

Finite Abstractions of Stochastic Max-Plus-Linear Systems*

Dieky Adzkiya¹, Sadegh Esmail Zadeh Soudjani¹, and Alessandro Abate^{2,1}

¹ Delft Center for Systems and Control
TU Delft, Delft University of Technology, The Netherlands
{D.Adzkiya,S.EsmailZadehSoudjani}@tudelft.nl

² Department of Computer Science
University of Oxford, United Kingdom
Alessandro.Abate@cs.ox.ac.uk

Abstract. This work investigates the use of finite abstractions to study the finite-horizon probabilistic invariance problem over Stochastic Max-Plus-Linear (SMPL) systems. SMPL systems are probabilistic extensions of discrete-event MPL systems that are widely employed in the engineering practice for timing and synchronisation studies. We construct finite abstractions by re-formulating the SMPL system as a discrete-time Markov process, then tailoring formal abstraction techniques in the literature to generate a finite-state Markov Chain (MC), together with precise guarantees on the level of the introduced approximation. This finally allows to probabilistically model check the obtained MC against the finite-horizon probabilistic invariance specification. The approach is practically implemented via a dedicated software, and elucidated in this work over numerical examples.

Keywords: Max-plus-linear systems, Max-plus algebra, Discrete-time stochastic processes, Continuous state spaces, Abstractions, Approximate probabilistic bisimulations.

1 Introduction

Max-Plus-Linear (MPL) systems are a class of discrete-event systems [1, 2] with a continuous state space characterising the timing of the underlying sequential discrete events. MPL systems are predisposed to describe the timing synchronisation between interleaved processes, under the assumption that timing events are dependent linearly (within the max-plus algebra) on previous event occurrences. MPL systems are widely employed in the analysis and scheduling of infrastructure networks, such as communication and railway systems [3], production and manufacturing lines [4, 5], or biological systems [6].

* The first two authors have equally contributed to this work. This work has been supported by the European Commission via STREP project MoVeS 257005, Marie Curie grant MANTRAS 249295, and IAPP project AMBI 324432.

Stochastic Max-Plus-Linear (SMPL) systems [7–9] are MPL systems where the delays between successive events (in the examples above, the processing or transportation times) are now characterised by random quantities. In practical applications SMPL systems are more realistic than simple MPL ones: for instance in a model for a railway network, train running times depend on driver behaviour, on weather conditions, and on passenger numbers at stations: they can arguably be more suitably modelled by random variables.

Only a few approaches have been developed in the literature to study the steady-state behaviour of SMPL systems, for example employing Lyapunov exponents and asymptotic growth rates [10–15]. The Lyapunov exponent of an SMPL system is analogous to the max-plus eigenvalue for an autonomous MPL system. The Lyapunov exponent of SMPL systems under some assumptions has been studied in [10], and later extended to approximate computations under other technical assumptions in [11, p. 251]. The application of model predictive control and system identification to SMPL systems is studied in [16, 17]. In contrast, our work focuses on one-step properties of SMPL systems and is based on developing finite-state abstractions: this is parallel to the approach in [18] for (deterministic) MPL systems. To the best of our knowledge, this contribution represents the first work on finite-state abstractions of SMPL systems.

Verification techniques and tools for deterministic, discrete-time, finite-state systems have been widely investigated and developed in the past decades [19]. The application of formal methods to stochastic models is typically limited to discrete-state structures, either in continuous or in discrete time [20, 21]. Continuous-space models on the other hand require the use of finite abstractions, as it is classically done for example with finite bisimulations of timed automata, which can be computed via the known region construction [22]. With focus on stochastic models, numerical schemes based on Markov Chain (MC) approximations of stochastic systems have been introduced in [23, 24], and applied to the approximate study of probabilistic reachability or invariance in [25, 26], however these finite abstractions do not come with explicit error bounds. On the contrary in [27], a technique has been introduced to instead provide formal abstractions of discrete-time, continuous-space Markov models [28], with the objective of investigating their probabilistic invariance by employing probabilistic model checking over a finite MC. In view of scalability and of generality, the approach has been improved and optimised in [29]. Interestingly the procedure has been shown [30] to introduce an approximate probabilistic bisimulation of the concrete model [31].

The aim of this work is to characterise and to compute the approximate solution of the finite-time invariance problem over SMPL systems: more precisely, for any allowable initial event time, we determine the probability that the time associated to the occurrence of N consecutive events will remain close to a given deterministic N -step schedule (cf. Section 2.2). The probabilistic invariance problem can be regarded as the dual of a reachability problem [28], and can be computed by constructing finite abstractions of the SMPL system,

which are quantifiably close to the concrete model [27]. More precisely, our approach works as follows. We first formulate the given SMPL system as a discrete-time Markov process, as suggested by [8, 9]. Then we adapt the techniques in [27, 29] to the structure of the SMPL system, in order to generate a finite-state MC, together with guarantees on the level of approximation introduced in the process. The invariance property over the obtained MC can then be analysed via probabilistic model checking [20] and computed by existing software [32, 33]. The result obtained from the model checking software is then combined with the approximation guarantees, in order to provide an overall assessment of the probability that the concrete SMPL system satisfies the given property.

The article is structured as follows. Initially, Section 2.1 introduces the SMPL formalism, whereas Section 2.2 presents the probabilistic invariance problem. Section 3 discusses the formal abstraction of an SMPL system as an MC. Furthermore, with focus on the probabilistic invariance problem, the quantification of the abstraction error and some numerical examples are presented in Section 4. Finally, Section 5 concludes the presentation of this work.

2 Preliminaries

This section introduces the basics of max-plus algebra and of autonomous SMPL systems, and discusses the probabilistic invariance problem, which is to be further elaborated throughout the paper.

2.1 Modelling: Stochastic Max-Plus-Linear Systems

The notations \mathbb{N} and \mathbb{N}_n represent the whole positive integers $\{1, 2, \dots\}$ and the first n positive integers $\{1, 2, \dots, n\}$, respectively. We use the bold letters for vectors and usual letters with the same name and index for the elements of the vector, for instance $\mathbf{x} = [x_1, \dots, x_n]^T$. Furthermore we define \mathbb{R}_ε , ε and e respectively as $\mathbb{R} \cup \{\varepsilon\}$, $-\infty$ and 0 . For $\alpha, \beta \in \mathbb{R}_\varepsilon$, introduce the two operations

$$\alpha \oplus \beta = \max\{\alpha, \beta\} \quad \text{and} \quad \alpha \otimes \beta = \alpha + \beta ,$$

where the element ε is considered to be absorbing w.r.t. \otimes [12, Definition 3.4], namely $\alpha \otimes \varepsilon = \varepsilon$ for all $\alpha \in \mathbb{R}_\varepsilon$. The rules for the order of evaluation of the max-algebraic operators correspond to those in the conventional algebra: max-algebraic multiplication has a higher precedence than max-algebraic addition [12, Sect. 3.1].

The basic max-algebraic operations are extended to matrices as follows. If $A, B \in \mathbb{R}_\varepsilon^{m \times n}$; $C \in \mathbb{R}_\varepsilon^{m \times p}$; $D \in \mathbb{R}_\varepsilon^{p \times n}$; and $\alpha \in \mathbb{R}_\varepsilon$, then

$$[\alpha \otimes A]_{ij} = \alpha \otimes A_{ij} , \quad [A \oplus B]_{ij} = A_{ij} \oplus B_{ij} , \quad [C \otimes D]_{ij} = \bigoplus_{k=1}^p C_{ik} \otimes D_{kj} ,$$

for each $i \in \mathbb{N}_m$ and $j \in \mathbb{N}_n$. Notice the analogy between \oplus , \otimes and respectively $+$, \times for matrix and vector operations in the conventional algebra. In this paper, the following notation is adopted for reasons of convenience. A vector with each component being equal to 0 (resp., $-\infty$) is also denoted by e (resp., ε). Furthermore, for practical reasons, the state space is taken to be \mathbb{R}^n (rather than \mathbb{R}_ε^n).

An autonomous SMPL system is defined as:

$$\mathbf{x}(k+1) = A(k) \otimes \mathbf{x}(k) , \quad (1)$$

where $\mathbf{x}(k) = [x_1(k), \dots, x_n(k)]^T \in \mathbb{R}^n$; $\{A_{ij}(\cdot)\}$ are discrete-time stationary random processes¹ taking values in \mathbb{R}_ε ; further $A_{ij}(k)$ are independent for all $k \in \mathbb{N} \cup \{0\}$ and $i, j \in \mathbb{N}_n$. We assume each random variable has fixed support [7, Definition 1.4.1], i.e. the probability of ε is either 0 or 1. The random sequence $\{A_{ij}(\cdot)\}$ is then characterised by a given density function $t_{ij}(\cdot)$ and corresponding distribution function $T_{ij}(\cdot)$ (cf. Theorem 1). The independent variable k denotes an increasing occurrence index, whereas the state variable $\mathbf{x}(k)$ defines the (continuous) time of the k -th occurrence of the discrete events. The state component $x_i(k)$ denotes the time of the k -th occurrence of the i -th event. Since this article is based exclusively on autonomous (that is, not non-deterministic) SMPL systems, the adjective will be dropped for simplicity.

Example 1. Consider the following SMPL system representing a simple railway network between two connected stations. The state variables $x_i(k)$ for $i = 1, 2$ denote the time of the k -th departure at station i :

$$\begin{aligned} \mathbf{x}(k+1) = A(k) \otimes \mathbf{x}(k), \quad A(k) &= \begin{bmatrix} 2 + e_{11}(k) & 5 + e_{12}(k) \\ 3 + e_{21}(k) & 3 + e_{22}(k) \end{bmatrix} \quad \text{or equivalently,} \\ \begin{bmatrix} x_1(k+1) \\ x_2(k+1) \end{bmatrix} &= \begin{bmatrix} \max\{2 + e_{11}(k) + x_1(k), 5 + e_{12}(k) + x_2(k)\} \\ \max\{3 + e_{21}(k) + x_1(k), 3 + e_{22}(k) + x_2(k)\} \end{bmatrix} , \end{aligned}$$

where we have assumed that $e_{11}(\cdot) \sim \text{Exp}(1)$, $e_{12}(\cdot) \sim \text{Exp}(2/5)$, $e_{21}(\cdot) \sim \text{Exp}(2/3)$, and $e_{22}(\cdot) \sim \text{Exp}(2/3)$, and $\text{Exp}(\lambda)$ represents the exponential distribution with rate λ . Notice that $A_{ij}(\cdot)$ denotes the traveling time from station j to station i and amounts to a deterministic constant plus a delay modelled by the random variable $e_{ij}(\cdot)$. A few sample trajectories of the SMPL system, initialised at $\mathbf{x}(0) = [1, 0]^T$, are displayed in Figure 1. Note that when all random delays are assumed to be equal to zero, the above deterministic system admits the unique solution $\mathbf{x}(k) = \mathbf{x}(0) + dk = [1 + 4k, 4k]^T$, where $d = 4$ is the max-plus eigenvalue of matrix A , and $[1, 0]^T$ is the corresponding eigenvector of the deterministic MPL system [12, Sect. 3.7]. Such a periodic trajectory can be used as a regular schedule for the train departures (cf. Section 2.2). \square

¹ Notice that, for deterministic MPL systems, matrix A is instead given and time-invariant.




Fig. 1. Sample trajectories of the SMPL system in Example 1 for 50 discrete steps (horizontal axis) and both coordinates (vertical axis)

2.2 Problem: Probabilistic Invariance

Let us consider events that are scheduled to occur regularly, that is let us select a time between consecutive events that is a positive given constant, say d . We call this a *regular schedule* and assume that it does not affect the time of occurrence of all events, e.g. any event may occur ahead of the regular schedule. In this work, we consider an N -step finite-horizon probabilistic invariance problem w.r.t. a regular schedule: more specifically, for each possible time of initial occurrence of all events $(x_i(0), i \in \mathbb{N}_n)$, we are interested in determining the probability that the time of k -th occurrence of all events $(\mathbf{x}(k))$ remains close to the corresponding time of the regular schedule, for $k \in \mathbb{N}_N \cup \{0\}$. For instance, we may want to determine the probability that the time of occurrence of all events is at least 5 time units ahead of the given regular schedule, as well as at most 5 time units behind it. The safe set is then defined as the desired time of occurrence of all events w.r.t. the regular schedule.

The techniques in [27, 29], developed to provide the characterisation and the computation of the quantity of interest over general Markov processes, can be directly applied to the SMPL system (1). However, in order to prevent the growth of the safe set as the event horizon N increases (which in general leads to a decrease in computational performance), we reformulate the SMPL system based on the given regular schedule, so that a fixed safe set is obtained. Since we are interested in the delay of event occurrences with respect to the given schedule, we introduce new variables defined as the difference between the states of the original SMPL system and the regular schedule. More precisely, first we define a vector $\mathbf{s}(\cdot)$ that characterises the regular schedule. The dynamics of $\mathbf{s}(\cdot)$ are determined by the time duration $d \in \mathbb{R}$ between consecutive events² and the arbitrary initial condition $\mathbf{s}(0) \in \mathbb{R}^n$, i.e. $\mathbf{s}(k+1) = d \otimes \mathbf{s}(k)$. As mentioned, new

² Our results can be generalised to event-dependent time durations. In this case the Markov process becomes inhomogeneous, which will greatly increase the computational complexity of the procedure.

states $\mathbf{z}(\cdot)$ are defined as the difference between the states of the original SMPL system (1) and the regular schedule $\mathbf{s}(\cdot)$, i.e. $\mathbf{z}(k) = \mathbf{x}(k) - \mathbf{s}(k)$ for $k \in \mathbb{N} \cup \{0\}$. The dynamics of the newly introduced SMPL system are then given by

$$\mathbf{z}(k+1) = [A(k) + D] \otimes \mathbf{z}(k) , \quad (2)$$

where $D = [d_{ij}]_{i,j} \in \mathbb{R}^{n \times n}$ (i.e. d_{ij} is the entry of matrix D at row i and column j), $d_{ij} = s_j(0) - s_i(0) - d$, and $\mathbf{z}(k) = [z_1(k), \dots, z_n(k)]^T \in \mathbb{R}^n$. Notice that $A_{ij}(k) \otimes d_{ij}$ are independent for all $k \in \mathbb{N} \cup \{0\}$ and $i, j \in \mathbb{N}_n$. The density (resp., distribution) function of $A_{ij}(k) \otimes d_{ij}$ corresponds to the density (resp., distribution) function of $A_{ij}(k)$ shifted forward of d_{ij} units. The independent variable k again denotes an increasing occurrence index, whereas the state variable $\mathbf{z}(k)$ defines the delay w.r.t. the schedule of k -th occurrence of all events: in particular the state component $z_i(k)$ denotes the delay w.r.t. the schedule of k -th occurrence of the i -th event. Notice that if the delay is negative then the event occurs ahead of schedule, whereas if the delay is positive then the event occurs behind schedule. The next theorem shows that, much like the original model in (1), the new SMPL system can be described as a discrete-time homogeneous Markov process.

Theorem 1. *The SMPL system in (2) is fully characterised by the following conditional density function*

$$t_z(\bar{\mathbf{z}}|\mathbf{z}) = \prod_{i=1}^n t_i(\bar{z}_i|\mathbf{z}) \quad \text{where} \quad (3)$$

$$t_i(\bar{z}_i|\mathbf{z}) = \sum_{j=1}^n \left[t_{ij}(\bar{z}_i - d_{ij} - z_j) \prod_{k=1, k \neq j}^n T_{ik}(\bar{z}_i - d_{ik} - z_k) \right] \quad i \in \mathbb{N}_n .$$

Employing the introduced SMPL system (2), the problem can be formulated as the following N -step invariance probability

$$P_{z_0}(\mathcal{A}) = \Pr\{\mathbf{z}(k) \in \mathcal{A} \text{ for all } k \in \mathbb{N}_N \cup \{0\} | \mathbf{z}(0) = \mathbf{z}_0\} ,$$

where \mathcal{A} is called the safe set and is assumed to be Borel measurable. The next proposition provides a theoretical framework to study the problem.

Proposition 1 ([28, Lemma 1]). *Consider value functions $V_k : \mathbb{R}^n \rightarrow [0, 1]$, for $k \in \mathbb{N}_N \cup \{0\}$, computed through the following backward recursion:*

$$V_k(\mathbf{z}) = \mathbb{1}_{\mathcal{A}}(\mathbf{z}) \int_{\mathcal{A}} V_{k+1}(\bar{\mathbf{z}}) t_z(\bar{\mathbf{z}}|\mathbf{z}) d\bar{\mathbf{z}} \quad \text{for all } \mathbf{z} \in \mathbb{R}^n ,$$

initialised with $V_N(\mathbf{z}) = \mathbb{1}_{\mathcal{A}}(\mathbf{z})$ for all $\mathbf{z} \in \mathbb{R}^n$. Then $P_{z_0}(\mathcal{A}) = V_0(\mathbf{z}_0)$.

For any $k \in \mathbb{N}_N \cup \{0\}$, notice that $V_k(\mathbf{z})$ represents the probability that an execution of the SMPL system (2) remains within the safe set \mathcal{A} over the residual event horizon $\{k, \dots, N\}$, starting from \mathbf{z} at event step k . This result characterises the finite-horizon probabilistic invariance problem as a dynamic programming problem. Since an explicit analytical solution to the problem is generally impossible to be found, we leverage the techniques developed in [27, 29] to provide a numerical computation with exact associated error bounds. This is elaborated in the next section.

3 Abstraction by a Finite State Markov Chain

We tailor the abstraction procedure presented in [27, Sect. 3.1] towards the goal of generating a finite-state MC (\mathcal{P}, T_p) from a given SMPL system and a safe set \mathcal{A} , and employ it to approximately compute the probabilistic invariance of interest.

Let $\mathcal{P} = \{\phi_1, \dots, \phi_{m+1}\}$ be a set of finitely many discrete states and $T_p : \mathcal{P} \times \mathcal{P} \rightarrow [0, 1]$ a related transition probability matrix, such that $T_p(\phi_i, \phi_j)$ characterises the probability of transitioning from state ϕ_i to state ϕ_j and thus induces a conditional discrete probability distribution over the finite space \mathcal{P} . Given a safe set \mathcal{A} , Algorithm 1 provides a procedure to abstract an SMPL system by a finite-state MC. The set $\mathcal{A}_p = \{\phi_1, \dots, \phi_m\}$ denotes the discrete safe set. In Algorithm 1, $\Xi : \mathcal{A}_p \rightarrow 2^{\mathcal{A}}$ represents the concretisation function, i.e. a set-valued map that associates to any discrete state (point) $\phi_i \in \mathcal{A}_p$ the corresponding continuous partition set $\mathcal{A}_i \subset \mathcal{A}$. Furthermore the abstraction function $\xi : \mathcal{A} \rightarrow \mathcal{A}_p$ associates to any point $z \in \mathcal{A}$ on the SMPL state space, the corresponding discrete state in \mathcal{A}_p . Additionally, notice that an absorbing discrete state ϕ_{m+1} is added to the state space of the MC in order to render the transition probability matrix T_p stochastic: the absorbing discrete state ϕ_{m+1} represents the complement of the safe set \mathcal{A} for the SMPL system, namely $\mathbb{R}^n \setminus \mathcal{A}$, and accounts for the associated dynamics.

Algorithm 1. Generation of a finite-state MC from an SMPL system and a safe set

Input: An SMPL system in (2) and a safe set \mathcal{A}

Output: A finite-state MC (\mathcal{P}, T_p)

1. Select a finite partition of set \mathcal{A} of cardinality m , as $\mathcal{A} = \cup_{i=1}^m \mathcal{A}_i$
2. For each \mathcal{A}_i , select a single representative point $z_i \in \mathcal{A}_i$
3. Define $\mathcal{A}_p = \{\phi_i, i \in \mathbb{N}_m\}$ and take $\mathcal{P} = \mathcal{A}_p \cup \{\phi_{m+1}\}$ as the finite state-space of the MC (ϕ_{m+1} is an absorbing state, as explained in the text)
4. Compute the transition probability matrix T_p as

$$T_p(\phi_i, \phi_j) = \begin{cases} \int_{\Xi(\phi_j)} t_z(\bar{z}|z_i) d\bar{z} , & \text{if } 1 \leq j \leq m \text{ and } 1 \leq i \leq m , \\ 1 - \sum_{\bar{\phi} \in \mathcal{A}_p} \int_{\Xi(\bar{\phi})} t_z(\bar{z}|z_i) d\bar{z} , & \text{if } j = m+1 \text{ and } 1 \leq i \leq m , \\ 1 , & \text{if } j = i = m+1 , \\ 0 , & \text{if } 1 \leq j \leq m \text{ and } i = m+1 , \end{cases}$$

Remark 1. The bottleneck of Algorithm 1 lies in the computation of transition probability matrix T_p , due to the integration of kernel t_z . This integration can be circumvented if the distribution functions $T_{ij}(\cdot)$ for all $i, j \in \mathbb{N}_n$ have explicit analytical form, e.g. an exponential distribution.

The procedure in Algorithm 1 has been shown [30] to introduce an approximate probabilistic bisimulation of the concrete model [31].

Algorithm 1 can be applied to abstract an SMPL system as a finite-state MC, regardless of the particular safe set \mathcal{A} . However the quantification of the abstraction error in Section 4 requires that the safe set \mathcal{A} is bounded. \square

Considering the obtained finite-state, discrete-time MC (\mathcal{P}, T_p) and the discretised safe set $\mathcal{A}_p \subset \mathcal{P}$, the probabilistic invariance problem amounts to evaluating the probability that a finite execution associated with the initial condition $\phi_0 \in \mathcal{P}$ remains within the discrete safe set \mathcal{A}_p during the given event horizon. This can be stated as following probability:

$$p_{\phi_0}(\mathcal{A}_p) = \Pr\{\phi(k) \in \mathcal{A}_p \text{ for all } k \in \mathbb{N}_N \cup \{0\} | \phi(0) = \phi_0\} ,$$

where $\phi(k)$ denotes the discrete state of the MC at step k .

The solution of this finite-horizon probabilistic invariance problem over the MC abstraction can be determined via a discrete version of Proposition 1.

Proposition 2. *Consider value functions $V_k^p : \mathcal{P} \rightarrow [0, 1]$, for $k \in \mathbb{N}_N \cup \{0\}$, computed through the following backward recursion:*

$$V_k^p(\phi) = \mathbb{1}_{\mathcal{A}_p}(\phi) \sum_{\bar{\phi} \in \mathcal{P}} V_{k+1}^p(\bar{\phi}) T_p(\phi, \bar{\phi}) \quad \text{for all } \phi \in \mathcal{P} ,$$

initialised with $V_N^p(\phi) = \mathbb{1}_{\mathcal{A}_p}(\phi)$ for all $\phi \in \mathcal{P}$. Then $p_{\phi_0}(\mathcal{A}_p) = V_0^p(\phi_0)$.

For any $k \in \mathbb{N}_N \cup \{0\}$, notice that $V_k^p(\phi)$ represents the probability that an execution of the finite-state MC remains within the discrete safe set \mathcal{A}_p over the residual event horizon $\{k, \dots, N\}$, starting from ϕ at event step k . The quantities in Proposition 2 can be easily computed via linear algebra. It is of interest to provide a quantitative comparison between the discrete outcome obtained by Proposition 2 and the continuous solution that results from Proposition 1: in other words, we are interested in deriving bounds on the abstraction error. The following section accomplishes this goal.

4 Quantification of the Abstraction Error

This section starts by precisely defining the error related to the abstraction procedure, which is due to the approximation of a continuous concrete model with a finite discrete one. Then a bound of the approximation error in [29] is recalled, and applied to the probabilistic invariance problem under some structural assumptions, namely in the case of Lipschitz continuous density functions, or alternatively of piecewise Lipschitz continuous density functions.

The approximation error is defined as the maximum difference between the outcomes obtained by Propositions 1 and 2 for any pair of initial conditions $\mathbf{z}_0 \in \mathcal{A}$ and $\xi(\mathbf{z}_0) \in \mathcal{A}_p$. Since an exact computation of this error is not possible in general, we resort to determining an upper bound of the approximation error, which is denoted as E . More formally, we are interested in quantifying E that satisfies

$$|P_{\mathbf{z}_0}(\mathcal{A}) - p_{\xi(\mathbf{z}_0)}(\mathcal{A}_p)| \leq E \quad \text{for all } \mathbf{z}_0 \in \mathcal{A} . \quad (4)$$

We raise the following assumption on the SMPL system. Recall that the density function of $A_{ij}(k) \otimes d_{ij}$ in (2) corresponds to the density function of $A_{ij}(k)$ in (1) shifted d_{ij} units forward.

Assumption 3. *The density functions $t_{ij}(\cdot)$ for $i, j \in \mathbb{N}_n$ are bounded:*

$$t_{ij}(z) \leq M_{ij} \quad \text{for all } z \in \mathbb{R} .$$

Assumption 3 implies the distribution functions $T_{ij}(\cdot)$ for $i, j \in \mathbb{N}_n$ are Lipschitz continuous. Recall that the (global) Lipschitz constant of a one-dimensional function can be computed as the maximum of the absolute value of the first derivative of the function. Thus

$$|T_{ij}(z) - T_{ij}(z')| \leq M_{ij}|z - z'| \quad \text{for all } z, z' \in \mathbb{R} .$$

For computation of the bound on approximation error, we use the following result based on [29], which has inspired most of this work.

Proposition 4 ([29, pp. 933-934]). *Suppose Assumption 3 holds and the density function $t_z(\bar{z}|\mathbf{z})$ satisfies the condition*

$$\int_{\mathcal{A}} |t_z(\bar{z}|\mathbf{z}) - t_z(\bar{z}|\mathbf{z}')| d\bar{z} \leq H\|\mathbf{z} - \mathbf{z}'\| \quad \text{for all } \mathbf{z}, \mathbf{z}' \in \mathcal{A} ,$$

then an upper bound on the approximation error in (4) is $E = NH\delta$, where N is the event horizon and δ is the partition diameter.

The partition diameter δ is defined in [27, Sect. 3.1]. We first determine the constant H for Lipschitz continuous density functions, then generalise the result to piecewise Lipschitz continuous density functions.

4.1 Lipschitz Continuous Density Functions

Assumption 5. *The density functions $t_{ij}(\cdot)$ for $i, j \in \mathbb{N}_n$ are Lipschitz continuous, namely there exist finite and positive constants h_{ij} , such that*

$$|t_{ij}(z) - t_{ij}(z')| \leq h_{ij}|z - z'| \quad \text{for all } z, z' \in \mathbb{R} .$$

Under Assumptions 3 and 5, the conditional density function $t_z(\bar{z}|\mathbf{z})$ is Lipschitz continuous. This opens up the application of the results in [27, 29] for the approximate solution of the probabilistic invariance problem. Notice that the Lipschitz constant of $t_z(\bar{z}|\mathbf{z})$ may be large, which implies a rather conservative upper bound on the approximation error. To improve this bound, we can instead directly use Proposition 4 presented before – an option also discussed in [29]. In particular we present three technical lemmas that are essential for the computation of the constant H . After the derivation of the improved bound, the obtained results are applied to a numerical example.

Lemma 1. *Any one-dimensional continuous distribution function $T(\cdot)$ satisfies the inequality*

$$\int_{\mathbb{R}} |T(\bar{z} - z) - T(\bar{z} - z')| d\bar{z} \leq |z - z'| \quad \text{for all } z, z' \in \mathbb{R} .$$

Lemma 2. *Suppose the random vector \bar{z} can be organised as $\bar{z} = [\bar{z}_1^T, \bar{z}_2^T]^T$, so that its conditional density function is the multiplication of conditional density functions of \bar{z}_1, \bar{z}_2 as:*

$$f(\bar{z}|\mathbf{z}) = f_1(\bar{z}_1|\mathbf{z})f_2(\bar{z}_2|\mathbf{z}) .$$

Then for a given set $\mathcal{A} \in \mathcal{B}(\mathbb{R}^n)$ it holds that

$$\int_{\mathcal{A}} |f(\bar{z}|\mathbf{z}) - f(\bar{z}|\mathbf{z}')| d\bar{z} \leq \sum_{i=1}^2 \int_{\Pi_i(\mathcal{A})} |f_i(\bar{z}_i|\mathbf{z}) - f_i(\bar{z}_i|\mathbf{z}')| d\bar{z}_i ,$$

where $\Pi_i(\cdot)$ represents the projection operator on the i -th axis.

Lemma 3. *Suppose the vector \mathbf{z} can be organised as $\mathbf{z} = [\mathbf{z}_1^T, \mathbf{z}_2^T]^T$, and that the density function of the conditional random variable $(\bar{z}|\mathbf{z})$ is of the form*

$$f(\bar{z}|\mathbf{z}) = f_1(\bar{z}, \mathbf{z}_1)f_2(\bar{z}, \mathbf{z}_2) ,$$

where $f_1(\bar{z}, \mathbf{z}_1), f_2(\bar{z}, \mathbf{z}_2)$ are bounded non-negative functions with $M_1 = \sup f_1(\bar{z}, \mathbf{z}_1)$ and $M_2 = \sup f_2(\bar{z}, \mathbf{z}_2)$. Then for a given set $\mathcal{C} \in \mathcal{B}(\mathbb{R})$:

$$\begin{aligned} & \int_{\mathcal{C}} |f(\bar{z}|\mathbf{z}_1, \mathbf{z}_2) - f(\bar{z}|\mathbf{z}'_1, \mathbf{z}'_2)| d\bar{z} \\ & \leq M_2 \int_{\mathcal{C}} |f_1(\bar{z}, \mathbf{z}_1) - f_1(\bar{z}, \mathbf{z}'_1)| d\bar{z} + M_1 \int_{\mathcal{C}} |f_2(\bar{z}, \mathbf{z}_2) - f_2(\bar{z}, \mathbf{z}'_2)| d\bar{z} . \end{aligned}$$

Theorem 2. *Under Assumptions 3 and 5, the constant H in Proposition 4 is*

$$H = \sum_{i,j=1}^n H_{ij} + (n-1)M_{ij} ,$$

where $H_{ij} = \mathcal{L}_i h_{ij}$, and where the constant $\mathcal{L}_i = \mathcal{L}(\Pi_i(\mathcal{A}))$ is the Lebesgue measure of the projection of the safe set onto the i -th axis.

We now elucidate the above results on a case study, and select a beta distribution to characterise delays. A motivation for employing a beta distribution is that its density function has bounded support. Thus by scaling and shifting the density function, we can construct a distribution taking positive real values within an interval. Recall that this distribution is used to model processing or transportation time, and as such it can only take positive values. Furthermore, the beta distribution can be used to approximate the normal distribution with arbitrary accuracy.

Definition 1 (Beta Distribution). *The general formula for the density function of the beta distribution is*

$$t(x; \alpha, \beta, a, b) = \frac{(x-a)^{\alpha-1}(b-x)^{\beta-1}}{B(\alpha, \beta)(b-a)^{\alpha+\beta-1}} \quad \text{if } a \leq x \leq b ,$$

and 0 otherwise, where $\alpha, \beta > 0$ are the shape parameters; $[a, b]$ is the support of the density function; and $B(\cdot, \cdot)$ is the beta function. A random variable X characterised by this distribution is denoted by $X \sim \text{Beta}(\alpha, \beta, a, b)$.

The case where $a = 0$ and $b = 1$ is called the standard beta distribution. Let us remark that the density function of the beta distribution is unbounded if any of the shape parameters belongs to the interval $(1, 2)$. We remark that if the shape parameters are positive integers, the beta distribution has a piecewise polynomial density function, which has been used for system identification of SMPPL systems in [17, Sect. 4.3].

Example 2. We apply the results in Theorem 2 to the following two-dimensional SMPPL system (1), where $A_{ij}(\cdot) \sim \text{Beta}(\alpha_{ij}, \beta_{ij}, a_{ij}, b_{ij})$,

$$\begin{bmatrix} \alpha_{11} & \alpha_{12} \\ \alpha_{21} & \alpha_{22} \end{bmatrix} = \begin{bmatrix} 2 & 4 \\ 2 & 2 \end{bmatrix}, \quad \begin{bmatrix} \beta_{11} & \beta_{12} \\ \beta_{21} & \beta_{22} \end{bmatrix} = \begin{bmatrix} 5 & 2 \\ 2 & 4 \end{bmatrix}, \quad \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} = \begin{bmatrix} 0 & 2 \\ 2 & 0 \end{bmatrix}, \quad \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} = \begin{bmatrix} 7 & 6.5 \\ 4 & 9 \end{bmatrix}.$$

Skipping the details of the direct calculations, the supremum and the Lipschitz constant of the density functions are respectively

$$\begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix} = \begin{bmatrix} 1536/4375 & 15/32 \\ 3/4 & 15/64 \end{bmatrix}, \quad \begin{bmatrix} h_{11} & h_{12} \\ h_{21} & h_{22} \end{bmatrix} = \begin{bmatrix} 30/49 & 80/81 \\ 3/2 & 20/81 \end{bmatrix}.$$

Considering a regular schedule with $\mathbf{s}(0) = [0, 0]^T$ and $d = 4$, selecting safe set $\mathcal{A} = [-5, 5]^2$, and event horizon $N = 5$, according to Theorem 2 we obtain an error $E = 176.4\delta$. In order to obtain an approximation error bounded by $E = 0.1$, we would need to discretise set \mathcal{A} uniformly with 24942 bins per each dimension (step 1 of Algorithm 1). The representative points have been selected at the centre of the squares obtained by uniform discretisation (step 2). The obtained finite-state MC has $24942^2 + 1$ discrete states (step 3). The procedure to construct transition probability matrix (step 4) works as follows. For each $i, j \in \{1, \dots, 24942^2 + 1\}$, we compute $T_p(\phi_i, \phi_j)$ which consists of four possible cases. If $1 \leq i, j \leq 24942^2$, then $T_p(\phi_i, \phi_j)$ is defined as the probability of transitioning from the i -th representative point \mathbf{z}_i to the j -th partition set \mathcal{A}_j . If $1 \leq i \leq 24942^2$ and $j = 24942^2 + 1$, then $T_p(\phi_i, \phi_j)$ is defined as the probability of transitioning from the i -th representative point \mathbf{z}_i to the complement of the safe set $\mathbb{R}^n \setminus \mathcal{A}$. Since the discrete state ϕ_{24942^2+1} is absorbing, then $T_p(\phi_{24942^2+1}, \phi_j) = 1$ if $j = 24942^2 + 1$, and is equal to 0 otherwise. The solution of the invariance problem obtained over the abstract model (cf. Proposition 2) is computed via the software tool FAUST² [34] and is depicted in Figure 2 (left panel). \square

4.2 Piecewise Lipschitz Continuous Density Functions

It is clear that the structural assumptions raised in the previous section pose limitations on the applicability of the ensuing results. For the sake of generality, we now extend the previous results to the more general case encompassed by the following requirement.

Assumption 6. *The density functions $t_{ij}(\cdot)$ for $i, j \in \mathbb{N}_n$ are piecewise Lipschitz continuous, namely there exist partitions $\mathbb{R} = \cup_{k=1}^{m_{ij}} D_{ij}^k$ and corresponding finite and positive constants h_{ij}^k , such that*

$$\begin{aligned} t_{ij}(z) &= \sum_{k=1}^{m_{ij}} t_{ij}^k(z) \mathbb{1}_{D_{ij}^k}(z) && \text{for all } z \in \mathbb{R} \text{ ,} \\ |t_{ij}^k(z) - t_{ij}^k(z')| &\leq h_{ij}^k |z - z'| && \text{for all } k \in \mathbb{N}_{m_{ij}} \text{ and } z, z' \in D_{ij}^k \text{ .} \end{aligned}$$

The notation k used in Assumption 6 is not a power and is not an occurrence index (1), but it denotes the index of a set in the partition of cardinality $\sum_{i,j} m_{ij}$. Notice that if Assumption 6 holds and the density functions are Lipschitz continuous, then Assumption 5 is automatically satisfied with $h_{ij} = \max_k h_{ij}^k$. In other words, with Assumption 6 we allow relaxing Assumption 5 to hold only within arbitrary sets partitioning the state space of the SMPL system. In fact, we could limit the assumptions to the safe set.

Under Assumptions 3 and 6, we now present a result extending Theorem 2 for the computation of the constant H .

Theorem 3. *Under Assumptions 3 and 6, the constant H in Proposition 4 is*

$$H = \sum_{i,j=1}^n H_{ij} + (n-1)M_{ij} \text{ ,}$$

where $H_{ij} = \mathcal{L}_i \max_k h_{ij}^k + \sum_k |J_{ij}^k|$ and $\mathcal{L}_i = \mathcal{L}(\Pi_i(\mathcal{A}))$. The notation $J_{ij}^k = \lim_{z \downarrow c_{ij}^k} t_{ij}(z) - \lim_{z \uparrow c_{ij}^k} t_{ij}(z)$ denotes the jump distance of the density function $t_{ij}(\cdot)$ at the k -th discontinuity point c_{ij}^k .

The constants H_{ij} in Theorem 3 are chosen for the satisfaction of the following inequalities

$$\int_{\Pi_i(\mathcal{A})} |t_{ij}(\bar{z}_i - d_{ij} - z_j) - t_{ij}(\bar{z}_i - d_{ij} - z'_j)| d\bar{z}_i \leq H_{ij} |z_j - z'_j| \text{ .} \quad (5)$$

In some cases, it is possible to obtain a smaller value for H_{ij} by substituting the density function directly into the inequality in (5). Furthermore H_{ij} may be independent of the size of the safe set. For instance, if the delay is modelled by an exponential distribution as in Example 1, then $A_{ij}(\cdot)$ for all $i, j \in \mathbb{N}_n$ follows a shifted exponential distribution, i.e. $A_{ij}(\cdot) \sim \text{SExp}(\lambda_{ij}, s_{ij})$. In this case, $H_{ij} = \lambda_{ij} + \lambda_{ij}^2 \mathcal{L}_i$, as per Theorem 3. However if we compute directly the left-hand side of (5), we get the quantity $H_{ij} = 2\lambda_{ij}$, which is independent of the shape of the safe set. This fact is now proven in general, for a class of distribution functions, in Theorem 4. Let us first introduce the following definition.

Definition 2 (Shifted Exponential Distribution). *The density function of an exponential distribution shifted by ς is given by*

$$t(x; \lambda, \varsigma) = \lambda \exp\{-\lambda(x - \varsigma)\} \theta(x - \varsigma) \ ,$$

where $\theta(\cdot)$ is the unit step function. A random variable X characterised by this distribution is denoted by $X \sim \text{SExp}(\lambda, \varsigma)$.

Theorem 4. *Any random sequence $A_{ij}(\cdot) \sim \text{SExp}(\lambda_{ij}, \varsigma_{ij})$ satisfies inequality (5) with $H_{ij} = 2\lambda_{ij}$.*

Given the previous result, the bound related to the invariance-related abstraction error over SMPL systems with $A_{ij}(\cdot) \sim \text{SExp}(\lambda_{ij}, \varsigma_{ij})$ can be improved and explicitly shown as follows. The maximum value of the density function $t_{ij}(\cdot)$ equals λ_{ij} , i.e. $M_{ij} = \lambda_{ij}$ for all $i, j \in \mathbb{N}_n$. By Theorem 3 and Proposition 4, the bound of the approximation error is then

$$E = N\delta(n+1) \sum_{i,j} \lambda_{ij} \ .$$

Let us go back to Example 2 and adapt according to Definition 2 and Theorem 4.

Example 3. Consider the following two-dimensional SMPL system (1), where $A_{ij}(\cdot) \sim \text{SExp}(\lambda_{ij}, \varsigma_{ij})$ and

$$\begin{bmatrix} \lambda_{11} & \lambda_{12} \\ \lambda_{21} & \lambda_{22} \end{bmatrix} = \begin{bmatrix} 1/2 & 1/3 \\ 1 & 1/3 \end{bmatrix} \ , \quad \begin{bmatrix} \varsigma_{11} & \varsigma_{12} \\ \varsigma_{21} & \varsigma_{22} \end{bmatrix} = \begin{bmatrix} 0 & 2 \\ 2 & 0 \end{bmatrix} \ .$$

Considering a regular schedule with $\mathbf{s}(0) = [0, 0]^T$ and $d = 4$, selecting safe set $\mathcal{A} = [-5, 5]^2$, and event horizon $N = 5$, we get $E = 32.5\delta$. In order to obtain a desired error $E = 0.1$, we need to use 4597 bins per dimension on a uniform discretisation of the set \mathcal{A} . The solution of the invariance problem over the abstract model is presented in Figure 2 (right panel).

Let us now validate this outcome. We have computed 1000 sample trajectories, with an initial condition that has been uniformly generated from the level set corresponding to the probability 0.3, namely within the set $\{\mathbf{z} : P_{\mathbf{z}}(\mathcal{A}) \geq 0.3\}$. Practically, this means we have sampled the initial condition on points corresponding to colours warmer than the ‘‘orange line.’’ Given the error bound $E = 0.1$, we would expect that the trajectories are invariant with a likelihood greater than 0.2. Among the cohort, we have found that 374 trajectories stay inside the safe set for the given 5 steps, which is aligned with the guarantee we have derived.

Furthermore we have compared the approximate solution against the following empirical approach: for each representative point, we generate 1000 sample trajectories starting from it and compute ratio of the number of trajectories that stay in the safe set for 5 steps to the total number of trajectories (1000).

The maximum absolute difference between the approximate solution and the empirical approach for all representative points is 0.0565, which aligns with the error bound of 0.1.

We have also done these two comparisons for the SMPL system in Example 2. The results are quite analogous to the ones obtained in this example. \square




Fig. 2. The left and right plots show solution of the finite-horizon probabilistic invariance problem for two-dimensional SMPL systems with beta (Example 2) and exponential (Example 3) distributions, respectively. The plots have been obtained by computing the problem over finite abstractions obtained by uniform discretisation of the set of interest and selection of central representative points.

5 Conclusions and Future Work

This work has employed finite abstractions to study the finite-horizon probabilistic invariance problem over Stochastic Max-Plus-Linear (SMPL) systems. We have assumed that each random variable has a fixed support, which implies that the topology of the SMPL system is fixed over time. Along this line, we are interested to relax this assumption in order to obtain results that are robust against small topological changes. Furthermore, we are interested in considering extensions of the probabilistic invariance problem. Computationally, we are interested in improving the software and integrating it with FAUST² [34]. Finally, we have been exploring the existence of distributions associated to an analytical solution to the finite-horizon probabilistic invariance problem.

References

1. Baccelli, F., Cohen, G., Gaujal, B.: Recursive equations and basic properties of timed Petri nets. *Discrete Event Dynamic Systems: Theory and Applications* 1(4), 415–439 (1992)
2. Hillion, H., Proth, J.: Performance evaluation of job-shop systems using timed event graphs. *IEEE Trans. Autom. Control* 34(1), 3–9 (1989)

3. Heidergott, B., Olsder, G., van der Woude, J.: *Max Plus at Work—Modeling and Analysis of Synchronized Systems: A Course on Max-Plus Algebra and Its Applications*. Princeton University Press (2006)
4. Roset, B., Nijmeijer, H., van Eekelen, J., Lefeber, E., Rooda, J.: Event driven manufacturing systems as time domain control systems. In: Proc. 44th IEEE Conf. Decision and Control and European Control Conf. (CDC-ECC 2005), pp. 446–451 (December 2005)
5. van Eekelen, J., Lefeber, E., Rooda, J.: Coupling event domain and time domain models of manufacturing systems. In: Proc. 45th IEEE Conf. Decision and Control (CDC 2006), pp. 6068–6073 (December 2006)
6. Brackley, C.A., Broomhead, D.S., Romano, M.C., Thiel, M.: A max-plus model of ribosome dynamics during mRNA translation. *Journal of Theoretical Biology* 303, 128–140 (2012)
7. Heidergott, B.: *Max-Plus Linear Stochastic Systems and Perturbation Analysis (The International Series on Discrete Event Dynamic Systems)*. Springer-Verlag New York, Inc., Secaucus (2006)
8. Olsder, G., Resing, J., De Vries, R., Keane, M., Hooghiemstra, G.: Discrete event systems with stochastic processing times. *IEEE Trans. Autom. Control* 35(3), 299–302 (1990)
9. Resing, J., de Vries, R., Hooghiemstra, G., Keane, M., Olsder, G.: Asymptotic behavior of random discrete event systems. *Stochastic Processes and their Applications* 36(2), 195–216 (1990)
10. van der Woude, J.W., Heidergott, B.: Asymptotic growth rate of stochastic max-plus systems that with a positive probability have a sunflower-like support. In: Proc. 8th Int. Workshop Discrete Event Systems, pp. 451–456 (July 2006)
11. Goverde, R., Heidergott, B., Merlet, G.: A coupling approach to estimating the Lyapunov exponent of stochastic max-plus linear systems. *European Journal of Operational Research* 210(2), 249–257 (2011)
12. Baccelli, F., Cohen, G., Olsder, G., Quadrat, J.P.: *Synchronization and Linearity, An Algebra for Discrete Event Systems*. John Wiley and Sons (1992)
13. Baccelli, F., Hong, D.: Analytic expansions of max-plus Lyapunov exponents. *The Annals of Applied Probability* 10(3), 779–827 (2000)
14. Gaubert, S., Hong, D.: Series expansions of Lyapunov exponents and forgetful monoids. Technical Report RR-3971, INRIA (July 2000)
15. Merlet, G.: Cycle time of stochastic max-plus linear systems. *Electronic Journal of Probability* 13(12), 322–340 (2008)
16. Farahani, S., van den Boom, T., van der Weide, H., De Schutter, B.: An approximation approach for model predictive control of stochastic max-plus linear systems. In: Proc. 10th Int. Workshop Discrete Event Systems, Berlin, DE, pp. 386–391 (August/September 2010)
17. Farahani, S., van den Boom, T., De Schutter, B.: Exact and approximate approaches to the identification of stochastic max-plus-linear systems. In: *Discrete Event Dynamic Systems: Theory and Applications*, pp. 1–25 (April 2013)
18. Adzkiya, D., De Schutter, B., Abate, A.: Finite abstractions of max-plus-linear systems. *IEEE Trans. Autom. Control* 58(12), 3039–3053 (2013)
19. Kurshan, R.P.: *Computer-Aided Verification of Coordinating Processes: The Automata-Theoretic Approach*. Princeton Series in Computer Science. Princeton University Press (1994)
20. Baier, C., Katoen, J.-P., Hermanns, H.: Approximate symbolic model checking of continuous-time Markov chains. In: Baeten, J.C.M., Mauw, S. (eds.) *CONCUR 1999*. LNCS, vol. 1664, pp. 146–162. Springer, Heidelberg (1999)

21. Kwiatkowska, M., Norman, G., Segala, R., Sproston, J.: Verifying quantitative properties of continuous probabilistic timed automata. In: Palamidessi, C. (ed.) CONCUR 2000. LNCS, vol. 1877, pp. 123–137. Springer, Heidelberg (2000)
22. Alur, R., Dill, D.: A theory of timed automata. *Theoretical Computer Science* 126(2), 183–235 (1994)
23. Kushner, H.J., Dupuis, P.: *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer, New York (2001)
24. Chaput, P., Danos, V., Panangaden, P., Plotkin, G.: Approximating markov processes by averaging. *Journal of the ACM* 61(1), 5:1–5:45 (2014)
25. Koutsoukos, X., Riley, D.: Computational methods for reachability analysis of stochastic hybrid systems. In: Hespanha, J.P., Tiwari, A. (eds.) HSCC 2006. LNCS, vol. 3927, pp. 377–391. Springer, Heidelberg (2006)
26. Prandini, M., Hu, J.: Stochastic reachability: Theory and numerical approximation. In: Cassandras, C., Lygeros, J. (eds.) *Stochastic Hybrid Systems. Automation and Control Engineering Series*, vol. 24, pp. 107–138. Taylor & Francis Group/CRC Press (2006)
27. Abate, A., Katoen, J.P., Lygeros, J., Prandini, M.: Approximate model checking of stochastic hybrid systems. *European Journal of Control* 16(6), 624–641 (2010)
28. Abate, A., Prandini, M., Lygeros, J., Sastry, S.: Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica* 44(11), 2724–2734 (2008)
29. Esmail Zadeh Soudjani, S., Abate, A.: Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes. *SIAM Journal on Applied Dynamical Systems* 12(2), 921–956 (2013)
30. Abate, A.: Approximation metrics based on probabilistic bisimulations for general state-space Markov processes: a survey. *Electronic Notes in Theoretical Computer Sciences*, 3–25 (2014)
31. Desharnais, J., Edalat, A., Panangaden, P.: Bisimulation for labelled Markov processes. *Information and Computation* 179(2), 163–193 (2002)
32. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 585–591. Springer, Heidelberg (2011)
33. Katoen, J.P., Khattri, M., Zapreev, I.S.: A Markov reward model checker. In: Proc. 2nd Int. Conf. Quantitative Evaluation of Systems (QEST 2005), pp. 243–244. IEEE Computer Society, Los Alamos (2005)
34. Esmail Zadeh Soudjani, S., Gevaerts, C., Abate, A.: FAUST²: Formal Abstractions of Uncountable-STate STochastic processes arXiv:1403.3286 (2014)