

Probabilistic Model Checking of Labelled Markov Processes via Finite Approximate Bisimulations

Alessandro Abate¹, Marta Kwiatkowska¹, Gethin Norman², and David Parker³

¹ Department of Computer Science, University of Oxford, UK

² School of Computing Science, University of Glasgow, UK

³ School of Computer Science, University of Birmingham, UK

Abstract. This paper concerns labelled Markov processes (LMPs), probabilistic models over uncountable state spaces originally introduced by Prakash Panangaden and colleagues. Motivated by the practical application of the LMP framework, we study its formal semantics and the relationship to similar models formulated in control theory. We consider notions of (exact and approximate) probabilistic bisimulation over LMPs and, drawing on methods from both formal verification and control theory, propose a simple technique to compute an approximate probabilistic bisimulation of a given LMP, where the resulting abstraction is characterised as a finite-state labelled Markov chain (LMC). This construction enables the application of automated quantitative verification and policy synthesis techniques over the obtained abstract model, which can be used to perform approximate analysis of the concrete LMP. We illustrate this process through a case study of a multi-room heating system that employs the probabilistic model checker PRISM.

1 Introduction

Labelled Markov processes (LMPs) are a celebrated class of models encompassing concurrency, interaction and probability over uncountable state spaces, originally introduced and studied by Prakash Panangaden and colleagues [14,37]. LMPs evolve sequentially (i.e., in discrete time) over an uncountably infinite state space, according to choices from a finite set of available actions (called labels). They also allow for the possible rejection of a selected action, resulting in termination. LMPs can be viewed as a generalisation of labelled transition systems, allowing state spaces that might be uncountable and that include discrete state spaces as a special case. LMPs also extend related discrete-state probabilistic models from the literature, e.g. [34], and are related to uncountable-state Markov decision processes [38].

The formal semantics of LMPs has been actively studied in the past (see the Related Work section below). One of the earliest contributions is the notion of exact probabilistic bisimulation in [14], obtained as a generalisation of its discrete-state counterpart [34] and used to characterise the LMP model semantics. Exact bisimulation is in general considered a very conservative requirement,

and approximate notions have been consequently developed [15,19], which are based on relaxing the notion of process equivalence or on distance (pseudo-)metrics. These metrics encode exact probabilistic bisimulation, in that the distance between a pair of states is zero if and only if the states are bisimilar. While the exact notion of probabilistic bisimulation can be characterised via a Boolean logic, these approximate notions of probabilistic bisimilarity can be encompassed by real-valued logics, e.g. [30]. In view of their underlying uncountable state spaces, the analysis of LMPs is not tractable, and approximate bisimulation notions can serve as a means to derive abstractions of the original LMPs. Such abstractions, if efficiently computable and finite, can provide a formal basis for approximate verification of LMPs.

Separately from the above work rooted in semantics and logic, models that are closely related to LMPs have also been defined and studied in decision and control [7,28,36]. Of particular interest is the result that quantitative finite abstractions of uncountable-space stochastic processes [2,3] are related to the original, uncountable-state models by notions of approximate probabilistic bisimulations [41]. These notions are characterised via distances between probability measures. Alternatively these formal relations between abstract and concrete models can be established via metrics over trajectories, which are obtained using Lyapunov-like functionals as proposed in [1,31,39], or by randomisation techniques as done in [4]. There is an evident connection between approximation notions and metrics proposed for LMPs and for related models in decision and control, and it is at this intersection that the present contribution unfolds.

In this paper, we build upon existing work on LMPs, with the aim of developing automated verification, as well as optimal policy synthesis, for these models against specifications given in quantitative temporal logic. Drawing on results from the decision and control literature, we give an explicit interpretation of the formal semantics of LMPs. We consider notions of (exact and approximate) probabilistic bisimulation over LMPs, and propose a simple technique to compute an approximate probabilistic bisimulation of a given LMP, where the resulting abstraction is characterised as a finite-state labelled Markov chain (LMC). This enables the direct application of automated quantitative verification techniques over the obtained abstract model by means of the probabilistic model checker PRISM [32], which supports a number of (finite-state) probabilistic models [32,26], including LMCs. We implement an algorithm for computing abstractions of LMPs represented as LMCs and, thanks to the established notion of approximate probabilistic bisimulation, the analysis of the abstraction corresponds to an approximate analysis of the concrete LMP. We illustrate the techniques on a case study of a multi-room heating system, performing both quantitative verification and policy synthesis against a step-bounded variant of the probabilistic temporal logic PCTL [27]. We thus extend the capability of PRISM to also provide analysis methods for (uncountable-state) LMPs, which was not possible previously.

Related Work. Approximation techniques for LMPs can be based on metrics [15,19,42] and coalgebras [43,44]. Approximate notions of probabilistic bisimilarity are formally characterised and computed for finite-state labelled Markov processes in [17]. Metrics are also discussed and employed in [16] and applied to *weak* notions of bisimilarity for finite-state processes, and in [23,24,25] for (finite and infinite) Markov decision processes – in particular, [25] looks at models with uncountable state spaces. The work in [23] is extended by on-the-fly techniques in [12] over finite-state Markov decision processes. LMP approximations are also investigated in [13] and, building on the basis of [17,23], looked at from a different perspective (that of Markov processes as transformers of functions) in [11]. Along the same lines, [33] considers a novel logical characterisation of notions of bisimulations for Markov decision processes. The relationship between exact bisimulation and (CSL) logic is explored in [18] over a continuous-time version of LMPs. Abstractions that are related to Panangaden’s finite-state approximants are studied over PCTL properties in [29]. In control theory, the goal of [2,3] is to enable the verification of step-bounded PCTL-like properties [21], as well as time-bounded [41] or unbounded [40] linear-temporal specifications. It is then shown that these approximations are related to the original, uncountable-state models by notions of approximate probabilistic bisimulations [41]. Regarding algorithms for computing abstractions, [9] employs Monte-Carlo techniques for the (approximate) computation of the concepts in [15,19] which relates to the randomised techniques in [4].

Organisation of the Paper. The paper is structured as follows. Section 2 introduces LMPs and discusses two distinct perspectives to their semantic definition. Section 3 discusses notions of exact and approximate probabilistic bisimulations from the literature, with an emphasis on their computability aspects. Section 4 proposes an abstraction procedure that approximates an LMP with an LMC and formally relates the two models. Section 5 describes PRISM model checking of the LMC as a way to study properties of the original LMP. Finally, Section 6 illustrates the technique over a case study.

2 Labelled Markov Processes: Model and Semantics

We consider probabilistic processes defined over uncountable spaces [36], which we assume to be homeomorphic to a Borel subset of a Polish space, namely a metrizable space that is complete (i.e., where every Cauchy sequence converges) and separable (i.e., which contains a countable dense subset). Such a space is endowed with a Borel σ -algebra, which consists of sets that are Borel measurable. The reference metric can be reduced to the Euclidean one.

The uncountable state space is denoted by \mathcal{S} , and the associated σ -algebra by $\mathcal{B}(\mathcal{S})$. We also introduce a space of labels (or actions) \mathcal{U} , which is assumed to be finite (that is, elements taken from a finite alphabet). For later reference, we extend state and action/label spaces with the additional elements e and \bar{u} , respectively, letting $\mathcal{S}^e = \mathcal{S} \cup \{e\}$ and $\mathcal{U}^e = \mathcal{U} \cup \{\bar{u}\}$. We assume a finite set of atomic propositions AP, a function $L : \mathcal{S} \rightarrow 2^{\text{AP}}$ which labels states with the

propositions that hold in that state, and a reward structure $\mathbf{r} : \mathcal{S} \times \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$, which assigns rewards to state-label pairs over the process.

Processes will evolve in discrete time over the finite interval $[0, N]$ over a sample space $\Omega_{N+1} = \mathcal{S}^{N+1}$, equipped with the canonical product σ -algebra $\mathcal{B}(\Omega_{N+1})$. The selection of labels at each time step depends on a policy (or strategy), which can base its choice on the previous evolution of the process. Formally a policy is a function $\sigma : \{\mathcal{S}^i \mid 1 \leq i \leq N\} \rightarrow \text{dist}(\mathcal{U})$, where $\text{dist}(\mathcal{U})$ is the set of probability distributions over \mathcal{U} and $\sigma(s_0, \dots, s_k) = \mu$ represents the fact that the policy selects the label u_k in state s_k at time instant k with probability $\mu(u_k)$, given that the states at the previous time instances were s_0, \dots, s_{k-1} .

Under a fixed policy the process is fully probabilistic and we can then reason about the likelihood of events. However, due to the uncountable state space this is not possible for all policies. Following [10], we restrict our attention to so called measurable policies, for which we can define a probability measure, denoted \mathbb{P}_s^σ , over the sample space Ω_{N+1} when the initial state of the process equals s .

The following definition is taken from [14,15,19] (these contributions mostly deal with analytic spaces that represent a generalisation of the Borel measurable space we focus on).

Definition 1 (Labelled Markov Process). *A labelled Markov process (LMP) \mathcal{S} is a structure:*

$$(\mathcal{S}, s_0, \mathcal{B}(\mathcal{S}), \{\tau_u \mid u \in \mathcal{U}\})$$

where \mathcal{S} is the state space, $s_0 \in \mathcal{S}$ is the initial state, $\mathcal{B}(\mathcal{S})$ is the Borel σ -field on \mathcal{S} , \mathcal{U} is the set of labels, and for each $u \in \mathcal{U}$:

$$\tau_u : \mathcal{S} \times \mathcal{B}(\mathcal{S}) \longrightarrow [0, 1]$$

is a sub-probability transition function, namely, a set-valued function $\tau_u(s, \cdot)$ that is a sub-probability measure on $\mathcal{B}(\mathcal{S})$ for all $s \in \mathcal{S}$, and such that the function $\tau_u(\cdot, S)$ is measurable for all $S \in \mathcal{B}(\mathcal{S})$. \square

In this work, we will often assume that the initial state s_0 can be any element of \mathcal{S} and thus omit it from the definition. Furthermore, we will implicitly assume that the state space is a standard Borel space, so the LMP \mathcal{S} will often be referred to simply as the pair $(\mathcal{S}, \{\tau_u \mid u \in \mathcal{U}\})$.

It is of interest to explicitly elucidate the underlying semantics of the model that is syntactically characterised in Definition 1. The semantics hinges on how the sub-probability measures are dealt with in the model: we consider two different options, the first drawn from the literature on testing [34], and the second originating from models of decision processes [38]. Recall that we consider finite traces over the discrete domain $[0, N]$ (this is because of the derivation of abstractions that we consider below – an extension of the semantics to infinite traces follows directly). The model is initialised at time $k=0$ at state s_0 , which is deterministically given or obtained by sampling a given probability distribution π_0 , namely $s_0 \sim \pi_0$. At any (discrete) time $0 \leq k \leq N-1$, given a state $s_k \in \mathcal{S}$ and selecting a discrete action $u_k \in \mathcal{U}$, this action is accepted with a probability

$\int_{\mathcal{S}} \tau_{u_k}(s_k, dx)$, whereas it is rejected with probability $1 - \int_{\mathcal{S}} \tau_{u_k}(s_k, dx)$. If the action u_k is rejected, then the model can exhibit two possible behaviours:

1. (Testing process) the dynamics stops, that is, the value s_{k+1} is undefined and the process returns the finite trace

$$((s_0, u_0), (s_1, u_1), \dots, (s_k, u_k));$$

2. (Decision process) a default action $u \in \mathcal{U}^e$ is selected and the process continues its evolution, returning the sample $s_{k+1} \sim \tau_u(s_k, \cdot)$. The default action can, for instance, coincide with the label selected (and accepted) at time $k-1$, i.e. $u = u_{k-1} \in \mathcal{U}$, or with the additional label, i.e. $u = \bar{u}$. At time instant $N-1$, the process completes its course and further returns the trace

$$((s_0, u_0), (s_1, u_1), \dots, (s_k, u_k), (s_{k+1}, u) \dots, (s_{N-1}, u), s_N).$$

Note that the above models can also be endowed with a set of output or observable variables, which are defined over an “observation space” \mathcal{O} via an observation map $h : \mathcal{S} \times \mathcal{U} \rightarrow \mathcal{O}$. In the case of “full observation,” the map h can simply correspond to the identity and the observation space coincides with the domain of the map. The testing literature often employs a map $h : \mathcal{U} \rightarrow \mathcal{O}$, whereas in the decision literature it is customary to consider a map $h : \mathcal{S} \rightarrow \mathcal{O}$. That is, the emphasis in the testing literature is on observing actions/labels, whereas in the decision and control field the focus is on observing variables (and thus on the corresponding underlying dynamics of the model).

We elucidate the discussion above by two examples.

Example 1 (Testing process). Consider a fictitious reactive system that takes the shape of a slot or a vending machine, outputting a chosen label and endowed with an internal state with one n -bit memory register retaining a random number. For simplicity, we select a time horizon $N < 2^{n-1}$, so that the internal state never under- or overflows. The state of the machine is $s_k \in \{-2^{n-1}, \dots, 0, \dots, 2^{n-1}\}$, where the index k is a discrete counter initialised at zero. At its k -th use, an operator pushes one of $\mathcal{U} = \{0, 1, 2, \dots, M\}$ buttons u_k , to which the machine responds with probability $\frac{1}{1+e^{-s_k}}$ and, in such a case, resets its state to $s_{k+1} = s_k + u_k \xi_k$, where ξ_k is a fair Bernoulli random variable taking values in $\{-1, 1\}$. On the other hand, if the label/action is not accepted, then the machine gets stuck at state s_k .

Clearly, the dynamics of the process hinges on the external inputs provided by the user (the times of which are not a concern; what matters is the discrete counter k for the input actions). The process generates traces as long as the input actions are accepted. We may be interested in assessing if a given periodic input policy applied within the finite time horizon (for instance, the periodic sequence $(M-1, M, M-1, M, \dots)$) is accepted with probability greater than a given threshold over the model initialised at a given state, where this probability depends on the underlying model. Alternatively, we may be interested in generating a policy that is optimal with respect to a specification over the space of possible labels (for example, we might want the one that minimises the occurrence of choices within the set $\{0, 1, 2\} \subseteq \mathcal{U}$). \square

Example 2 (Decision process). Let us consider a variant of the temperature control model presented in [3,5] and which will be further elaborated upon in Section 6. The temperature of a room is controlled at the discrete time instants $t_k = 0, \delta, 2\delta, \dots, N\delta$, where $\delta \in \mathbb{R}^+$ represents a given fixed sampling time. The temperature is affected by the heat inflow generated by a heater that is controlled by a thermostat, which at each time instant t_k can either be switched off or set to a level between 1 and M . This choice between heater settings is represented by the labels $\mathcal{U} = \{u_0, u_1, \dots, u_M\}$ of the LMP, where $u_0 = 0$ and $0 < u_1 < \dots < u_M$. The (real-valued) temperature s_{k+1} at time t_{k+1} depends on that at time t_k as follows:

$$s_{k+1} = s_k + h(s_k - s_a) + h_{u_k} \zeta(s_k, u_k) + \xi_k,$$

where

$$\zeta(s_k, u_k) = \begin{cases} u_k & \text{w.p. } 1 - s_k \cdot \frac{u_k}{u_M} \cdot \alpha \\ \bar{u} & \text{else} \end{cases}$$

$\xi_k \sim \mathcal{N}[0, 1]$, s_a represents the constant ambient temperature outside the room, the coefficient h denotes the heat loss, h_{u_k} is the heat inflow when the heater setting corresponds to the label u_k , and α is a normalisation constant.

The quantity $\zeta(s_k, u_k)$ characterises an accepted action (u_k) with a probability that decreases both as the temperature increases (we suppose increasing the temperature has a negative affect through heat-related noise on the correct operation of the thermostat) and as the heater level increases (increasing the heater level puts more stress on the heater, which is then more likely to fail), and conversely provides a default value if the action is not accepted. The default action could feature the heater in the OFF mode (u_0), or the heater stuck to the last viable control value (u_{k-1}). In other words, once an action/label is rejected, the dynamics progresses by adhering to the default action.

We stress that, unlike in the previous example, here the dynamics proceeds regardless of whether the action is accepted or not, since the model variable (s_k) describes a physical quantity with its own dynamics that simply cannot be physically stopped by whatever input choice. Given this model, we may be interested in assessing whether the selected policy satisfies a desired property with a specified probability (similar to the testing case), or in synthesising a policy that maximises the probability of a given specification – say, to maintain the room temperature within a certain comfort interval. Policy synthesis problems appear to be richer for models of decision processes, since the dynamics play a role in a more explicit manner. \square

The second semantics (related to a decision process) leads to a reinterpretation of the LMP as a special case of an (infinite-space) MDP [38]. Next, we aim at leveraging this interpretation for *both* semantics: in other words, the two semantics of the LMP can be interpreted in a consistent manner by extending the dynamics and by properly “completing” the sub-stochastic kernels, by means of new absorbing states and additional labels, as described next. This connection has been qualitatively discussed in [37] and is now expounded in detail and newly applied at a semantical level over the different models. Let us start with

the testing process. Given a state $s_k \in \mathcal{S}$ and an action $u_k \in \mathcal{U}$, we introduce the binary random variable taking values in the set $\{u_k, \bar{u}\}$ with probability $\{\int_{\mathcal{S}} \tau_{u_k}(s_k, dx), 1 - \int_{\mathcal{S}} \tau_{u_k}(s_k, dx)\}$, respectively. Consider the extended spaces \mathcal{S}^e and \mathcal{U}^e . Here e is an absorbing state, namely e is such that $\int_{\mathcal{S}} \tau_u(e, dx) = 0$ for all $u \in \mathcal{U}^e$ (any action selected at state e is rejected), and such that $\tau_{\bar{u}}(s, dx) = \delta_e(dx)$ for all $s \in \mathcal{S}^e$, where $\delta_e(\cdot)$ denotes the Dirac delta function over \mathcal{S}^e . We obtain:

$$s_{k+1} \sim \begin{cases} \int_{\mathcal{S}} \tau_{u_k}(s_k, dx) \tau_{u_k}(s_k, \cdot) & \text{if the action is accepted} \\ (1 - \int_{\mathcal{S}} \tau_{u_k}(s_k, dx)) \tau_{\bar{u}}(s_k, \cdot) & \text{if the action is rejected.} \end{cases} \quad (1)$$

The labelling map $h : \mathcal{U} \rightarrow \mathcal{O}$ is inherited and extended to \mathcal{U}^e , so that $h(\bar{u}) = \emptyset$.

Let us now focus on the decision process. Similarly to the testing case, at time k and state s_k , label/action u_k is chosen and this value accepted with a certain probability, else a default value u is given. In the negative instance, the actual value of u depends on the context (see the discussion in the example above) and can correspond to an action within \mathcal{U} (say, the last accepted action) or to the additional action \bar{u} outside this finite set but in \mathcal{U}^e . Then, as in the testing case, s_{k+1} is selected according to the probability laws in (1). However, we impose the following condition: once an action is rejected and label u is selected, the very same action is retained deterministically for the remaining part of the time horizon, namely $u_j = u$ for all $k \leq j \leq N-1$. Essentially, it is as if, for any time instant $k \leq j \leq N-1$, the action space collapsed into the singleton set $\{u\}$. Notice that, in the decision case, the state space \mathcal{S}^e does not need to be extended; however, the kernel $\tau_{\bar{u}}, \bar{u} \in \mathcal{U}^e \setminus \mathcal{U}$, should be defined and indeed have a non-trivial dynamical meaning if the action \bar{u} is used. Finally, the labelling map $h : \mathcal{S} \rightarrow \mathcal{O}$ is inherited from above.

Let us emphasise that, according to the completion procedure described above, LMPs (in general endowed with sub-probability measures) can be considered as special cases of MDPs, which allows connecting with the rich literature on the subject [7,28].

3 Exact and Approximate Probabilistic Bisimulations

We now recall the notions of exact and approximate probabilistic bisimulation for LMPs [14,17]. We also extend these definitions to incorporate the labelling and reward functions introduced in Section 2. We emphasise that both concepts are to be regarded as *strong* notions – we do not consider hidden actions or internal nondeterminism in this work, and thus refrain from dealing with *weak* notions of bisimulation.

Definition 2 ((Exact) Probabilistic Bisimulation). *Consider an LMP $\mathcal{S} = (\mathcal{S}, \{\tau_u \mid u \in \mathcal{U}\})$. An equivalence relation R on \mathcal{S} is a probabilistic bisimulation if, whenever $s_1 R s_2$ for $s_1, s_2 \in \mathcal{S}$, then $L(s_1) = L(s_2)$, $\mathbf{r}(s_1, u) = \mathbf{r}(s_2, u)$ for all $u \in \mathcal{U}$ and, for any $u \in \mathcal{U}$ and set $\tilde{S} \in \mathcal{S}/R$ (which is Borel measurable), it holds that*

$$\tau_u(s_1, \tilde{S}) = \tau_u(s_2, \tilde{S}).$$

A pair of states $s_1, s_2 \in \mathcal{S}$ are said to be probabilistically bisimilar if there exists a probabilistic bisimulation R such that $s_1 R s_2$. \square

Observe that the autonomous case of general Markov chains with sub-probability measures, which is characterised by a trivial labels set with a single element, can be obtained as a special case of the above definition.

Let R be a relation on a set A . A set $\tilde{A} \subseteq A$ is said to be R -closed if $R(\tilde{A}) = \{t \mid s R t \wedge s \in \tilde{A}\} \subseteq \tilde{A}$. This notion will be employed shortly – for the moment, note that Definition 2 can be equivalently given by considering the condition on the transition kernel to hold over R -closed measurable sets $\tilde{S} \subseteq \mathcal{S}$.

The exact bisimulation relation given above directly extends the corresponding notions for finite Markov chains and Markov decision processes (that is, models characterised by discrete state spaces). However, although intuitive, it can be quite conservative when applied over uncountable state spaces, and procedures to compute such relations over these models are in general deemed to be undecidable. Furthermore, the concept does not appear to accommodate computational robustness [20,45], arguably limiting its applicability to real-world models in engineering and science. These considerations lead to a notion of approximate probabilistic bisimulation with level ε , or simply ε -probabilistic bisimulation [17], as described next.

Definition 3 (Approximate Probabilistic Bisimulation). Consider an LMP $\mathcal{S} = (\mathcal{S}, \{\tau_u \mid u \in \mathcal{U}\})$. A relation R_ε on \mathcal{S} is an ε -probabilistic bisimulation relation if, whenever $s_1 R_\varepsilon s_2$ for $s_1, s_2 \in \mathcal{S}$, then $L(s_1) = L(s_2)$, $\mathbf{r}(s_1, u) = \mathbf{r}(s_2, u)$ for all $u \in \mathcal{U}$ and, for any $u \in \mathcal{U}$ and R_ε -closed set $\tilde{S} \subseteq \mathcal{S}$, it holds that

$$|\tau_u(s_1, \tilde{S}) - \tau_u(s_2, \tilde{S})| \leq \varepsilon. \quad (2)$$

In this case we say that the two states are ε -probabilistically bisimilar. \square

Unlike the equivalence relation R in the exact case, in general, the relation R_ε does not satisfy the transitive property (the triangle inequality does not hold: each element of a pair of states may be close to a common third element, but map to very different transition measures among each other), and as such is not an equivalence relation [17]. Hence, it induces a cover of the state space \mathcal{S} , but not necessarily a partition.

The above notions can be used to relate or compare two separate LMPs, say \mathcal{S}_1 and \mathcal{S}_2 , with the same action space \mathcal{U} by considering an LMP \mathcal{S}^+ characterised as follows [37]. The state space \mathcal{S}^+ is given by the direct sum of the state spaces of the two processes (i.e. the disjoint union of \mathcal{S}_1 and of \mathcal{S}_2), where the associated σ -algebra is given by $\mathcal{B}(\mathcal{S}_1) \cup \mathcal{B}(\mathcal{S}_2)$. The labelling and reward structure combine those for the separate processes, using the fact that the state space is the directed sum of these processes. The transition kernel $\tau_u^+ : \mathcal{S}^+ \times \mathcal{B}(\mathcal{S}^+) \rightarrow [0, 1]$ is such that, for any $u \in \mathcal{U}$, $1 \leq i \leq 2$, $s_i \in \mathcal{S}_i$, $S^+ \subseteq \mathcal{B}(\mathcal{S}_1) \cup \mathcal{B}(\mathcal{S}_2)$ we have $\tau_u^+(s_i, S^+) = \tau_u^i(s_i, S^+ \cap \mathcal{S}_i)$. The initial states of the composed model are characterised by considering those of the two generating processes with equal likelihood. In the instance of the exact notion, we have the following definition.

Definition 4. Consider two LMPs $\mathcal{S}_i = (\mathcal{S}_i, \{\tau_u^i \mid u \in \mathcal{U}\})$ where $i = 1, 2$, endowed with the same action space \mathcal{U} , and their direct sum \mathcal{S}^+ . An equivalence relation R on \mathcal{S}^+ is a probabilistic bisimulation relation between \mathcal{S}_1 and \mathcal{S}_2 if, whenever $s_1 R s_2$ for $s_1 \in \mathcal{S}_1, s_2 \in \mathcal{S}_2$, then $L(s_1) = L(s_2)$, $\mathbf{r}(s_1, u) = \mathbf{r}(s_2, u)$ for all $u \in \mathcal{U}$ and, for any given $u \in \mathcal{U}$ and R -closed set $S^+ \in \mathcal{B}(\mathcal{S}_1) \cup \mathcal{B}(\mathcal{S}_2)$, it holds that

$$\tau_u^+(s_1, \tilde{S}^+) = \tau_u^1(s_1, \tilde{S}^+ \cap \mathcal{S}_1) = \tau_u^2(s_2, \tilde{S}^+ \cap \mathcal{S}_2) = \tau_u^+(s_2, \tilde{S}^+).$$

A pair of states $(s_1, s_2) \in \mathcal{S}_1 \times \mathcal{S}_2$ is said to be probabilistically bisimilar if there exists a relation R such that $s_1 R s_2$. Two LMPs \mathcal{S}_i are probabilistically bisimilar if their initial states are. \square

The inequality in (2) can be considered as a correspondence between states in the pair (s_1, s_2) that could result from the existence of a (pseudo-)metric over probability distributions on the state space. This approach has been taken up by a number of articles in the literature, which have introduced metrics as a means to relate two models. Such metrics have been defined based on logical characterizations [15,19,33], categorical notions [43,44], games [17], normed distances over process trajectories [1,31], as well as distances between probability measures [42].

3.1 Computability of Approximate Probabilistic Bisimulations

While for processes over discrete and finite state spaces there exist algorithmic procedures to compute exact [6,34] and approximate [17] probabilistic bisimulations, the computational aspects related to these notions for processes over uncountable state spaces appear to be much harder to deal with. We are of course interested in characterising computationally finite relations, which will be the goal pursued in the next section. Presently, only a few approaches exist to approximate uncountable-space processes with finite-state ones: LMPs [9,11,13,19], (infinite-state) MDPs [33], general Markov chains [29] and stochastic hybrid systems (SHSs) [2,3].

Alternative approaches to check the existence of an approximate probabilistic bisimulations between two models, which hinge on the computation of a function relating the trajectories of the two processes [1,31,39], are limited to models that are both defined over an uncountable space, and do not appear to allow for the constructive synthesis of approximate models from concrete ones. Computation of abstract models, quantitatively related to corresponding concrete ones, is investigated in [4], which leverages randomised approaches and, as such, can enforce similarity requirements only up to a confidence level. Recent work on symbolic abstractions [46] refer to approximation notions over (higher-order) moments on the distance between the trajectory of the abstract model and the solution of the concrete one.

In conclusion, analysing the precision, quality, and scalability properties of constructive approximation techniques for uncountable-state stochastic processes is a major goal with relevant applications that we deem worthwhile pursuing.

4 From LMPs to Finite Labelled Markov Chains

In this section, we introduce an abstraction procedure to approximate a given LMP as a finite-state labelled Markov chain (LMC). The abstraction procedure is based on a discretisation of the state space of the LMP (recall that the space of labels (actions) is finite and as such requires no discretisation) and is inspired by the early work in [3] over (fully-probabilistic) SHS models. It is now extended to account for the presence of actions and to accommodate the LMP framework (with sub-probability measures). The relationship between abstract and concrete models as an approximate probabilistic bisimulation is drawn from results in [41].

Let us start from an LMP $\mathcal{S} = (\mathcal{S}, s_0, \mathcal{B}(\mathcal{S}), \{\tau_u \mid u \in \mathcal{U}\})$, represented via its extended dynamics independently from its actual semantics¹. We consider a finite partition of the space $\mathcal{S} = \cup_{i=1}^Q S_i$ such that $S_i \cap S_j = \emptyset$ for all $1 \leq i \neq j \leq Q$. In addition, we assume that states in the same element of the partition have the same labelling and reward values, that is, for any $1 \leq i \leq Q$ we have $L(s) = L(s')$ and $\mathbf{r}(s, u) = \mathbf{r}(s', u)$ for all $s, s' \in S_i$ and $u \in \mathcal{U}$. Let us associate to this partition a finite σ -algebra corresponding to $\sigma(S_1, \dots, S_Q)$. The finiteness of the introduced σ -algebra, in particular, implies that, for any $1 \leq i \leq Q$, states $s_1, s_2 \in S_i$, measurable set $S \in \sigma(S_1, \dots, S_Q)$ and label $u \in \mathcal{U}$, we have:

$$\tau_u(s_1, S) = \tau_u(s_2, S).$$

This follows from the finite structure of the σ -algebra and the definition of measurability. Let us now select for each $1 \leq i \leq Q$ a single fixed state $s_i \in S_i$. Using these states, for any label $u \in \mathcal{U}$ we then approximate the kernel τ_u by the matrix $p_u \in [0, 1]^Q \times [0, 1]^Q$, where for any $1 \leq i, j \leq Q$:

$$p_u(i, j) = \tau_u(s_i, S_j).$$

Observe that, for any $u \in \mathcal{U}$ and $1 \leq i \leq Q$, we have $\sum_{j=1}^Q p_u(i, j) \leq 1$. The structure resulting from this procedure is called a finite labelled Markov chain (LMC). Note that, in general, LMCs do not correspond to (finite-state) MDPs: this correspondence holds only if we have abstracted an LMP that has been “completed” with the procedure described in Section 2, and which as such can be reinterpreted as an uncountable-state MDP.

Let us comment on the procedure above. We have started from the LMP $\mathcal{S} = (\mathcal{S}, \{\tau_u \mid u \in \mathcal{U}\})$, endowed with an uncountable state space \mathcal{S} with the corresponding (uncountable) Borel σ -algebra $\mathcal{B}(\mathcal{S})$. We have partitioned the space \mathcal{S} into a finite quotient made up of uncountable sets S_i , and associated to this finite quotient a finite σ -algebra. We call this intermediate model \mathcal{S}^f : the obtained model is still defined over an uncountable state space, but its probabilistic structure is simpler, being characterised by a finite σ -algebra $\sigma(S_1, \dots, S_Q)$ and piecewise constant kernels – call them $\tau_u^f(s, \cdot)$ – for completeness $\mathcal{S}^f = (\mathcal{S}, s_0, \sigma(S_1, \dots, S_Q), \{\tau_u^f \mid u \in \mathcal{U}\})$. This latter feature has allowed

¹ With slight abuse of notation but for simplicity sake, we avoid referring to extended state and/or action spaces as more proper for “completed” LMP models.

us, in particular, to select an arbitrary state s_i for each of the partition sets S_i , which has led to a finite state space $\mathcal{S}^d = \{s_1, \dots, s_Q\}$. For each label $u \in \mathcal{U}$ we have introduced the (sub-)probability transition matrix p_u . The new model $\mathcal{S}^d = (\mathcal{S}^d, s_0^d, \sigma(S_1, \dots, S_Q), \{p_u \mid u \in \mathcal{U}\})$ is an LMC. Here s_0^d is the discrete state in \mathcal{S}^d that corresponds to the quotient set, including the concrete initial condition s_0 of \mathcal{S} .

Let us emphasise that, whilst the structure of the state spaces of \mathcal{S}^f and of \mathcal{S}^d are not directly comparable, their probabilistic structure is equivalent and finite – that is, the probability associated to any set in \mathcal{S}^f (the quotient of \mathcal{S}) for \mathcal{S}^f is matched by that defined over the finite set of states in \mathcal{S}^d for \mathcal{S}^d . The model \mathcal{S}^f allows us to formally relate the concrete, uncountable state-space model \mathcal{S} with the discrete and finite abstraction \mathcal{S}^d .

The formal relationship between the concrete and the abstract models can be derived under the following assumption on the regularity of the kernels of \mathcal{S} .

Assumption 1. Consider the LMP $\mathcal{S} = (\mathcal{S}, \{\tau_u \mid u \in \mathcal{U}\})$. For any label $u \in \mathcal{U}$ and states $s', s'', t \in \mathcal{S}$, there exists a positive and finite constant $k(u)$, such that

$$|T_u(s', t) - T_u(s'', t)| \leq k(u) \|s' - s''\|$$

where T_u is the density associated to the kernel τ_u , which is assumed to admit an integral form so that $\int T_u(s, t) dt = \int \tau_u(s, dt)$ for all $s \in \mathcal{S}$ and $u \in \mathcal{U}$. \square

Consider the concrete LMP \mathcal{S} , and recall the finite partition for its state space, $\mathcal{S} = \cup_{i=1}^Q S_i$. Let R be the relation over \mathcal{S} such that $s' R s''$ if and only if the states are in the same element of the partition, i.e. there exists $1 \leq i \leq Q$ such that $s', s'' \in S_i$. The relation R is trivially symmetric and reflexive. Furthermore, if $s' R s''$, then for any $S \in \{S_1, \dots, S_Q\}$:

$$\begin{aligned} |\tau_u(s', S) - \tau_u(s'', S)| &= \left| \int_S \tau_u(s', dx) - \int_S \tau_u(s'', dx) \right| \\ &= \left| \int_S T_u(s', x) dx - \int_S T_u(s'', x) dx \right| \\ &\leq \int_S |T_u(s', x) - T_u(s'', x)| dx \\ &\leq \int_S k(u) \|s' - s''\| dx \\ &\leq \mathcal{L}(S) k(u) \delta(S) \end{aligned} \tag{3}$$

where $\delta(S) = \sup_{s', s'' \in S} \|s' - s''\|$ denotes the diameter of the partition set S and $\mathcal{L}(S)$ denotes the Lebesgue measure of the set S . By virtue of the inequality established in (3) and Definition 3, we obtain the following result.

Theorem 1. Consider the LMP $\mathcal{S} = (\mathcal{S}, s_0, \sigma(S_1, \dots, S_Q), \{\tau_u \mid u \in \mathcal{U}\})$. The introduced relation R is an (approximate) ε -probabilistic bisimulation over \mathcal{S} where

$$\varepsilon = \max_{u \in \mathcal{U}} \max_{1 \leq i \leq Q} \mathcal{L}(S_i) k(u) \delta(S_i). \quad \square$$

From this point on, we assume that we are interested in the dynamics of the LMP over a bounded set \mathcal{S} , which allows us to conclude that ε is finite (since its volume $\mathcal{L}(\mathcal{S})$ and its diameter $\delta(\mathcal{S})$ are). More specifically, the approximation level ε can be tuned by reducing the quantity $\delta(\cdot)$, the diameter of the partitions of \mathcal{S} . Likewise, better bounds based on local Lipschitz continuity (rather than global, as per Assumption 1) can improve the error, as further elaborated in [21].

We now introduce the model \mathcal{S}^f , with its corresponding finite σ -algebra and piecewise constant kernel functions τ_u^f . Working with the same relation R as above, using (3) we have that if $s R s^f$ and $S \in \{S_1, \dots, S_Q\}$, then

$$|\tau_u(s, S) - \tau_u^f(s^f, S)| = |\tau_u(s, S) - \tau_u(s_i, S)| \leq \mathcal{L}(S)k(u)\delta(S).$$

This leads us to conclude, via Definition 4, that the LMPs \mathcal{S} and \mathcal{S}^f are ε -probabilistically bisimilar, where ε is taken from Theorem 1. Notice that Definition 4 can be used to relate LMPs with different structures, since it does not require the LMPs to have the same state or probability spaces – the only requirement is that the processes share the same action space. Having argued that the probabilistic structure of \mathcal{S}^f and of \mathcal{S}^d are the same, we proceed now by comparing the LMP \mathcal{S} with the LMC \mathcal{S}^d . Consider their direct sum \mathcal{S}^+ and relation R where, for $s \in \mathcal{S}$ and $s_i \in \mathcal{S}^d$, we have $s R s_i$ if and only if $s \in S_i$. Now, any R -closed set S^+ is such that $S^+ \cap \mathcal{S}_d = S_j$ and $S^+ \cap \mathcal{S} = S_j$ for any $1 \leq j \leq Q$. It therefore follows that

$$\begin{aligned} |\tau_u^+(s, S^+) - \tau_u^+(s_i, S^+)| &= |\tau_u(s, S_j) - p_u(i, j)| \\ &= |\tau_u(s, S_j) - \tau_u(s_i, S_j)| \\ &\leq \mathcal{L}(S_j)k(u)\delta(S_i) \end{aligned}$$

which leads to the following result.

Theorem 2. *Models \mathcal{S} and \mathcal{S}^d are ε -probabilistically bisimilar.* □

Remark 1. The above theorem establishes a formal relationship between \mathcal{S} and \mathcal{S}^d by way of comparing \mathcal{S} with \mathcal{S}^f over the same state space. Unlike most of the mentioned approaches in the LMPs approximations literature, the result comes with a simple procedure to compute the finite abstraction \mathcal{S}^d , with a quantitative relationship between the finite abstraction and the original LMP model [3,21]. Thanks to the dependence of the error on the (max) diameter among the partition sets, the approximation level ε can be tuned by selecting a more refined partition of the state space \mathcal{S} . Of course this entails obtaining a partition set with larger cardinality by employing smaller partitions. □

5 Model Checking Labelled Markov Chains with PRISM

The previous section has described a procedure to approximate an infinite-state LMP by a finite-state LMC. In this section, we show how probabilistic model checking over this finite abstract model can be used to verify properties of the original, concrete uncountable-space LMP.

Probabilistic model checking is a powerful and efficient technique for formally analysing a large variety of quantitative properties of probabilistic models. The properties are specified as formulae in a (probabilistic) temporal logic. In this paper, we use a time-bounded fragment of the logic PCTL [8,27] for discrete-time models, augmented with an operator to reason about costs and rewards [26], although the relationship established in the previous section between LMPs and LMCs in fact preserves a more general class of linear-time properties over a bounded horizon [41].

We will explain our use of probabilistic model checking in the context of (finite-state) LMCs, and subsequently explain the relationship with LMPs. We use logical properties defined according to Φ in the following syntax:

$$\begin{aligned}\Phi &::= P_{\sim p}[\phi \text{ U}^{\leq K} \psi] \mid \mathbf{R}_{\sim x}^{\mathbf{r}}[\mathbf{C}^{\leq K}] \\ \phi &::= \mathbf{true} \mid a \mid \phi \wedge \psi \mid \neg \phi\end{aligned}$$

where $\sim \in \{<, \leq, >, \geq\}$ is a binary comparison operator, $p \in [0, 1]$ is a probability bound, $x \in \mathbb{R}_{\geq 0}$ is a reward bound, $K \in \mathbb{N}$ is a time bound, \mathbf{r} is a reward structure and a is an atomic proposition. A property $P_{\sim p}[\phi \text{ U}^{\leq K} \psi]$ asserts that the probability of ψ becoming true within K time steps, and ϕ remaining true up until that point, satisfies $\sim p$. In standard fashion, we can also reason about (bounded) probabilistic reachability and invariance:

$$\begin{aligned}P_{\sim p}[\diamond^{\leq K} \phi] &\stackrel{\text{def}}{=} P_{\sim p}[\mathbf{true} \text{ U}^{\leq K} \phi] \\ P_{\geq p}[\square^{\leq K} \phi] &\stackrel{\text{def}}{=} P_{\leq 1-p}[\diamond^{\leq K} \neg \phi]\end{aligned}$$

A property $\mathbf{R}_{\sim x}^{\mathbf{r}}[\mathbf{C}^{\leq K}]$ asserts that the expected amount of reward (from reward structure \mathbf{r}) accumulated over K steps satisfies $\sim x$. State formulae ϕ can identify states according to the atomic propositions that label them, and can be combined by Boolean operations on these propositions.

We define satisfaction of a logical formulae Φ with respect to a state s^d and policy σ^d of an LMC \mathcal{S}^d . We write $\mathcal{S}^d, s^d, \sigma^d \models \Phi$ to denote that, starting from state s^d of \mathcal{S}^d , and under the control of σ^d , Φ is satisfied. We can then treat the analysis of a formula Φ against a model \mathcal{S}^d in two distinct ways. We can *verify* that a formula Φ is satisfied under all policies of \mathcal{S}^d , or we can *synthesize* a single policy that satisfies Φ . In fact, in practice, whichever kind of analysis is required, the most practical solution is to compute the minimum or maximum value, over all policies, for the required property. For example, for an until property $\phi \text{ U}^{\leq K} \psi$, we might compute the maximum probability of satisfaction when the initial state is s^d :

$$P_{\max=?}[\phi \text{ U}^{\leq K} \psi] \stackrel{\text{def}}{=} \max_{\sigma^d} P_{s^d}^{\sigma^d}(\phi \text{ U}^{\leq K} \psi)$$

where $P_{s^d}^{\sigma^d}(\phi \text{ U}^{\leq K} \psi)$ is the probability under the policy σ^d when the initial state is s^d of ψ becoming true within K time steps, and ϕ remaining true up until that point.

Computing minimum or maximum probabilities or rewards (and thus checking a property Φ against an LMC) can be performed using existing probabilistic model checking algorithms for Markov decision processes [8,26]. These methods

are implemented in the PRISM model checker, which we use for the case study in the next section. When computing optimal values, a corresponding policy (strategy) that achieves them can also be synthesised.

Finally, we discuss how probabilistic model checking of an LMC obtained from an LMP allows us to analyse the original, concrete LMP.

Theorem 3. *Consider a concrete LMP \mathcal{S} and an abstract LMC \mathcal{S}^d which are ε -probabilistic bisimilar. For two ε -probabilistically bisimilar states $s \in \mathcal{S}, s^d \in \mathcal{S}^d$ and until property $\phi \text{ U}^{\leq K} \psi$ we have:*

- for any (measurable) policy σ of \mathcal{S} there exists a policy σ^d of \mathcal{S}^d such that

$$| \mathbf{P}_s^\sigma(\phi \text{ U}^{\leq K} \psi) - \mathbf{P}_{s^d}^{\sigma^d}(\phi \text{ U}^{\leq K} \psi) | \leq \varepsilon K$$

- for any policy σ^d of \mathcal{S}^d there exists a (measurable) policy σ of \mathcal{S} such that

$$| \mathbf{P}_{s^d}^{\sigma^d}(\phi \text{ U}^{\leq K} \psi) - \mathbf{P}_s^\sigma(\phi \text{ U}^{\leq K} \psi) | \leq \varepsilon K.$$

Furthermore, the above bounds apply to the case of optimal policy synthesis, for instance (in the case of maximisation) considering policies σ, σ^d within the same class for \mathcal{S} and \mathcal{S}^d , respectively, it holds that

$$| \max_{\sigma^d} \mathbf{P}_{s^d}^{\sigma^d}(\phi \text{ U}^{\leq K} \psi) - \max_{\sigma} \mathbf{P}_s^\sigma(\phi \text{ U}^{\leq K} \psi) | \leq \varepsilon K. \quad \square$$

The above theorem also generalises to expected reward properties and general linear-time properties over a finite horizon, such as bounded linear-time temporal logic (BLTL) or properties expressed as deterministic finite automata.

6 Case Study

This section presents a case study of the multi-room heating benchmark introduced in [22], based on a model proposed by [35] and already discussed in Section 2. The objective is to evaluate the usefulness of probabilistic model checking for the (approximate) verification (and optimisation) of an LMP. The model is an extension of that presented in [5], in that the control set is richer than the binary one considered in the reference, and is also related to that in [3].

We study a model for the control of the temperature evolution of two adjacent rooms. Each room is equipped with a heater and there is a single control which can switch the heaters between $M=10$ different levels of heat flow, with 0 corresponding to the heaters being OFF and 10 to the heaters being ON at full power. The uncountable state space is \mathbb{R}^2 , modelling the temperature evolution in the two rooms.

As in Section 2, the average temperature of a room evolves according to a stochastic difference equation during the finite time horizon $[0, N]$. As there are now two rooms, following [22] we also include the heat transfer between the rooms in the equations. Letting $\mathbf{s}_k \in \mathbb{R}^2$ denote the temperatures in the rooms

at time instant t_k , we have that the equation for room $i \in 1, 2$ (assuming j is the other room) is given by:

$$\mathbf{s}_{k+1}(i) = \mathbf{s}_k(i) + b_i(\mathbf{s}_k(i) - x_a) + a(\mathbf{s}_k(j) - \mathbf{s}_k(i)) + h_{u_k} \zeta(\mathbf{s}_k(i), u_k) + \boldsymbol{\xi}_k(i) \quad (4)$$

where x_a represents the ambient temperature (assumed to be constant) and a the heat transfer rate between the rooms. The quantity b_i is a non-negative constant representing the average heat transfer rate from room i to the ambient and h_{u_k} denotes the heat rate supplied to room i by the corresponding heater at time k . The quantity $\zeta(\mathbf{s}_k(\cdot), u_k)$ characterises an accepted action (u_k) with a probability that, as in Section 2, decreases both as the temperature increases and as the heater level increases. The disturbances $\langle \boldsymbol{\xi}_k(i) \rangle_{0 \leq k \leq N-1}$ affecting the temperature evolution in room i are assumed to be a sequence of independent identically distributed Gaussian random variables with zero mean and variance ν^2 . We also assume that the disturbances affecting the temperature in the two rooms are independent.

The continuous transition kernel τ_u describing the evolution of the uncountable state $\mathbf{s} = (s(1), s(2))$ can easily be derived from (4). Let $\mathcal{N}(\cdot; \mu, V)$ denote the Gaussian measure over $(\mathbb{R}^2, \mathcal{B}(\mathbb{R}^2))$, with mean $\mu \in \mathbb{R}^2$ and covariance matrix $V \in \mathbb{R}^{2 \times 2}$. Then, $\tau_u : \mathcal{B}(\mathbb{R}^2) \times \mathcal{S} \rightarrow [0, 1]$ can be expressed as:

$$\tau_u(\cdot | s) = \mathcal{N}(\cdot; s + Bs + C, \nu^2 I) \quad (5)$$

where $B \in \mathbb{R}^{2 \times 2}$ with $[B]_{ii} = b_i - a$ and $[B]_{ij} = a$, and $C \in \mathbb{R}^2$ with $[C]_i = -b_i a + u$. With reference to the semantic classification in Section 2, we are dealing here with a decision process.

Let us select a time horizon of $N=180$ time steps. We are interested in the optimal probability and the corresponding policy that the model dynamics stay within a given “safe” temperature interval in both rooms, say $\mathcal{I} = [17.5, 22.5] \subset \mathcal{S}$ degrees Celsius, and also the optimal expected time and associated policy that the dynamics stay within the interval. We assume that the process is initialised with the temperature in each room being at the mid-point of this interval (if it is initialised outside it, then the associated probability is trivially equal to zero).

We proceed by abstracting the model as a labelled Markov chain [3] as follows. The set \mathcal{I} is partitioned uniformly into $B=5$ bins or sub-intervals. The labels of the model correspond to choosing the heat-flow level of the heaters for the next time instant. Regarding the atomic propositions and labelling function over the abstract LMC (and concrete LMP), we assign the atomic proposition *safe* to those states where the temperature is within the interval. In addition, to allow the analysis of the time spent within the temperature interval, we use the structure reward \mathbf{r} which assigns the reward 1 to states-label pairs (both of the LMC and LMP) for which the temperature in the state is within the interval and 0 otherwise.

We use PRISM to obtain the minimum and maximum probability of remaining within the *safe* temperature interval over the time horizon, and the minimum and maximum expected time spent in the safe interval up to the horizon. The properties used are $\mathbf{P}_{\max=?} [\square^{\leq K} \text{safe}]$ and $\mathbf{R}_{\max=?}^{\mathbf{r}} [C^{\leq K}]$, as well as the




Fig. 1. PCTL model checking for the case study

corresponding properties for minimum, rather than maximum, values (see the previous section for details of the notation). The results are presented in Fig. 1(a) and Fig. 1(b), respectively.

The graph plots demonstrate that the minimum probability quickly reaches zero, and that the minimum expected time stabilises as the time horizon increases. Examining with PRISM the policies that obtain these minimum values, we see that the policies coincide and correspond to never switching the heaters on (i.e. setting the heating level to be 0 at each step up until the time horizon). Although at first this may seem the obvious policy for minimising these values, an alternative policy could be to keep the heaters on full at each step (i.e. setting the heating level to 10), as it may be quicker to heat the rooms to above the temperature interval, as opposed to letting the rooms cool to below the interval.

Using PRISM, we find that this alternative approach is actually far less effective in the expected time case, and for small time horizons when considering the probabilistic invariance property. This is due to the fact that it takes much longer to heat a room to above the temperature interval than it does to reach the lower bound by keeping the heaters off. For example for a time bound of 10 minutes, the probability of remaining within the interval equals $1.68e-15$ when the heaters are kept off, while if the heaters are on full the probability of remaining within the interval is 0.01562. The fact that there is a chance that the heaters fail at each time step only increases the difference between these policies with regards to remaining within the temperature interval, as it is clearly detrimental to keep the heaters on full, but has no influence when the heaters are kept off. This can be seen in the expected time graph (see Fig. 1(b)), where the expected time of remaining within the temperature interval for the “full” policy keeps increasing while the minimum policy levels off.

In the case of the maximum values for the properties in Fig. 1, we see that for small time horizons there is a very high chance that we can remain within the temperature interval, but as the horizon increases the chance of remaining within the interval drops off. Consider the maximum expected time spent within the interval; this keeps increasing as the horizon increases, but at a lesser rate.

The reason for this behaviour is due to the fact that the heaters can fail and, once a heater fails, there is nothing we can do to stop the temperature in the corresponding room decreasing. Regarding the corresponding policies, we see that, while the heaters are working, the optimal approach is to initially set the heaters to be on full and then lower the heater level as one approaches the upper bound of the temperature interval. In addition, if the temperature in the rooms starts to drop, then the policy repeats the process by setting the heaters to high and then reducing as the temperature nears the upper bound of the interval.

7 Conclusions

This paper has put forward a computable technique to derive finite abstractions of labelled Markov processes (LMPs) in the form of labelled Markov Chains (LMCs), a probabilistic model related to Markov decision processes. The abstract LMC models are shown to correspond to the concrete LMPs via the notion of approximate probabilistic bisimulation. The technique enables the use of PRISM for probabilistic model checking and optimal policy synthesis over the abstract LMCs, extending its current capability to uncountable-state space models. The usefulness of the approach is demonstrated by means of a case study.

Acknowledgments. The authors are supported in part by the ERC Advanced Grant VERIWARE, the EU FP7 project HIERATIC, the EU FP7 project MoVeS, the EU FP7 Marie Curie grant MANTRAS, the EU FP7 project AMBI, and by the NWO VENI grant 016.103.020.

References

1. Abate, A.: A contractivity approach for probabilistic bisimulations of diffusion processes. In: Proc. 48th IEEE Conf. Decision and Control, Shanghai, China, pp. 2230–2235 (December 2009)
2. Abate, A., D’Innocenzo, A., Di Benedetto, M.D.: Approximate abstractions of stochastic hybrid systems. *IEEE Transactions on Automatic Control* 56(11), 2688–2694 (2011), doi:10.1109/TAC.2011.2160595
3. Abate, A., Katoen, J.-P., Lygeros, J., Prandini, M.: Approximate model checking of stochastic hybrid systems. *European Journal of Control* 16, 624–641 (2010)
4. Abate, A., Prandini, M.: Approximate abstractions of stochastic systems: A randomized method. In: 2011 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC), pp. 4861–4866 (2011)
5. Abate, A., Prandini, M., Lygeros, J., Sastry, S.: Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica* 44(11), 2724–2734 (2008)
6. Baier, C., Katoen, J.-P.: Principles of model checking. The MIT Press (2008)
7. Bertsekas, D.P., Shreve, S.E.: Stochastic optimal control: The discrete time case, vol. 139. Academic Press (1978)
8. Bianco, A., de Alfaro, L.: Model checking of probabilistic and nondeterministic systems. In: Thiagarajan, P.S. (ed.) FSTTCS 1995. LNCS, vol. 1026, pp. 499–513. Springer, Heidelberg (1995)

9. Bouchard-Cote, A., Ferns, N., Panangaden, P., Precup, D.: An approximation algorithm for labelled Markov processes: towards realistic approximation. In: Proc. 2nd Int. Conf. Quantitative Evaluation of Systems, QEST 2005 (2005)
10. Cattani, S., Segala, R., Kwiatkowska, M., Norman, G.: Stochastic transition systems for continuous state spaces and non-determinism. In: Sassone, V. (ed.) FOSSACS 2005. LNCS, vol. 3441, pp. 125–139. Springer, Heidelberg (2005)
11. Chaput, P., Danos, V., Panangaden, P., Plotkin, G.: Approximating markov processes by averaging. In: Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikolettseas, S., Thomas, W. (eds.) ICALP 2009, Part II. LNCS, vol. 5556, pp. 127–138. Springer, Heidelberg (2009)
12. Comanici, G., Panangaden, P., Precup, D.: On-the-fly algorithms for bisimulation metrics. In: Proc. 9th Int. Conf. Quantitative Evaluation of Systems (QEST 2012), pp. 681–692 (2012)
13. Danos, V., Desharnais, J., Panangaden, P.: Labelled markov processes: Stronger and faster approximations. *Electr. Notes Theor. Comput. Sci.* 87, 157–203 (2004)
14. Desharnais, J., Edalat, A., Panangaden, P.: Bisimulation for labelled Markov processes. *Information and Computation* 179(2), 163–193 (2002)
15. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labelled Markov processes. *Theoretical Computer Science* 318(3), 323–354 (2004)
16. Desharnais, J., Jagadeesan, R., Gupta, V., Panangaden, P.: The metric analogue of weak bisimulation for probabilistic processes. In: Proc. 17th Annual IEEE Symp. Logic in Computer Science (LICS 2002), pp. 413–422 (2002)
17. Desharnais, J., Laviolette, F., Tracol, M.: Approximate analysis of probabilistic processes: logic, simulation and games. In: Proc. 5th Int. Conf. Quantitative Evaluation of SysTems (QEST 2008), pp. 264–273 (September 2008)
18. Desharnais, J., Panangaden, P.: Continuous stochastic logic characterizes bisimulation of continuous-time Markov processes. *The Journal of Logic and Algebraic Programming* 56(1-2), 99–115 (2003)
19. Desharnais, J., Panangaden, P., Jagadeesan, R., Gupta, V.: Approximating labeled Markov processes. In: Proc. 15th Annual IEEE Symp. Logic in Computer Science (LICS 2000), pp. 95–105 (2000)
20. D’Innocenzo, A., Abate, A., Katoen, J.-P.: Robust PCTL model checking. In: Proc. 15th ACM Int. Conf. Hybrid Systems: computation and control, Beijing, PRC, pp. 275–285 (April 2012)
21. Esmail Zadeh Soudjani, S., Abate, A.: Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes. *SIAM Journal on Applied Dynamical Systems* 12(2), 921–956 (2013)
22. Fehnker, A., Ivančić, F.: Benchmarks for hybrid systems verification. In: Alur, R., Pappas, G.J. (eds.) HSCC 2004. LNCS, vol. 2993, pp. 326–341. Springer, Heidelberg (2004)
23. Ferns, N., Panangaden, P., Precup, D.: Metrics for finite Markov decision processes. In: Proc. 20th Conf. Uncertainty in Artificial Intelligence (UAI 2004), Banff, Canada, pp. 162–169 (2004)
24. Ferns, N., Panangaden, P., Precup, D.: Metrics for Markov decision processes with infinite state spaces. In: Proc. 21st Conf. Uncertainty in Artificial Intelligence, UAI 2005 (2005)
25. Ferns, N., Panangaden, P., Precup, D.: Bisimulation metrics for continuous Markov decision processes. *SIAM Journal of Computing* 60(4), 1662–1724 (2011)
26. Forejt, V., Kwiatkowska, M., Norman, G., Parker, D.: Automated verification techniques for probabilistic systems. In: Bernardo, M., Issarny, V. (eds.) SFM 2011. LNCS, vol. 6659, pp. 53–113. Springer, Heidelberg (2011)

27. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. *Formal Aspects of Computing* 6(5), 512–535 (1994)
28. Hernández-Lerma, O., Lasserre, J.B.: *Discrete-time Markov control processes, Applications of Mathematics*, vol. 30. Springer, New York (1996)
29. Huth, M.: On finite-state approximants for probabilistic computation tree logic. *Theoretical Computer Science* 346(1), 113–134 (2005)
30. Huth, M., Kwiatkowska, M.Z.: Quantitative analysis and model checking. In: *12th Annual IEEE Symp. Logic in Computer Science*, pp. 111–122. IEEE Computer Society (1997)
31. Julius, A.A., Pappas, G.J.: Approximations of stochastic hybrid systems. *IEEE Transactions on Automatic Control* 54(6), 1193–1203 (2009)
32. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) *CAV 2011. LNCS*, vol. 6806, pp. 585–591. Springer, Heidelberg (2011)
33. Larsen, K.G., Mardare, R., Panangaden, P.: Taking it to the limit: Approximate reasoning for markov processes. In: Rovan, B., Sassone, V., Widmayer, P. (eds.) *MFCS 2012. LNCS*, vol. 7464, pp. 681–692. Springer, Heidelberg (2012)
34. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. *Information and Computation* 94, 1–28 (1991)
35. Malhame, R., Chong, C.-Y.: Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *IEEE Transactions on Automatic Control* 30(9), 854–860 (1985)
36. Meyn, S.P., Tweedie, R.L.: *Markov chains and stochastic stability. Communications and Control Engineering Series*. Springer, London (1993)
37. Panangaden, P.: *Labelled Markov Processes*. Imperial College Press (2009)
38. Puterman, M.: *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, Inc. (1994)
39. Tkachev, I., Abate, A.: On infinite-horizon probabilistic properties and stochastic bisimulation functions. In: *2011 50th IEEE Conf. Decision and Control and European Control Conference (CDC-ECC)*, pp. 526–531 (2011)
40. Tkachev, I., Abate, A.: Characterization and computation of infinite horizon specifications over Markov processes. *Theoretical Computer Science* 515, 1–18 (2013)
41. Tkachev, I., Abate, A.: Formula-free finite abstractions for linear temporal verification of stochastic hybrid systems. In: *Proc. 16th ACM Int. Conf. Hybrid Systems: Computation and Control*, pp. 283–292 (2013)
42. Tkachev, I., Abate, A.: On approximation metrics for linear temporal model-checking of stochastic systems. In: *Proc. 17th ACM Int. Conf. Hybrid Systems: Computation and Control* (2014)
43. van Breugel, F., Worrell, J.: An algorithm for quantitative verification of probabilistic transition systems. In: Larsen, K.G., Nielsen, M. (eds.) *CONCUR 2001. LNCS*, vol. 2154, pp. 336–350. Springer, Heidelberg (2001)
44. van Breugel, F., Worrell, J.: Towards quantitative verification of probabilistic transition systems. In: Orejas, F., Spirakis, P.G., van Leeuwen, J. (eds.) *ICALP 2001. LNCS*, vol. 2076, pp. 421–432. Springer, Heidelberg (2001)
45. Wimmer, R., Becker, B.: Correctness issues of symbolic bisimulation computation for Markov chains. In: *Proc. 15th Int. GI/ITG Conf. Measurement, Modelling, and Evaluation of Computing Systems and Dependability and Fault Tolerance (MMB-DFT)*, pp. 287–301. Springer, Heidelberg (2010)
46. Zamani, M., Esfahani, P.M., Majumdar, R., Abate, A., Lygeros, J.: Symbolic control of stochastic systems via approximately bisimilar finite abstractions. *arXiv:1302.3868* (2013)