \alpha_j \cdot \text{Prob}_{\mu_j}^{\mathcal{T}_1} \left(\bigcup_{j' \in \mathbb{N}} \bigcup_{(j_2, \dots, j_n) \in \mathbb{N}_{\geq k}^{n-1}} \text{Cyl}(B_{[j']}^c, A'_1, B_{[j_2]}^c, \dots, B_{[j_n]}^c, A'_1, B_{[k]}^c) \right), \end{aligned}$$

where for each $0 \leq j \leq m$, $0 < \alpha_j < 1$ and $\mu_j \in \text{Dist}(S_1)$ are given by Lemma 5, where α_j corresponds to:

$$\alpha_j = \text{Prob}_\nu^{\mathcal{T}_1} (\text{Cyl}((B^c \setminus A'_1)_{[j]}, A'_1, B_{[k]}^c)).$$

Note that this is possible due to the fact that we look at the union of all $j_1 \geq k$. Using the induction hypothesis and this last equality, we get that

$$\begin{aligned} \text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \bigcup_{(j_1, \dots, j_n) \in \mathbb{N}_{\geq k}^n} \text{Cyl}(B_{[j]}^c, A'_1, B_{[j_1]}^c, A'_1, B_{[j_2]}^c, \dots, B_{[j_n]}^c, A'_1, B_{[k]}^c) \right) \\ \leq (1-p)^n \cdot \text{Prob}_\nu^{\mathcal{T}_1} \left(\bigcup_{j=0}^m \text{Cyl}(B_{[j]}^c, A'_1, B_{[k]}^c) \right) \\ \leq (1-p)^{n+1} \end{aligned}$$

where the last inequality stands from what we have done in case $n = 0$. Through the limit over m , we can thus deduce that (9) is still true for $n + 1$.

Finally coming back to (8), through the limit over n this time, we conclude that

$$\text{Prob}_\mu^{\mathcal{T}_1} (\mathbf{G F} A'_1 \wedge \mathbf{G} \alpha^{-1}(B^c)) \leq \lim_{n \rightarrow \infty} (1-p)^n = 0.$$

This concludes the proof. ◀

► **Proposition 37.** Let \mathcal{T}_2 be a finite Markov chain such that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 . Fix $\mu \in \text{Dist}(S_1)$, and assume that \mathcal{T}_1 is μ -fair w.r.t. every α -closed set. Then \mathcal{T}_1 is μ -decisive w.r.t. every α -closed set.

Proof. As \mathcal{T}_2 is a finite Markov chain, it can be viewed as a graph. We can therefore speak of the *bottom strongly connected components (BSCC)* of \mathcal{T}_2 (a BSCC is a subset $C \subseteq S_2$ such that for all $s, s' \in C$, if s' is reachable from s , then s is reachable from s' as well). We write $\text{BSCC}(\mathcal{T}_2)$ for the set of BSCCs of \mathcal{T}_2 . We define $\mathcal{C} = \{s \in S_2 \mid \exists C \in \text{BSCC}(\mathcal{T}_2), s \in C\}$. We first prove that $\text{Prob}_\mu^{\mathcal{T}_1} (\mathbf{F} \alpha^{-1}(\mathcal{C})) = 1$. In order to establish this, we show that for each $s \in S_2$, $\text{Prob}_\mu^{\mathcal{T}_1} (\mathbf{G F} \alpha^{-1}(s)) > 0$ implies that $s \in \mathcal{C}$. Indeed, pick $s \in S_2$ such that:

$$\text{Prob}_\mu^{\mathcal{T}_1} (\mathbf{G F} \alpha^{-1}(\{s\})) > 0.$$

We can state that for each $k \geq 1$ and for each $s_0, s_1, \dots, s_k \in S_2$ with $s_0 = s$ and such that for each $0 \leq i < k$, $\kappa_2(s_i, s_{i+1}) > 0$, it holds that

$$\text{Prob}_\mu^{\mathcal{T}_1} (\mathbf{G F} \alpha^{-1}(s_k) \mid \mathbf{G F} \alpha^{-1}(s)) = 1.$$

We prove this by induction over k . First fix $k = 1$ and let $s_1 \in S_2$ such that $\kappa_2(s, s_1) > 0$. Then for every $\nu \in \text{Dist}(\alpha^{-1}(s))$, $\text{Prob}_\nu^{\mathcal{T}_1}(\text{Cyl}(\alpha^{-1}(\{s\}), \alpha^{-1}(\{s_1\}))) > 0$. Hence $\alpha^{-1}(s) \in \text{PreProb}^{\mathcal{T}}(\{\alpha^{-1}(s_1)\})$. And since \mathcal{T}_1 is fair w.r.t. α -closed sets, we get that

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(\{s_1\}) \mid \mathbf{GF} \alpha^{-1}(\{s\})) = 1.$$

Now fix $k > 1$ and assume that for each $1 \leq j < k$ and for each $s_0, \dots, s_j \in S_2$ with $s_0 = s$ and such that for each $0 \leq i < j$, $\kappa_2(s_i, s_{i+1}) > 0$, it holds that

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(s_j) \mid \mathbf{GF} \alpha^{-1}(s)) = 1.$$

We want to show that it is still the case for k . Fix $s_0, s_1, \dots, s_k \in S_2$ satisfying all the desired hypotheses. Using the induction hypothesis, we know that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(s_{k-1}) \mid \mathbf{GF} \alpha^{-1}(s_0)) = 1$ and $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(s_k) \mid \mathbf{GF} \alpha^{-1}(s_{k-1})) = 1$. We can then compute:

$$\begin{aligned} & \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(s_k) \mid \mathbf{GF} \alpha^{-1}(s_0)) \\ &= \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(s_k) \wedge \mathbf{GF} \alpha^{-1}(s_{k-1}) \mid \mathbf{GF} \alpha^{-1}(s_0)) \\ &= \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(s_k) \mid \mathbf{GF} \alpha^{-1}(s_{k-1}) \wedge \mathbf{GF} \alpha^{-1}(s_0)) \cdot \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(s_{k-1}) \mid \mathbf{GF} \alpha^{-1}(s_0)) \\ &= 1 \end{aligned}$$

from the induction hypothesis. This shows that for every state s' which is reachable from s in \mathcal{T}_2 ,

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(\{s'\}) \mid \mathbf{GF} \alpha^{-1}(\{s\})) = 1.$$

Then fix s' reachable from s in \mathcal{T}_2 . We can show that s is also reachable from s' . Towards a contradiction, assume that it is not the case. It follows that

$$\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(\{s'\}) \wedge \mathbf{GF} \alpha^{-1}(\{s\})) = 0$$

which is a contradiction with $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(\{s'\}) \mid \mathbf{GF} \alpha^{-1}(\{s\})) = 1$ and $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(\{s\})) > 0$. We deduce thus that s belongs to a BSCC of \mathcal{T}_2 .

We can now prove that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(\mathcal{C})) = 1$. Indeed observe first that from the finiteness of \mathcal{T}_2 , it holds that for every paths $\rho = t_0 t_1 t_2 \dots \in \text{Paths}(\mathcal{T}_1)$, there is $s \in S_2$ such that $\{i \in \mathbb{N} \mid t_i \in \alpha^{-1}(s)\}$ is infinite. Keeping this in mind, we write $S_2 = \{s_1, \dots, s_k, s_{k+1}, \dots, s_n\}$ where $k \geq 1$ and $\{s_1, \dots, s_k\} = \mathcal{C}$. Then we can write

$$\begin{aligned} \text{Paths}(\mathcal{T}_1) &= \text{Ev}_{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(s_1)) \cup \text{Ev}_{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(s_2) \wedge \mathbf{FG} \neg \alpha^{-1}(s_1)) \\ &\quad \cup \dots \cup \text{Ev}_{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(s_n) \wedge \bigwedge_{i=1}^{n-1} \mathbf{FG} \neg \alpha^{-1}(s_i)). \end{aligned}$$

From what we have shown previously, we now get that for each $j \geq k + 1$,

$$0 = \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(s_j)) \geq \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(s_j) \wedge \bigwedge_{i=1}^{j-1} \mathbf{FG} \neg \alpha^{-1}(s_i)).$$

And we conclude that

$$\begin{aligned} 1 &= \text{Prob}_\mu^{\mathcal{T}_1}(\text{Paths}(\mathcal{T}_1)) \\ &= \sum_{j=1}^k \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{GF} \alpha^{-1}(s_j) \wedge \bigwedge_{i=1}^{j-1} \mathbf{FG} \neg \alpha^{-1}(s_i)) \\ &\leq \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(\mathcal{C})). \end{aligned}$$

We are now able to prove that \mathcal{T}_1 is $\text{Dec}(\mu, \mathcal{B})$. Fix $B \subseteq S_2$, we want to show that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B})) = 1$. We have that

$$\begin{aligned} & \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B})) \\ &= \sum_{\substack{C \in \text{BSCC}(\mathcal{T}_2) \text{ s.t.} \\ \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(C)) > 0}} \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(C)) \cdot \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B}) \mid \mathbf{F} \alpha^{-1}(C)). \end{aligned}$$

Now we fix some $C \in \text{BSCC}(\mathcal{T}_2)$ such that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(C)) > 0$. There are two cases:

- first if there is $s \in C$ such that $s \in B$, then $\alpha^{-1}(s) \subseteq \alpha^{-1}(B)$ and thus $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B}) \mid \mathbf{F} \alpha^{-1}(C)) = 1$;
- or for each $s \in C$, $s \in \widetilde{B}$ which implies that $\alpha^{-1}(C) \subseteq \alpha^{-1}(\widetilde{B})$ and in that case again $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B}) \mid \mathbf{F} \alpha^{-1}(C)) = 1$.

We finally conclude that

$$\begin{aligned} \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(B) \vee \mathbf{F} \alpha^{-1}(\widetilde{B})) &= \sum_{\substack{C \in \text{BSCC}(\mathcal{T}_2) \text{ s.t.} \\ \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(C)) > 0}} \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(C)) \\ &= \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \alpha^{-1}(C)) = 1. \end{aligned}$$

◀

D Technical results of Section 5

► **Lemma 41.** *Assume that A is an attractor for \mathcal{T} . Then $A \times Q$ is an attractor for $\mathcal{T} \times \mathcal{M}$. Furthermore, if A is finite, then so is $A \times Q$.*

We first prove the following lemma.

► **Lemma 95.** *Fix $\mu \in \text{Dist}(S)$ and assume that $A \in \Sigma$ is a μ -attractor for \mathcal{T} . Then for each $q \in Q$, $A \times Q$ is a $(\mu \times \delta_q)$ -attractor for $\mathcal{T} \times \mathcal{M}$.*

Proof. Fix $\mu \in \text{Dist}(S)$ and $A \in \Sigma$ such that $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} A) = 1$. Fix $q \in Q$. We know that

$$\text{Ev}_{\mathcal{T} \times \mathcal{M}}(\mathbf{F} A \times Q) = \text{Ev}_{\mathcal{T} \times \mathcal{M}}\left(\bigcup_{n \in \mathbb{N}} \overbrace{\text{Cyl}(S', \dots, S', A \times Q)}^{n \text{ times}}\right).$$

Then from Lemma 85, we know that for each $n \in \mathbb{N}$

$$\begin{aligned} \text{Prob}_{\mu \times \delta_q}^{\mathcal{T} \times \mathcal{M}}(\overbrace{\text{Cyl}(S', \dots, S', A \times Q)}^{n \text{ times}}) &= \sum_{u_1, \dots, u_n \in 2^{\text{AP}}} \text{Prob}_\mu^{\mathcal{T}}(\text{Cyl}(\mathcal{L}^{-1}(u_1), \dots, \mathcal{L}^{-1}(u_n), A)) \\ &= \text{Prob}_\mu^{\mathcal{T}}(\overbrace{\text{Cyl}(S, \dots, S, A \times Q)}^{n \text{ times}}). \end{aligned}$$

As this holds true for each $n \geq 0$, we thus get that $\text{Prob}_{\mu \times \delta_q}(\mathbf{F} A \times Q) = \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} A) = 1$ from the hypothesis. This concludes the proof. ◀



Proof of Lemma 41. Fix $A \in \Sigma$ such that for each $\mu \in \text{Dist}(S)$, $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} A) = 1$. We want to prove that for each $\nu \in \text{Dist}(S \times Q)$, $\text{Prob}_\nu^{\mathcal{T} \times \mathcal{M}}(\mathbf{F} A \times Q) = 1$. Fix $\nu \in \text{Dist}(S \times Q)$ and compute:

$$\text{Prob}_\nu^{\mathcal{T} \times \mathcal{M}}(\mathbf{F} A \times Q) = \sum_{q \in Q} \nu(S \times \{q\}) \cdot \text{Prob}_{\nu_{S \times \{q\}}}^{\mathcal{T} \times \mathcal{M}}(\mathbf{F} A \times Q).$$

Note that $\nu_{S \times \{q\}}$ induces a distribution $\nu_q \in \text{Dist}(S)$ as follows: for each $B \in \Sigma$, $\nu_q(B) = \nu_{S \times \{q\}}(B \times \{q\})$. Writing $\mu = \nu_q$ it then holds that $\nu_{S \times \{q\}} = \mu \times \delta_q$. We then get, from the hypothesis and Lemma 95, that $\text{Prob}_{\nu_{S \times \{q\}}}^{\mathcal{T} \times \mathcal{M}}(\mathbf{F} A \times Q) = 1$ for each $q \in Q$. Hence, $\text{Prob}_\nu^{\mathcal{T} \times \mathcal{M}}(\mathbf{F} A \times Q) = \sum_{q \in Q} \nu(S \times \{q\}) = 1$ which concludes the proof. \blacktriangleleft

► **Lemma 47.** Let $\alpha_{\mathcal{M}} : S_1 \times Q \rightarrow S_2 \times Q$ be the lifting of α such that $\alpha_{\mathcal{M}}(s, q) = (\alpha(s), q)$. If \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 , then $\mathcal{T}_2 \times \mathcal{M}$ is an $\alpha_{\mathcal{M}}$ -abstraction of $\mathcal{T}_1 \times \mathcal{M}$. Furthermore, if $\mathcal{T}_1 \times \mathcal{M}$ is $\text{Dec}(\mathcal{B})$ where $\mathcal{B} = \{\alpha_{\mathcal{M}}^{-1}(B) \mid B \in \Sigma'_2\}$, then $\mathcal{T}_2 \times \mathcal{M}$ is a sound $\alpha_{\mathcal{M}}$ -abstraction of $\mathcal{T}_1 \times \mathcal{M}$.

Proof. We first show that $\mathcal{T}_2 \times \mathcal{M}$ is an $\alpha_{\mathcal{M}}$ -abstraction of $\mathcal{T}_1 \times \mathcal{M}$. It suffices to show that for each $\mu \in \text{Dist}(S_1)$, for each $q, q' \in Q$ and for each $B_{q'} \in \Sigma_2$,

$$\text{Prob}_{\mu \times \delta_q}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Cyl}(S_1 \times Q, \alpha_{\mathcal{M}}^{-1}(B_{q'} \times \{q'\}))) > 0 \Leftrightarrow \text{Prob}_{(\alpha_{\mathcal{M}})_{\#}(\mu \times \delta_q)}^{\mathcal{T}_2 \times \mathcal{M}}(\text{Cyl}(S_2 \times Q, B_{q'} \times \{q'\})) > 0. \quad (11)$$

Fix $\mu \in \text{Dist}(S_1)$, $q, q' \in Q$ and $B_{q'} \in \Sigma_2$. Write $u \in 2^{\text{AP}}$ for the unique label such that $(q, u, q') \in E$. In order to prove (11), we will use the fact that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_2 . And in order to make the link with the wanted equivalence, we will use Lemma 85. We can establish that $(\alpha_{\mathcal{M}})_{\#}(\mu \times \delta_q) = \alpha_{\#}(\mu) \times \delta_q$. Indeed given $p \in Q$ and $C_p \in \Sigma_2$, it holds that

$$\begin{aligned} (\alpha_{\mathcal{M}})_{\#}(\mu \times \delta_q)(C_p \times \{p\}) &= (\mu \times \delta_q)(\alpha^{-1}(C_p)\{p\}) \\ &= \mu(\alpha^{-1}(C_p)) \cdot \delta_q(p) \\ &= \alpha_{\#}(\mu)(\alpha^{-1}(C_p)) \cdot \delta_q(p) = (\alpha_{\#}(\mu) \times \delta_q)(C_p \times \{p\}). \end{aligned}$$

Hence we get that

$$\begin{aligned} \text{Prob}_{(\alpha_{\mathcal{M}})_{\#}(\mu \times \delta_q)}^{\mathcal{T}_2 \times \mathcal{M}}(\text{Cyl}(S_2 \times Q, B_{q'} \times \{q'\})) > 0 &\Leftrightarrow \text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\text{Cyl}(\mathcal{L}_2^{-1}(u), B_{q'})) > 0 \\ &\Leftrightarrow \text{Prob}_{\mu}^{\mathcal{T}_1}(\text{Cyl}(\mathcal{L}_1^{-1}(u), \alpha^{-1}(B_{q'}))) > 0 \\ &\Leftrightarrow \text{Prob}_{\mu \times \delta_q}^{\mathcal{T}_1 \times \mathcal{M}}(\text{Cyl}(S_1 \times Q, \alpha_{\mathcal{M}}^{-1}(B_{q'} \times \{q'\}))) > 0 \end{aligned}$$

where the first and third equivalences hold from Lemma 85, and the second equivalence holds from the fact that \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 .

Finally, since $\mathcal{T}_1 \times \mathcal{M}$ is decisive w.r.t $\alpha_{\mathcal{M}}^{-1}(B)$ for each $B \in \Sigma'_2$ and since $\mathcal{T}_2 \times \mathcal{M}$ is an $\alpha_{\mathcal{M}}$ -abstraction of $\mathcal{T}_1 \times \mathcal{M}$, Proposition 40 allows us to conclude that $\mathcal{T}_2 \times \mathcal{M}$ is a sound $\alpha_{\mathcal{M}}$ -abstraction of $\mathcal{T}_1 \times \mathcal{M}$. \blacktriangleleft

We give here the (partial) counter-example mentioned in Remark 48.

► **Example 96.** We illustrate Remark 48 by exhibiting an example where soundness (w.r.t. a fixed distribution) as well as decisiveness properties do not transfer to the product with a deterministic Muller automaton.

Consider the DMC \mathcal{T}_1 depicted on the left of Figure 15 which corresponds to the random walk over \mathbb{N} from Example 2, when $p = 2/3$. Consider also the finite MC \mathcal{T}_2 on the right of the same figure. Clearly enough, \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 for the mapping $\alpha : \mathbb{N} \rightarrow \{s_0, s_1, s_2\}$ defined as follows: $\alpha(0) = s_0$, $\alpha(1) = s_1$ and $\alpha(i) = s_2$ for any $i \geq 2$.

Define $\mu = \delta_0$ as the initial distribution in \mathcal{T}_1 . For any $B \subseteq \mathbb{N}$, $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} B) = 1$ and it follows that \mathcal{T}_2 is a μ -sound α -abstraction of \mathcal{T}_1 . It should be noted that it is however not sound when considering $\mu' = \delta_1$ as initial distribution. Indeed, $\text{Prob}_{\mu'}^{\mathcal{T}_1}(\mathbf{F} \{0\}) < 1$ though $\text{Prob}_{\delta_{s_1}}^{\mathcal{T}_2}(\mathbf{F} \{s_0\}) = 1$ (and $\delta_{s_1} = \alpha_\#(\mu')$).



■ **Figure 15** Left, \mathcal{T}_1 a random walk over \mathbb{N} and right, its sound finite abstraction \mathcal{T}_2 .

Consider now the Muller automaton of Section 2 on the left of Figure 2. As stated in Lemma 47, it holds that $\mathcal{T}_2 \times \mathcal{M}$ is an $\alpha_{\mathcal{M}}$ -abstraction of $\mathcal{T}_1 \times \mathcal{M}$ where for each $n \in \mathbb{N}$ and each $q \in Q$, $\alpha_{\mathcal{M}}((n, q)) = (\alpha(n), q)$. Consider $\mu \times \delta_{q_0} = \delta_{(0, q_0)}$ and $B = \{(s_0, q_2)\}$. It then holds that $(\alpha_{\mathcal{M}})_\#(\mu \times \delta_{q_0}) = \delta_{(s_0, q_0)}$ and that $\alpha_{\mathcal{M}}^{-1}(B) = \{(0, q_2)\}$. It is easily observed that starting in state $(0, q_0)$ (resp. (s_0, q_0)) in $\mathcal{T}_1 \times \mathcal{M}$ (resp. $\mathcal{T}_2 \times \mathcal{M}$), then if we visit in the future a state $(0, q)$ (resp. (s_0, q)) we will necessarily get that $q = q_2$. Keeping this in mind, one can see that $\text{Prob}_{\delta_{(s_0, q_0)}}^{\mathcal{T}_2 \times \mathcal{M}}(\mathbf{F} B) = 1$ while

$$\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \alpha_{\mathcal{M}}^{-1}(B)) = \text{Prob}_{\delta_{(1, q_1)}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \alpha_{\mathcal{M}}^{-1}(B)) = \text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} \{0\}) < 1$$

where the first equality holds from Lemma 5 and the second equality holds from Lemma 85. This proves that $\mathcal{T}_2 \times \mathcal{M}$ is not $(\mu \times \delta_{q_0})$ -sound for $\mathcal{T}_1 \times \mathcal{M}$.

Now, observe that \mathcal{T}_1 is decisive w.r.t. any set of states $B \subseteq \mathbb{N}$ from μ as we have seen that $\text{Prob}_\mu^{\mathcal{T}_1}(\mathbf{F} B) = 1$ for any set of states B . It should be noted that \mathcal{T}_1 is not decisive by considering μ' as the initial distribution and $B = \{0\}$. In this case, $\widetilde{\{0\}} = \emptyset$ and thus $\text{Prob}_{\mu'}^{\mathcal{T}_1}(\mathbf{F} \{0\}) \vee \mathbf{F} \widetilde{\{0\}} = \text{Prob}_{\mu'}^{\mathcal{T}_1}(\mathbf{F} \{0\}) < 1$. Consider now $\mathcal{T}_1 \times \mathcal{M}$, we have already shown that $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \{(0, q_2)\}) < 1$. It can be established that $\widetilde{\{(0, q_2)\}} = (2\mathbb{N} + 1) \times \{q_0, q_2\} \cup 2\mathbb{N} \times \{q_1\}$ which are states not reachable from $(0, q_0)$. We deduce that $\text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \{(0, q_2)\}) \vee \mathbf{F} \widetilde{\{(0, q_2)\}} = \text{Prob}_{\mu \times \delta_{q_0}}^{\mathcal{T}_1 \times \mathcal{M}}(\mathbf{F} \{(0, q_2)\}) < 1$. This shows that $\mathcal{T}_1 \times \mathcal{M}$ is not decisive w.r.t. $\{(0, q_2)\}$ from $\mu \times \delta_{q_0}$.

E Technical results of Section 6

E.1 Additional technical results for Subsection 6.1

► **Lemma 97.** For every $\mu \in \text{Dist}(S)$

- (i) $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B \wedge (\neg B \mathbf{U} \widetilde{B})) = 0$;
- (ii) $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B \wedge \mathbf{F} \widetilde{B}) = 0$.

Proof. We first prove point (i). Since B cannot be reached while we are in $\neg B$, it holds that

$$\text{Prob}_\mu^\mathcal{T}(\mathbf{F} B \wedge (\neg B \mathbf{U} \widetilde{B})) = \text{Prob}_\mu^\mathcal{T}(\neg B \mathbf{U} (\widetilde{B} \wedge \mathbf{F} B)).$$

Relaxing the constraint on the until, we get $\text{Prob}_\mu^\mathcal{T}(\neg B \mathbf{U} (\widetilde{B} \wedge \mathbf{F} B)) \leq \text{Prob}_\mu^\mathcal{T}(\mathbf{F} (\widetilde{B} \wedge \mathbf{F} B))$, and the latter is null by definition of \widetilde{B} . This proves the first item.

Point (ii) is straightforward from the definition of \widetilde{B} by observing that $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \wedge \mathbf{F} \widetilde{B}) \leq \text{Prob}_\mu^\mathcal{T}(\mathbf{F} (\widetilde{B} \wedge \mathbf{F} B)) = 0$. \blacktriangleleft

► **Lemma 98.** For every $\mu \in \text{Dist}(S)$, if \mathcal{T} is $\text{PersDec}(\mu, B)$, then $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{B} \wedge \mathbf{F} \widetilde{\widetilde{B}}) = 0$.

Proof. Assume that \mathcal{T} is $\text{PersDec}(\mu, B)$, i.e. for each $p \geq 0$, $\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq p} B \vee \mathbf{F}_{\geq p} \widetilde{B}) = 1$. Towards a contradiction, we suppose that $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{B} \wedge \mathbf{F} \widetilde{\widetilde{B}}) > 0$. Since

$$\text{Ev}_\mathcal{T}(\mathbf{F} \widetilde{B} \wedge \mathbf{F} \widetilde{\widetilde{B}}) = \bigcup_{n \geq 0} \bigcup_{m \geq 0} \text{Ev}_\mathcal{T}(\mathbf{F}_{=n} \widetilde{B}) \cap \text{Ev}_\mathcal{T}(\mathbf{F}_{=m} \widetilde{\widetilde{B}}),$$

we deduce that there are $n, m \geq 0$ such that $\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{=n} \widetilde{B} \wedge \mathbf{F}_{=m} \widetilde{\widetilde{B}}) > 0$. We write e for the event $e = \text{Ev}_\mathcal{T}(\mathbf{F}_{=n} \widetilde{B} \wedge \mathbf{F}_{=m} \widetilde{\widetilde{B}})$. We can show that $\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq n} B \mid e) = 0$ and $\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq m} \widetilde{B} \mid e) = 0$. Indeed we get that:

$$\begin{aligned} \text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq n} B \mid e) &= \frac{\text{Prob}_\mu^\mathcal{T}((\mathbf{F}_{\geq n} B) \wedge e)}{\text{Prob}_\mu^\mathcal{T}(e)} \\ &\leq \frac{\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq n} B \wedge \mathbf{F}_{=n} \widetilde{B})}{\text{Prob}_\mu^\mathcal{T}(e)} \\ &= 0 \end{aligned}$$

from the definition of \widetilde{B} . The equality $\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq m} \widetilde{B} \mid e) = 0$ is proved similarly. Writing $q = \max(m, n)$, it follows that

$$\text{Prob}_\mu^\mathcal{T}(\mathbf{F}_{\geq q} B \vee \mathbf{F}_{\geq q} \widetilde{B} \mid e) = 0.$$

And since $\text{Prob}_\mu^\mathcal{T}(e) > 0$, this contradicts the fact that \mathcal{T} is $\text{PersDec}(\mu, B)$, which concludes the proof. \blacktriangleleft

E.2 Qualitative analysis of simple properties

► **Proposition 99.** Let $\mu \in \text{Dist}(S)$. Then we have the following implications, yielding various characterizations for the qualitative analysis of STSs (under specified assumptions):

Almost-sure reachability

- if $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} B) = 1$ then $\text{Prob}_\mu^\mathcal{T}(\neg B \mathbf{U} \widetilde{B}) = 0$;
- if \mathcal{T} is $\text{Dec}(\mu, B)$ and $\text{Prob}_\mu^\mathcal{T}(\neg B \mathbf{U} \widetilde{B}) = 0$, then $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} B) = 1$.

Almost-sure repeated reachability

- if $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B) = 1$ then $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{B}) = 0$;
- if \mathcal{T} is $\text{StrDec}(\mu, B)$ and $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{B}) = 0$, then $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B) = 1$.

Positive repeated reachability

- if \mathcal{T} is $\text{Dec}(\mu, \widetilde{B})$ and if $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B) > 0$, then $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{B}) > 0$;
- if \mathcal{T} is $\text{PersDec}(\mu, B)$ and if $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{B}) > 0$, then $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B) > 0$.

Proof. We start with almost-sure reachability. We start with the first implication. Since the event $\text{Ev}_\mathcal{T}(\mathbf{F} B)$ is almost-sure, we have

$$\text{Prob}_\mu^\mathcal{T}(\neg B \mathbf{U} \widetilde{B}) = \text{Prob}_\mu^\mathcal{T}((\neg B \mathbf{U} \widetilde{B}) \wedge \mathbf{F} B)$$

and then it is straightforward from point (i) of Lemma 97.

In order to prove the other implication, we need the assumption that \mathcal{T} is $\text{Dec}(\mu, B)$. We have that:

$$\begin{aligned} 1 &= \text{Prob}_\mu^\mathcal{T}(\mathbf{F} B \vee \mathbf{F} \widetilde{B}) = \text{Prob}_\mu^\mathcal{T}(\mathbf{F} B \vee (\neg B \mathbf{U} \widetilde{B})) \quad \text{from Lemma 14 (fifth item)} \\ &= \text{Prob}_\mu^\mathcal{T}(\mathbf{F} B) + \text{Prob}_\mu^\mathcal{T}(\neg B \mathbf{U} \widetilde{B}) \quad \text{from Lemma 97 (point (i))}. \end{aligned}$$

Then from $\text{Prob}_\mu^\mathcal{T}(\neg B \mathbf{U} \widetilde{B}) = 0$, we derive that $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} B) = 1$.

We now consider almost-sure repeated reachability. Since the event $\text{Ev}_\mathcal{T}(\mathbf{G} \mathbf{F} B)$ is almost-sure, we have

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

and then it is straightforward from point (ii) of Lemma 97.

In order to prove the second item, we assume that \mathcal{T} is $\text{StrDec}(\mu, B)$, i.e. $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$. By assumption, the event $\text{Ev}_\mathcal{T}(\mathbf{F} \widetilde{B})$ has probability 0, and thus $\text{Ev}_\mathcal{T}(\mathbf{G} \mathbf{F} B)$ is almost-sure.

We now consider positive repeated reachability. For the first item, we only require \mathcal{T} to be $\text{Dec}(\mu, \widetilde{B})$, that is $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{B} \vee \mathbf{F} \widetilde{B}) = 1$. Since the event $\text{Ev}_\mathcal{T}(\mathbf{F} \widetilde{B} \vee \mathbf{F} \widetilde{B})$ is almost-sure, we derive the equality:

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

Now from point (ii) of Lemma 97, we get that $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \wedge (\mathbf{F} \widetilde{B} \vee \mathbf{F} \widetilde{B})) = \text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \wedge \mathbf{F} \widetilde{B})$. Therefore $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \wedge \mathbf{F} \widetilde{B}) = \text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B) > 0$, and thus $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{B}) > 0$.

Assume now that \mathcal{T} is $\text{PersDec}(\mu, B)$ and that $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{B}) > 0$. Lemma 98 implies that $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} \widetilde{B}) < 1$. Since $\text{PersDec}(\mu, B)$ implies $\text{StrDec}(\mu, B)$, it follows that $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B \vee \mathbf{F} \widetilde{B}) = 1$ and thus, $\text{Prob}_\mu^\mathcal{T}(\mathbf{G} \mathbf{F} B) > 0$. ◀

F Technical results of Section 7

► **Proposition 57** (Approximation scheme for reachability properties). If \mathcal{T} is $\text{Dec}(\mu, B)$, then the two sequences $(p_n^{\text{Yes}})_n$ and $(1 - p_n^{\text{No}})_n$ are adjacent^a and converge to $\text{Prob}_\mu^\mathcal{T}(\mathbf{F} B)$.

^a Recall that two sequences $(a_n)_{n \in \mathbb{N}}$ and $(b_n)_{n \in \mathbb{N}}$ are said *adjacent* if w.l.o.g. (a_n) is non-decreasing, (b_n) is non-increasing and the sequence $(a_n - b_n)_{n \in \mathbb{N}}$ converges to 0.

Proof. We have that:

$$\begin{aligned}
 \lim_{n \rightarrow +\infty} p_n^{\text{Yes}} + p_n^{\text{No}} &= \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B) + \text{Prob}_\mu^{\mathcal{T}}(\neg B \mathbf{U} \widetilde{B}) \\
 &= \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B \vee (\neg B \mathbf{U} \widetilde{B})) \quad \text{from point (i) of Lemma 97} \\
 &= \text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} B \vee \mathbf{F} \widetilde{B}) \quad \text{from Lemma 14 (fifth item)} \\
 &= 1 .
 \end{aligned}$$

The last equality comes from the decisiveness assumption. ◀

► **Proposition 58** (Approximation scheme for repeated reachability). If \mathcal{T} is $\text{PersDec}(\mu, B)$ and $\text{Dec}(\mu, \widetilde{B})$, then the two sequences $(q_n^{\text{Yes}})_n$ and $(1 - q_n^{\text{No}})_n$ are adjacent and converge to $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G} \mathbf{F} B)$.

Proof. Since \mathcal{T} is $\text{Dec}(\mu, \widetilde{B})$, it holds that $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F} \widetilde{B} \vee \mathbf{F} \widetilde{\widetilde{B}}) = 1$. Since \mathcal{T} is $\text{PersDec}(\mu, B)$, one derives from Lemma 98 that

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

We can now show that

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

It comes from the fact that $\text{PersDec}(\mu, B)$ is equivalent to $\text{StrDec}(\mu, B)$ and from point (ii) of Lemma 97. This proves the first part of the corollary.

Finally, we can directly establish from Lemma 98 and from the hypothesis $\text{Dec}(\mu, \widetilde{B})$, that $\lim_{n \rightarrow +\infty} q_n^{\text{Yes}} + q_n^{\text{No}} = 1$. ◀

G

 Technical results of Section 8

In this section, we argue why GSMPs with no cycle of immediate events are almost-surely non-zero. We call *immediate event* a fixed-delay event with delay 0.

► **Lemma 100.** Let $\mathcal{G} = (Q, \mathcal{E}, \ell, u, f, \mathbf{E}, \text{Succ})$ be a GSMP with no cycle with immediate fixed-delay events. Fix $q_0 \in Q$ an initial state, and μ the measure assigning probability 1 to q_0 . Then:

$$\text{Prob}_\mu^{\mathcal{T}_{\mathcal{G}}}(\{\rho \in \text{Paths}(\mathcal{T}_{\mathcal{G}}) \mid \rho \text{ is zero}\}) = 0$$

Sketch. Let $d > 0$ be smaller than any constant appearing in the non-immediate events of \mathcal{G} . There is $\lambda_0 > 0$ such that for every non-immediate event e , $\int_{t=d}^{\infty} f_e(t) dt \geq \lambda_0$.

We consider a non-stochastic interpretation of \mathcal{G} , where delays of events are selected non-deterministically in the supports of the distributions. Pick a finite run ρ that can be generated that way from an initial configuration, and let $\gamma = (q, \nu)$ be its last configuration. In any fireable sequence of transitions $q \xrightarrow{E_1} q_1 \dots \xrightarrow{E_N} q_N$ of length $N > |Q| \cdot |\mathcal{E}|$ from γ , there is an event which is newly enabled along that sequence, and there is $1 \leq i < k \leq N$ with $e \in \mathbf{E}(q_j)$ for every $i \leq j < k$ and $e \in E_k$.

Towards a contradiction assume it is not the case, then this means that each event e in E_i for some $1 \leq i \leq N$ is either an immediate event or an event in $\mathbf{E}(q) \cap \bigcap_{j < i} \mathbf{E}(q_j) \cap \bigcap_{j < i} E_j^c$ (that is, e was already enabled in q , it is fired by E_i , and was not disabled inbetween). There can be at most $|\mathbf{E}(q)| \leq |\mathcal{E}|$ such events which are not immediate. Furthermore, by

assumption, there is no cycle with only immediate events. Hence as soon as $N > |Q| \cdot |\mathcal{E}|$, this is not possible. Hence, this implies the above claim.

Hence with probability lower-bounded by λ_0 , the duration of a continuation of ρ along that sequence of edges will be larger than d . Hence, providing more details here, we deduce that almost-surely, runs will diverge. ◀

