

Safety Verification of Continuous-Space Pure Jump Markov Processes

Sadegh Esmail Zadeh Soudjani¹, Rupak Majumdar², and Alessandro Abate¹

¹ Department of Computer Science
University of Oxford, United Kingdom
{Sadegh.Soudjani,Alessandro.Abate}@cs.ox.ac.uk
² Max Planck Institute for Software Systems
rupak@mpi-sws.org

Abstract. We study the probabilistic safety verification problem for pure jump Markov processes, a class of models that generalizes continuous-time Markov chains over continuous (uncountable) state spaces. Solutions of these processes are piecewise constant, right-continuous functions from time to states. Their jump (or reset) times are realizations of a Poisson process, characterized by a jump rate function that can be both time- and state-dependent. Upon jumping in time, the new state of the solution process is specified according to a (continuous) stochastic conditional kernel. After providing a full characterization of safety properties of these processes, we describe a formal method to abstract the process as a finite-state discrete-time Markov chain; this approach is formal in that it provides a-priori error bounds on the precision of the abstraction, based on the continuity properties of the stochastic kernel of the process and of its jump rate function. We illustrate the approach on a case study of thermostatically controlled loads.

1 Introduction

Stochastic processes evolving in continuous time are used to model many phenomena in science and engineering. In recent years, there has been a lot of work in the algorithmic analysis and formal verification of such models with respect to quantitative temporal specifications. For example, the verification of continuous-time Markov chains over finite state spaces has been widely addressed in the literature against properties expressed in temporal logics such as CSL [5–7], MTL [13], and timed-automata specifications [14], and there exist efficient software tools [24, 27].

In this paper, we extend this line of work and study the class of continuous-space, pure jump Markov processes (cPJMP, for short). A cPJMP evolves in continuous time. The process starts at state $X_{t_0} = x_0$ at time $t = t_0$ and waits until a random time $t = T_1$, governed by a Poisson process depending on x_0 and possibly time-inhomogeneous, when it makes a jump to a new state $X_{T_1} = x_1$ based on a transition kernel that is conditional on the jumping time and on x_0 . Then it waits until time $t = T_2$, when it makes another jump to state $X_{T_2} = x_2$

with probability that depends on the current time and on x_1 , and so on. The states take values over a continuous domain, hence the transition kernel induces continuous measures.

cPJMPs generalize continuous-time, finite-state Markov chains (CTMCs) by allowing time-inhomogeneous behaviors (the waiting times and transition probabilities can depend on time) and allowing for general, continuous state spaces. Correspondingly, non-deterministic extensions of cPJMPs (not explicitly discussed in this work, but directly obtainable from cPJMPs) extend general-space MDPs [10] and LMPs [30] by allowing a random inter-arrival time in between stochastic resets over their continuous state space. cPJMPs can be employed in the definition and analysis of jump-diffusion processes [25]: of interest to this work, the jump component can capture event-driven uncertainties, such as corporate defaults, operational failures, or insured events [31]. It is likewise possible to obtain a cPJMP by random time sampling of a general stochastic differential equation (SDE) – indeed cPJMPs can be as well thought of as SDEs with jumps, with drift and diffusion terms that are equal to zero. This connection with diffusions driven by Wiener processes renders cPJMP relevant to areas including financial and economic modeling [31], systems biology [4], physics and chemistry [34].

We study the problem of approximately computing the bounded-time safety probability of a cPJMP by generalizing the corresponding algorithms for CTMCs. First, we show that a cPJMP can be embedded into a discrete-time, continuous-space Markov process (DTMP). In this process, we “compile away” the time inhomogeneity of the process by explicitly modeling the time as an additional state variable. Second, we characterize the bounded-time safety probability of the discrete-time Markov process as the least fixed point solution of a system of integral equations that generalize the Bellman equations for CTMCs. Finally, under Lipschitz continuity assumptions on the jump rate function and on the jump measure of the cPJMP, we show how the continuous-space discrete-time Markov process can be approximated by a *finite*-state discrete-time Markov chain (DTMC), up to any desired degree of precision. Our technical result shows a *guaranteed* upper bound on the error incurred in computing the bounded-time safety probability introduced by the finite-state approximation.

While we focus on bounded-time safety probability computation, our algorithms can be generalized to provide approximate model checking algorithms for more expressive temporal logics such as continuous-time stochastic logic (CSL) [5, 9]. We demonstrate our results on a case study from energy systems, modeling thermostatically-controlled loads as a cPJMP.

2 Pure Jump Markov Processes in Continuous Time

2.1 Model Definition - Syntax and Semantics

Let (K, \mathcal{K}) be a measurable space, where K is the (not necessarily finite) state space and \mathcal{K} is a sigma-algebra on K . Let Ω be a sample space. Let $\mathbb{R}_{\geq 0}$ be the

set of non-negative reals. We consider stochastic processes $X : \Omega \times \mathbb{R}_{\geq 0} \rightarrow K$ in continuous time. For any $t \in \mathbb{R}_{\geq 0}$, the function $X(\cdot, t) : \Omega \rightarrow K$ is a random variable, which we denote by X_t . For every $I \subseteq \mathbb{R}_{\geq 0}$ we write $\mathcal{F}_I = \sigma(X_t, t \in I)$ for the sigma-algebra on Ω generated by the stochastic process X restricted to the index set I . We suppose that for every $t \in \mathbb{R}_{\geq 0}$ and $x \in K$, a probability $\mathbb{P}^{t,x}$ is given on $(\Omega, \mathcal{F}_{[t,\infty)})$. The stochastic process $X : \Omega \times \mathbb{R}_{\geq 0} \rightarrow K$ is a (*pure*) *jump Markov process* if the following conditions hold:

- (a) \mathcal{K} contains all one-point sets and $\mathbb{P}^{t,x}(X_t = x) = 1$ for every $t \in \mathbb{R}_{\geq 0}, x \in K$.
- (b) For every $0 \leq t \leq s$ and $A \in \mathcal{K}$ the function $x \mapsto \mathbb{P}^{t,x}(X_s \in A)$ is \mathcal{K} -measurable.
- (c) **[Markov property]** For every $0 \leq u \leq t \leq s$, $A \in \mathcal{K}$ we have $\mathbb{P}^{u,x}(X_s \in A | \mathcal{F}_{[u,t]}) = \mathbb{P}^{t,X_t}(X_s \in A)$, $\mathbb{P}^{u,x}$ -a.s.
- (d) **[Pure Jump property]** For every $\omega \in \Omega$ and $t \geq 0$ there exists $\delta > 0$ such that $X_s(\omega) = X_t(\omega)$ for $s \in [t, t + \delta]$; this is equivalent to requiring that all the trajectories of X are càdlàg [11] when K is given the discrete topology (where all subsets are open).
- (e) **[Non-explosive property]** For every $\omega \in \Omega$ the number of jumps of the trajectory $t \mapsto X_t(\omega)$ is finite on every bounded interval.

Condition (a) enables us to assign probabilities to any points $x \in K$. In particular, the probability measure $\mathbb{P}^{t,x}$ assigns probability 1 to x , so that the process is initialized deterministically at x at time t . Condition (b) is essential for transporting any probability measure on X_t to the events $X_s \in A$, $A \in \mathcal{K}$, for any $t \leq s$.

Intuitively, a Markov process $X : \Omega \times \mathbb{R}_{\geq 0} \rightarrow K$ in continuous time is a pure jump process if, starting from any point $x \in K$, the process is right continuous, admits constant trajectories except at isolated jumps, and allows only for a finite number of isolated jumps within any bounded interval. A cPJMP is described by means of the joint law Q of the first jump time T_1 and of the corresponding position X_{T_1} . To proceed formally, we first fix $t \geq 0$ and $x \in K$ and define the first jump time

$$T_1(\omega) = \inf\{s > t : X_s(\omega) \neq X_t(\omega)\}, \quad (1)$$

with the convention that $T_1(\omega) = \infty$ if the indicated set is empty. Clearly, the value of T_1 depends on t . Its associated probability measure also depends on x through $\mathbb{P}^{t,x}$. Allowing this jump time to be equal to infinity requires extending the definition of the process X as follows. Take an extra dummy point $\Delta \notin K$ and redefine $X : \Omega \times \mathbb{R}_{\geq 0} \cup \{\infty\} \rightarrow K \cup \{\Delta\}$ such that $X(\omega, \infty) = X_\infty(\omega) = \Delta$ for all $\omega \in \Omega$. Then $\bar{X}_{T_1} : \Omega \rightarrow K \cup \{\Delta\}$ is well defined. Note that X_{T_1} is associated with a probability measure first through the random variable T_1 (the first jump time) and then through the process X conditioned on knowing this jump time.

On the extended space $S := (\mathbb{R}_{\geq 0} \times K) \cup \{(\infty, \Delta)\}$ we consider the smallest sigma-algebra, denoted by \mathcal{S} , containing $\{(\infty, \Delta)\}$ and all sets of $\mathcal{B}(\mathbb{R}_{\geq 0}) \otimes \mathcal{K}$ (here and in the following $\mathcal{B}(A)$ denotes the Borel sigma-algebra of a topological space A , and $Y \otimes Z$ is the product sigma-algebra of two sigma-algebras Y, Z ,

that is the smallest sigma-algebra generated by subsets of the form $A_1 \times A_2$, $A_1 \in \mathcal{Y}, A_2 \in \mathcal{Z}$). Note that this sigma-algebra \mathcal{S} is smaller than the product of two sigma-algebras defined on $\mathbb{R}_{\geq 0} \cup \{\infty\}$ and $K \cup \{\Delta\}$. The extended process X ensures that \mathcal{S} is sufficient to contain the associated probability measure of (T_1, X_{T_1}) . With these definitions, (T_1, X_{T_1}) is a random variable with values in (S, \mathcal{S}) , and its law under $\mathbb{P}^{t,x}$ is denoted by $Q(t, x, \cdot)$.

We first construct $Q(t, x, \cdot)$ for the continuous part of \mathcal{S} and later discuss how to assign probabilities to the single point (∞, Δ) . We will assume that Q is constructed starting from a given transition measure from $\mathbb{R}_{\geq 0} \times K$ to \mathcal{K} , called *rate measure* and denoted by $\nu(t, x, A)$, $t \in \mathbb{R}_{\geq 0}, x \in K, A \in \mathcal{K}$. We require that $A \mapsto \nu(t, x, A)$ is a positive measure on \mathcal{K} for all $t \in \mathbb{R}_{\geq 0}$ and $x \in K$, and that $(t, x) \mapsto \nu(t, x, A)$ is $\mathcal{B}(\mathbb{R}_{\geq 0}) \otimes \mathcal{K}$ -measurable for all $A \in \mathcal{K}$. We also assume that the rate measure ν satisfies the two conditions

- (f) $\sup\{\nu(t, x, K) | t \in \mathbb{R}_{\geq 0}, x \in K\} < \infty$ and
- (g) $\nu(t, x, \{x\}) = 0$ for all $t \in \mathbb{R}_{\geq 0}, x \in K$.

The condition (f) implies a finite number of jumps in a bounded interval, which satisfies the non-explosive condition (e) raised above. The condition (g) enforces no jump from a state to itself, which is in accordance with the definition of jump time in (1). Define

$$\lambda(t, x) = \nu(t, x, K), \quad \pi(t, x, A) = \begin{cases} \frac{\nu(t, x, A)}{\lambda(t, x)}, & \text{if } \lambda(t, x) > 0, \\ \mathbf{1}_A(x), & \text{if } \lambda(t, x) = 0, \end{cases}$$

where $\mathbf{1}_A(\cdot)$ is the indicator function of any set A . Therefore λ is a nonnegative bounded measurable function and π is a transition probability on K satisfying

$$\pi(t, x, \{x\}) = \begin{cases} 0, & \text{if } \lambda(t, x) > 0, \\ \delta_x, & \text{if } \lambda(t, x) = 0, \end{cases}$$

where δ_x is the Dirac measure at x . Function λ is called the *jump rate function*, and π the *jump measure*. Note that we have $\nu(t, x, A) = \lambda(t, x)\pi(t, x, A)$, $\forall t \in \mathbb{R}_{\geq 0}, x \in K, A \in \mathcal{K}$. Given the rate measure ν , we require that for the Markov process X we have, for $0 \leq t \leq a < b \leq \infty, x \in K, A \in \mathcal{K}$,

$$Q(t, x, (a, b) \times A) = \int_a^b \pi(s, x, A) \lambda(s, x) \exp \left[- \int_t^s \lambda(r, x) dr \right] ds, \quad (2)$$

where Q was described above as the law of (T_1, X_{T_1}) under $\mathbb{P}^{t,x}$. Note that (2) completely specifies the probability measure $Q(t, x, \cdot)$ on (S, \mathcal{S}) : indeed simple computations show that

$$\begin{aligned} \mathbb{P}^{t,x}(T_1 = \infty) &= Q(t, x, (\infty, \Delta)) \\ &:= 1 - Q(t, x, (t, \infty) \times K) = \exp \left[- \int_t^\infty \lambda(r, x) dr \right], \end{aligned} \quad (3)$$

$$\mathbb{P}^{t,x}(T_1 \in (s, \infty]) = 1 - Q(t, x, (t, s] \times K) = \exp \left[- \int_t^s \lambda(r, x) dr \right], \quad (4)$$

for all $s \geq t$ and we clearly have $\mathbb{P}^{t,x}(T_1 \leq t) = Q(t, x, [0, t] \times K) = 0$. Note that (3) assigns probability to the single point (∞, Δ) , which completes the definition of $Q(t, x, \cdot)$ on (S, \mathcal{S}) .

We may interpret (4) as the statement that T_1 has exponential distribution on $[t, \infty]$ with variable rate $\lambda(r, x), r \geq t$. Moreover, the probability $\pi(s, x, A)$ can be interpreted as the conditional probability that X_{T_1} is in $A \in \mathcal{K}$, given that the jump time $T_1 = s$, or more precisely,

$$\mathbb{P}^{t,x}(X_{T_1} \in A, T_1 < \infty | T_1) = \pi(T_1, x, A) \mathbf{1}_{T_1 < \infty}, \quad \mathbb{P}^{t,x} - \text{a.s.}$$

2.2 Examples and Related Models

Example 1. Poisson-driven differential equation [31, Section 1.7]. Let the process $\{\mathcal{N}_t \mid t \geq 0\}$ represent a standard Poisson process with homogeneous rate λ .³ Consider a pure jump process $X = \{X_t, t \in \mathbb{R}_{\geq 0}, X_t \in \mathbb{R}\}$, driven by the Poisson process \mathcal{N}_t , where its value X_t at time t satisfies the SDE

$$dX_t = c(t, X_{t-})d\mathcal{N}_t \quad \forall t \in \mathbb{R}_{\geq 0},$$

with the deterministic initial value $X_0 \in \mathbb{R}$. The function $c : \mathbb{R}_{\geq 0} \times \mathbb{R} \rightarrow \mathbb{R}$ is called the *jump coefficient*. For the case of $c(t, x) = c_0 x$ with the constant $c_0 \in \mathbb{R}_{\geq 0}$ and initial value $X_0 > 0$, the process has an explicit representation⁴

$$X_t = X_0(c_0 + 1)^{\mathcal{N}_t}, \quad \text{for } t \in \mathbb{R}_{\geq 0}.$$

From this explicit representation, we can compute properties of the process X_t , such as the probability that the process does not exceed $\alpha_{\mathfrak{h}} > 0$ within the time interval $[0, T]$. This probability is analytically computable for the above simple process: defining $\beta_{\mathfrak{h}} = (\ln \alpha_{\mathfrak{h}} - \ln X_0) / \ln(c_0 + 1)$, this probability is $\sum_{n=0}^{n \leq \beta_{\mathfrak{h}}} e^{-\lambda T} (\lambda T)^n / n!$. \square

Example 2. Compound Poisson processes [31, Section 1.1] represent a generalization of Poisson processes, with exponential waiting times between jumps but where jump sizes, rather than being deterministic, follow an arbitrary distribution. Let $\{y_n\}_{n \geq 1}$ be a sequence of independent random variables with distribution μ for all $n \geq 1$ and assume that the standard Poisson process $\{\mathcal{N}_t \mid t \geq 0\}$ with parameter $\lambda > 0$ is independent of $\{y_n\}_{n \geq 1}$. The compound Poisson process X_t is represented in the form $X_t = \sum_{n=1}^{\mathcal{N}_t} y_n$. A typical application of compound Poisson processes is to model the aggregate claim up to time t generated by a portfolio of insurance policies, where the individual claims are distributed according to μ . Let us assume the gamma distribution $y_n \sim \Gamma(a, b)$ for the individual claims [35] and answer the same safety question as in the previous example: what

³ Recall that a (homogeneous) Poisson process $\{\mathcal{N}_t \mid t \geq 0\}$ with rate λ is a Lévy process with $\mathcal{N}_0 = 0$ and $\mathbb{P}\{\mathcal{N}_t = n\} = \frac{(\lambda t)^n}{n!} e^{-\lambda t}$.

⁴ The solution can be derived observing that the process satisfies the recursive equation $X_{\tau_{n+1}} - X_{\tau_n} = c_0 X_{\tau_n}$, where the jumps occur at τ_n , $n = 1, 2, 3, \dots$ according to \mathcal{N}_t .

is the probability that the aggregate claim does not exceed $\alpha_{\mathfrak{h}} > 0$ in the time interval $[0, T]$? This probability is also analytically computable, and results in

$$e^{-\lambda T} + e^{-\lambda T} \sum_{n=1}^{\infty} \frac{\gamma(na, \alpha_{\mathfrak{h}}/b)}{\Gamma(na)} \frac{(\lambda T)^n}{n!},$$

where $\Gamma(\cdot)$ is the gamma function, and $\gamma(\cdot, \cdot)$ is the lower incomplete gamma function. \square

Notice that the safety probability is expressible analytically in the above two examples. This is first because the trajectories of the solution are always non-decreasing, and secondly since the distribution of the solution process conditioned on the value of the underlying Poisson process is computable analytically. Unfortunately in general trajectories of cPJMPs cannot be derived explicitly, and as such the safety probability is not analytically expressible. In Section 3 we provide a general characterization of the solution of the probabilistic safety problem. In Section 4 we also work out a formal approximation method to numerically compute the solution.

Example 3. Continuous-time Markov chains [8]. The class of cPJMP we consider includes, as special cases, all the time-homogeneous, nonexplosive, jump Markov processes: these correspond to a function ν not depending on the time variable t . Within this time-homogeneous case we need to retain the boundedness assumption in (f) for the rate function. Assuming further that K is a finite or countably infinite set, we obtain the class of continuous-time Markov chains characterized by the transition rates matrix $\nu(x, \{y\})_{x, y \in K}$, namely

$$\mathbb{P}^{t,x}(X_{T_1} = x', T_1 < \infty | T_1) = \frac{\nu(x, \{x'\})}{E(x)} \left[1 - e^{-E(x)t} \right] \mathbf{1}_{T_1 < \infty},$$

where $E(x) = \sum_{x' \in K} \nu(x, \{x'\})$. The probability that the system stays within a set $A \subseteq K$ in the interval $[0, T]$ can be expressed as the solution of a system of integral equations [7], which is a special case of the Bellman fixed-point equation developed in Section 3 for cPJMPs, but not expressible in closed form. \square

Example 4. cPJMP defined by dynamical systems. Consider a process X with piecewise-constant trajectories, which resets (or jumps) at time t over a space K according to a vector field $f : K \times \mathbb{R}^n \times \mathbb{R}_{\geq 0} \rightarrow K$, so that

$$x(t^+) = f(x(t^-), \zeta(t), t), \tag{5}$$

where $\zeta(\cdot)$ is a continuous-time stationary process with a given, time-independent density function. The resets for the process follow a Poisson process $\mathcal{N}_t, t \geq 0, \mathcal{N}_0 = 0$, with a rate λ depending on time t and on the continuous state of the process $x(t)$. Notice that the dependence of the vector field f on time is in accordance with [25]. The map f , together with the distribution of the process $\zeta(\cdot)$, uniquely defines a jump measure $\pi(t, x, A)$, which gives the probability of

jumping from any state x at time t to any (measurable) subset of the state space $A \subseteq K$ [26, Proposition 7.6]:

$$\pi(t, x, A) = T_\zeta(\zeta \in \mathbb{R}^n : f(x, \zeta, t) \in A),$$

where T_ζ is the distribution of the random vector $\zeta(0)$ (in fact, of any $\zeta(t)$ since the process is stationary and time-independent). \square

2.3 Embedded Discrete-Time Markov Process of a cPJMP

We have defined a cPJMP on a measurable space (K, \mathcal{K}) through the transition measure ν . The trajectories of a cPJMP are piecewise constant, which makes them fully representable by their jump times and corresponding values. It is worth studying the properties of the random variables $(T_n(w), X_{T_n}(w))$, $n \in \mathbb{N} := \{0, 1, 2, \dots\}$, where T_n is the n^{th} jump time and X_{T_n} is the corresponding value of the process. The ensuing Theorem 1 states that $(T_n, X_{T_n})_{n \in \mathbb{N}}$ can be considered as a discrete-time Markov process (DTMP) by slight extension of the definition of Q . The discrete time is indexed by nonnegative natural numbers $n \in \mathbb{N}$, as opposed to continuous time indexed by $t \in \mathbb{R}_{\geq 0}$.

Definition 1. A discrete-time Markov process $(Y_n)_{n \in \mathbb{N}}$ is uniquely defined by a triple $\mathfrak{D} = (E_\eta, \mathcal{E}_\eta, P_\eta)$, where $(E_\eta, \mathcal{E}_\eta)$ is a measurable space and $P_\eta : E_\eta \times \mathcal{E}_\eta \rightarrow [0, 1]$ is a transition kernel such that for any $y \in E_\eta$ and $A \in \mathcal{E}_\eta$, $P_\eta(y, A)$ gives the probability that $Y_{n+1} \in A$, conditioned on $Y_n = y$. E_η is called the state space of the DTMP \mathfrak{D} and the elements of E_η are the states of \mathfrak{D} . The process is time-inhomogeneous if P_η depends also on the time index n .

We adapt the following result from [23, Chapter III, Section 1, Theorem 2].

Theorem 1. Starting from $T_0 = t$ define inductively $T_{n+1} = \inf\{s > T_n : X_s \neq X_{T_n}\}$, with the convention that $T_{n+1} = \infty$ if the indicated set is empty. Under the probability $\mathbb{P}^{t,x}$, the sequence $(T_n, X_{T_n})_{n \in \mathbb{N}}$ is a DTMP in (S, \mathcal{S}) with transition kernel Q , provided we extend the definition of Q making the state (∞, Δ) absorbing, by defining $Q(\infty, \Delta, \mathbb{R}_{\geq 0} \times K) = 0$, $Q(\infty, \Delta, \{(\infty, \Delta)\}) = 1$. Note that $(T_n, X_{T_n})_{n \in \mathbb{N}}$ is time-homogeneous, although in general X is not.

Theorem 1 states that given the stochastic process $X : \Omega \times \mathbb{R}_{\geq 0} \rightarrow K$ with probability measure $\mathbb{P}^{t,x}$ as defined in Section 2, we can construct a DTMP on (S, \mathcal{S}) with the extended transition kernel Q , whose state includes jump times and jump values of X . The inverse is also true, as described next, which allows for a simple description of the process X . Suppose one starts with a DTMP $(\tau_n, \xi_n)_{n \in \mathbb{N}}$ in S with transition probability kernel Q and a given starting point $(t, x) \in \mathbb{R}_{\geq 0} \times K$. One can then define a process Z in K setting $Z_t = \sum_{n=0}^{N_\eta} \xi_n \mathbf{1}_{[\tau_n, \tau_{n+1})}(t)$, where $N_\eta := \sup\{n \in \mathbb{N} : \tau_n < \infty\}$. Then Z has the same law as the process X under $\mathbb{P}^{t,x}$.

Example 5. For a CTMC defined by its transition rate matrix $\nu(x, \{x'\})$, we get that $\pi(x, \{y\})_{x,y \in K}$ is the stochastic transition matrix of the corresponding embedded discrete-time Markov chain (DTMC). \square




Fig. 1. Transition system for the safety problem.

3 Bounded-Time Safety Probability for cPJMPs

In this section, we characterize the bounded-time safety probability for a cPJMP X , that is the quantity

$$p_{\mathcal{A}}(t_0, x_0, T) = \mathbb{P}^{t_0, x_0} \{X_u \in \mathcal{A}, \text{ for all } u \in [t_0, T] | X_{t_0} = x_0\}, \quad (6)$$

for a given initial time $t_0 \in [0, T]$, $T < \infty$, and initial state $X_{t_0} = x_0$.⁵ Note that in this setup we must account for the initial time t_0 alongside the initial state x_0 because the process is time-inhomogeneous.

For the characterization of $p_{\mathcal{A}}(t_0, x_0, T)$, we first construct the DTMP $\mathfrak{M} = (S, \mathcal{S}, Q)$ with state $s_n = (\tau_n, x_n) \in S$ according to Theorem 1. In order to formulate the safety problem over the new process \mathfrak{M} , we introduce a transition system with a set of states $\mathcal{Q} = \{\mathcal{S}, \mathcal{U}, \mathcal{W}\}$ representing *Safe*, *Unsafe*, and *Wait*. The transition system is initialized at \mathcal{W} or \mathcal{U} depending on whether the initial state of the process is in the safe set or not. A transition from \mathcal{W} to \mathcal{S} is activated if the next jump time τ^+ is outside the interval $\mathcal{T} := [0, T]$ and the next state $x^+ \in K$. A transition $\mathcal{W} \rightarrow \mathcal{U}$ is activated if $\tau^+ \in \mathcal{T}$ and the next state is outside the safe set. Finally, the self loops at all the states of \mathcal{Q} characterize all other dynamics of the transition system.

Based on the transition system in Figure 1, the quantity $p_{\mathcal{A}}(t_0, x_0, T)$ can be characterized as the probability of reaching the state \mathcal{S} in the transition system under the dynamics of \mathfrak{M} , which is equal to the likelihood associated to the set of words $\{\mathcal{W}^+\mathcal{S}\} = \{\mathcal{W}\mathcal{S}, \mathcal{W}^2\mathcal{S}, \mathcal{W}^3\mathcal{S}, \dots\}$ (we have denoted by \mathcal{W}^+ the Kleene star without ϵ). This can be written as the infinite series $p_{\mathcal{A}}(t_0, x_0, T) = \sum_{n=1}^{\infty} \mathbb{P}\{\mathcal{W}^n\mathcal{S}\}$, which equals to

$$p_{\mathcal{A}}(t_0, x_0, T) = \sum_{n=1}^{\infty} \mathbb{P}\{(s_0, s_1, \dots, s_n) \in \mathcal{G}^n \mathcal{H} | s_0 = (t_0, x_0)\}, \quad (7)$$

where $\mathcal{G} := \mathcal{T} \times \mathcal{A}$ and $\mathcal{H} := (T, \infty) \times K \cup \{(\infty, \Delta)\}$. Note that the non-explosive condition posed in (e) and reinforced by assumption (f) on $\nu(\cdot)$ implies that $\lim_{n \rightarrow \infty} \mathbb{P}\{\mathcal{W}^n\mathcal{S}\} = 0$, which is a necessary condition for the series (7) to converge.

⁵ A slight modification of the approach presented in this paper allows for verifying more general quantitative questions such as $\mathbb{P}_{\sim p}(\Phi \mathbf{U}^{(T_1, T_2)} \Psi)$, defined over any state labels Ψ, Φ and over any (possibly unbounded) time interval (T_1, T_2) – an adaptation is required on the construction of sets \mathcal{G}, \mathcal{H} in (7).

We show in the rest of this section that the infinite series (7) converges and is approximately computable via its partial sums under a mild assumption on the jump rate function (a bound on the integral of $\lambda(\cdot)$ over the interval \mathcal{T}).

The reformulation of $p_{\mathcal{A}}(t_0, x_0, T)$ as (7) indicates its close relationship with the infinite-horizon probabilistic reach-avoid specification over DTMPs. This problem is studied in [32, 33], which formulate the solution as a Bellman equation and describe convergence properties of the series based on contractivity of the stochastic operator associated to the DTMP. The next theorem can be seen as an extension of [33, Section 3.1] and presents a Bellman equation for the characterization of the safety probability $p_{\mathcal{A}}(t_0, x_0, T)$, which is an equation for the infinite-horizon reach-avoid problem over the DTMP \mathfrak{M} with the safe set $(\mathcal{G} \cup \mathcal{H}) \in \mathcal{S}$ and target set $\mathcal{H} \in \mathcal{S}$.

Theorem 2. *The solution of the probabilistic safety problem defined in (6) can be characterized as $p_{\mathcal{A}}(t_0, x_0, T) = V(t_0, x_0) - \mathbf{1}_{\mathcal{H}}(t_0, x_0)$, where the value function $V : S \rightarrow [0, 1]$ is the least solution of the fixed-point Bellman equation*

$$V(s) = \mathbf{1}_{\mathcal{H}}(s) + \mathbf{1}_{\mathcal{G}}(s) \int_S V(\bar{s})Q(s, d\bar{s}), \quad \forall s = (t, x) \in S. \quad (8)$$

In order to characterize the solution of the fixed-point equation (8), we consider the value functions $V_n : S \rightarrow [0, 1]$, $k \in \mathbb{N}$, for the finite-horizon reach-avoid probability $V_n(s) := \mathbb{P}\{\mathbb{S}, \mathbb{W}\mathbb{S}, \mathbb{W}^2\mathbb{S}, \dots, \mathbb{W}^n\mathbb{S}\}$. These functions satisfy the Bellman recursion

$$V_{n+1}(s) = \mathbf{1}_{\mathcal{H}}(s) + \mathbf{1}_{\mathcal{G}}(s) \int_S V_n(\bar{s})Q(s, d\bar{s}), \quad V_0(s) = \mathbf{1}_{\mathcal{H}}(s). \quad (9)$$

Then we have that $V(s) = \lim_{n \rightarrow \infty} V_n(s)$, where the limit is point-wise non-decreasing [33, Section 3.1]. Equation (9) indicates that the support of the value functions $V_n(\cdot)$ is bounded by the set $\mathcal{G} \cup \mathcal{H}$. These value functions are equal to one over the set \mathcal{H} and satisfy the following recursion for any $s = (t, x) \in \mathcal{G}$:

$$V_{n+1}(s) = g(s) + \int_{\mathcal{G}} V_n(\bar{s})Q(s, d\bar{s}), \quad g(t, x) := \exp \left[- \int_t^T \lambda(r, x) dr \right]. \quad (10)$$

In the following we provide an operator perspective to (10), show that the associated operator is contractive, and quantify an upper bound for the quantity $\|V - V_n\|$ as a function of n .

Let \mathbb{B} denote the space of all real-valued, bounded and measurable functions on \mathcal{G} . Then \mathbb{B} is a Banach space with a norm given by $\|f\| := \sup_{s \in \mathcal{G}} |f(s)|$ for $f \in \mathbb{B}$. An operator $\mathcal{J} : \mathbb{B} \rightarrow \mathbb{B}$ is called linear if

$$\mathcal{J}(\alpha_1 f_1 + \alpha_2 f_2) = \alpha_1 \mathcal{J}(f_1) + \alpha_2 \mathcal{J}(f_2), \quad \forall f_1, f_2 \in \mathbb{B}, \forall \alpha_1, \alpha_2 \in \mathbb{R}.$$

The quantity $\|\mathcal{J}\| = \sup_{\|f\| \leq 1} \|\mathcal{J}(f)\|$ is called the norm of the linear operator \mathcal{J} . We say that a linear operator \mathcal{J} is a contraction whenever it holds that $\|\mathcal{J}\| < 1$. We define the linear operator $\mathcal{I}_{\mathcal{G}} f(s) := \int_{\mathcal{G}} f(\bar{s})Q(s, d\bar{s})$, which is associated with equation (10). The following lemma raises assumptions on the jump rate function λ to render the operator $\mathcal{I}_{\mathcal{G}}$ contractive.

Lemma 1. For a given set $\mathcal{G} = \mathcal{T} \times \mathcal{A}$, with $\mathcal{T} = [0, T]$ and bounded safe set \mathcal{A} , suppose there exists a finite constant $\kappa \geq \sup \left\{ \int_0^T \lambda(r, x) dr, x \in \mathcal{A} \right\}$. Then the invariance operator $\mathcal{I}_{\mathcal{G}}$ is contractive with the norm $\|\mathcal{I}_{\mathcal{G}}\| \leq 1 - e^{-\kappa}$.

Theorem 3. Under the assumption of Lemma 1 the sequence $\{V_n\}_{n \in \mathbb{N}}$ satisfies

$$V_{n+1}(s) = g(s) + \mathcal{I}_{\mathcal{G}} V_n(s), \quad \forall s \in \mathcal{G},$$

and converges uniformly to $V(\cdot)$. Moreover, $\|V - V_n\| \leq (1 - e^{-\kappa})^n$ for all $n \in \mathbb{N}$.

The previous result allows us to select a sufficiently large n in order to make the difference between V and V_n smaller than a predefined threshold. For a given threshold, say ϵ_1 , one can select $N \geq \ln \epsilon_1 / \ln(1 - e^{-\kappa})$ and compute V_N . Theorem 3 then guarantees that $|V(s) - V_N(s)| \leq \epsilon_1$ for all $s \in S$. The next section is devoted to the precise computation of V_N over the uncountable state space S , for a preselected N .

4 Finite DTMCs as Formal Approximations of cPJMPs

In the previous sections we have shown that the bounded-time safety verification of the given cPJMP can be approximated by a step-bounded reach-avoid verification of a DTMP, with guaranteed error bounds. Due to lack of analytical solutions, the verification of DTMPs against PCTL specifications (amongst which reach-avoid) is studied in the literature via finite abstractions [1, 16], which result in the PCTL verification of discrete time, finite space Markov chains (DTMCs) [17, 18]. In other words, the goal of the DTMC abstraction is to provide a discrete and automated computation of the reach-avoid probability. The approach is formal in that it allows for the computation of explicit bounds on the error associated with the abstraction.

The DTMC is obtained by state-space partitioning of the DTMP: equation (10) indicates that we only need to partition the bounded set \mathcal{G} . The abstraction procedure, presented in Algorithm 1, generates a DTMC (S_a, P_a) with finite state space S_a and transition probability matrix P_a . Over this DTMC we compute $p_a(s_i, N)$, which is the probability of reaching target state s_{m+1} , while avoiding the unsafe state s_{m+2} , during the step horizon $0, \dots, N$, as a function of the initial state $s_i \in S_a$. This is obtained via a discrete version of equation (10), which boils down to via matrix manipulations [8].

We now introduce some regularity assumptions on the jump rate function $\lambda(\cdot)$ (Assumption 1) and on the jump measure $\pi(\cdot)$ (Assumption 2), which are needed to quantify the abstraction error resulting from the DTMC (S_a, P_a) .

Assumption 1 Assume the space K is endowed with a metric $\rho : K \times K \rightarrow \mathbb{R}$. Suppose the jump rate function $\lambda(\cdot)$ is bounded and Lipschitz-continuous, namely that there are finite constants Λ and h_λ such that $\lambda(t, x) \leq \Lambda$ and

$$|\lambda(t, x) - \lambda(t, x')| \leq h_\lambda \rho(x, x'),$$

for all $(t, x), (t, x') \in \mathcal{G}$.

Algorithm 1 Finite-state abstraction of the DTMP \mathfrak{M}

Require: DTMP $\mathfrak{M} = (S, \mathcal{S}, Q)$, the sets $\mathcal{G} = [0, T] \times \mathcal{A}$, $\mathcal{H} = (T, \infty) \times K \cup \{(\infty, \Delta)\}$

- 1: Select an arbitrary finite partition of the set $\mathcal{G} = \cup_{i=1}^m \mathcal{D}_i$ (\mathcal{D}_i are non-overlapping)
- 2: Define $\mathcal{D}_{m+1} := \mathcal{H}$, $\mathcal{D}_{m+2} := S \setminus (\mathcal{G} \cup \mathcal{H})$, to obtain a partition of $S = \cup_{i=1}^{m+2} \mathcal{D}_i$
- 3: For each \mathcal{D}_i , select one representative point $s_i \in \mathcal{D}_i$
- 4: Introduce DTMC (S_a, P_a) , with state space $S_a = \{s_1, s_2, \dots, s_{m+2}\}$, and transition matrix P_a :

$$P_a(i, j) = \begin{cases} Q(s_i, \mathcal{D}_j) & 1 \leq i \leq m, 1 \leq j \leq m+2 \\ 1 & i = j \in \{m+1, m+2\} \\ 0 & \text{otherwise} \end{cases}$$

5: **return** DTMC (S_a, P_a)

Assumption 1 implies the Lipschitz continuity of $g(\cdot)$.

Lemma 2. *Under Assumption 1, the function $g(\cdot)$ in (10) is Lipschitz continuous, namely for all $s = (t, x), s' = (t', x') \in \mathcal{G}$,*

$$|g(t, x) - g(t', x')| \leq Th_\lambda \rho(x, x') + \Lambda |t - t'|.$$

The next assumption is on the regularity of the jump measure $\pi(\cdot)$ through its associated density function.

Assumption 2 *Let \mathcal{K} be the Borel sigma-algebra on K . Assume that the jump measure π on (K, \mathcal{K}) given $(\mathbb{R}_{\geq 0} \times K, \mathcal{B}(\mathbb{R}_{\geq 0}) \otimes \mathcal{K})$ is an integral kernel, i.e. that there exists a sigma-finite basis measure μ on (K, \mathcal{K}) and a jointly measurable function $p : \mathbb{R}_{\geq 0} \times K \times K \rightarrow \mathbb{R}_{\geq 0}$ such that $\pi(t, x, dy) = p(t, x, y) \mu(dy)$, i.e. $\pi(t, x, A) = \int_A p(t, x, y) \mu(dy)$ for any $(t, x) \in \mathbb{R}_{\geq 0} \times K, A \in \mathcal{K}$. Suppose further that the density function $p(\tau, x, y)$ is Lipschitz-continuous, namely that there exists a finite constant h_p , such that*

$$|p(\tau, x, y) - p(\tau, x', y)| \leq h_p \rho(x, x'), \quad \forall x, x', y \in \mathcal{A}, \tau \in \mathcal{T}.$$

Example 6. The density function $p(t, x, y)$ is computable for the dynamical system representation (5) in Example 4 under suitable assumptions on vector field f given the density function of $\zeta(\cdot)$ [19, 21]. \square

Remark 1. Assumption 2 enables us to specify the conditional density function of the DTMP $(T_n, X_{T_n})_{n \in \mathbb{N}}$ as

$$t_s(\tau, y|t, x) = p(\tau, x, y) \lambda(\tau, x) \exp \left[- \int_t^\tau \lambda(r, x) dr \right] \mathbf{1}_{[t, \infty)}(\tau),$$

which gives the integral representation of the stochastic kernel of the process as $Q(t, x, (a, b), A) = \int_a^b \int_A t_s(\tau, y|t, x) \mu(dy) d\tau$. \square

Using Assumptions 1, 2, and its consequences Theorem 3 and Lemmas 1, 2, we finally establish the following result for the error computation of the abstraction.

Theorem 4. *Under Assumptions 1 and 2, the following inequality holds:*

$$|p_{\mathcal{A}}(t_0, x_0, T) - p_a(s_{\tau}, N)| \leq (1 - e^{-\kappa})^N + N(h_x \delta_x + h_t \delta_t), \quad \forall (t_0, x_0) \in \mathcal{G},$$

where $h_x = h_p \mu(\mathcal{A}) + 3Th_\lambda$, $h_t = 3\Lambda$, whereas κ is defined in Lemma 1. The constants δ_x, δ_t denote the partition diameters of state-space and time, namely

$$\begin{aligned} \delta_x &= \sup\{\rho(x, x'), \quad \forall (\tau, x), (\tau, x') \in \mathcal{D}_i, i = 1, 2, \dots, m\}, \\ \delta_t &= \sup\{|\tau - \tau'|, \quad \forall (\tau, x), (\tau', x) \in \mathcal{D}_i, i = 1, 2, \dots, m\}. \end{aligned}$$

In the inequality above, s_{τ} is the representative point of the partition set to which the state (t_0, x_0) belongs, and $p_a(s_{\tau}, N)$ is the reach-avoid probability computed over the DTMC (S_a, P_a) with finite step-horizon N .

Notice that there are two terms contributing to the error in Theorem 4. The first term is caused by replacing the discrete infinite-step reach-avoid problem with an N -step one. The second term results from the DTMC abstraction. Augmenting the number of steps N decreases the first term exponentially and increases the second term linearly: as such, this upper bound on the error can be tuned by selecting a sufficiently large step-horizon N , and accordingly small partition diameters δ_t, δ_x .

5 Case Study: Thermostatically Controlled Loads

Thermostatically Controlled Loads (TCLs) have shown potential to be engaged in power system services such as load shifting, peak shaving, and demand response programs. Recent studies have focused on the development of models for aggregated populations of TCLs [12, 20, 28]. Formal abstraction techniques have also been employed to verify properties of TCL models [2, 20]. We employ the model of a TCL as the case study in this paper. The model describes the continuous-time evolution of the temperature in a TCL by a linear SDE. The value of the temperature is available to a thermostat for regulation via a network of independent asynchronous sensors [3, 29]. We recast this model as a cPJMP and quantitatively verify user comfort as a probabilistic safety problem.

Dynamical Model for the Case Study The continuous-time evolution of the temperature $\theta = \{\theta_t, t \in \mathbb{R}_{\geq 0}\}$ in a *cooling* TCL can be specified by the following linear SDE:

$$d\theta_t = \frac{dt}{RC}(\theta_a - q_t R P_{rate} - \theta_t) + \sigma dW_t, \quad (11)$$

where $\{W_t, t \in \mathbb{R}_{\geq 0}\}$ is the standard Brownian motion, θ_a is the ambient temperature, C and R indicate the thermal capacitance and resistance, P_{rate} is the rate of energy transfer, and σ is standard deviation of the noise term. The process $\{q_t, t \in \mathbb{R}_{\geq 0}\}$ represents the state of the thermostat at time t , $q_t \in \{0, 1\}$ for

OFF and ON modes (the latter meaning that the cooler is functioning), respectively. For a given temperature θ_t at time t and a fixed mode q_t , the temperature at time $s \geq t$ is characterized by the solution of (11), namely

$$\theta_s = a\theta_t + (1 - a)(\theta_a - q_t RP_{rate}) + w_s,$$

where $a = \exp[-(s - t)/RC]$ and $w_s \sim \mathcal{N}(0, \frac{1}{2}\sigma^2 RC(1 - a^2))$.

We assume the value of temperature is available to the thermostat via a network of sensors at possibly non-uniform time samples $\{\tau_n, n \in \mathbb{N}\}$. For a network of independent and asynchronous sensors, the time between two consecutive available values of temperature $(\tau_{n+1} - \tau_n)$, when the number of sensors is large, can be approximated by an exponential distribution [3, 29]. We assume that the associated rate depends on temperature, $\lambda(\theta_{\tau_n})$, where θ_{τ_n} is the latest available temperature (at time τ_n).

The temperature of the cooling TCL is regulated by updating the thermostat mode via the equation $q_{\tau_{n+1}} = f(q_{\tau_n}, \theta_{\tau_{n+1}})$, which is based on discrete switching

$$f(q, \theta) = \begin{cases} 0, & \theta < \theta_s - \delta_d/2 := \theta_- \\ 1, & \theta > \theta_s + \delta_d/2 := \theta_+ \\ q, & \text{else,} \end{cases} \quad (12)$$

where θ_s denotes a given temperature set-point and δ_d a dead-band, and together characterize the temperature operating range. Then the mode q_t is a piecewise-constant and right-continuous function of time, which can change value from q_{τ_n} to $q_{\tau_{n+1}}$ at time τ_{n+1} according to the logic in (12).

cPJMP for the Case Study The values of temperature and the mode of the thermostat evolve over the *hybrid* state space $K = \{0, 1\} \times \mathbb{R}$, namely a space made up of discrete *and* continuous components [2]. The temperature space \mathbb{R} is endowed with the Euclidean metric and with the Borel sigma-algebra. The jump measure of the process is an integral kernel (Assumption 2 is valid), with μ being the Lebesgue measure and with the density function

$$p(\tau - t, q, \theta, \bar{q}, \bar{\theta}) = \delta_d [\bar{q} - f(q, \bar{\theta})] \phi(\bar{\theta}; m_{\eta}(\tau - t, q, \theta), \sigma_{\eta}^2(\tau - t)),$$

where $\delta_d[\cdot]$ is the Kronecker delta function, $\phi(\cdot; \bar{m}, \bar{\sigma}^2)$ is the Gaussian density function with mean \bar{m} and variance $\bar{\sigma}^2$, and

$$m_{\eta}(u, q, \theta) = a(u)\theta + (1 - a(u))(\theta_a - qRP_{rate}), \\ \sigma_{\eta}^2(u) = 2\sigma^2 RC(1 - a(u)^2), \quad a(u) = \exp[-u/RC].$$

We are interested in quantifying a proxy for user comfort: we quantify whether the likelihood of having the temperature inside a dead-band $[\theta_-, \theta_+]$ during the time interval $[0, T]$ is greater than a given threshold. This problem can be mathematically formulated as computing the safety probability of the model over the safe set $\mathcal{A} = \{0, 1\} \times [\theta_-, \theta_+]$.

Note that the density function $\pi(\cdot)$ is slightly different from the general formulation of cPJMPs in Section 2 in that it depends on $(\tau - t)$ (through $a(\cdot)$), instead of just the jump time τ . This difference requires a slight modification of the abstraction error, which is presented next.

Computation of Probabilistic Safety We consider a jump rate function $\lambda(t, \theta) = \lambda_0 e^{-\alpha t} \cosh[2\beta(\theta - \theta_s)]$ with positive constants λ_0, α , and β . The term $e^{-\alpha t}$ models the reduction of the sampling rate of the sensors in time. The cosine hyperbolic function $\cosh[2\beta(\theta - \theta_s)]$ shows that more frequent temperature measurements are provided by the sensors for larger deviation of the temperature from the set-point. The assumption raised on the jump rate function in Lemma 1 holds with constant $\kappa = \lambda_0 \cosh(\beta\delta_d)/\alpha$, whereas Assumption 1 holds with $h_\lambda = 2\lambda_0\beta \sinh(\beta\delta_d)$ and $\Lambda = \lambda_0 \cosh(\beta\delta_d)$. The application of the abstraction technique presented in this paper to the case study leads to the error

$$E = (1 - e^{-\kappa})^N + N(h_1\delta_\theta + h_2\delta_t + h_3\sqrt{\delta_t}), \quad (13)$$

with constants h_1, h_2, h_3 defined as

$$h_1 := 3Th_\lambda + \frac{\Lambda}{2\sigma}\sqrt{\pi RC}, \quad h_2 := 3\Lambda + \frac{\Lambda\theta_+\sqrt{\pi}}{2\sigma\sqrt{RC}} + \frac{4\Lambda}{\sqrt{2\pi}}, \quad h_3 := \frac{8\Lambda\sqrt{RC}}{\sqrt{\pi}}.$$

The additional terms contributing to the error, in comparison with the results of Theorem 4, are due to the dependence of the mean and variance of the Gaussian density function ϕ from the current time t . We use the values in Figure 2 (left) for the parameters in the numerical simulation. The standard deviation of the process noise is $\sigma = 0.1 [^\circ C s^{-1/2}]$. The time bound for the safety specification is $T = 1[h]$. The parameters of the jump rate functions are $\alpha = 1, \beta = 1, \lambda_0 = 1$, which means if the TCL is initialized at the set-point, the rate of temperature observations is 20 times higher than the decay rate of the TCL ($1/RC$).

We have implemented Algorithm 1 for the abstraction and computation of safety probability over the model using the software tool FAUST² [22]. Figure 2 (right) shows the error bound from (13), as a function of numbers of partition bins for the temperature n_θ and the time n_t , with a fixed step-horizon $N = 8$. One can see that for instance the abstraction algorithm guarantees an error bound of 0.23 by selecting $n_\theta = n_t = 4 \times 10^3$ ($\delta_\theta = 1.25 \times 10^{-4}, \delta_t = 2.5 \times 10^{-4}$), which generates a DTMC with 3.2×10^7 states. This indicates that meaningful error bounds (less than one) may lead to large DTMCs.

The derived error bounds can be in general conservative. To demonstrate the conservativeness of bounds, we perform the analysis with partition diameters $\delta_\theta = \delta_t = 0.0125$ ($n_\theta = 40, n_t = 80$), which result in a DTMC with 6400 states for which the error bounds are not meaningful. Figure 3 (top row) shows the computed safety probabilities as a function of initial temperature θ_0 at initial time t_0 , the left plot for ON mode and the right plot for OFF mode. Figure 3, bottom row, shows the safety likelihood estimated via Monte Carlo simulations with 1000 runs initialized at the representative points used in Algorithm 1. The computation and the estimation are very close to each other with a maximum relative difference of 12%. The results suggest that the error bounds can be reduced by employing advanced gridding techniques [18, 21].

Parameter	Interpretation	Value
θ_s	temperature set-point	20 [$^{\circ}C$]
δ_d	dead-band width	0.5 [$^{\circ}C$]
θ_a	ambient temperature	32 [$^{\circ}C$]
R	thermal resistance	2 [$^{\circ}C/kW$]
C	thermal capacitance	10 [$kWh/^{\circ}C$]
P_{rate}	power	14 [kW]




Fig. 2. Values of parameters for the TCL case study [20] (left). Error as a function of numbers of partition sets for temperature n_{θ} and time n_t (right).




Fig. 3. Safety probabilities as a function of initial temperature θ_0 and initial time t_0 . Left and right columns for ON $q_0 = 1$ and OFF $q_0 = 0$ modes, respectively. First and second rows are computed via abstraction approach in this paper and via Monte Carlo simulations, respectively.

6 Conclusions

We have presented an abstraction-based safety verification procedure for pure jump Markov processes with continuous states. While the focus of the work has been on the study of probabilistic safety, the technique can be extended to verify richer temporal properties. The errors can be sharpened via adaptive, non-uniform schemes [18, 19]. cPJMP are a generalization of CTMC with an assumption of constant values in between jumps: we plan to investigate the challenging problem of non-constant dynamics between jumps [15].

References

1. 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.
2. 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, 2008.
3. Y. Ait-Sahalia and P.A. Mykland. The effects of random and discrete sampling when estimating continuous-time diffusions. *Econometrica*, 71(2):483–549, 2003.
4. L.J.S. Allen. *An Introduction to Stochastic Processes with Applications to Biology*. Pearson/Prentice Hall, 2003.
5. A. Aziz, K. Sanwal, V. Singhal, and R. Brayton. Verifying continuous time Markov chains. In *Proceedings of the 8th International Conference on Computer Aided Verification*, pages 269–276, London, UK, 1996. Springer.
6. C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen. Model checking continuous-time Markov chains by transient analysis. In *Computer Aided Verification*, volume 1855 of *LNCS*, pages 358–372. Springer Berlin Heidelberg, 2000.
7. C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen. Model-checking algorithms for continuous-time Markov chains. *IEEE Transactions on Software Engineering*, 29(6):524–541, June 2003.
8. C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
9. C. Baier, J.-P. Katoen, and H. Hermanns. Approximate symbolic model checking of continuous-time Markov chains. In *Proceedings of the 10th International Conference on Concurrency Theory, CONCUR '99*, pages 146–161, London, UK, UK, 1999. Springer-Verlag.
10. D.P. Bertsekas and S.E. Shreve. *Stochastic Optimal Control: The Discrete-Time Case*. Athena Scientific, 1996.
11. P. Billingsley. *Convergence of Probability Measures*. Wiley, New York, 1999.
12. D.S. Callaway. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energy Conversion and Management*, 50(5):1389–1400, 2009.
13. T. Chen, M. Diciolla, M. Kwiatkowska, and A. Mereacre. Time-bounded verification of CTMCs against real-time specifications. In *Formal Modeling and Analysis of Timed Systems*, volume 6919 of *LNCS*, pages 26–42. Springer, 2011.
14. T. Chen, T. Han, J.-P. Katoen, and A. Mereacre. Model Checking of Continuous-Time Markov Chains Against Timed Automata Specifications. *Logical Methods in Computer Science*, Volume 7, Issue 1, March 2011.
15. M.H.A. Davis. *Markov Models and Optimization*. Chapman & Hall/CRC Press, London, 1993.
16. S. Esmaeil Zadeh Soudjani and A. Abate. Adaptive gridding for abstraction and verification of stochastic hybrid systems. In *QEST*, pages 59–69, 2011.
17. S. Esmaeil Zadeh Soudjani and A. Abate. Higher-order approximations for verification of stochastic hybrid systems. In *ATVA*, volume 7561 of *LNCS*, pages 416–434. Springer, 2012.
18. S. Esmaeil 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.
19. S. Esmaeil Zadeh Soudjani and A. Abate. Precise approximations of the probability distribution of a Markov process in time: an application to probabilistic invariance. In *TACAS*, volume 8413 of *LNCS*, pages 547–561. Springer, 2014.

20. S. Esmail Zadeh Soudjani and A. Abate. Aggregation and control of populations of thermostatically controlled loads by formal abstractions. *Control Systems Technology, IEEE Transactions on*, 23(3):975–990, May 2015.
21. S. Esmail Zadeh Soudjani and A. Abate. Quantitative approximation of the probability distribution of a Markov process by formal abstractions. *Logical Methods in Computer Science*, pages 1–29, 2015.
22. S. Esmail Zadeh Soudjani, C. Gevaerts, and A. Abate. FAUST²: Formal abstractions of uncountable-state stochastic processes. In *TACAS*, volume 9035 of *LNCS*, pages 272–286. Springer, 2015.
23. I.I. Gihman and A.V. Skorokhod. *The Theory of Stochastic Processes: II*, volume 218 of *Die Grundlehren der mathematischen Wissenschaften*. Springer, 1975.
24. A. Hinton, M. Kwiatkowska, G. Norman, and D. Parker. PRISM: A tool for automatic verification of probabilistic systems. In *Tools and Algorithms for the Construction and Analysis of Systems*, volume 3920 of *LNCS*, pages 441–444. Springer, 2006.
25. J. Jacod and A. Shiryaev. *Limit Theorems for Stochastic Processes*. Grundlehren der mathematischen Wissenschaften. Springer, 2010.
26. O. Kallenberg. *Foundations of Modern Probability*. Probability and its Applications. Springer Verlag, New York, 2002.
27. J.-P. Katoen, M. Khattri, and I. S. Zapreev. A Markov reward model checker. In *QEST*, pages 243–244. IEEE, 2005.
28. J.L. Mathieu, S. Koch, and D.S. Callaway. State estimation and control of electric loads to manage real-time energy imbalance. *IEEE Transactions on Power Systems*, 28(1):430–440, 2013.
29. M. Micheli and M. Jordan. Random sampling of a continuous-time stochastic dynamical system. In *Proc. of the 15th International Symposium on the Mathematical Theory of Networks and Systems (MTNS)*, pages 1–15, 2002.
30. P. Panangaden. *Labelled Markov Processes*. Imperial College Press, 2009.
31. E. Platen and N. Bruti-Liberati. *Numerical Solution of Stochastic Differential Equations with Jumps in Finance*. Stochastic Modelling and Applied Probability. Springer, 2010.
32. 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.
33. I. Tkachev and A. Abate. Characterization and computation of infinite-horizon specifications over Markov processes. *Theoretical Computer Science*, 515(0):1–18, 2014.
34. N.G. Van Kampen. *Stochastic Processes in Physics and Chemistry*. North-Holland Personal Library. Elsevier Science, 2011.
35. Y. Ziai. *Statistical Models of Claim Amount Distributions in General Insurance*. PhD thesis, School of Engineering & Maths, City University London, 1979.