

QUANTITATIVE APPROXIMATION OF THE PROBABILITY DISTRIBUTION OF A MARKOV PROCESS BY FORMAL ABSTRACTIONS

SADEGH ESMAEIL ZADEH SOUDJANI AND ALESSANDRO ABATE

Delft Center for Systems and Control, Technische Universiteit Delft, The Netherlands
e-mail address: S.EsmaeilZadehSoudjani@tudelft.nl

Department of Computer Science, University of Oxford, United Kingdom
e-mail address: Alessandro.Abate@cs.ox.ac.uk

ABSTRACT. The goal of this work is to formally abstract a Markov process evolving in discrete time over a general state space as a finite-state Markov chain, with the objective of precisely approximating its state probability distribution in time, which allows for its approximate, faster computation by that of the Markov chain. The approach is based on formal abstractions and employs an arbitrary finite partition of the state space of the Markov process, and the computation of average transition probabilities between partition sets. The abstraction technique is formal, in that it comes with guarantees on the introduced approximation that depend on the diameters of the partitions: as such, they can be tuned at will. Further in the case of Markov processes with unbounded state spaces, a procedure for precisely truncating the state space within a compact set is provided, together with an error bound that depends on the asymptotic properties of the transition kernel of the original process. The overall abstraction algorithm, which practically hinges on piecewise constant approximations of the density functions of the Markov process, is extended to higher-order function approximations: these can lead to improved error bounds and associated lower computational requirements. The approach is practically tested to compute probabilistic invariance of the Markov process under study, and is compared to a known alternative approach from the literature.

1. INTRODUCTION

Verification techniques and tools for deterministic, discrete time, finite-state systems have been available for many years [20]. Formal methods in the stochastic context are typically limited to finite-state structures, either in continuous or in discrete time [5, 23]. Stochastic

1998 ACM Subject Classification: PROBABILITY AND STATISTICS, Markov processes, Distribution functions.

Key words and phrases: Continuous-Space Markov Processes, Discrete-time Stochastic Systems, Formal Abstractions, PCTL Verification, Probabilistic Invariance, Higher-Order Approximations.

This article generalises and completes the results presented in [14] and specifically benefits from extensions first discussed in [11]. This work has been supported by the European Commission via STREP project MoVeS 257005 and IAPP project AMBI 324432, and by the John Fell Oxford University Press Research Fund.

processes evolving over continuous (uncountable) spaces are often related to undecidable problems (the exception being when they admit analytical solutions). It is thus of interest to resort to formal approximation techniques that allow solving decidably corresponding problems over finite discretisations of the original models. In order to formally relate the computable approximate solutions to the original problems, it is of interest to come up with explicit bounds on the error introduced by the approximations. The use of formal approximations techniques over complex models can be looked at from the perspective of research on abstractions, which are of wide use in formal verification.

Successful numerical schemes based on Markov chain approximations of general stochastic systems in continuous time have been introduced in the literature [21]. However, the finite abstractions are only related to the original models asymptotically (at the limit, that is weakly), with no explicit error bounds. This approach has been applied to the approximate study of probabilistic reachability or safety of stochastic hybrid models in [19, 27]. An alternative line of work on approximations of continuous-space, discrete-time Markov processes is pursued in [8, 9], where the discussed approximation scheme generates a finite-state model. In [7] the idea of approximating by averaging is introduced, where the conditional expectation is used to compute the approximation, and is later extended in [6]. The weak point of these contributions is the fact that the approximation errors that are essential in assessing the quality of the approximation are not computed.

As an alternative to qualitative approximations, in [2] a technique has been introduced to provide formal abstractions of discrete-time, continuous-space Markov models [4], with the objective of investigating their probabilistic invariance (safety) via model checking procedures over a finite Markov chain. In view of computational scalability, the approach has been improved and optimised in [10, 13], extended to a wider class of processes [12, 15], and practically implemented as a software tool [16]. These abstraction techniques hinge on piecewise-constant approximations of the kernels of the Markov process. Linear projection operators are employed in [11] to generalise these techniques via higher-order interpolations that provide improved error bounds on the approximation level.

In this work we show that the approach in [2, 13] can be successfully employed to approximately compute the statistics in time of a stochastic process over a continuous state space. We first provide a forward recursion for the approximate computation of the state distribution in time of the Markov process. This computation is based on a partitioning of the state space, and on the abstraction of the Markov process as a finite-state Markov chain. Further, a higher-order approximation method is presented, as a generalisation of the approach above, and an upper bound on the error related to the new approximation is formally derived. Based on the information gained from the state distribution in time, we show that the method can be used as an alternative to [2, 11, 12, 13] to approximately compute probabilistic invariance (safety) for discrete-time stochastic systems over general state spaces. Probabilistic invariance (safety) is the dual problem to probabilistic reachability. Over deterministic models reachability and safety have been vastly studied in the literature, and computational algorithms and tools have been developed based on both forward and backward recursions. Similarly, for the probabilistic models under study, we compare the presented approach (based on forward computations) with existing approaches in the literature [2, 11, 12, 13] (which hinge on backward computations), particularly in terms of the introduced approximation error.

The Markov chain abstraction applied to the forward/backward computation of probabilistic invariance can be generalised to other specifications expressed as non-nested PCTL

formulae or to reward-based criteria characterised via value function recursions. Moreover, the constructed Markov chain can be shown to represent an *approximate probabilistic bisimulation* of the original process [1, 29, 3].

The article is structured as follows. Section 2 introduces the model under study and discusses some structural assumptions needed for the abstraction procedure. The procedure comprises two separate parts: Section 3 describes the truncation of the dynamics of the model, whereas Section 4 details the abstraction of the dynamics (approximation of the transition kernel) – both parts contribute to the associated approximation error. Section 5 considers higher-order approximation schemes and quantifies the introduced approximation error. Section 6 specialises these higher-order schemes to explicit algorithms for low-dimensional models using known interpolation bases. Section 7 discusses the application of the procedure to the computation of probabilistic invariance, and compares it against an alternative approach in the literature.

In this article we use $\mathbb{N} \doteq \{1, 2, 3, \dots\}$ to indicate the natural numbers, $\mathbb{N}_m \doteq \{1, 2, \dots, m\}$ and $\mathbb{Z}_m \doteq \{0, 1, 2, \dots, m\}$ for any $m \in \mathbb{N}$.

2. MODELS, PRELIMINARIES, AND GOALS OF THIS WORK

We consider a discrete time Markov process $\mathcal{M}_{\mathfrak{s}}$ defined over a general state space, which is characterised by a pair $(\mathcal{S}, T_{\mathfrak{s}})$, where \mathcal{S} is a continuous state space that we assume endowed with a metric and be separable¹. We denote by $(\mathcal{S}, \mathcal{B}(\mathcal{S}), \mathcal{P})$ the probability structure on \mathcal{S} , with $\mathcal{B}(\mathcal{S})$ being the Borel σ -algebra² in \mathcal{S} and \mathcal{P} a probability measure to be characterised shortly. $T_{\mathfrak{s}}$ is a stochastic kernel that assigns to each point $s \in \mathcal{S}$ a probability measure $T_{\mathfrak{s}}(\cdot|s)$, so that for any measurable set $A \in \mathcal{B}(\mathcal{S})$, $\mathcal{P}(s(1) \in A|s(0) = s) = T_{\mathfrak{s}}(A|s)$. We assume that the stochastic kernel $T_{\mathfrak{s}}$ admits a density function $t_{\mathfrak{s}}$, namely $T_{\mathfrak{s}}(A|s) = \int_A t_{\mathfrak{s}}(\bar{s}|s) d\bar{s}$.

Given the measurable space $(\mathcal{S}, \mathcal{B}(\mathcal{S}), \mathcal{P})$, we set up the product space \mathcal{S}^{t+1} containing elements $\mathbf{s}(t) = [s(0), s(1), \dots, s(t)]$, where the bold typeset is used in the sequel to indicate vector quantities. Suppose that the initial state of the Markov process $\mathcal{M}_{\mathfrak{s}}$ is distributed according to the density function $\pi_0 : \mathcal{S} \rightarrow \mathbb{R}^{\geq 0}$. Then the multi-variate density function $\pi_0(s_0)t_{\mathfrak{s}}(s_1|s_0)t_{\mathfrak{s}}(s_2|s_1) \dots t_{\mathfrak{s}}(s_t|s_{t-1})$ is a probability measure \mathbb{P} on the product space \mathcal{S}^{t+1} . On the other hand the state distribution of $\mathcal{M}_{\mathfrak{s}}$ at time $t \in \mathbb{N}$ is characterised by a density function $\pi_t : \mathcal{S} \rightarrow \mathbb{R}^{\geq 0}$, which fully describes the statistics of the process at time t and is in particular such that, for all $A \in \mathcal{B}(\mathcal{S})$,

$$\mathbb{P}(s(t) \in A) = \int_A \pi_t(s) ds,$$

where the symbol \mathbb{P} is used to indicate the probability associated to events over the product space \mathcal{S}^{t+1} (note that the event $s(t) \in A$ is equivalent to $\mathbf{s}(t) \in \mathcal{S}^t \times A$ on their corresponding probability spaces).

The state density functions $\pi_t(\cdot)$ can be characterised recursively, as follows:

$$\pi_{t+1}(\bar{s}) = \int_{\mathcal{S}} t_{\mathfrak{s}}(\bar{s}|s) \pi_t(s) ds \quad \forall \bar{s} \in \mathcal{S}. \quad (2.1)$$

¹A metric space \mathcal{S} is called separable if it has a countable dense subset.

²The Borel σ -algebra $\mathcal{B}(\mathcal{S})$ is the smallest σ -algebra in \mathcal{S} that contains all open subsets of \mathcal{S} . For a separable metric space \mathcal{S} , $\mathcal{B}(\mathcal{S})$ equals the σ -algebra generated by the open (or closed) balls of \mathcal{S} .

In practice the forward recursion in (2.1) rarely yields a closed form for $\pi_t(\cdot)$. A special instance where this is the case is represented by a linear dynamical system perturbed by a Gaussian process noise: due to the closure property of Gaussian distributions over addition and multiplication by a constant, it is possible to explicitly write recursive formulae for the mean and the variance of the distribution, and thus express in a closed form the distribution in time of the solution process. In more general cases, it is necessary to numerically (hence, approximately) compute this density function in time.

This article provides a numerical approximation of the density function of \mathcal{M}_s in time as the probability mass function (pmf) of a finite-state Markov chain \mathcal{M}_p . The Markov chain \mathcal{M}_p is obtained as an abstraction of the concrete Markov process \mathcal{M}_s . The abstraction is associated with a guaranteed and tunable error bound, and algorithmically it leverages a state-space partitioning procedure. The procedure is comprised of two steps:

- (1) since the state space \mathcal{S} is generally unbounded, it is first properly truncated;
- (2) subsequently, a partition of the truncated dynamics is introduced.

Section 3 discusses the error generated by the state-space truncation, whereas Section 4 describes the construction of the Markov chain by state-space partitioning. The discussed Markov chain abstraction is based on a piecewise-constant approximation of the density functions. In order to improve the efficiency and the precision of the approximation, we generalise the abstraction method in Section 5 utilizing higher-order approximations of the density functions. We employ the following example throughout the article as a running case study.

Example 2.1. Consider the one-dimensional stochastic dynamical system

$$s(t+1) = as(t) + b + \sigma w(t),$$

where the parameters $a, \sigma > 0$, whereas $b \in \mathbb{R}$, and $w(\cdot)$ is a process comprised of independent, identically distributed random variables with a standard normal distribution. The initial state of the process is selected uniformly within the bounded interval $[\beta_0, \gamma_0] \subset \mathbb{R}$. The solution of the model is a Markov process, evolving over the state space $\mathcal{S} = \mathbb{R}$, and fully characterised by the conditional density function

$$t_s(\bar{s}|s) = \phi_\sigma(\bar{s} - as - b), \quad \text{where} \quad \phi_\sigma(u) = \frac{1}{\sigma\sqrt{2\pi}}e^{-u^2/2\sigma^2}.$$

□

We raise the following assumptions in order to relate the state density function of \mathcal{M}_s to the probability mass function of \mathcal{M}_p . More precisely, these assumptions are employed for the computation of the approximation error, but the abstraction approach proposed in this article can be applied without raising them.

Assumption 2.2. For given sets $\Gamma \subset \mathcal{S}^2$ and $\Lambda_0 \subset \mathcal{S}$, there exist positive constants ϵ and ϵ_0 , such that $t_s(\bar{s}|s)$ and $\pi_0(s)$ satisfy the following conditions:

$$t_s(\bar{s}|s) \leq \epsilon, \quad \forall (s, \bar{s}) \in \mathcal{S}^2 \setminus \Gamma, \quad \text{and} \quad \pi_0(s) \leq \epsilon_0, \quad \forall s \in \mathcal{S} \setminus \Lambda_0. \quad (2.2)$$

Assumption 2.3. The density functions $\pi_0(s)$ and $t_s(\bar{s}|s)$ are (globally) Lipschitz continuous, namely there exist finite constants λ_0, λ_f , such that the following Lipschitz continuity

conditions hold:

$$|\pi_0(s) - \pi_0(s')| \leq \lambda_0 \|s - s'\|, \quad \forall s, s' \in \Lambda_0, \quad (2.3)$$

$$|t_s(\bar{s}|s) - t_s(\bar{s}'|s)| \leq \lambda_f \|\bar{s} - \bar{s}'\|, \quad \forall s, \bar{s}, \bar{s}' \in \mathcal{S}. \quad (2.4)$$

Moreover, there exists a finite constant M_f such that

$$M_f = \sup \left\{ \int_{\mathcal{S}} t_s(\bar{s}|s) ds \mid \bar{s} \in \mathcal{S} \right\}. \quad (2.5)$$

The Lipschitz constants λ_0, λ_f are practically computed by taking partial derivatives of the density functions $\pi_0(\cdot), t_s(\cdot|s)$ and maximising their norm. The sets Λ_0 and Γ will be used to truncate the support of the density functions $\pi_0(\cdot)$ and $t_s(\cdot|s)$, respectively. Assumption 2.2 enables the precise study of the behaviour of density functions $\pi_t(\cdot)$ over the truncated state space. Furthermore, the Lipschitz continuity conditions in Assumption 2.3 are essential to derive error bounds related to the abstraction of the Markov process over the truncated state space. In order to compute these error bounds, we assign the infinity norm to the space of bounded measurable functions over the state space \mathcal{S} , namely

$$\|f(s)\|_\infty = \sup_{s \in \mathcal{S}} |f(s)|, \quad \forall f \in \mathbb{B}(\mathcal{S}) \doteq \{g : \mathcal{S} \rightarrow \mathbb{R}, g \text{ bounded and measurable}\}.$$

In the sequel the function $\mathbb{1}_A(\cdot)$ denotes the indicator function of a set $A \subseteq \mathcal{S}$, namely $\mathbb{1}_A(s) = 1$, if $s \in A$; else $\mathbb{1}_A(s) = 0$.

Example 2.1 (Continued). Select the interval $\Lambda_0 = [\beta_0, \gamma_0]$ and define the set Γ by the linear inequality

$$\Gamma = \{(s, \bar{s}) \in \mathbb{R}^2 \mid |\bar{s} - as - b| \leq \alpha\sigma\}.$$

The initial density function π_0 of the process can be represented by the function

$$\pi_0(s) = \mathbb{1}_{[\beta_0, \gamma_0]}(s) / (\gamma_0 - \beta_0).$$

Then Assumption 2.2 is valid with constants $\epsilon = \phi_1(\alpha)/\sigma$ and $\epsilon_0 = 0$. The constant M_f in Assumption 2.3 is equal to $1/a$. Lipschitz continuity, as per (2.3) and (2.4), holds for the constants $\lambda_0 = 0$ and $\lambda_f = 1/(\sigma^2 \sqrt{2\pi e})$. \square

3. STATE-SPACE TRUNCATION PROCEDURE

We limit the support of the density functions π_0, t_s to the sets Λ_0, Γ respectively, and recursively compute support sets Λ_t , as in (3.2), that are associated to the density functions π_t . Then we employ the quantities ϵ, ϵ_0 in Assumption 2.2 to compute bounds ϵ_t , as in (3.1), on the error incurring in disregarding the value of the density functions π_t outside the sets Λ_t . Finally we truncate the original, unbounded state space to the set $\Upsilon = \cup_{t=0}^N \Lambda_t$.

As intuitive, the error related to the spatial truncation depends on the behaviour of the conditional density function t_s over the eliminated regions of the state space. Suppose that sets Γ, Λ_0 are selected such that Assumption 2.2 is satisfied with constants ϵ, ϵ_0 : then Theorem 3.3 provides an upper bound on the error obtained from manipulating the density functions in time $\pi_t(\cdot)$ exclusively over the truncated regions of the state space.

Theorem 3.1. *Under Assumption 2.2 the functions π_t satisfy the bound*

$$0 \leq \pi_t(s) \leq \epsilon_t, \quad \forall s \in \mathcal{S} \setminus \Lambda_t,$$

where the quantities $\{\varepsilon_t, t \in \mathbb{Z}_N\}$ are defined recursively by

$$\varepsilon_{t+1} = \varepsilon + M_{\mathfrak{f}}\varepsilon_t, \quad (3.1)$$

whereas the support sets $\{\Lambda_t, t \in \mathbb{Z}_N\}$ are computed as

$$\Lambda_{t+1} = \Pi_{\bar{s}}(\Gamma \cap (\Lambda_t \times \mathcal{S})), \quad (3.2)$$

where $\Pi_{\bar{s}}$ denotes the projection map along the second set of coordinates³.

Remark 3.2. Notice that if the shape of the sets Γ and Λ_0 is computationally manageable (e.g., if these sets are polytopes), then it is possible to precisely implement the computation of the recursion in (3.2) by available software tools, such as the MPT toolbox [22]. Further, if for some t_0 , $\Lambda_{t_0+1} \supset \Lambda_{t_0}$, then for all $t \geq t_0$, $\Lambda_{t+1} \supset \Lambda_t$. Similarly, we have that

- if for some t_0 , $\Lambda_{t_0+1} \subset \Lambda_{t_0}$, then for all $t \geq t_0$, $\Lambda_{t+1} \subset \Lambda_t$.
- if for some t_0 , $\Lambda_{t_0+1} = \Lambda_{t_0}$, then for all $t \geq t_0$, $\Lambda_t = \Lambda_{t_0}$.

In order to clarify the role of Γ in the computation of Λ_t , we emphasize that $\Lambda_{t+1} = \cup_{s \in \Lambda_t} \Xi(s)$, where Ξ depends only on Γ and is defined by the set-valued map

$$\Xi : \mathcal{S} \rightarrow 2^{\mathcal{S}}, \quad \Xi(s) = \{\bar{s} \in \mathcal{S} \mid (s, \bar{s}) \in \Gamma\}.$$

Figure 1 provides a visual illustration of the recursion in (3.2). □

Let us introduce the quantity $\kappa(t, M_{\mathfrak{f}})$, which plays a role in the solution of (3.1) and will be frequently used shortly:

$$\kappa(t, M_{\mathfrak{f}}) = \begin{cases} \frac{1-M_{\mathfrak{f}}^t}{1-M_{\mathfrak{f}}}, & M_{\mathfrak{f}} \neq 1 \\ t, & M_{\mathfrak{f}} = 1. \end{cases} \quad (3.3)$$

The following theorem provides a truncation procedure, valid over a finite time horizon $\mathbb{Z}_N = \{0, 1, \dots, N\}$, which reduces the state space \mathcal{S} to the set $\Upsilon = \bigcup_{t=0}^N \Lambda_t$. The theorem also formally quantifies the associated truncation error.

Theorem 3.3. *Suppose that the state space of the process $\mathcal{M}_{\bar{s}}$ has been truncated to the set $\Upsilon = \bigcup_{t=0}^N \Lambda_t$. Let us introduce the following recursion to compute functions $\mu_t : \mathcal{S} \rightarrow \mathbb{R}^{\geq 0}$ as an approximation of the density functions π_t :*

$$\mu_{t+1}(\bar{s}) = \mathbb{1}_{\Upsilon}(\bar{s}) \int_{\mathcal{S}} t_{\bar{s}}(\bar{s}|s) \mu_t(s) ds, \quad \mu_0(s) = \mathbb{1}_{\Lambda_0}(s) \pi_0(s), \quad \forall \bar{s} \in \mathcal{S}. \quad (3.4)$$

Then the introduced approximation error is $\|\pi_t - \mu_t\|_{\infty} \leq \varepsilon_t$, for all $t \in \mathbb{Z}_N$.

Recapitulating, Theorem 3.3 leads to the following procedure to approximate the density functions π_t of $\mathcal{M}_{\bar{s}}$ over an unbounded state space \mathcal{S} :

- (1) truncate π_0 in such a way that μ_0 has a bounded support Λ_0 ;
- (2) truncate the conditional density function $t_{\bar{s}}(\cdot|s)$ over a bounded set for all $s \in \mathcal{S}$, then quantify $\Gamma \subset \mathcal{S}^2$ as the support of the truncated density function;
- (3) leverage the recursion in (3.2) to compute the support sets Λ_t ;
- (4) use the recursion in (3.4) to compute the approximate density functions μ_t over the set $\Upsilon = \bigcup_{t=0}^N \Lambda_t$. Note that the recursion in (3.4) is effectively computed over the set Υ , since $\mu_t(s) = 0$ for all $s \in \mathcal{S} \setminus \Upsilon$.

³Recall that both Γ and $\Lambda \times \mathcal{S}$ are defined over $\mathcal{S}^2 = \mathcal{S} \times \mathcal{S}$.




Figure 1: Graphical representation of the recursion in (3.2) for sets Λ_t .

Note that we could as well handle the support of $\mu_t(\cdot)$ over the time-varying sets Λ_t , by adapting the recursion in (3.4) with $\mathbb{1}_{\Lambda_{t+1}}$ instead of $\mathbb{1}_{\Upsilon}$. However, while employing the (larger) set Υ may lead to an increase in memory requirements at each stage, it will considerably simplify the computations of the state-space partitioning and of the Markov chain abstraction: indeed, employing time-varying sets Λ_t would render the partitioning procedure also time-dependent, and the obtained Markov chain would be time-inhomogeneous. We therefore opt to work directly with set Υ in order to avoid these difficulties.

Example 2.1 (Continued). We easily obtain a closed form for the sets $\Lambda_t = [\beta_t, \gamma_t]$, via

$$\beta_{t+1} = a\beta_t + b - \alpha\sigma, \quad \gamma_{t+1} = a\gamma_t + b + \alpha\sigma.$$

Set Υ is the union of intervals $[\beta_t, \gamma_t]$. The error of the state-space truncation over Υ is

$$\|\pi_t - \mu_t\|_\infty \leq \varepsilon_t = \kappa(t, M_f) \frac{\phi_1(\alpha)}{\sigma}, \quad M_f = \frac{1}{a}.$$

□

The recursion in (3.4) indicates that the support of functions μ_t are always contained in the set Υ (namely, they are equal to zero over the complement of Υ). We are thus only interested in computing these functions over the set Υ , which allows simplifying the recursion in (3.4) as follows:

$$\mu_{t+1}(\bar{s}) = \int_{\Upsilon} t_{\bar{s}}(\bar{s}|s) \mu_t(s) ds, \quad \forall \bar{s} \in \Upsilon, \quad t \in \mathbb{Z}_{N-1}. \quad (3.5)$$

4. MARKOV CHAIN ABSTRACTION VIA STATE-SPACE PARTITIONING

In this section we assume that sets Γ, Λ_0 have been properly selected so that $\Upsilon = \cup_{t=0}^N \Lambda_t$ is bounded. In order to formally abstract process \mathcal{M}_s as a finite Markov chain \mathcal{M}_p and to approximate its state density functions, we select a finite partition of the bounded set Υ as $\Upsilon = \cup_{i=1}^n \mathcal{A}_i$, where sets \mathcal{A}_i have non-trivial measure. We then complete the partition over the whole state space $\mathcal{S} = \cup_{i=1}^{n+1} \mathcal{A}_i$ by additionally including set $\mathcal{A}_{n+1} = \mathcal{S} \setminus \Upsilon$. This results in a finite Markov chain \mathcal{M}_p with $n+1$ discrete abstract states in the set $\mathbb{N}_{n+1} = \{1, 2, \dots, n, n+$

1}, and characterised by the transition probability matrix $P = [P_{ij}] \in \mathbb{R}^{(n+1) \times (n+1)}$, where the probability of jumping from any pair of states i to j (P_{ij}) is computed as

$$\begin{aligned} P_{ij} &= \frac{1}{\mathcal{L}(\mathcal{A}_i)} \int_{\mathcal{A}_j} \int_{\mathcal{A}_i} t_{\mathfrak{s}}(\bar{s}|s) ds d\bar{s}, \quad \forall i \in \mathbb{N}_n, \\ P_{(n+1)j} &= \delta_{(n+1)j}, \end{aligned} \quad (4.1)$$

for all $j \in \mathbb{N}_{n+1}$, and where $\delta_{(n+1)j}$ is the Kronecker delta function (the abstract state $n+1$ of $\mathcal{M}_{\mathfrak{p}}$ is absorbing), and $\mathcal{L}(\cdot)$ denotes the Lebesgue measure of a set (i.e., its volume). The quantities in (4.1) are well-defined since the set Υ is bounded and the measures $\mathcal{L}(\mathcal{A}_i), i \in \mathbb{N}_n$, are finite and non-trivial. Notice that matrix P for the Markov chain $\mathcal{M}_{\mathfrak{p}}$ is stochastic, namely

$$\begin{aligned} \sum_{j=1}^{n+1} P_{ij} &= \sum_{j=1}^{n+1} \frac{1}{\mathcal{L}(\mathcal{A}_i)} \int_{\mathcal{A}_j} \int_{\mathcal{A}_i} t_{\mathfrak{s}}(\bar{s}|s) ds d\bar{s} = \frac{1}{\mathcal{L}(\mathcal{A}_i)} \int_{\mathcal{A}_i} \left(\sum_{j=1}^{n+1} \int_{\mathcal{A}_j} t_{\mathfrak{s}}(\bar{s}|s) d\bar{s} \right) ds \\ &= \frac{1}{\mathcal{L}(\mathcal{A}_i)} \int_{\mathcal{A}_i} \int_{\mathcal{S}} t_{\mathfrak{s}}(\bar{s}|s) d\bar{s} ds = \frac{1}{\mathcal{L}(\mathcal{A}_i)} \int_{\mathcal{A}_i} ds = 1. \end{aligned}$$

The initial distribution of $\mathcal{M}_{\mathfrak{p}}$ is the pmf $\mathbf{p}_0 = [p_0(1), p_0(2), \dots, p_0(n+1)]$, and it is obtained from π_0 as $p_0(i) = \int_{\mathcal{A}_i} \pi_0(s) ds, \forall i \in \mathbb{N}_{n+1}$. Then the pmf associated to the state distribution of $\mathcal{M}_{\mathfrak{p}}$ at time t can be computed as $\mathbf{p}_t = \mathbf{p}_0 P^t$.

It is intuitive that the discrete pmf \mathbf{p}_t of the Markov chain $\mathcal{M}_{\mathfrak{p}}$ approximates the continuous density function π_t of the Markov process $\mathcal{M}_{\mathfrak{s}}$. In the rest of the section we show how to formalise this relationship: \mathbf{p}_t is used to construct an approximate function, denoted by ψ_t , of the density π_t . Theorem 4.2 shows that ψ_t is a piecewise constant approximation (with values in its codomain that are the entries of the pmf \mathbf{p}_t , normalised by the Lebesgue measure of the associated partition set) of the original density function π_t . Moreover, under the continuity assumption in (2.4) (ref. Lemma 4.1) we can establish the Lipschitz continuity of π_t , which enables the quantification (in Theorem 4.2) of the error related to its piecewise constant approximation ψ_t .

Lemma 4.1. *Suppose that the inequality in (2.4) holds. Then the state density functions $\pi_t(\cdot)$ are globally Lipschitz continuous with constant $\lambda_{\mathfrak{f}}$ for all $t \in \mathbb{N}$:*

$$|\pi_t(s) - \pi_t(s')| \leq \lambda_{\mathfrak{f}} \|s - s'\|, \quad \forall s, s' \in \mathcal{S}.$$

Theorem 4.2. *Under Assumptions 2.2 and 2.3, the functions $\pi_t(\cdot)$ can be approximated by piecewise constant functions $\psi_t(\cdot)$, defined as*

$$\psi_t(s) = \sum_{i=1}^n \frac{p_t(i)}{\mathcal{L}(\mathcal{A}_i)} \mathbb{1}_{\mathcal{A}_i}(s), \quad \forall t \in \mathbb{N}, \quad (4.2)$$

where $\mathbb{1}_B(\cdot)$ is the indicator function of a set $B \subset \mathcal{S}$. The approximation error is upper-bounded by the quantity

$$\|\pi_t - \psi_t\|_{\infty} \leq \varepsilon_t + E_t, \quad \forall t \in \mathbb{N}, \quad (4.3)$$

where E_t can be recursively computed as

$$E_{t+1} = M_{\mathfrak{f}} E_t + \lambda_{\mathfrak{f}} \delta, \quad E_0 = \lambda_0 \delta, \quad (4.4)$$

and δ is an upper bound on the diameters of the partition sets $\{\mathcal{A}_i\}_{i=1}^n$, namely $\delta = \sup \{\|s - s'\|, s, s' \in \mathcal{A}_i, i \in \mathbb{N}_n\}$.

Note that the functions ψ_t are defined over the whole state space \mathcal{S} , but (4.2) implies that they are equal to zero outside the set Υ .

Corollary 4.3. *The recursion in (4.4) admits the explicit solution*

$$E_t = [\kappa(t, M_{\mathfrak{f}})\lambda_{\mathfrak{f}} + M_{\mathfrak{f}}^t \lambda_0] \delta,$$

where $\kappa(t, M_{\mathfrak{f}})$ is introduced in (3.3).

Underlying Theorem 4.2 is the fact that $\psi_t(\cdot)$ are in general sub-stochastic density functions:

$$\begin{aligned} \int_{\mathcal{S}} \psi_t(s) ds &= \int_{\mathcal{S}} \sum_{i=1}^n \frac{p_t(i)}{\mathcal{L}(\mathcal{A}_i)} \mathbb{1}_{\mathcal{A}_i}(s) ds = \sum_{i=1}^n \frac{p_t(i)}{\mathcal{L}(\mathcal{A}_i)} \int_{\mathcal{S}} \mathbb{1}_{\mathcal{A}_i}(s) ds \\ &= \sum_{i=1}^n \frac{p_t(i)}{\mathcal{L}(\mathcal{A}_i)} \mathcal{L}(\mathcal{A}_i) = \sum_{i=1}^n p_t(i) = 1 - p_t(n+1) \leq 1. \end{aligned}$$

This is clearly due to the fact that we are operating on the dynamics of $\mathcal{M}_{\mathfrak{s}}$ truncated over the set Υ . It is thus intuitive that the approximation procedure and the derived error bounds are also valid for the case of sub-stochastic density functions [3], namely

$$\int_{\mathcal{S}} t_{\mathfrak{s}}(\bar{s}|s) d\bar{s} \leq 1, \quad \forall s \in \mathcal{S}, \quad \int_{\mathcal{S}} \pi_0(s) ds \leq 1,$$

the only difference being that the obtained Markov chain $\mathcal{M}_{\mathfrak{p}}$ is as well sub-stochastic.

Further, whenever the Lipschitz continuity requirement on the initial density function, as per (2.3) in Assumption 2.3, does not hold, (for instance, this is the case when the initial state of the process is deterministic) we can relax this continuity assumption on the initial distribution of the process by starting the discrete computation from the time step $t = 1$. In this case we define the pmf $\mathbf{p}_1 = [p_1(1), p_1(2), \dots, p_1(n+1)]$, where

$$p_1(i) = \int_{\mathcal{A}_i} \int_{\mathcal{S}} t_{\mathfrak{s}}(\bar{s}|s) \pi_0(s) ds d\bar{s}, \quad \forall i \in \mathbb{N}_{n+1}, \quad (4.5)$$

and derive $\mathbf{p}_t = \mathbf{p}_1 P^{t-1}$ for all $t \in \mathbb{N}$. Theorem 4.2 follows along similar lines, except for equation (4.4), where the initial error is set to $E_0 = 0$ and the time-dependent terms E_t can be derived as $E_t = \kappa(t, M_{\mathfrak{f}})\lambda_{\mathfrak{f}}\delta$.

It is important to emphasise the practical computability of the derived errors, and the fact that they can be tuned by selecting a finer partition of set Υ that relates to a smaller global parameter δ . Further, in order to attain abstractions that are practically useful, it is imperative to seek improvements on the derived error bounds: in particular, the approximation errors can be computed locally (under corresponding local Lipschitz continuity assumptions), following the procedures discussed in [13].

Example 2.1 (Continued). The error of the Markov chain abstraction can be expressed as

$$\|\pi_t - \psi_t\|_{\infty} \leq \kappa(t, M_{\mathfrak{f}}) \left[\frac{\delta}{\sigma^2 \sqrt{2\pi e}} + \frac{\phi_1(\alpha)}{\sigma} \right], \quad M_{\mathfrak{f}} = \frac{1}{a}. \quad (4.6)$$

The error can be tuned in two distinct ways:

- (1) by selecting larger values for α , which on the one hand leads to a less narrow truncation, but on the other requires the partition of a larger interval;
- (2) by reducing partitions diameter δ , which of course results in a larger cardinality of the partition sets.

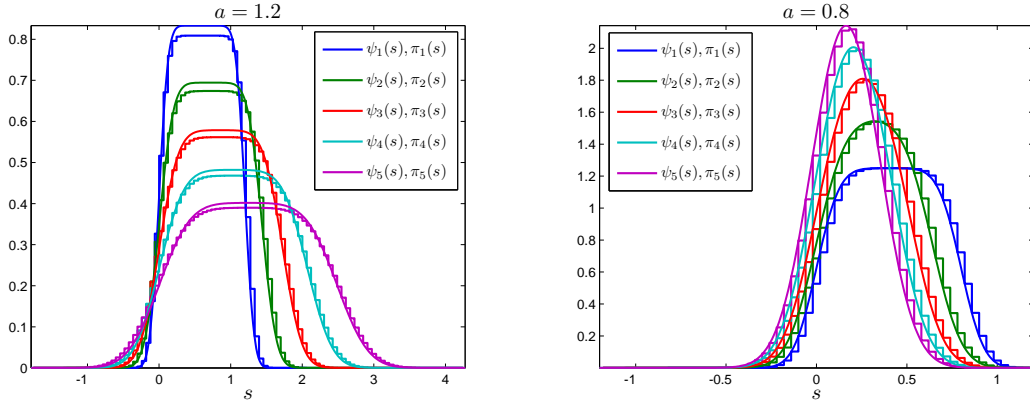


Figure 2: Piecewise constant approximation $\psi_t(\cdot)$ of the state density function $\pi_t(\cdot)$ (derived analytically), for parameters $a = 1.2$ (left) and $a = 0.8$ (right).

Let us select values $b = 0, \beta_0 = 0, \gamma_0 = 1, \sigma = 0.1$, and time horizon $N = 5$. For $a = 1.2$ we need to partition the interval $\Upsilon = [-0.75\alpha, 2.49 + 0.75\alpha]$, which results in the error $\|\pi_t - \psi_t\|_\infty \leq 86.8\delta + 35.9\phi_1(\alpha)$ for all $t \in \mathbb{Z}_N$. For $a = 0.8$ we need to partition the smaller interval $\Upsilon = [-0.34\alpha, 0.33 + 0.34\alpha]$, which results in the error $\|\pi_t - \psi_t\|_\infty \leq 198.6\delta + 82.1\phi_1(\alpha)$ for all $t \in \mathbb{Z}_N$. Notice that in the case of $a = 1.2$, we partition a larger interval and obtain a smaller error, while for $a = 0.8$ we partition a smaller interval with correspondingly a larger error. It is obvious that the parameters δ, α can be chosen properly to ensure that a certain error precision is met. This simple model admits a solution in closed form, and its state density functions can be obtained as the convolution of a uniform distribution (the contribution of initial state) and a zero-mean Gaussian distribution with time-dependent variance (the contributions of the process noise). This leads to the plots in Figure 2, which display the original and the approximated state density functions for the set of parameters $\alpha = 2.4, \delta = 0.05$. \square

5. HIGHER-ORDER APPROXIMATION SCHEMES

In the previous section we have shown that a Markov chain abstraction can be employed to formally approximate the density function of a Markov process in time. This abstraction is interpreted as a piecewise-constant approximation of the density function of the Markov model. In this section we argue that this procedure can be extended to approximation techniques based on higher-order interpolations.

With focus on the truncated region of the state space, let us denote with $\mathbb{B}(\Upsilon)$ the space of bounded and measurable functions $f : \Upsilon \rightarrow \mathbb{R}$, equipped with the infinity norm $\|f\|_\infty = \sup\{|f(s)|, s \in \Upsilon\}$, for all $f \in \mathbb{B}(\Upsilon)$. The linear operator \mathcal{R}_Υ , defined over $\mathbb{B}(\Upsilon)$ by

$$\mathcal{R}_\Upsilon(f)(\bar{s}) = \int_\Upsilon t_s(\bar{s}|s)f(s)ds, \quad \forall \bar{s} \in \Upsilon, \quad \forall f \in \mathbb{B}(\Upsilon), \quad (5.1)$$

characterises the solution of the recursion in (3.5) as $\mu_t(s) = \mathcal{R}_\Upsilon^t(\mu_0)(s)$, for any $t \in \mathbb{N}_N$. While in Section 4 we have proposed approximations of functions $\mu_t(\cdot)$ by piecewise-constant

functions $\psi_t(\cdot)$ with an explicit quantification of the associated error, in this section we are interested in considering approximations via higher-order interpolations.

5.1. Quantification of the Error of a Projection Over a Function Space. Consider a set of basis functions $\Phi = \{\phi_1(s), \phi_2(s), \dots, \phi_h(s)\}$, $h \in \mathbb{N}$, the function space $\Psi = \text{span } \Phi$ generated by this set as a subset of $\mathbb{B}(\Upsilon)$, and a linear operator $\Pi_\Upsilon : \mathbb{B}(\Upsilon) \rightarrow \Psi$, which projects any function $f \in \mathbb{B}(\Upsilon)$ onto the function space Ψ . Theorem 5.1 provides a theoretical result for approximating the solution of (3.5): the following section provides details on turning this result into a useful tool for approximations.

Theorem 5.1. *Assume that a linear projection operator $\Pi_\Upsilon : \mathbb{B}(\Upsilon) \rightarrow \Psi$ satisfies the inequality*

$$\|\Pi_\Upsilon(t_s(\cdot|s)) - t_s(\cdot|s)\|_\infty \leq \mathcal{E}^h, \quad \forall s \in \Upsilon, \quad (5.2)$$

and that there exists a finite constant M_f^h , such that

$$\int_\Upsilon |\Pi_\Upsilon(t_s(\bar{s}|s))| ds \leq M_f^h, \quad \forall \bar{s} \in \Upsilon. \quad (5.3)$$

Define the functions $\psi_t^h(\cdot)$ as approximations of $\mu_t(\cdot)$ (cf. (5.1)), by

$$\psi_t^h = (\Pi_\Upsilon \mathcal{R}_\Upsilon)^t(\mu_0), \quad t \in \mathbb{Z}_N. \quad (5.4)$$

Then it holds that

$$\|\mu_t - \psi_t^h\|_\infty \leq E_t^h, \quad t \in \mathbb{N}_N, \quad (5.5)$$

where the error E_t^h satisfies the difference equation

$$E_{t+1}^h = M_f^h E_t^h + \mathcal{E}^h, \quad E_0^h = 0.$$

Corollary 5.2. *Under the assumptions raised in (5.2)-(5.3), the error E_t^h can be alternatively expressed explicitly as $E_t^h = \mathcal{E}^h \kappa(t, M_f^h)$.*

The error E_t^h formulated in Theorem 5.1 is comparable with the quantity E_t computed in Theorem 4.2. Both E_t^h, E_t represent bounds on the approximation error introduced by $\mu_t(\cdot)$, the density function obtained after state space truncation. The difference is in the initialisation of the corresponding recursions, where we have $E_0^h = 0$ because $\psi_0^h = \mu_0$, but $E_0 = \lambda_0 \delta$ since ψ_0 is a piecewise constant approximation of μ_0 in (4.2). As we mentioned before, the quantities in (4.5) can be alternatively employed as the starting values of the computation to relax the continuity assumption on π_0 , which results in an initial error $E_0 = 0$, thus providing a complete similarity between E_t and E_t^h .

5.2. Construction of the Projection Operator. In the ensuing sections we focus, for the sake of simplicity, on a Euclidean domain, namely $\Upsilon \subset \mathcal{S} = \mathbb{R}^d$, where d denotes a finite dimension. We discuss a general form for an interpolation operator related to the discussed projection operation. Let $\phi_j : \mathcal{D} \subset \mathbb{R}^d \rightarrow \mathbb{R}, j \in \mathbb{N}_h$, be independent functions defined over a generic set \mathcal{D} . The interpolation operator $\Pi_{\mathcal{D}}$ is defined as a projection map onto the function space $\Psi = \text{span}\{\phi_1(s), \phi_2(s), \dots, \phi_h(s)\}$, which projects any function $f : \mathcal{D} \rightarrow \mathbb{R}$ to a unique function $\Pi_{\mathcal{D}}(f) = \sum_{j=1}^h \alpha_j \phi_j$, using a finite set of data $\{(s_j, f(s_j)) | s_j \in \mathcal{D}, j \in \mathbb{N}_h\}$ and such that $\Pi_{\mathcal{D}}(f)(s_j) = f(s_j)$. The projection coefficients $\alpha_j, j \in \mathbb{N}_h$, satisfy the linear

equation $\mathbf{f} = \mathcal{Q}\boldsymbol{\alpha}$, where $\mathbf{f} = [f(s_i)]_{i \in \mathbb{N}_h}$ and $\boldsymbol{\alpha} = [\alpha_j]_{j \in \mathbb{N}_h}$ are h -dimensional column vectors, and $\mathcal{Q} = [\phi_j(s_i)]_{i,j}$ is the associated $(h \times h)$ -dimensional interpolation matrix.

Let us now shift the focus to the recursion in (3.5) discussed in the previous section and tailor the operators above accordingly. Let us select a partition $\{\mathcal{A}_i, i \in \mathbb{N}_n\}$ for the set Υ , with finite cardinality n . Selecting a basis $\{\phi_{ij}, j \in \mathbb{N}_h\}$ for each partition, let us introduce the interpolation operators $\Pi_{\mathcal{A}_i}$ for the projection over each partition set \mathcal{A}_i , which is done as described above by replacing the domain \mathcal{D} with \mathcal{A}_i . Finally, let us introduce the (global) linear operator Π_{Υ} , acting on a function $f : \Upsilon \rightarrow \mathbb{R}$ by

$$\Pi_{\Upsilon}(f) = \sum_{i=1}^n \mathbb{1}_{\mathcal{A}_i} \Pi_{\mathcal{A}_i}(f|_{\mathcal{A}_i}), \quad (5.6)$$

where $f|_{\mathcal{A}_i}$ represents the restriction of the domain of function f to the partition set \mathcal{A}_i .

5.3. Approximation Algorithm. An advantage of the interpolation operator in (5.6) is that $\Pi_{\Upsilon}(f)$ is fully characterised by the interpolation coefficients α_{ij} , since

$$\Pi_{\Upsilon}(f) = \sum_{i=1}^n \sum_{j=1}^h \alpha_{ij} \phi_{ij} \mathbb{1}_{\mathcal{A}_i}.$$

The set of interpolation coefficients α_{ij} is computable by matrix multiplication based on the data set $\{f(s_{ij}), i \in \mathbb{N}_n, j \in \mathbb{N}_h\}$. More precisely, we have $[f(s_{iv})]_{v \in \mathbb{N}_h} = \mathcal{Q}_i[\alpha_{ij}]_{j \in \mathbb{N}_h}$ with the interpolation matrices $\mathcal{Q}_i = [\phi_{ij}(s_{iv})]_{v,j \in \mathbb{N}_h}$. These matrices depend solely on the interpolation points s_{ij} and on the basis functions ϕ_{ij} evaluated at these points and can be computed off-line (see step 4 in Algorithm 1, to be discussed shortly). Moreover, values of the function f only at the interpolation points s_{ij} are sufficient for the computation of α_{ij} .

Let us now focus on the recursion in (5.4), namely $\psi_{t+1}^h = \Pi_{\Upsilon} \mathcal{R}_{\Upsilon}(\psi_t^h)$, given the initialisation $\psi_0^h = \mu_0$, for the approximate computation of the value functions. This recursion indicates that the approximate functions $\psi_t^h, t \in \mathbb{N}_N$, belong to the image of the operator Π_{Υ} , and as such can be expressed as

$$\psi_t^h = \sum_{i=1}^n \sum_{j=1}^h \alpha_{ij}^t \phi_{ij} \mathbb{1}_{\mathcal{A}_i},$$

where α_{ij}^t denote the interpolation coefficients referring to function ψ_t^h (at step t). This suggests that we need to store and update the coefficients α_{ij}^t for each iteration in (5.4).

Writing the recursion in the form $\psi_{t+1}^h = \Pi_{\Upsilon}(\mathcal{R}_{\Upsilon}(\psi_t^h))$ indicates that the function ψ_{t+1}^h is in the range of the projection Π_{Υ} . Therefore, it is sufficient to evaluate the function $\mathcal{R}_{\Upsilon}(\psi_t^h)$ over the interpolation points in order to compute the coefficients α_{ij}^{t+1} . In the following expressions, the pair i, u indicates the indices of related partition sets, namely $\mathcal{A}_i, \mathcal{A}_u$, whereas the pair of indices j, v show the ordering positions within partition sets. For an arbitrary interpolation point s_{uv} we have:

$$\mathcal{R}_{\Upsilon}(\psi_t^h)(s_{uv}) = \int_{\Upsilon} t_{\mathfrak{s}}(s_{uv}|s) \psi_t^h(s) ds = \sum_{i=1}^n \sum_{j=1}^h \alpha_{ij}^t \int_{\mathcal{A}_i} t_{\mathfrak{s}}(s_{st}|s) \phi_{ij}(s) ds.$$

Introducing the following quantities

$$P_{ij}^{uv} = \int_{\mathcal{A}_i} t_{\mathfrak{s}}(s_{uv}|s) \phi_{ij}(s) ds,$$

we can succinctly express

$$\mathcal{R}_{\Upsilon}(\psi_t^{\mathfrak{h}})(s_{uv}) = \sum_{i=1}^n \sum_{j=1}^h \alpha_{ij}^t P_{ij}^{uv} \doteq \beta_{uv}^{t+1}.$$

Algorithm 1 provides a general procedure for the discrete computation of the interpolation coefficients and of the approximate value functions.

Algorithm 1 Approximate computation of the functions $\psi_t^{\mathfrak{h}}$

Require: Density function $t_{\mathfrak{s}}(\bar{s}|s)$, set Υ

- 1: Select a finite n -dimensional partition of the set $\Upsilon = \cup_{i=1}^n \mathcal{A}_i$ (\mathcal{A}_i are non-overlapping)
- 2: For each \mathcal{A}_i , select interpolation basis functions ϕ_{ij} and points $s_{ij} \in \mathcal{A}_i$, where $j \in \mathbb{N}_h$
- 3: Compute $P_{ij}^{uv} = \int_{\mathcal{A}_i} t_{\mathfrak{s}}(s_{uv}|s) \phi_{ij}(s) ds$, where $i, u \in \mathbb{N}_n$ and $j, v \in \mathbb{N}_h$
- 4: Compute a matrix representation for the operators $\Pi_{\mathcal{A}_i}$, namely $\mathcal{Q}_i = [\phi_{ij}(s_{iv})]_{v,j}$
- 5: Set $t = 1$ and $\beta_{ij}^1 = \int_{\Upsilon} t_{\mathfrak{s}}(s_{ij}|s) \mu_0(s) ds$, for all i, j
- 6: **if** $t \leq N$ **then**
- 7: Compute interpolation coefficients α_{ij}^t ; based on equation $\mathcal{Q}_i[\alpha_{ij}^t]_{j \in \mathbb{N}_h} = [\beta_{iv}^t]_{v \in \mathbb{N}_h}$, given β_{ij}^t and matrices \mathcal{Q}_i in step 4
- 8: Compute values β_{uv}^{t+1} as $\beta_{uv}^{t+1} = \sum_i \sum_j \alpha_{ij}^t P_{ij}^{uv}$, for all u, v
- 9: $t = t + 1$
- 10: **end if**

Ensure: Approximate functions $\psi_t^{\mathfrak{h}} = \sum_i \sum_j \alpha_{ij}^t \phi_{ij} \mathbb{1}_{\mathcal{A}_i}$, $t \in \mathbb{N}_N$

It is possible to simplify Algorithm 1 when the interpolation matrices \mathcal{Q}_i are nonsingular. Let us transform the basis $\{\phi_{ij}, j \in \mathbb{N}_h\}$ to its equivalent basis using matrix $(\mathcal{Q}_i^T)^{-1}$. The interpolation matrices corresponding to the new basis will be the identity matrix. In other words, the new basis functions admit $\mathcal{Q}_i = \mathbb{I}_h$ and thus step 4 can be skipped, and that the main update (steps 7 and 8) can be simplified as follows:

$$\alpha_{uv}^{t+1} = \sum_{i=1}^n \sum_{j=1}^h \alpha_{ij}^t P_{ij}^{uv}, \quad \forall u \in \mathbb{N}_n, v \in \mathbb{N}_h.$$

In Algorithm 1, the interpolation points s_{ij} are in general pair-wise distinct. By extending the domain of interpolation \mathcal{A}_i to its closure $\bar{\mathcal{A}}_i$, it is legitimate to use boundary points as interpolation points, which can lead to a reduction of the number of integrations required in Algorithm 1. In the ensuing sections, we will exploit this feature by specifically selecting equally spaced interpolation points.

6. SPECIAL FORMS OF THE PROJECTION OPERATOR

In this section we leverage known interpolation theorems for the construction of the projection operator Π_{Υ} : this should both yield useful schemes for a number of standard models, and further help with the understanding of the details discussed in the previous section.

6.1. Piecewise Constant Approximations. We focus on the special case of the approximation of a function by a piecewise constant one, which has inspired Section 4. Let us select the basis functions $\phi_{ij}(s) = 1$ for all $i \in \mathbb{N}_n, j \in \mathbb{N}_1$ – the cardinality of these sets of basis functions is simply equal to $h = 1$ (we eliminate the corresponding indices when appropriate). In this case the matrix operators $\mathcal{Q}_i, i \in \mathbb{N}_n$, (cf. step 4 in Algorithm 1) correspond to the identity matrix, and the projection operator Π_Υ becomes

$$\Pi_\Upsilon(f) = \sum_{i=1}^n f(s_i) \mathbf{1}_{\mathcal{A}_i}, \quad \forall f \in \mathbb{B}(\Upsilon), \quad (6.1)$$

where the quantities P_{ij}^{uv} (cf. step 3 in Algorithm 1) form a square matrix (see step 3 in Algorithm 2). The procedure is detailed in Algorithm 2, while the associated error is formally quantified in Theorem 6.1.

Algorithm 2 Piecewise constant computation of the functions ψ_t^h

Require: Density function $t_s(\bar{s}|s)$, set Υ

- 1: Select a finite n -dimensional partition of the set $\Upsilon = \cup_{i=1}^n \mathcal{A}_i$ (\mathcal{A}_i are non-overlapping)
- 2: For each \mathcal{A}_i , select one representative point $s_i \in \mathcal{A}_i$
- 3: Compute matrix $P = [P(i, j)]_{i,j}$ with entries $P(i, j) = \int_{\mathcal{A}_i} t_s(s_j|s) ds$, where $i, j \in \mathbb{N}_n$
- 4: Set $t = 1$ and $\alpha_1(i) = \int_{\Upsilon} t_s(s_i|s) \mu_0(s) ds$, for all i
- 5: **if** $t < N$ **then**
- 6: Compute the row vector $\alpha_{t+1} = [\alpha_{t+1}(i)]_i$ based on $\alpha_{t+1} = \alpha_t P$
- 7: $t = t + 1$
- 8: **end if**

Ensure: Approximate functions $\psi_t^h = \sum_{i=1}^n \alpha_t(i) \mathbf{1}_{\mathcal{A}_i}, t \in \mathbb{N}_N$

Theorem 6.1. *Suppose the density function $t_s(\cdot|s)$ satisfies the Lipschitz continuity assumption (2.4) with constant λ_f . Then the projection operator (6.1) satisfies the inequality*

$$\|\Pi_\Upsilon(t_s(\cdot|s)) - t_s(\cdot|s)\|_\infty \leq \lambda_f \delta, \quad \forall s \in \Upsilon,$$

where $\delta = \max_i \delta_i$ is the partition diameter of $\cup_{i=1}^n \mathcal{A}_i = \Upsilon$, with $\delta_i = \sup\{\|s - s'\| : s, s' \in \mathcal{A}_i\}$. Theorem 5.1 ensures that the approximation error of Algorithm 2 is upper bounded by the quantity

$$\|\mu_t - \psi_t^h\|_\infty \leq E_t^h = \lambda_f \delta \kappa(t, M_f), \quad t \in \mathbb{N}_N,$$

with the constant M_f defined in Assumption 2.3.

Notice that the error E_t^h of Theorem 6.1 reduces to E_t in Theorem 4.2 when employing the quantities in (4.5) to relax the continuity assumption on π_0 .

Let us compare Algorithms 1 and 2 in terms of their computational complexity. Algorithm 1 requires $nh(nh + 1)$ integrations in the marginalisation steps (3 and 5), whereas $n(n + 1)$ integrations are required in Algorithm 2. Furthermore, steps 4 and 7 in Algorithm 1 can be skipped by using proper equivalent basis functions, whereas these steps are not needed at all in Algorithm 2. As a bottom line, higher interpolation orders increase the computational complexity of the approximation procedure, however this can as well lead to a lower global approximation error. From a different perspective, since the global approximation error depends on the local partitioning sets (their diameter and the local

continuity of the density function), for a given error higher interpolation procedures may require less partitions sets.

As a final note, comparing the transition probabilities of (4.1) with quantities $P(i, j)$ in step 3 of Algorithm 2 reveals that the Markov chain abstraction presented in Section 4 is a special case of Algorithm 2. More precisely, the *mean value theorem* for integration ensures the existence of representative points s_i such that $P(i, j)$ of Algorithm 2 is equal to P_{ij} in (4.1).

6.2. Higher-order Approximations for One-Dimensional Systems. We study higher-order interpolations over the real axis, where the partition sets \mathcal{A}_i are real-valued intervals. We use this simple setting to quantify the error related to the approximate computation of the functions μ_t . We select equally spaced points as the interpolation points and employ polynomial basis functions within each interval.

Consider a one dimensional Markov process, $\mathcal{S} = \mathbb{R}$, with a partitioning of $\Upsilon = \cup_{i=1}^n \mathcal{A}_i$ which is such that $\mathcal{A}_i = [a_i, b_i]$. Define the interpolation operator Π_Υ of (5.6) over the polynomial basis functions $\phi_{ij}(s) = s^{j-1}$, $i \in \mathbb{N}_n, j \in \mathbb{N}_h, h \geq 2$, (or their equivalent Lagrange polynomials [24]) using equally spaced interpolation points $s_{ij} \in \mathcal{A}_i$,

$$s_{ij} = a_i + (j-1) \frac{b_i - a_i}{h-1}, \quad j \in \mathbb{N}_h.$$

The following result can be adapted from [24].

Theorem 6.2. *Assume that the density function $t_s(\cdot|s)$ is h -times differentiable and define the constant*

$$\mathcal{M}_h = \max_{s, \bar{s} \in \Upsilon} \left| \frac{\partial^h t_s(\bar{s}|s)}{\partial \bar{s}^h} \right|.$$

The interpolation operator Π_Υ , constructed with polynomial basis functions and equally spaced interpolation points, satisfies the inequality

$$\|\Pi_\Upsilon(t_s(\cdot|s)) - t_s(\cdot|s)\|_\infty \leq \mathcal{E}^h = \frac{\mathcal{M}_h}{4h} \left(\frac{\delta}{h-1} \right)^h, \quad \forall s \in \Upsilon,$$

where $\delta = \max_i \delta_i$, with $\delta_i = b_i - a_i$, $i \in \mathbb{N}_n$, and where h is the cardinality of the set of basis functions.

Theorem 6.2 provides the necessary ingredients for Theorem 5.1, leading to the quantification of the approximation error: employing Algorithm 1 with equally spaced points and polynomial basis functions of degree less than h , the approximation error is upper bounded by the quantity

$$\|\mu_t - \psi_t^h\|_\infty \leq E_t^h = \mathcal{E}^h \kappa(t, M_f^h), \quad t \in \mathbb{N}_N,$$

with the constant M_f^h defined in (5.3) and computed for this particular choice of basis functions and points.

It is worth highlighting that, unlike the piecewise constant case of Section 4, with higher-order approximation approaches the global error is a nonlinear function of the partition size δ , namely it depends on a power of the partition size contingent on the order of the selected interpolation operator. As such, its convergence speed, as δ is decreased, increases over that of the piecewise constant case.




Figure 3: Comparison of the first-order (affine) approximation $\psi_N^h(\cdot)$ versus the Markov chain abstraction (zero-order approximation, constant) $\psi_N(\cdot)$ of the state density function $\pi_N(\cdot)$ (derived analytically), for $N = 5$ and parameters $a = 1.2$ (left) and $a = 0.8$ (right).

Example 2.1 (Continued). We partition the set $\Upsilon = \cup_{i=1}^n \mathcal{A}_i$ for the one dimensional system of Example 2.1 with the intervals $\mathcal{A}_i = [a_i, b_i]$. We select interpolation points $\{a_i, a_{i+1}\}$ with polynomial basis functions $\{1, s\}$, leading to piecewise affine approximations (namely, first-order interpolation with $h = 2$) of the density functions $\pi_t(\cdot)$. This set of basis functions can be equivalently transformed to

$$\Phi = \left\{ \phi_{i1} = \frac{b_i - s}{b_i - a_i}, \phi_{i2} = \frac{s - a_i}{b_i - a_i} \right\},$$

to obtain $\mathcal{Q}_i = \mathbb{I}_2$. The constant M_f^h has the same value as $M_f = 1/a$ and the quantity \mathcal{M}_2 in Theorem 6.2 is $\mathcal{M}_2 = 1/\sigma^3\sqrt{2\pi}$. The error related to this first-order approximation can be upper bounded as

$$\|\pi_t - \psi_t^h\|_\infty \leq \kappa(t, M_f) \left[\frac{\delta^2}{\sigma^3 8\sqrt{2\pi}} + \frac{\phi_1(\alpha)}{\sigma} \right]. \quad (6.2)$$

Notice that the first part of the error in (6.2), which specifically relates to the first-order approximation, is proportional to δ^2 – this improves the error bound computed in (4.6). Algorithm 1 is implemented for this linear system with the aforementioned parameters $b = 0, \beta_0 = 0, \gamma_0 = 1, \sigma = 0.1$, and the time horizon $N = 5$. The errors corresponding to the values $a = 1.2$ and $a = 0.8$ are analytically upper-bounded, for any $t \in \mathbb{N}_N$, as $\|\pi_t - \psi_t^h\|_\infty \leq 179\delta^2 + 35.9\phi_1(\alpha)$ and as $\|\pi_t - \psi_t^h\|_\infty \leq 409.3\delta^2 + 82.1\phi_1(\alpha)$, respectively. The plots in Figure 3 display the first- and zero-order approximations of the density function $\pi_N(\cdot)$ and compare it with the analytic solution for two different values $a = 1.2$ (left) and $a = 0.8$ (right). The partition size $n = 25$ and parameter $\alpha = 2.4$ have been selected in order to illustrate the differences of the two approximation methods in Figure 3, but may be increased at will in order to decrease the related error bound to match a desired value.

6.3. Bilinear Interpolation for Two-Dimensional Systems. We directly tailor the results of Section 5 to a general two-dimensional Markov process, where $s = (s_1, s_2) \in \mathcal{S} =$

\mathbb{R}^2 . Assume that set Υ is replaced by a superset that is comprised of a finite union of rectangles: this replacement does not violate the bound on the truncation error formulated in Section 3. Consider a uniform partition (using squared partition sets of size δ) for the set Υ . We employ a bilinear interpolation within each partition set $\mathcal{A}_i = [a_{i1}, b_{i1}] \times [a_{i2}, b_{i2}]$ with basis $\{\phi_{i1}(s) = 1, \phi_{i2}(s) = s_1, \phi_{i3}(s) = s_2, \phi_{i4}(s) = s_1 s_2\}$, $i \in \mathbb{N}_n$ (or their equivalent Lagrange polynomials [11]). Assume that the density function $t_s(\cdot|s)$ is partially differentiable and define the following bounds on its derivatives

$$\left| \frac{\partial^2 t_s(\bar{s}|s)}{\partial \bar{s}_k^2} \right| \leq \mathcal{M}_2^k, \quad \left| \frac{\partial^3 t_s(\bar{s}|s)}{\partial \bar{s}_k^2 \partial \bar{s}_{3-k}} \right| \leq \mathcal{M}_3^k, \quad k \in \mathbb{N}_2, \forall s, \bar{s} \in \Upsilon.$$

The operator Π_Υ in (5.6), constructed with bilinear interpolation within each partition set, satisfies the inequality

$$\|\Pi_\Upsilon(t_s(\cdot|s)) - t_s(\cdot|s)\|_\infty \leq \mathcal{E}^h = \frac{\delta^2}{16} (\mathcal{M}_2^1 + \mathcal{M}_2^2) + \frac{\delta^3}{8\sqrt{2}} (\mathcal{M}_3^1 + \mathcal{M}_3^2), \quad \forall s \in \Upsilon,$$

where $\delta = \max_i \delta_i$, with $\delta_i = [(b_{i1} - a_{i1})^2 + (b_{i2} - a_{i2})^2]^{1/2}$, $i \in \mathbb{N}_n$. We implement Algorithm 1 for two-dimensional processes using bilinear interpolation: the approximation error is upper-bounded by the quantity

$$\|\mu_t - \psi_t^h\|_\infty \leq E_t^h = \mathcal{E}^h \kappa(t, M_f^h), \quad t \in \mathbb{N}_N,$$

with the constant M_f^h defined in (5.3) and computed for this particular choice of basis functions and points. It can be proved that for bilinear interpolation basis functions, the constant M_f^h of (5.3) is upper bounded by M_f of Assumption 2.3, and thus can be replaced by this quantity in the error computation.

6.4. Trilinear Interpolation for Three-Dimensional Systems. We now apply the results of Section 5 to a general three-dimensional Markov process, where $s = (s_1, s_2, s_3) \in \mathcal{S} = \mathbb{R}^3$. Again we replace the set Υ by a superset that is comprised of a finite union of boxes, without violating the bound on the truncation error formulated in Section 3. Consider a uniform partition (using cubic sets of size δ) for the set Υ . We employ a trilinear interpolation within each partition set with basis functions

$$\Phi = \{1, s_1, s_2, s_3, s_1 s_2, s_2 s_3, s_3 s_1, s_1 s_2 s_3\}.$$

Assume that the density function $t_s(\cdot|s)$ is partially differentiable and define the following bounds on its derivatives

$$\left| \frac{\partial^2 t_s(\bar{s}|s)}{\partial \bar{s}_i^2} \right| \leq \mathcal{M}_2^i, \quad \left| \frac{\partial^3 t_s(\bar{s}|s)}{\partial \bar{s}_i^2 \partial \bar{s}_j} \right| \leq \mathcal{M}_3^{ij}, \quad \left| \frac{\partial^3 t_s(\bar{s}|s)}{\partial \bar{s}_1 \partial \bar{s}_2 \partial \bar{s}_3} \right| \leq \mathcal{M}_3, \quad i, j \in \mathbb{N}_3, i \neq j.$$

We implement Algorithm 1 for three-dimensional processes using trilinear interpolation in the operator Π_Υ (5.6). The approximation error is then upper bounded by the quantity

$$\|\mu_t - \psi_t^h\|_\infty \leq \mathcal{E}^h \kappa(t, M_f^h) \quad t \in \mathbb{N}_N,$$

with the constant

$$\mathcal{E}^h = \frac{\delta^2}{24} (\mathcal{M}_2^1 + \mathcal{M}_2^2 + \mathcal{M}_2^3) + \frac{\delta^3}{12\sqrt{3}} (\mathcal{M}_3^{12} + \mathcal{M}_3^{21} + \mathcal{M}_3^{23} + \mathcal{M}_3^{32} + \mathcal{M}_3^{13} + \mathcal{M}_3^{31} + 3\mathcal{M}_3).$$

Similar to the bilinear interpolation case, the constant M_f^h can be replaced by M_f of Assumption 2.3 in the error computation.

7. APPLICATION OF THE FORMAL APPROXIMATION PROCEDURE TO THE PROBABILISTIC INVARIANCE PROBLEM

The problem of probabilistic invariance (or, equivalently, safety) for general Markov processes has been theoretically characterised in [4] and further investigated computationally in [2, 10, 11, 12]. With reference to a discrete-time Markov process \mathcal{M}_s over a continuous state space \mathcal{S} , and to a safe set $\mathcal{A} \in \mathcal{B}(\mathcal{S})$, the goal is to quantify the probability

$$p_s^N(\mathcal{A}) = \mathbb{P}\{s(t) \in \mathcal{A}, \text{ for all } t \in \mathbb{Z}_N | s(0) = s\}.$$

More generally, it is of interest to quantify the probability $p_{\pi_0}^N(\mathcal{A})$, where the initial condition of the process $s(0)$ is a random variable characterised by the density function $\pi_0(\cdot)$. In Section 7.1 we present a forward computation of probabilistic invariance by application of the approximation procedure above, then review results on backward computation [2, 10, 11, 12] in Section 7.2. We conclude in Section 7.3 with a comparison of the two approaches.

7.1. Forward Computation of Probabilistic Invariance. The technique for approximating the density function of a process in time can be easily employed for the approximate computation of probabilistic invariance. Define functions $W_t : \mathcal{S} \rightarrow \mathbb{R}^{\geq 0}$, characterised as

$$W_{t+1}(\bar{s}) = \mathbf{1}_{\mathcal{A}}(\bar{s}) \int_{\mathcal{S}} W_t(s) t_s(\bar{s}|s) ds, \quad W_0(\bar{s}) = \mathbf{1}_{\mathcal{A}}(\bar{s}) \pi_0(\bar{s}), \quad \forall \bar{s} \in \mathcal{S}. \quad (7.1)$$

Then the solution of the problem is obtained as $p_{\pi_0}^N(\mathcal{A}) = \int_{\mathcal{S}} W_N(s) ds$. A comparison of the recursions in (7.1) and in (3.4) reveals how probabilistic invariance can be computed as a special case of the general approximation procedure in this work. In applying the procedure, the only difference consists in replacing set Υ by the safe set \mathcal{A} , and in restricting Assumption 2.3 to hold over the safe set – the solution over the complement of this set is trivially known, as such the error related to the truncation of the state space can be disregarded. The procedure consists in partitioning the safe set, in constructing the Markov chain \mathcal{M}_p as per (4.1), and in computing $\psi_t(\cdot)$ as an approximation of $W_t(\cdot)$ based on (4.2). The error of this approximation is $\|W_t - \psi_t\|_{\infty} \leq E_t$, which results in the following:

$$\left| p_{\pi_0}^N(\mathcal{A}) - \int_{\mathcal{A}} \psi_t(s) ds \right| \leq E_N \mathcal{L}(\mathcal{A}) = \kappa(N, M_f) \lambda_f \delta \mathcal{L}(\mathcal{A}) \doteq E_f.$$

Note that the sub-density functions satisfy the inequalities

$$1 \geq \int_{\mathcal{S}} W_0(s) ds \geq \int_{\mathcal{S}} W_1(s) ds \geq \dots \geq \int_{\mathcal{S}} W_N(s) ds \geq 0.$$

7.2. Backward Computation of Probabilistic Invariance. The contributions in [2, 10, 11, 12] have characterised specifications in PCTL with an alternative formulation based on backward recursions. In particular, the computation of probabilistic invariance can be obtained via the value functions $V_t : \mathcal{S} \rightarrow [0, 1]$, which are characterised as

$$V_t(s) = \mathbf{1}_{\mathcal{A}}(s) \int_{\mathcal{S}} V_{t+1}(\bar{s}) t_s(\bar{s}|s) d\bar{s}, \quad V_N(s) = \mathbf{1}_{\mathcal{A}}(s), \quad \forall s \in \mathcal{S}. \quad (7.2)$$

The desired probabilistic invariance is expressed as $p_{\pi_0}^N(\mathcal{A}) = \int_{\mathcal{S}} V_0(s) \pi_0(s) ds$. The value functions always map the state space to the interval $[0, 1]$ and they are non-increasing, namely $V_t(s) \leq V_{t+1}(s)$ for any fixed $s \in \mathcal{S}$. The contributions in [2, 10, 11, 12] discuss efficient algorithms for the approximate computation of the quantity $p_{\pi_0}^N(\mathcal{A})$, relying on

different assumptions on the model under study. The easiest and most straightforward procedure is based on the following assumption [2].

Assumption 7.1. *The conditional density function of the process is globally Lipschitz continuous with respect to the conditional state within the safe set. Namely, there exists a finite constant λ_b , such that*

$$|t(\bar{s}|s) - t(\bar{s}|s')| \leq \lambda_b \|s - s'\|, \quad \forall s, s', \bar{s} \in \mathcal{A}.$$

A finite constant M_b is introduced as $M_b = \sup_{s \in \mathcal{A}} \int_{\mathcal{A}} t_s(\bar{s}|s) d\bar{s} \leq 1$.

The procedure introduces a partition of the safe set $\mathcal{A} = \cup_{i=1}^n \mathcal{A}_i$ and extends it to $\mathcal{S} = \cup_{i=1}^{n+1} \mathcal{A}_i$, with $\mathcal{A}_{n+1} = \mathcal{S} \setminus \mathcal{A}$. Then it selects arbitrary representative points $s_i \in \mathcal{A}_i$ and constructs a finite-state Markov chain \mathcal{M}_b over the finite state space $\{s_1, s_2, \dots, s_{n+1}\}$, endowed with transition probabilities

$$P(s_i, s_j) = \int_{\mathcal{A}_j} t_s(\bar{s}|s_i) d\bar{s}, \quad P(s_{n+1}, s_j) = \delta_{(n+1)j}, \quad (7.3)$$

for all $i \in \mathbb{N}_n, j \in \mathbb{N}_{n+1}$. The error of such an approximation is [11]:

$$E_b \doteq \kappa(N, M_b) \lambda_b \delta \mathcal{L}(\mathcal{A}),$$

where δ is the max partitions diameter, and $\mathcal{L}(\mathcal{A})$ is the Lebesgue measure of set \mathcal{A} .

7.3. Comparison of the Two Approaches. We first compare the two constructed Markov chains. The Markov chain \mathcal{M}_p obtained with the abstraction from the forward approach is a special case of the Markov chain \mathcal{M}_b from the backward approach: in the latter case in fact the representative points can be selected intelligently to determine the average probability of jumping from one partition set to another. More specifically, the quantities (4.1) are a special case of those in (7.3) (based on the mean value theorem for integration). We will show that this leads to a less conservative (smaller) error bound for the approximation.

The forward computation is in general more informative than the backward computation since it provides not only the solution of the safety problem in time, but also the state distribution over the safe set. Further the forward approach may provide some insight to the solution of the infinite-horizon safety problem [28, 30] for a given initial distribution. As discussed in [30], solution of the infinite-horizon safety problem depends on the existence of absorbing subsets of the safe set. The outcome of the forward approach can provide evidence on the non existence of such subsets. Finally, the forward approach presented in Sections 3-6 for approximating density functions can be used to approximate the value functions in the recursion (7.1) over *unbounded* safe sets since we do not require the state space (thus also the safe set) to be bounded, while boundedness of the safe set is required in all the results in the literature that are based on backward computations.

Next, we compare errors and related assumptions. The error computations rely on two different assumptions: the Lipschitz continuity of the conditional density function with respect to the current state or to the next state, respectively. Further, the constants M_f and M_b are generally different and play an important role in the form of the error. M_b represents the maximum probability of remaining within a given set, while M_f is an indication of the maximum concentration of the process evolution towards one state, over a single time-step. M_b is always less than or equal to one, while M_f could be any finite positive number.

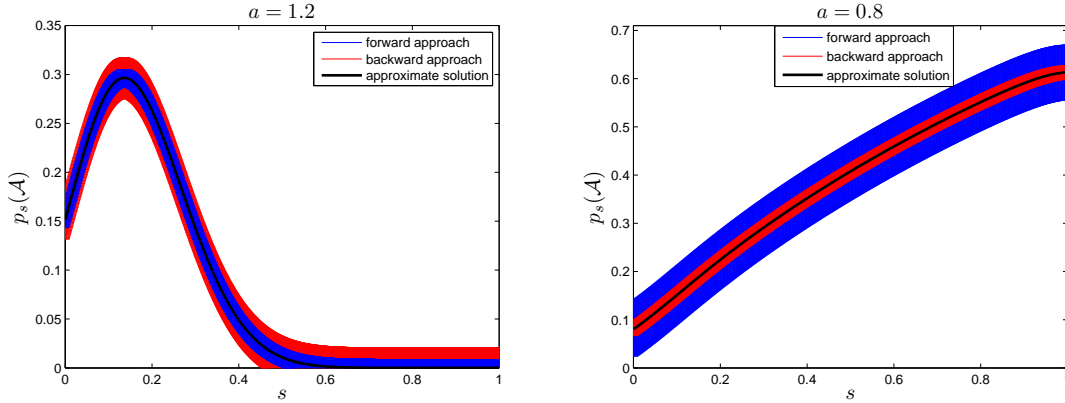


Figure 4: Approximate solution of the probabilistic invariance problem (thin black line), together with error intervals of forward (blue band) and backward (red band) approaches, for $a = 1.2$ (left) and $a = 0.8$ (right).

Example 2.1 (Continued). The constants λ_f , M_f and λ_b , M_b for the one dimensional dynamical system of Example 2.1 are

$$\lambda_f = \frac{1}{\sigma^2 \sqrt{2\pi e}}, \quad \lambda_b = a\lambda_f, \quad M_f = \frac{1}{a}, \quad M_b \leq 1.$$

If $0 < a < 1$, the system trajectories converge to an equilibrium point (in expected value). In this case the model solution has higher chances of ending up in a neighbourhood of the equilibrium in time, and the backward recursion provides a better error bound. If $a > 1$, the system trajectories tend to diverge with time. In this case the forward recursion provides a much better error bound, compared to the backward recursion.

For the numerical simulation we select a safety set $\mathcal{A} = [0, 1]$, a noise level $\sigma = 0.1$, and a time horizon $N = 10$. The solution of the safety problem for the two cases $a = 1.2$ and $a = 0.8$ is plotted in Figure 4. We have computed constants $\lambda_f = 24.20$, $M_b = 1$ (in both cases), while $\lambda_b = 29.03$, $M_f = 0.83$ for the first case and $\lambda_b = 19.36$, $M_f = 1.25$ for the second case. We have selected the center of the partition sets (distributed uniformly over the set \mathcal{A}) as representative points for the Markov chain \mathcal{M}_b . In order to compare the two approaches, we have assumed the same computational effort (related to the same partition size of $\delta = 0.7 \times 10^{-4}$), and have obtained an error $E_f = 0.008$, $E_b = 0.020$ for $a = 1.2$ and $E_f = 0.056$, $E_b = 0.014$ for $a = 0.8$. The simulations show that the forward approach works better for $a = 1.2$, while the backward approach is better suitable for $a = 0.8$. Note that the approximate solutions provided by the two approaches are very close: the difference of the transition probabilities computed via the Markov chains $\mathcal{M}_f, \mathcal{M}_b$ are in the order of 10^{-8} , and the difference in the approximate solutions (black curve in Figure 4) is in the order of 10^{-6} . This has been due to the selection of very fine partition sets that have resulted in small abstraction errors. \square

Remark 7.2. Over deterministic models, [25] compares forward and backward reachability analysis and provides insights on their differences: the claim is that for systems with significant contraction, forward reachability is more effective than backward reachability because of numerical stability issues. On the other hand, for the probabilistic models under study,

the result indicates that under Lipschitz continuity of the transition kernel the backward approach is more effective in systems with convergence in the state distribution. If we treat deterministic systems as special (limiting) instances of stochastic systems, our result is not contradicting with [25] since the Lipschitz continuity assumption on the transition kernels of probabilistic models does not hold over deterministic ones. \square

Motivated by the previous example, we study how the convergence properties of a Markov process are related to the constant M_f .

Theorem 7.3. *Assume that the initial density function $\pi_0(s)$ is bounded and that the constant M_f is finite and $M_f < 1$. If the state space is unbounded, the sequence of density functions $\{\pi_t(s) | t \geq 0\}$ uniformly exponentially converges to zero. The sequence of probabilities $\mathbb{P}\{s(t) \in \mathcal{A}\}$ and the corresponding solution of the safety problem for any compact safe set \mathcal{A} exponentially converge to zero.*

Theorem 7.3 indicates that under the invoked assumptions the probability “spreads out” over the unbounded state space as time progresses. Moreover, the theorem ensures the absence of absorbing sets [28, 30], which are indeed known to characterise the solution of infinite-horizon properties. Example 7.4 studies the relationship between constant M_f and the stability of linear stochastic difference equations.

Example 7.4. Consider the stochastic linear difference equations

$$s(t + 1) = As(t) + w(t), \quad s(\cdot), w(\cdot) \in \mathbb{R}^d,$$

where $w(\cdot)$ are i.i.d. random vectors with known distributions. For such systems $M_f = 1/|\det A|$, then the condition $M_f < 1$ implies instability of the system in expected value. Equivalently, mean-stability of the system implies $M_f \geq 1$. Note that for this class of systems $M_f > 1$ does not generally imply stability, since $\det A$ is only the product of the eigenvalues of the system. \square

The Lipschitz constants λ_f and λ_b have a different nature, as clarified in Example 7.5.

Example 7.5. Consider the dynamical system

$$s(t + 1) = f(s(t), w(t)), \quad s(\cdot), w(\cdot) \in \mathbb{R}^d,$$

where $w(\cdot)$ are i.i.d. with known distribution $t_w(\cdot)$. Suppose that the vector field $f : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is continuously differentiable and that the matrix $\frac{\partial f}{\partial w}$ is invertible. Then the *implicit function theorem* guarantees the existence and uniqueness of a function $g : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that $w(t) = g(s(t + 1), s(t))$. The conditional density function of the system in this case is [26]:

$$t_{\bar{s}}(\bar{s} | s) = \left| \det \left[\frac{\partial g}{\partial \bar{s}}(\bar{s}, s) \right] \right| t_w(g(\bar{s}, s)).$$

The Lipschitz constants λ_f, λ_b are specified by the dependence of function $g(\bar{s}, s)$ from the variables \bar{s}, s , respectively. As a special case the invertibility of $\frac{\partial f}{\partial w}$ is guaranteed for systems with additive process noise, namely $f(s, w) = f_a(s) + w$. Then $g(\bar{s}, s) = \bar{s} - f_a(s)$, λ_f is the Lipschitz constant of $t_w(\cdot)$, while λ_b is the multiplication of the Lipschitz constant of $t_w(\cdot)$ and of $f_a(\cdot)$. \square

8. CONCLUSIONS

This contribution has put forward new algorithms, based on Markov chain abstractions, for the efficient computation of approximate solutions of the state distribution function in time of Markov processes evolving over continuous state spaces. A higher-order function approximation method has also been presented, with a formal derivation of an upper bound on the associated error. The approach has been applied to the verification of a particular non-nested PCTL formula (expressing probabilistic safety or invariance), and compared with an alternative computational approach from the literature.

The authors plan to integrate the presented procedures within the software tool **FAUST**², which is developed for the formal abstraction and verification of uncountable-state stochastic processes [16]. The software enables the user to automatically export the finite-state abstracted model to existing probabilistic model checking tools, such as PRISM and MRMC [17, 18], for further quantitative analysis, verification, and synthesis objectives.

REFERENCES

- [1] A. Abate. Approximation metrics based on probabilistic bisimulations for general state-space markov processes: a survey. *Electron. Notes Theor. Comput. Sci.*, 2012.
- [2] A. Abate, J.-P. Katoen, J. Lygeros, and M. Prandini. Approximate model checking of stochastic hybrid systems. *European Journal of Control*, 6:624–641, 2010.
- [3] A. Abate, M. Kwiatkowska, G. Norman, and D. Parker. Probabilistic model checking of labelled markov processes via finite approximate bisimulations. In F. van Breugel, E. Kashefi, C. Palamidessi, and J. Rutten, editors, *Horizons of the Mind – P. Panangaden Festschrift*, Lecture Notes in Computer Science 8464, pages 40–58. Springer Verlag, 2014.
- [4] A. Abate, M. Prandini, J. Lygeros, and S. Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11):2724–2734, November 2008.
- [5] C. Baier, J.-P. Katoen, and H. Hermanns. Approximate symbolic model checking of continous-time Markov chains. In J.C.M. Baeten and S. Mauw, editors, *Concurrency Theory*, volume 1664 of *Lecture Notes in Computer Science*, pages 146–162. Springer Verlag, Berlin Heidelberg, 1999.
- [6] P. Chaput, V. Danos, P. Panangaden, and G. Plotkin. Approximating Markov processes by averaging. *J. ACM*, 61(1):5:1–5:45, January 2014.
- [7] V. Danos, J. Desharnais, and P. Panangaden. Conditional expectation and the approximation of labelled Markov processes. In Roberto Amadio and Denis Lugiez, editors, *CONCUR 2003 - Concurrency Theory*, volume 2761 of *Lecture Notes in Computer Science*, pages 477–491. Springer Berlin Heidelberg, 2003.
- [8] V. Danos, J. Desharnais, and P. Panangaden. Labelled Markov processes: Stronger and faster approximations. *Electr. Notes Theor. Comput. Sci.*, 87:157–203, 2004.
- [9] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Approximating labelled Markov processes. *Information and Computation*, 184(1):160 – 200, 2003.
- [10] S. Esmail Zadeh Soudjani and A. Abate. Adaptive gridding for abstraction and verification of stochastic hybrid systems. In *Proceedings of the 8th International Conference on Quantitative Evaluation of Systems*, pages 59–69, September 2011.
- [11] S. Esmail Zadeh Soudjani and A. Abate. Higher-Order Approximations for Verification of Stochastic Hybrid Systems. In S. Chakraborty and M. Mukund, editors, *Automated Technology for Verification and Analysis*, volume 7561 of *Lecture Notes in Computer Science*, pages 416–434. Springer Verlag, Berlin Heidelberg, 2012.
- [12] S. Esmail Zadeh Soudjani and A. Abate. Probabilistic invariance of mixed deterministic-stochastic dynamical systems. In *ACM Proceedings of the 15th International Conference on Hybrid Systems: Computation and Control*, pages 207–216, Beijing, PRC, April 2012.
- [13] S. Esmail Zadeh Soudjani and A. Abate. Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes. *SIAM Journal on Applied Dynamical Systems*, 12(2):921–956, 2013.

- [14] S. Esmail Zadeh Soudjani and A. Abate. Precise approximations of the probability distribution of a Markov process in time: an application to probabilistic invariance. In E. Abraham and K. Havelund, editors, *International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 8413 of *Lecture Notes in Computer Science*, pages 547–561. Springer Verlag, 2014.
- [15] S. Esmail Zadeh Soudjani and A. Abate. Probabilistic reach-avoid computation for partially-degenerate stochastic processes. *IEEE Transactions on Automatic Control*, 59(2):528–534, 2014.
- [16] S. Esmail Zadeh Soudjani, C. Gevaerts, and A. Abate. **FAUST²**: Formal abstractions of uncountable-state stochastic processes. In C. Baier and C. Tinelli, editors, *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 9035 of *Lecture Notes in Computer Science*, pages 272–286. Springer Verlag, 2015.
- [17] A. Hinton, M. Kwiatkowska, G. Norman, and D. Parker. PRISM: A tool for automatic verification of probabilistic systems. In H. Hermanns and J. Palsberg, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, volume 3920 of *Lecture Notes in Computer Science*, pages 441–444. Springer Verlag, Berlin Heidelberg, 2006.
- [18] J.-P. Katoen, M. Khattri, and I. S. Zapreev. A Markov reward model checker. In *IEEE Proceedings of the International Conference on Quantitative Evaluation of Systems*, pages 243–244, Los Alamos, CA, USA, 2005.
- [19] K. Koutsoukos and D. Riley. Computational methods for reachability analysis of stochastic hybrid systems. In J. Hespanha and A. Tiwari, editors, *Hybrid Systems: Computation and Control*, volume 3927 of *Lecture Notes in Computer Science*, pages 377–391. Springer Verlag, Berlin Heidelberg, 2006.
- [20] R. P. Kurshan. *Computer-Aided Verification of Coordinating Processes: The Automata-Theoretic Approach*. Princeton Series in Computer Science. Princeton University Press, 1994.
- [21] H. J. Kushner and P.G. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, New York, 2001.
- [22] M. Kvasnica, P. Grieder, and M. Baotić. Multi-parametric toolbox (MPT), 2004.
- [23] M. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. Verifying quantitative properties of continuous probabilistic timed automata. In *CONCUR’00*, volume 1877 of *Lecture Notes in Computer Science*, pages 123–137. Springer Verlag, Berlin Heidelberg, 2000.
- [24] G. Mastroianni and G.V. Milovanovic. *Interpolation Processes: Basic Theory and Applications*. Springer Verlag, 2008.
- [25] I.M. Mitchell. Comparing forward and backward reachability as tools for safety analysis. In *Proceedings of the 10th international conference on Hybrid systems: computation and control*, HSCC’07, pages 428–443, Berlin, Heidelberg, 2007. Springer-Verlag.
- [26] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, 3rd edition, 1991.
- [27] M. Prandini and J. Hu. Stochastic reachability: Theory and numerical approximation. In C.G. Cassandras and J. Lygeros, editors, *Stochastic hybrid systems*, Automation and Control Engineering Series 24, pages 107–138. Taylor & Francis Group/CRC Press, 2006.
- [28] I. Tkachev and A. Abate. On infinite-horizon probabilistic properties and stochastic bisimulation functions. In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, pages 526–531, Orlando, FL, December 2011.
- [29] I. Tkachev and A. Abate. Formula-free Finite Abstractions for Linear Temporal Verification of Stochastic Hybrid Systems. In *Proceedings of the 16th International Conference on Hybrid Systems: Computation and Control*, pages 283–292, Philadelphia, PA, April 2013.
- [30] I. Tkachev and A. Abate. Characterization and computation of infinite-horizon specifications over Markov processes. *Theoretical Computer Science*, 515(0):1–18, 2014.

APPENDIX A. PROOF OF STATEMENTS

Proof of Theorem 3.1. We prove the theorem inductively. The statement is trivial for $t = 0$ based on Assumption 2.2. Suppose it is true for t , then we prove it for $t+1$. Take $\bar{s} \in \mathcal{S} \setminus \Lambda_{t+1}$, then

$$\pi_{t+1}(\bar{s}) = \int_{\mathcal{S}} t_{\bar{s}}(\bar{s}|s) \pi_t(s) ds = \int_{\mathcal{S} \setminus \Lambda_t} t_{\bar{s}}(\bar{s}|s) \pi_t(s) ds + \int_{\Lambda_t} t_{\bar{s}}(\bar{s}|s) \pi_t(s) ds.$$

The first integral is upper bounded by $\varepsilon_t M_{\mathfrak{f}}$. The domain of the second integral implies that $(s, \bar{s}) \in \Lambda_t \times \mathcal{S}$. Combining this with the definition (3.2) of Λ_{t+1} and $\bar{s} \notin \Lambda_{t+1}$ results in $(s, \bar{s}) \notin \Gamma$. Then

$$\int_{\Lambda_t} t_{\bar{s}}(\bar{s}|s) \pi_t(s) ds \leq \epsilon \int_{\Lambda_t} \pi_t(s) ds \leq \epsilon.$$

Then we obtain $\pi_{t+1}(\bar{s}) \leq \varepsilon_t M_{\mathfrak{f}} + \epsilon = \varepsilon_{t+1}$, for all $\bar{s} \in \mathcal{S} \setminus \Lambda_{t+1}$. \square

Proof of Theorem 3.3. The initial density function μ_0 satisfies the following inequality:

$$\|\pi_0 - \mu_0\|_{\infty} = \|\pi_0 - \mathbb{1}_{\Lambda_0} \pi_0\|_{\infty} = \|\mathbb{1}_{\mathcal{S} \setminus \Lambda_0} \pi_0\|_{\infty} = \sup \{\pi_0(s), s \in \mathcal{S} \setminus \Lambda_0\} \leq \varepsilon_0.$$

Suppose μ_t satisfies the inequality. We prove that it is also true for μ_{t+1} . Take any $\bar{s} \in \mathcal{S} \setminus \Upsilon$,

$$\mathcal{S} \setminus \Upsilon \subset \mathcal{S} \setminus \Lambda_{t+1} \Rightarrow |\pi_{t+1}(\bar{s}) - \mu_{t+1}(\bar{s})| = \pi_{t+1}(\bar{s}) \leq \varepsilon_{t+1}.$$

Take any $\bar{s} \in \Upsilon$, we have

$$|\pi_{t+1}(\bar{s}) - \mu_{t+1}(\bar{s})| \leq \int_{\mathcal{S}} t_{\bar{s}}(\bar{s}|s) |\pi_t(s) - \mu_t(s)| ds \leq \varepsilon_t \int_{\mathcal{S}} t_{\bar{s}}(\bar{s}|s) ds \leq \varepsilon_t M_{\mathfrak{f}} \leq \varepsilon_{t+1}.$$

\square

Proof of Lemma 4.1. For any $t \in \mathbb{N}$ and $\bar{s}, \bar{s}' \in \mathcal{S}$,

$$|\pi_t(\bar{s}) - \pi_t(\bar{s}')| \leq \int_{\mathcal{S}} \pi_{t-1}(s) |t_{\bar{s}}(\bar{s}|s) - t_{\bar{s}'}(\bar{s}'|s)| ds \leq \lambda_{\mathfrak{f}} \|\bar{s} - \bar{s}'\| \int_{\mathcal{S}} \pi_t(s) ds = \lambda_{\mathfrak{f}} \|\bar{s} - \bar{s}'\|.$$

\square

Proof of Theorem 4.2. We use the triangle inequality as

$$\|\pi_t - \psi_t\|_{\infty} \leq \|\pi_t - \mu_t\|_{\infty} + \|\mu_t - \psi_t\|_{\infty} \leq \varepsilon_t + \|\mu_t - \psi_t\|_{\infty}.$$

Define the set of Lebesgue integrable functions \mathcal{D} and its subset \mathcal{D}_1 :

$$\mathcal{D} = \left\{ f : \mathcal{S} \rightarrow \mathbb{R}^{\geq 0} \mid \int_{\mathcal{S}} f(s) ds < \infty \right\},$$

$$\mathcal{D}_1 = \left\{ f \in \mathcal{D} \mid \int_{\mathcal{S}} t_{\bar{s}}(\bar{s}|s) f(s) ds < \infty, \text{ for all } \bar{s} \in \mathcal{S} \right\},$$

and the operators

$$\mathcal{R} : \mathcal{D}_1 \rightarrow \mathcal{D}, \quad \mathcal{R}(f)(\bar{s}) = \mathbb{1}_{\Upsilon}(\bar{s}) \int_{\mathcal{S}} t_{\bar{s}}(\bar{s}|s) f(s) ds,$$

$$\mathcal{R}_a : \mathcal{D} \rightarrow \mathcal{D}_1, \quad \mathcal{R}_a(f)(\bar{s}) = \sum_{i=1}^n \frac{\int_{\mathcal{A}_i} f(s) ds}{\mathcal{L}(\mathcal{A}_i)} \mathbb{1}_{\mathcal{A}_i}(\bar{s}).$$

Then $\mu_t(s)$ is formulated as t -times application of the operator \mathcal{R} to the initial function $\pi_0(s)$,

$$\mu_{t+1} = \mathcal{R}(\mu_t) \Rightarrow \mu_t = \mathcal{R}^t(\pi_0) \quad \forall t \in \mathbb{N}.$$

Define functions ψ_t by the recursive equation:

$$\psi_{t+1} = (\mathcal{R}_a \mathcal{R})(\psi_t) \Rightarrow \psi_t = (\mathcal{R}_a \mathcal{R})^t(\psi_0) \quad \forall t \in \mathbb{N},$$

initialised with $\psi_0 = \mathcal{R}_a(\pi_0)$. Notice that the functions ψ_t , $t \in \mathbb{N}$, are all piecewise constant due to their recursive definition. We restrict our attention to the set Υ since the supports of both functions μ_t, ψ_t are included in Υ . The goal is to prove that ψ_t satisfies (4.2) and $\|\mu_t - \psi_t\|_\infty \leq E_t$. We achieve this goal using induction. The initial function ψ_0 is of the form

$$\psi_0(s) = \mathcal{R}_a(\pi_0)(s) = \sum_{i=1}^n \frac{\int_{\mathcal{A}_i} \pi_0(v) dv}{\mathcal{L}(\mathcal{A}_i)} \mathbb{1}_{\mathcal{A}_i}(s) = \sum_{i=1}^n \frac{p_0(i)}{\mathcal{L}(\mathcal{A}_i)} \mathbb{1}_{\mathcal{A}_i}(s).$$

Suppose the statement is true for t . Then for any $s \in \mathcal{A}_j$, $j \in \mathbb{N}_n$,

$$\begin{aligned} \psi_{t+1}(s) &= (\mathcal{R}_a \mathcal{R})(\psi_t)(s) = \frac{1}{\mathcal{L}(\mathcal{A}_j)} \int_{\mathcal{A}_j} \mathcal{R}(\psi_t)(v) dv \\ &= \frac{1}{\mathcal{L}(\mathcal{A}_j)} \int_{\mathcal{A}_j} \int_{\mathcal{S}} t_{\mathfrak{s}}(v|u) \psi_t(u) dudv = \frac{1}{\mathcal{L}(\mathcal{A}_j)} \sum_{i=1}^n \int_{\mathcal{A}_j} \int_{\mathcal{A}_i} t_{\mathfrak{s}}(v|u) \psi_t(u) dudv \\ &= \frac{1}{\mathcal{L}(\mathcal{A}_j)} \sum_{i=1}^n \frac{p_t(i)}{\mathcal{L}(\mathcal{A}_i)} \int_{\mathcal{A}_j} \int_{\mathcal{A}_i} t_{\mathfrak{s}}(v|u) dudv = \frac{1}{\mathcal{L}(\mathcal{A}_j)} \sum_{i=1}^n p_t(i) P_{ij} \\ &= \frac{1}{\mathcal{L}(\mathcal{A}_j)} p_{t+1}(j). \end{aligned}$$

We have proved (4.2). The approximation error for $t = 0$ is computed using (2.3) in Assumption 2.3. Take any state $s \in \mathcal{A}_i$, $i \in \mathbb{N}_n$:

$$|\mu_0(s) - \psi_0(s)| = \left| \pi_0(s) - \frac{\int_{\mathcal{A}_i} \pi_0(v) dv}{\mathcal{L}(\mathcal{A}_i)} \right| \leq \frac{\int_{\mathcal{A}_i} |\pi_0(s) - \pi_0(v)| dv}{\mathcal{L}(\mathcal{A}_i)} \leq \lambda_0 \delta = E_0.$$

For $t \geq 1$ we can write

$$\|\mu_{t+1} - \psi_{t+1}\|_\infty = \|\mathcal{R}(\mu_t) - (\mathcal{R}_a \mathcal{R})(\psi_t)\|_\infty \leq \|\mathcal{R}(\mu_t) - \mathcal{R}(\psi_t)\|_\infty + \|\mathcal{R}(\psi_t) - (\mathcal{R}_a \mathcal{R})(\psi_t)\|_\infty.$$

The error has two terms. The first term is upper bounded as follows

$$|\mathcal{R}(\mu_t)(\bar{s}) - \mathcal{R}(\psi_t)(\bar{s})| \leq \mathbb{1}_\Upsilon(\bar{s}) \int_{\mathcal{S}} t_{\mathfrak{s}}(\bar{s}|s) |\mu_t(s) - \psi_t(s)| ds \leq E_t \int_{\mathcal{S}} t_{\mathfrak{s}}(\bar{s}|s) ds \leq M_{\Upsilon} E_t.$$

Let us focus on the second term: for any arbitrary state $\bar{s} \in \mathcal{A}_j$, $j \in \mathbb{N}_n$, we have

$$\begin{aligned} |\mathcal{R}(\psi_t)(\bar{s}) - (\mathcal{R}_a \mathcal{R})(\psi_t)(\bar{s})| &= \left| \int_{\mathcal{S}} t_{\mathfrak{s}}(\bar{s}|u) \psi_t(u) du - \frac{1}{\mathcal{L}(\mathcal{A}_j)} \int_{\mathcal{A}_j} \int_{\mathcal{S}} t_{\mathfrak{s}}(v|u) \psi_t(u) dudv \right| \\ &\leq \int_{\mathcal{S}} \psi_t(u) \left| t_{\mathfrak{s}}(\bar{s}|u) - \frac{1}{\mathcal{L}(\mathcal{A}_j)} \int_{\mathcal{A}_j} t_{\mathfrak{s}}(v|u) dv \right| du \\ &\leq \int_{\mathcal{S}} \psi_t(u) \frac{1}{\mathcal{L}(\mathcal{A}_j)} \int_{\mathcal{A}_j} |t_{\mathfrak{s}}(\bar{s}|u) - t_{\mathfrak{s}}(v|u)| dv du \leq \int_{\mathcal{S}} \psi_t(u) \lambda_{\Upsilon} \delta du \leq \lambda_{\Upsilon} \delta. \end{aligned}$$

The summation of the two upper bounds leads to that in (4.4). \square

Proof of Theorem 5.1. The definition of $\psi_t^{\mathfrak{h}}$ implies that $\psi_{t+1}^{\mathfrak{h}} = \Pi_{\Upsilon} \mathcal{R}_{\Upsilon}(\psi_t^{\mathfrak{h}})$, with $\psi_0^{\mathfrak{h}} = \mu_0$. Then (5.5) is true for $t = 0$ with $E_0^{\mathfrak{h}} = 0$. Assume that $\|\mu_t - \psi_t^{\mathfrak{h}}\|_{\infty} \leq E_t^{\mathfrak{h}}$; then for any $\bar{s} \in \Upsilon$,

$$\begin{aligned} |\mu_{t+1}(\bar{s}) - \psi_{t+1}^{\mathfrak{h}}(\bar{s})| &= |\mathcal{R}_{\Upsilon}(\mu_t)(\bar{s}) - \Pi_{\Upsilon} \mathcal{R}_{\Upsilon}(\psi_t^{\mathfrak{h}})(\bar{s})| \\ &\leq |\mathcal{R}_{\Upsilon}(\mu_t)(\bar{s}) - \Pi_{\Upsilon} \mathcal{R}_{\Upsilon}(\mu_t)(\bar{s})| + |\Pi_{\Upsilon} \mathcal{R}_{\Upsilon}(\mu_t)(\bar{s}) - \Pi_{\Upsilon} \mathcal{R}_{\Upsilon}(\psi_t^{\mathfrak{h}})(\bar{s})|. \end{aligned}$$

The first term can be upper bounded based on the linearity of the operator Π_{Υ} as

$$\begin{aligned} |\mathcal{R}_{\Upsilon}(\mu_t)(\bar{s}) - \Pi_{\Upsilon} \mathcal{R}_{\Upsilon}(\mu_t)(\bar{s})| &= \left| \int_{\Upsilon} t_{\mathfrak{s}}(\bar{s}|s) \mu_t(s) ds - \int_{\Upsilon} \Pi_{\Upsilon}(t_{\mathfrak{s}}(\bar{s}|s)) \mu_t(s) ds \right| \\ &\leq \int_{\Upsilon} \mu_t(s) |t_{\mathfrak{s}}(\bar{s}|s) - \Pi_{\Upsilon}(t_{\mathfrak{s}}(\bar{s}|s))| ds \leq \mathcal{E}^{\mathfrak{h}} \int_{\Upsilon} \mu_t(s) ds \leq \mathcal{E}^{\mathfrak{h}}. \end{aligned}$$

On the other hand, the second term is upper bounded as follows:

$$\begin{aligned} |\Pi_{\Upsilon} \mathcal{R}_{\Upsilon}(\mu_t)(\bar{s}) - \Pi_{\Upsilon} \mathcal{R}_{\Upsilon}(\psi_t^{\mathfrak{h}})(\bar{s})| &\leq \int_{\Upsilon} |\Pi_{\Upsilon}(t_{\mathfrak{s}}(\bar{s}|s))| |\mu_t(s) - \psi_t^{\mathfrak{h}}(s)| ds \\ &\leq E_t^{\mathfrak{h}} \int_{\Upsilon} |\Pi_{\Upsilon}(t_{\mathfrak{s}}(\bar{s}|s))| ds \leq E_t^{\mathfrak{h}} M_{\mathfrak{f}}^{\mathfrak{h}}. \end{aligned}$$

The addition of the two bounds leads to the statement. \square

Proof of Theorem 7.3. Based on the recursion over the density functions, we have that

$$\begin{aligned} \pi_{t+1}(\bar{s}) &\leq \sup\{\pi_t(s), s \in \mathcal{S}\} \int_{\mathcal{S}} t_{\mathfrak{s}}(\bar{s}|s) ds \leq M_{\mathfrak{f}} \sup\{\pi_t(s), s \in \mathcal{S}\} \\ &\Rightarrow \pi_t(s) \leq \sup\{\pi_0(s), s \in \mathcal{S}\} M_{\mathfrak{f}}^t, \quad \forall t \in \mathbb{N}. \end{aligned}$$

Then $\{\pi_t\}$ converges to zero uniformly exponentially, with a rate $M_{\mathfrak{f}}$. With focus on the safety problem we have $\mathbb{P}\{s(u) \in \mathcal{A} \text{ for all } u \in \mathbb{N}_t\} \leq \mathbb{P}\{s(t) \in \mathcal{A}\}$, where

$$\mathbb{P}\{s(t) \in \mathcal{A}\} = \int_{\mathcal{A}} \pi_t(s) ds \leq \mathcal{L}(\mathcal{A}) \sup_{s \in \mathcal{S}} \pi_0(s) M_{\mathfrak{f}}^t.$$

Note that if $\lim_{t \rightarrow \infty} \mathbb{P}\{s(t) \in \mathcal{S}\} = \lim_{t \rightarrow \infty} 1 = 1$, then the state-space cannot be bounded under the assumptions. \square

APPENDIX B. LIST OF SYMBOLS

$\mathbb{N} = \{1, 2, 3, \dots\}$	the set of natural numbers
$\mathbb{N}_m = \{1, 2, \dots, m\}$	the finite set of natural numbers
$\mathbb{Z}_m = \{0, 1, 2, \dots, m\}$	the finite set of non-negative numbers
$\mathcal{M}_{\mathfrak{s}}$	discrete-time Markov process
\mathcal{S}	state space of the Markov process
$\mathcal{B}(\mathcal{S})$	Borel σ -algebra on the space \mathcal{S}
$T_{\mathfrak{s}}$	stochastic kernel
\mathcal{P}	probability measure on \mathcal{S} for one-step transition of the process
$t_{\mathfrak{s}}(\bar{s} s)$	conditional density function of the process
$\pi_0 : \mathcal{S} \rightarrow \mathbb{R}^{\geq 0}$	density function of the initial state
$\pi_t : \mathcal{S} \rightarrow \mathbb{R}^{\geq 0}$	density function of $s(t)$
\mathbb{P}	probability measure on the product space \mathcal{S}^{t+1}
$\mathbf{s}(t) = [s(0), s(1), \dots, s(t)]$	bold typeset is employed for vectors
$\mathcal{M}_{\mathfrak{p}}$	Markov chain for forward computation of the density function
a, b, σ, α	parameters of the example
$w(\cdot)$	process noise
$[\beta_0, \gamma_0] \subset \mathbb{R}$	support of the initial density function π_0 in the example
$\phi_{\sigma}(u)$	Gaussian density function with zero mean and standard deviation σ
$\Gamma \subset \mathcal{S}^2, \Lambda_0 \subset \mathcal{S}$	sets used in the truncation procedure
ϵ	threshold on truncation part of $t_{\mathfrak{s}}$: $t_{\mathfrak{s}}(\bar{s} s) \leq \epsilon$ for all $(s, \bar{s}) \in \mathcal{S}^2 \setminus \Gamma$
ϵ_0	threshold on truncation part of π_0 : $\pi_0(s) \leq \epsilon_0$ for all $s \in \mathcal{S} \setminus \Lambda_0$
ϵ_t	threshold on truncation part of π_t : $\pi_t(s) \leq \epsilon_t$ for all $s \in \mathcal{S} \setminus \Lambda_t$
λ_0	Lipschitz constants of $\pi_0(s)$
$\lambda_{\mathfrak{f}}$	Lipschitz constants of $t_{\mathfrak{s}}(\bar{s} s)$ with respect to \bar{s}
$M_{\mathfrak{f}}$	upper bound for the quantities $\int_{\mathcal{S}} t_{\mathfrak{s}}(\bar{s} s) ds, \bar{s} \in \mathcal{S}$
$\ \cdot\ _{\infty}$	infinity norm
$\mathbf{1}_A(\cdot)$	indicator function of set A
Λ_t	support set of $\pi_t(\cdot)$ obtained via a recursive procedure
$\Upsilon = \cup_{t=0}^N \Lambda_t$	truncated state space
$\Xi(s)$	set-valued map $\Xi : \mathcal{S} \rightarrow 2^{\mathcal{S}}$
$\kappa(t, M_{\mathfrak{f}})$	constant equal to t for $M_{\mathfrak{f}} = 1$ and to $(1 - M_{\mathfrak{f}}^t)/(1 - M_{\mathfrak{f}})$ for $M_{\mathfrak{f}} \neq 1$
$\mu_t(\cdot)$	approximate density function after state space truncation
$\Lambda_t = [\beta_t, \gamma_t]$	support set of $\pi_t(\cdot)$ in the example
$\Upsilon = \cup_{i=1}^n \mathcal{A}_i$	selected partition
n	partition size
δ	partition diameter
$\mathcal{A}_{n+1} = \mathcal{S} \setminus \Upsilon$	partition set associated with absorbing state of Markov chain
$P = [P_{ij}]$	transition probability matrix of the Markov chain
$\delta_{(n+1)j}$	Kronecker delta function, equal to one for $j = (n+1)$ and zero otherwise
$\mathcal{L}(\cdot)$	Lebesgue measure of a set
$\mathbf{p}_0 = [p_0(1), \dots, p_0(n+1)]$	pmf of a Markov chain at $t = 0$

$\mathbf{p}_t = [p_t(1), \dots, p_t(n+1)]$	pmf of a Markov chain at time t
ψ_t	approximation of the density function π_t after abstraction
E_t	abstraction error related to partitioning
$\mathbb{B}(\mathcal{S})$	space of bounded and measurable functions on \mathcal{S}
$\mathbb{B}(\Upsilon)$	space of bounded and measurable functions on Υ
\mathcal{R}_Υ	linear operator on $\mathbb{B}(\Upsilon)$ defined as $\mathcal{R}_\Upsilon(f)(\bar{s}) = \int_\Upsilon t_{\bar{s}}(\bar{s} s)f(s)ds$, for all $\bar{s} \in \Upsilon$
$\Phi = \{\phi_1(s), \dots, \phi_h(s)\}$	set of basis functions $\phi_j \in \mathbb{B}(\Upsilon)$
$\Psi = \text{span } \Phi$	function space generated by Φ
$\Pi_\Upsilon : \mathbb{B}(\Upsilon) \rightarrow \Psi$	projection operator
\mathcal{E}^h	upper bound for $\ \Pi_\Upsilon(t_{\bar{s}}(\cdot s)) - t_{\bar{s}}(\cdot s)\ _\infty$
M_j^h	upper bound for $\int_\Upsilon \Pi_\Upsilon(t_{\bar{s}}(\bar{s} s)) ds$, for all $\bar{s} \in \Upsilon$
ψ_t^h	approximation of $\pi_t(\cdot)$ computed via higher-order methods
E_t^h	error of higher-order approximation
d	dimension of the state space
$\Pi_{\mathcal{D}}$	interpolation operator projecting any function $f : \mathcal{D} \rightarrow \mathbb{R}$ to a unique function of Ψ such that $\Pi_{\mathcal{D}}(f) = \sum_{j=1}^h \alpha_j \phi_j$
$\mathbf{f} = [f(s_i)]_{i \in \mathbb{N}_h}$	h -dim. column vector containing values of f at the interpolation points
$\boldsymbol{\alpha} = [\alpha_j]_{j \in \mathbb{N}_h}$	h -dim. column vector containing interpolation coefficients
$\mathcal{Q} = [\phi_j(s_i)]_{i,j}$	h -dimensional interpolation matrix corresponding to $\Pi_{\mathcal{D}}$
$\{\phi_{ij}, j \in \mathbb{N}_h\}$	set of basis functions for partition set \mathcal{A}_i
$\{s_{ij}, j \in \mathbb{N}_h\}$	set of interpolation points in the partition set \mathcal{A}_i
$\{\alpha_{ij}, j \in \mathbb{N}_h\}$	set of interpolation coefficients in the partition set \mathcal{A}_i
$\mathcal{Q}_i = [\phi_{ij}(s_{iv})]_{v,j \in \mathbb{N}_h}$	matrix representation of interpolation inside partition set \mathcal{A}_i
\mathbb{I}_h	h -dimensional identity matrix
$f _{\mathcal{A}_i}$	function $f : \Upsilon \rightarrow \mathbb{R}$ with domain restricted to set $\mathcal{A}_i \subset \Upsilon$
α_{ij}^t	interpolation coefficients for ψ_t^h used in Algorithm 1
β_{uv}^t	defined as $\beta_{uv}^t \doteq \mathcal{R}_\Upsilon(\psi_{t-1}^h)(s_{uv})$ and used in Algorithm 1
P_{ij}^{uv}	defined as $P_{ij}^{uv} \doteq \int_{\mathcal{A}_i} t_{\bar{s}}(s_{uv} s)\phi_{ij}(s)ds$ and used in Algorithm 1
$P = [P(i,j)]_{i,j}$	matrix with entries $P(i,j) = \int_{\mathcal{A}_i} t_{\bar{s}}(s_j s)ds$ used for piecewise constant approximation of the density functions in Algorithm 2
$\boldsymbol{\alpha}_t = [\alpha_t(i)]_i$	a row vector used in Algorithm 2 for values of ψ_t^h
δ_i	diameter of the partition set \mathcal{A}_i , $\delta_i = \sup\{\ s - s'\ , s, s' \in \mathcal{A}_i\}$
$[a_i, b_i]$	partition set \mathcal{A}_i for one-dimensional systems
\mathcal{M}_h	an upper bound for the quantity $ \partial^h t_{\bar{s}}(\bar{s} s)/\partial \bar{s}^h $, for all $s, \bar{s} \in \Upsilon$
$[a_{i1}, b_{i1}] \times [a_{i2}, b_{i2}]$	partition set \mathcal{A}_i for two-dimensional systems
$\mathcal{M}_2^k, \mathcal{M}_3^k$	upper bounds on partial derivatives of $t_{\bar{s}}$ in 2-dim. systems
$\mathcal{M}_2^i, \mathcal{M}_3^{ij}, \mathcal{M}_3$	upper bounds on partial derivatives of $t_{\bar{s}}$ in 3-dim. systems
$p_s^N(\mathcal{A})$	safety probability over the set \mathcal{A} with time horizon N and initial state s
$p_{\pi_0}^N(\mathcal{A})$	safety probability over the set \mathcal{A} with the initial state admitting the density function π_0
$W_t : \mathcal{S} \rightarrow \mathbb{R}^{\geq 0}$	sub-density functions for forward computation of the safety probability

$V_t : \mathcal{S} \rightarrow [0, 1]$	value functions for backward computation of the safety probability
$E_f = \kappa(N, M_f)\lambda_f\delta\mathcal{L}(\mathcal{A})$	error of forward computation of the safety probability
$E_b = \kappa(N, M_b)\lambda_b\delta\mathcal{L}(\mathcal{A})$	error of backward computation of the safety probability
λ_b	Lipschitz constant of $t(\bar{s} s)$ with respect to s
M_b	upper bound for $\int_{\mathcal{A}} t_s(\bar{s} s)d\bar{s}$, for all $s \in \mathcal{A}$
\mathcal{M}_b	finite-state Markov chain obtained via the backward abstraction approach
$P(s_i, s_j)$	transition probabilities of the Markov chain \mathcal{M}_b