



Algorithmic Minimization of Uncertain Continuous-Time Markov Chains

Cardelli, Luca; Grosu, Radu; Larsen, Kim G.; Tribastone, Mirco; Tschaikowski, Max; Vandin, Andrea

Published in:
IEEE Transactions on Automatic Control

Link to article, DOI:
[10.1109/TAC.2023.3244093](https://doi.org/10.1109/TAC.2023.3244093)

Publication date:
2023

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Cardelli, L., Grosu, R., Larsen, K. G., Tribastone, M., Tschaikowski, M., & Vandin, A. (2023). Algorithmic Minimization of Uncertain Continuous-Time Markov Chains. *IEEE Transactions on Automatic Control*, 68(11), 6557-6572. <https://doi.org/10.1109/TAC.2023.3244093>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Algorithmic Minimization of Uncertain Continuous-Time Markov Chains

Luca Cardelli¹, Radu Grosu², Kim G. Larsen³, Mirco Tribastone⁴, Max Tschaikowski⁵ and Andrea Vandin⁶

Abstract—The assumption of perfect knowledge of rate parameters in continuous-time Markov chains (CTMCs) is undermined when confronted with reality, where they may be uncertain due to lack of information or because of measurement noise. Here we consider uncertain CTMCs (UCTMCs), where rates are assumed to vary non-deterministically with time from bounded continuous intervals. An uncertain CTMC can be therefore seen as a specific type of Markov decision process for which the analysis is computationally difficult. To tackle this, we develop a theory of minimization which generalizes the notion of lumpability for CTMCs. Our first result is a quantitative and logical characterization of minimization. Specifically, we show that the reduced UCTMC model has a macro-state for each block of a partition of the state space, which preserves value functions and logical formulae whenever rewards are equal within each block. The second result is an efficient minimization algorithm for UCTMCs by means of partition refinement. As application, we show that reductions in a number of CTMC benchmark models are robust with respect to uncertainties in original rates.

I. INTRODUCTION

Continuous-time Markov chains (CTMCs) play a central role for the stochastic modeling in a wide range of natural and engineering disciplines including chemistry [1], queuing theory [2], epidemiology [3], and biology [4]. A CTMC is typically characterized by a number of parameters such as arrival and service rates in a queuing system [2], transmission and recovery rates in epidemic processes [5], and kinetic rates in a chemical reaction network [6]. In essentially all practical situations, however, knowing the values of all parameters *precisely* is unlikely. This may be due to measurement noise when parameters are inferred from observations, as well as to our inability to accurately observe events at certain spatio-temporal scales—a well-known problem notably arising in computational systems biology (e.g., [7]). In addition, sometimes the modeler may wish to be deliberately imprecise about the value of certain parameters in order to account for the disagreement between the real system and its model.

¹ Luca Cardelli is with University of Oxford, UK: luca.a.cardelli@gmail.com

² Radu Grosu is with TU Wien, Austria: radu.grosu@tuwien.ac.at

³ Kim G. Larsen is with Aalborg University, Denmark: kgl@cs.aau.dk

⁴ Mirco Tribastone is with IMT Lucca, Italy: mirco.tribastone@imtlucca.it

⁵ Max Tschaikowski is with Aalborg University, Denmark: tschaikowski@cs.aau.dk

⁶ Andrea Vandin is with Sant'Anna School for Advanced Studies, Pisa, Italy and DTU Technical University of Denmark: andrea.vandin@santannapisa.it

L. Cardelli is supported by a Royal Society Research Professorship. The work has been partially supported by the ERC Advanced Grant LASSO, the Villum Investigator Grant S4OS, the EU-Ecsel project iDev40, the FFG project Adex, the PRIN project SEDUCE, no. 2017TWRCNB, the FWF project COCO no. M-2393-N32, the Poul Due Jensen Foundation grant no. 883901, and by the DFF RP1 project REDUCTO no. 9040-00224B.

These motivations call for frameworks where parameter uncertainty is a first-class citizen. In this paper we consider a model that we call *uncertain CTMCs* (UCTMCs). It extends CTMCs by defining bounded intervals for the values of the (time-varying) transition rates in order to model partial knowledge. The UCTMC dynamics depends on a scheduler that, at any point in time, nondeterministically picks a value within the interval to define the exponential rate at which the Markov chain makes a jump. Therefore, for instance, a time-inhomogeneous CTMC corresponds to the special case where the scheduler is deterministic but can be time-varying; a time-homogeneous CTMC corresponds to a deterministic scheduler that does not depend on time.

Overall, a UCTMC can be seen as a Markov decision process (MDP). Specifically, it can be alternatively defined as a time-inhomogeneous continuous-time MDP with an uncountable action space, which represents all values within the uncertainty intervals, see [8, Section 2.2]. In light of this connection, the analysis of a UCTMC corresponds to the optimization of some reward function. For example, the analogue to the transient probability distribution in a classic CTMC is the minimal and maximal probability of being in a state at a given time point across all schedulers. This can be obtained by optimizing with respect to a suitable reward. Therefore, it becomes apparent that, even for such a (relatively) straightforward analysis in the classic CTMC, its uncertain analogue is more challenging. Indeed, most existing results either consider finite actions spaces [9] or time-abstract policies [10].

Minimization, aka lumpability, is a natural way to cope with the analysis of computationally difficult problems. It seeks to find a more compact representation which can be related in some appropriate formal way to the original model for analysis purposes. For MDPs, the vast majority of minimization results has been obtained in the case of discrete-time and finite action spaces, by means of e.g., exact and approximate bisimulation, MAXQ, stochastic programming and homomorphism, see, e.g., [11], [12]. Since UCTMCs can be seen as continuous-time MDPs with infinite action spaces, the aforementioned approaches cannot be used in our context.

Here we develop a theory and an algorithm for the minimization for UCTMCs that extends the related notions of CTMC lumpability [2], [13] and stochastic bisimulation [14], [15]. Ordinary lumpability identifies a partition of the state space which induces a reduced (*lumped*) Markov chain where each macro-state represents a partition block; the probability of being in each macro-state at any time point is equal to the sum of the probabilities of the states of the original CTMC belonging to that block [13]. *Mutatis mutandis*, we present UCTMC lumpability as a partition of the state space which preserves any value function (whenever the rewards are equal for all states in each partition block). It turns out to be a conservative

extension of CTMC lumpability, in the sense that it collapses to the latter when all UCTMC transition rates are constants. For UCTMC we can prove a characterization result, namely that any partition of the state space that preserves value functions in the sense specified above must satisfy the conditions of UCTMC lumpability. This generalizes classic characterization results for CTMC lumpability from [13]. Additionally, we also establish that UCTMC lumpability enjoys a characterization by means of an extended continuous stochastic logic and thus generalize the classic logical characterization of lumpability [16], [17].


For a given CTMC, there trivially exists the *coarsest* ordinarily lumpable partition where all states collapse into a single block: this corresponds to the basic property that the whole probability mass is constant over time. Much more interestingly, for a given initial partition of states, it is possible to compute the coarsest *refinement* of such partition which satisfies the conditions of ordinary lumpability. In the literature of CTMCs, this is a well established result with algorithms that run in $\mathcal{O}(p \log n)$ time, where p is the number of transitions and n is the number of states of the CTMC [18], [19]. These algorithms are based on partition refinement, i.e., they iteratively split the blocks of the initial partition until the lumpability criteria are met.

Here we provide an efficient partition-refinement algorithm for UCTMCs. In particular, it uses the CTMC lumpability algorithm as an inner step: the coarsest UCTMC lumpable partition that refines a given initial partition of states is the coarsest refinement of the same initial partition in two time-homogeneous CTMCs. These are derived by choosing as rates the lower and upper bounds for all uncertainty intervals, respectively. The algorithm iterates through these two CTMCs until a fixed point: we prove that it takes $\mathcal{O}(pn \log n)$ time.

As an application, we consider the problem of analyzing the “robustness” of CTMC lumpability, i.e., to what extent the minimization depends on the specific choice of the parameters. Using a prototype implementation, we study whether adding uncertainty intervals around the constant values of the rates of several CTMC benchmark models preserves the original CTMC lumpability, demonstrating the scalability of our algorithm to models with millions of states.

Further work.: As indicated earlier, a UCTMC can be seen as a family of time-inhomogeneous CTMCs [8], [20], [17] which, in turn, can be interpreted as a time-inhomogeneous continuous-time MDP with an uncountable action space, representing the values within the uncertainty intervals. This model of uncertainty is different from the state of the art concerned with MDPs where the action space is finite and/or policies are time-independent (alternatively, untimed or time-invariant), see for instance [21], [22], [9], [10], [23], [24]. Another related model is that of parametric CTMCs and parametric MDPs [25], [26], [27], [28], [29], where certain transition probabilities have symbolic parameters. A parametric model underlies an (infinite) family of Markov models, one for each possible evaluation of the parameters. However, each member of this family is time-invariant because the instantiation of the parameters is assumed fixed throughout the time course evolution of the process. Parametric CTMCs whose parameters are subject to stochastic uncertainty have been considered

too [30] but come with statistical rather than deterministic guarantees. Most notions of lumpability and bisimulation for the aforementioned models of uncertainty impose constraints that must hold across all actions (in the case of MDPs [21], [26], [31]) or, analogously, for all parameter evaluations (for parametric Markov chains [26]). Instead, our notion of lumpability can aggregate states even when realizations of the uncertain transition rates make the resulting time-inhomogeneous Markov chain not lumpable. In order to clarify this difference, let us consider the simple graph structure in the right inset. If $q_{2,1}$ and $q_{3,1}$ are constant values, then the graph represents a continuous-time Markov chain. In this case, states 2 and 3 can be aggregated by



ordinary lumpability if $q_{2,1} = q_{3,1}$. In the case of a parametric Markov chain, $q_{2,1}$ and $q_{3,1}$ can be expressions over parameters; yet, parametric Markov chain lumpability requires these two expressions to be equal for all possible assignments of the parameters [26]—hence, each member of the family of Markov chains will be ordinarily lumpable. A similar remark applies to lumpability of parametric MDPs. Indeed, if $q_{i,j}(a)$ denotes the transition rate from state i into state j in the case of any action a , the lumpability condition requires that $q_{2,1}(a) = q_{3,1}(a)$. Instead, a UCTMC has bounded intervals as transitions. Applied to this simple example, our proposed notion of lumpability will require that the intervals of both transitions be equal; however, according to the semantic interpretation of a UCTMC, this model underlies behavior in the form of (time-varying) CTMCs which have different transition rates when the uncertainty is resolved. The closest notion to UCTMC lumpability is the alternating probabilistic bisimulation considered in [32] for discrete-time interval MDPs. Similarly to us, alternating probabilistic bisimulation: (i) does not require that realizations of the uncertain transition probabilities make the discrete-time Markov chain lumpable; (ii) can be computed in polynomial time; (iii) preserves quantitative and logical properties; however, in [32] it is not proved that the bisimulation is indeed necessary for the preservation of such properties.

A UCTMC can be seen as a family of time-inhomogeneous CTMCs where each time-inhomogeneous CTMC corresponds to a concrete realization of the nondeterminism. This closely aligns to open-loop control, allowing one to interpret a UCTMC as a bilinear control system (e.g., [33], [34]) with bounded controls accounting for uncertainties/disturbances [35], [36], [37], [38]. Due to this, UCTMC lumpability can be related to the bisimulation/abstraction of nonlinear dynamical systems [39], [33], [40], [41] where, for a given observation map and two copies of the same system, the largest abstraction gives rise to a reduced dynamical system which coincides with the original one up to the chosen observation map. Indeed, any UCTMC lumpability determines a linear observation map such that the lumped UCTMC is a (so-called consistent implementable) abstraction of the original one [42]. To the best of our knowledge, the computation of a (in a certain sense) minimal observation map has been investigated for linear dynamics with additive controls only [42]. Instead, UCTMC lumpability

considers the computation for a subclass of linear dynamics with multiplicative controls. As less closely related reduction approaches in control theory can be named [43] and [44] which consider the verification of quantitative and logical properties, respectively. Similarly to [32], a characterization is however not addressed and the time setting is discrete.

The present work extends [45] by characterizing the preservation of cost functionals rather than reachable probability sets. Moreover, we provide a closed-form solution for the value function and optimal control of a UCTMC by invoking Pontryagin's principle [46], allowing us in particular to avoid the computationally demanding DTMDP approximation from [45]. Moreover, using the closed-form solution, we demonstrate the speed-up due to UCTMC lumpability on benchmark models. All proofs are provided.

II. UNCERTAIN CONTINUOUS-TIME MARKOV CHAINS

Notation. We use ∂_t to denote the derivative with respect to time t , while x^T is the transpose of a vector x (in the following, we consider column vectors and use the transpose for row vectors). Pointwise equivalence of functions is denoted by \equiv , while $:=$ signifies a definition. Given two partitions \mathcal{H}_1 and \mathcal{H}_2 of a set \mathcal{S} , we say that \mathcal{H}_1 is a *refinement* of \mathcal{H}_2 if for any $H_1 \in \mathcal{H}_1$ there exists a (unique) $H_2 \in \mathcal{H}_2$ such that $H_1 \subseteq H_2$. We shall not distinguish among an equivalence relation and the partition induced by it.

We first introduce time-inhomogeneous (alternatively, time-varying) CTMCs.

Definition 1 (CTMC). A (time-varying) CTMC is a pair (\mathcal{S}, Q) where \mathcal{S} is a set of states $\mathcal{S} = \{1, \dots, n\}$ and $Q = (q_{i,j})_{i,j}$ is the transition rate matrix such that every transition rate function $q_{i,j} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is a piecewise continuous function with right limits when $i \neq j$ and $q_{i,i} = -\sum_{j \neq i} q_{i,j}$ for any i .

In what follows, quantities such as Q and π are time-dependent unless otherwise stated. To enhance readability, we will often suppress the explicit time dependence and write, for instance, π and Q instead of $\pi(t)$ and $Q(t)$, respectively. Explicit time dependence will be usually used in presence of different time points like $\pi(t)$ and $\pi(0)$. Abusing notation, we shall also write $Q \in [m; M]$ instead of $Q \in [m; M]^{\mathbb{R}_{\geq 0}}$ for any $m, M \in \mathbb{R}_{\geq 0}^S$ with $m \leq M$.

The following result relates the (transient) probability distributions of (\mathcal{S}, Q) to the Kolmogorov equations for time-varying transition rates [8, Section 2.2].

Theorem 1. Given a CTMC (\mathcal{S}, Q) and an initial probability distribution $\pi[0]$, the probability distributions $\pi(t)$ exist and satisfy, for all $t \in \mathbb{R}_{\geq 0}$, the Kolmogorov equations

$$\partial_t \pi(t)^T = \pi(t)^T Q(t), \quad \text{where } \pi(0) = \pi[0]. \quad (1)$$

For time-homogeneous CTMCs, ordinary lumpability requires a partition of the state space such that, for any two states in the same block, the total cumulative rate into any block is equal. The lumped CTMC has one state for each block, with transitions given by such cumulative rates. The Kolmogorov equations of the lumped CTMC preserve the sum of the probabilities in each block; the characterization result

states that if a partition preserves such probabilities for any initial condition, then it must satisfy the criteria for ordinary lumpability (e.g., [13]).

By Theorem 1, ordinary lumpability for time-varying CTMCs is a straightforward generalization which requires pointwise equivalence between cumulative rate functions.

Theorem 2 (Ordinary Lumpability). Given a CTMC (\mathcal{S}, Q) , a partition \mathcal{H} of the set of states \mathcal{S} is an ordinary lumpability if

$$\sum_{j \in H'} q_{i_1, j} \equiv \sum_{j \in H'} q_{i_2, j}, \quad \text{for all } H, H' \in \mathcal{H} \text{ and } i_1, i_2 \in H.$$

The lumped CTMC $(\hat{\mathcal{S}}, \hat{Q})$ is given by:

- states $\hat{\mathcal{S}} := \{i_H \mid H \in \mathcal{H}\}$, where $i_H \in H$ is an arbitrary representative of H ;
- transition rate matrix $\hat{Q} = (\hat{q}_{i_H, i_{H'}})_{H, H'}$, where

$$\hat{q}_{i_H, i_{H'}} := \sum_{j \in H'} q_{i_H, j} \quad \text{for all } H, H' \in \mathcal{H}.$$

If the initial probability distribution of $(\hat{\mathcal{S}}, \hat{Q})$ is defined by $\hat{\pi}[0]_{i_H} = \sum_{i \in H} \pi[0]_i$ for all $H \in \mathcal{H}$ and the transient probability distributions of $(\hat{\mathcal{S}}, \hat{Q})$ are denoted by $\hat{\pi}$, the following holds.

- If \mathcal{H} is an ordinary lumpability, then $\hat{\pi}_{i_H} \equiv \sum_{i \in H} \pi_i$ for all $H \in \mathcal{H}$ and $\pi[0]$.
- If \mathcal{H} is such that $\hat{\pi}_{i_H} \equiv \sum_{i \in H} \pi_i$ for all $H \in \mathcal{H}$ and $\pi[0]$, then \mathcal{H} is an ordinary lumpability.

Proof. Follows from the discrete-time case [13]. \square

We can now formally define a UCTMC. It allows transition rates to vary non-deterministically with time within bounded continuous intervals.

Definition 2 (Uncertain CTMC). An uncertain CTMC $(\mathcal{S}, m, M, r, \phi)$ is the following family of time-varying CTMCs:

- a set of states $\mathcal{S} = \{1, \dots, n\}$;
- lower bounds of transition rates $0 \leq m = (m_{i,j})_{i \neq j}$;
- upper bounds of transition rates $m \leq M = (M_{i,j})_{i \neq j}$;
- final reward function $\phi : \mathcal{S} \rightarrow \mathbb{R}$;
- running reward function $r : \mathbb{R}_{\geq 0} \times \mathcal{S} \rightarrow \mathbb{R}$.

The lower and upper bounds are numbers, while reward functions r and ϕ are assumed to be measurable. Any family member of a UCTMC is called a *realization* (of its nondeterminism).

A UCTMC $(\mathcal{S}, m, M, r, \phi)$ induces two extremal (time-homogeneous) CTMCs (\mathcal{S}, m) and (\mathcal{S}, M) by fixing all lower and upper bounds, respectively, for each transition rate. We call a UCTMC deterministic if $m = M$ because this yields a CTMC.

The value function of a UCTMC is essentially defined as the maximal (minimal) reward across all $Q \equiv (Q(t))_{t \geq 0} \in [m; M]$. It intuitively accounts for the uncertainty across all possible realizations. Reward functions are part of the UCTMC definition because, in general, they influence the extent to which the optimization problem induced by a family of CTMCs can be simplified.




Fig. 1: The $M/M/1/2$ queue with background- (b) and user-level (u) jobs as running example.

Definition 3 (UCTMC value function). Fix a UCTMC $\mathcal{U} = (\mathcal{S}, m, M, r, \phi)$ and a finite time horizon $T \geq 0$. The maximal value obtainable at $T \geq 0$ if starting in $i \in \mathcal{S}$ is

$$\begin{aligned} V^{\sup}(i, T) &= \sup_{Q \in [m; M]} V^Q(i, T) \\ &= \sup_{Q \in [m; M]} \mathbb{E} \left[\int_0^T r_{\mathcal{U}(t)}(t) dt + \phi_{\mathcal{U}(T)} \right] \\ &= \sup_{Q \in [m; M]} \left[\int_0^T \sum_{j \in \mathcal{S}} r_j(t) \pi_j^Q(t) dt + \sum_{j \in \mathcal{S}} \phi_j \pi_j^Q(T) \right], \end{aligned}$$

where $\pi_j^Q(0) = \mathbb{1}_i$ and $\pi_j^Q(t)$ denote the solution of (1) using the transition rate matrix $Q = (q_{i,j})_{i,j}$. The minimal value is defined in a similar way, that is

$$V^{\inf}(i, T) = \inf_{Q \in [m; M]} V^Q(i, T).$$

Remark 1. The probabilities with which the realizations of UCTMC $(\mathcal{S}, m, M, r, \phi)$ can reach a set $S \subseteq \mathcal{S}$ at $T \geq 0$ comprise the set of reachability probabilities $\{\sum_{i \in S} \pi_i^Q(T) \mid Q \in [m; M]\}$. The reachability probability can be expressed by the rewards $r = 0$ and $\phi = \mathbb{1}_S$, where $\mathbb{1}_S$ is the characteristic function of S .

Example. We will use as queuing system as a running example. This is essentially a birth/death process, a fundamental type of CTMCs (e.g., [47]). When used to model real systems such as computer applications, the assumptions made in a classical time-homogeneous setting may be restrictive [48]. Here we consider an $M/M/1/2$ queue: it has exponentially distributed arrival and service times, 1 server, and processes jobs of two classes according to a first-come first-served scheduling (FCFS). For illustrative purposes, we show the case of a buffer of size 2, i.e., the service facility can accept at most two jobs, e.g., background processes b and user-level processes u . Each state of the Markov chain tracks the current configuration of the queue. Real-world phenomena may lead to time- or class-dependent behaviors. For example, time-dependent arrival rates account for peak/off-peak variations (e.g., [49]); service rates may degrade as the

queue length increases, as would be identified by state-of-the-art learning methods for CTMCs [48]. If one wishes to account more precisely for these effects, arrival and service rates can be assumed to be contained in some intervals $[\underline{\lambda}; \bar{\lambda}]$ and $[\underline{\mu}; \bar{\mu}]$, respectively. The system can thus be modeled by the UCTMC depicted using standard notation in Figure 1. State ϵ denotes an empty queue; the other states show the queue configuration, e.g., state ub is the full queue where a user job is at the head, followed by a background job.

One of the most interesting performance metrics is the utilization, i.e., the probability that the queue is not empty [47]. In light of Remark 1, one can consider the final reward $\phi = \mathbb{1}_{\mathcal{S} \setminus \{\epsilon\}}$ which describes the probability of not being in state ϵ . Likewise, one could use the discounted reward $r = e^{-at} \mathbb{1}_{\mathcal{S} \setminus \{\epsilon\}}$ with discount factor $a > 0$. In the time-homogeneous deterministic case when $\underline{\lambda} = \bar{\lambda}$ and $\underline{\mu} = \bar{\mu}$, the resulting CTMC admits the ordinary lumpability consisting of blocks $\{\epsilon\}$, $\{b, u\}$, and $\{bb, bu, ub, uu\}$; in the lumped CTMC neither the job order nor the job class has to be tracked, and the model becomes the textbook $M/M/1/2$ single-class queue [47]. Notice that in the lumped CTMC the utilization can be recovered exactly because the state ϵ is not aggregated. Instead, we remark that, for nontrivial intervals of rates, the UCTMC admits time-varying behaviors that break the symmetries for CTMC ordinary lumpability by Theorem 2.

A. Closed-form Solutions of UCTMCs

Akin to the derivation of the famous LQR formula for linear control systems [46], we use Pontryagin's principle, a cornerstone of optimal control theory [46], to derive a closed-form solution for UCTMCs. To this end, we generalize the result from [50] which considered UCTMCs with final rewards only. In contrast to [50], we do not rely on the sufficient version of Pontryagin's principle [51] but use a direct argument.

Theorem 3. Assume that we are given a UCTMC $\mathcal{U} = (\mathcal{S}, m, M, r, \phi)$. Then, for a given finite time horizon $T \geq 0$, a minimizing $Q^* \in [m; M]$ can be obtained by solving the adjoint system of equations

$$\partial_t p_i(t) = - \sum_{j \neq i} (p_j - p_i) q_{i,j}^*(t) + r_i(t), \quad i \in \mathcal{S}, \quad (2)$$

where for all $i, j \in \mathcal{S}$ with $i \neq j$:

$$q_{i,j}^*(t) = \begin{cases} m_{i,j} & , p_j(t) - p_i(t) < 0 \\ M_{i,j} & , p_j(t) - p_i(t) \geq 0 \end{cases}$$

and $p_i(T) = -\phi_i$. That is, Q^* yields

$$V^{\inf}(i, T) = \inf_{Q \in [m; M]} V^Q(i, T) = V^{Q^*}(i, T), \quad i \in \mathcal{S}.$$

Instead, V^{\sup} is obtained if $p_i(T) = \phi_i$ and $r_i(t)$ is replaced with $-r_i(t)$ in (2) for all $i \in \mathcal{S}$.

Proof. See Appendix A. \square

Theorem 3 ensures that the value function of a UCTMC can be computed by solving a system of differential equations of size $2|\mathcal{S}|$.

B. UCTMCs as MDPs

We end the section by relating UCTMCs to MDPs. We begin by showing that a UCTMC can be interpreted as an instance of a time-inhomogeneous continuous-time MDP (CTMDP). To see this, we consider a CTMDP with the scheduler model as in [8, Section 2.2], which can be intuitively described as follows. For a sufficiently small time step $h > 0$, a CTMDP that is in state $i \in \mathcal{S}$ at some time $kh \geq 0$, where $k \geq 0$ is an integer, may choose an action a_i from $\mathcal{A}(i)$, the set of available actions in state i . With this, the CTMDP remains in state i on $[kh; kh + h)$, while at time $kh + h$ the state is:

- $j \neq i$, with probability $q_{i,j}(kh, a_i)h + o(h)$, where $o(h)$ refers to the standard small- o notation, while $q_{i,j}(kh, a_i)$ denotes the transition rate from state i into state j at time kh under action a_i ;
- i , with probability $1 + q_{i,i}(kh, a_i)h + o(h)$.

Note that $q_{i,j}(kh, a_i)h + o(h)$ and $1 + q_{i,i}(kh, a_i)h + o(h)$ can be interpreted as transition probabilities of the embedded discrete-time Markov chain [52] under action a_i at step k . Indeed, in the special case when the transition rates are time-invariant, the *time-homogeneous* CTMDP admits a characterization in terms of sojourn times and an embedded chain. According to this characterization, the choice of action $a_i \in \mathcal{A}(i)$ upon entering state $i \in \mathcal{S}$ gives a sojourn time in state i that is exponentially distributed with rate $-q_{i,i}(a_i)$, and the probability to move into a state $j \neq i$ equal to $-q_{i,j}(a_i)/q_{i,i}(a_i)$ (see Theorem 2.8.2 in [52]).

Under this model, the discussion in [8, Section 2.2] yields the following relationship between a UCTMC and a CTMDP, where, essentially the uncountable many actions of the latter encode the uncertainty intervals of the former.

Theorem 4. *For a given UCTMC $(\mathcal{S}, m, M, r, \phi)$, consider the CTMDP $(\mathcal{S}, \mathcal{A}, \mathcal{M}, r, \phi)$ where an action taken at time t in state i , denoted by $a_i(t)$, is a row vector such that each component $a_{i,j}(t)$ determines the transition rate from i into j at time t . More formally:*

- the set of actions in state $i \in \mathcal{S}$ be given by $\mathcal{A}(i) = \prod_{j \neq i} [m_{i,j}; M_{i,j}]$;
- the transition rate from state i to state j at time t under action $a_i \in \mathcal{A}(i)$ is denoted by $q_{i,j}(t, a_i)$ and is given by $a_{i,j} \in [m_{i,j}; M_{i,j}]$, where $a_{i,j}$ is the j -th entry of a_i ;
- the policies form the set \mathcal{M} and are given by piecewise continuous $a : [0; \infty) \rightarrow \prod_{i \in \mathcal{S}} \mathcal{A}(i)$.

If the value of CTMDP $(\mathcal{S}, \mathcal{A}, \mathcal{M}, r, \phi)$ arises by replacing V with W and $(q_{i,j})_{i,j}$ with $(a_{i,j})_{i,j}$ in Definition 3, then

$$V^*(i, T) = W^*(i, T)$$

for all $i \in \mathcal{S}$, $T \geq 0$ and $*$ \in $\{\inf, \sup\}$.

Proof. Follows from [8, Section 2.2]. \square

Thanks to Theorem 4, it is in principle possible to approximate UCTMCs by means of DTMDPs. While the corresponding approximation is computationally challenging compared to the closed-form solution of Theorem 3, it allows one to relate UCTMC lumpability to alternating probabilistic bisimulation, see Appendix B.




Fig. 2: Lumped UCTMC from Figure 1.

III. UCTMC LUMPABILITY

We define UCTMC lumpability in terms of lumpability of the two extremal CTMCs. Additionally we require that states within each block have the same rewards.

Definition 4 (UCTMC Lumpability). *A partition \mathcal{H} of \mathcal{S} is a UCTMC lumpability of UCTMC $(\mathcal{S}, m, M, r, \phi)$ if*

- \mathcal{H} is an ordinary lumpability of both CTMCs (\mathcal{S}, m) and (\mathcal{S}, M) ;
- r and ϕ are constant on \mathcal{H} , that is, $r_{i_1}(t) = r_{i_2}(t)$ and $\phi_{i_1} = \phi_{i_2}$ for all $H \in \mathcal{H}$, $i_1, i_2 \in H$, $t \geq 0$.

The lumped UCTMC is obtained as follows.

Definition 5 (Lumped UCTMC). *Let \mathcal{H} be a UCTMC lumpability of $(\mathcal{S}, m, M, r, \phi)$ and fix, for each $H \in \mathcal{H}$, some representative $i_H \in H$. The lumped UCTMC has:*

- set of states $\hat{\mathcal{S}} := \{i_H \mid H \in \mathcal{H}\}$,
- lower bounds $\hat{m}_{i_H, i_{H'}} := \sum_{j \in H'} m_{i_H, j}$,
- upper bounds $\hat{M}_{i_H, i_{H'}} := \sum_{j \in H'} M_{i_H, j}$,
- rewards $\hat{r}_{i_H}(t) := r_{i_H}(t)$ and $\hat{\phi}_{i_H} := \phi_{i_H}$,

where $i_H, i_{H'} \in \hat{\mathcal{S}}$.

Example. For the UCTMC from Figure 1, the UCTMC lumpability $\mathcal{H} = \{\{\epsilon\}, \{b, u\}, \{bb, bu, ub, uu\}\}$ induces the lumped UCTMC in Figure 2 for all rewards that are constant on \mathcal{H} . Each state is labeled with a representative of the corresponding partition block. We remark that the lumped UCTMC corresponds to an $M/M/1/2$ queue with FCFS scheduling and uncertain arrival and service rates where the identity of the jobs and their order is not tracked.

A. Characterization via Value Functions

The next result states that the value function of a UCTMC coincides with the value function of its lumped UCTMC. Moreover, it allows to retrieve from an optimal control of the lumped UCTMC an optimal control of the original UCTMC, thus circumventing the solution of the original optimization problem altogether.

Theorem 5 (Value preservation). *Assume that \mathcal{H} is a UCTMC lumpability of $\mathcal{U} = (\mathcal{S}, m, M, r, \phi)$. Then, for any $H \in \mathcal{H}$, $i \in H$ and $T \geq 0$, it holds that*

$$V_{\hat{\mathcal{U}}}^{\sup}(i_H, T) = V_{\mathcal{U}}^{\sup}(i, T) \quad \text{and} \quad V_{\hat{\mathcal{U}}}^{\inf}(i_H, T) = V_{\mathcal{U}}^{\inf}(i, T),$$

where $\hat{\mathcal{U}}$ refers to the lumped UCTMC induced by \mathcal{H} . Additionally, if $\hat{Q} \in [\hat{m}; \hat{M}]$ is minimizing the cost, that is

$$V_{\hat{\mathcal{U}}}^{\inf}(i_H, T) = \inf_{\hat{Q} \in [\hat{m}; \hat{M}]} V_{\hat{\mathcal{U}}}^{\hat{Q}}(i_H, T), \quad H \in \mathcal{H},$$

then the minimal cost of the original UCTMC is attained via

$$q_{i,j}(t) = m_{i,j} + \frac{M_{i,j} - m_{i,j}}{\sum_{j' \in H'} (M_{i,j'} - m_{i,j'})} (\hat{q}_{i_H, i_{H'}}(t) - \hat{m}_{i_H, i_{H'}}),$$

where $H, H' \in \mathcal{H}$ satisfy $i \in H$ and $j \in H'$. The statement remains true when minimization is replaced with maximization.

Proof. See Appendix C. \square

Example. Applied to the running example, this result states that the minimum and maximum utilization is preserved in the lumped UCTMC.

The next result is a converse of Theorem 5. Together, both results provide a characterization of UCTMC lumpability in terms of value functions.

Theorem 6 (Characterization). *Fix a triple (\mathcal{S}, m, M) , some partition \mathcal{H} of \mathcal{S} and a triple $(\hat{\mathcal{S}}, \hat{m}, \hat{M})$ such that $\hat{\mathcal{S}} = \{i_H \mid H \in \mathcal{H}\}$. Assume that for all reward functions $r : \mathbb{R}_{\geq 0} \times \mathcal{S} \rightarrow \mathbb{R}$ and $\phi : \mathcal{S} \rightarrow \mathbb{R}$ that are constant on \mathcal{H} , it holds that*

$$V_{\mathcal{U}}^*(i_H, T) = V_{\hat{\mathcal{U}}}^*(i, T), \text{ for all } H \in \mathcal{H}, T \geq 0,$$

where $\mathcal{U} = (\mathcal{S}, m, M, r, \phi)$ and $\hat{\mathcal{U}} = (\hat{\mathcal{S}}, \hat{m}, \hat{M}, \hat{r}, \hat{\phi})$ with $\hat{r} = r|_{\mathbb{R}_{\geq 0} \times \hat{\mathcal{S}}}$ and $\hat{\phi} = \phi|_{\hat{\mathcal{S}}}$. Then

- \mathcal{H} is a UCTMC lumpability of $(\mathcal{S}, m, M, r, \phi)$ for all rewards $r, \phi : \mathcal{S} \rightarrow \mathbb{R}$ that are constant on \mathcal{H} and;
- $(\hat{\mathcal{S}}, \hat{m}, \hat{M}, r|_{\mathbb{R}_{\geq 0} \times \hat{\mathcal{S}}}, \phi|_{\hat{\mathcal{S}}})$ is its lumped UCTMC.

Proof. See Appendix C. \square

We next note that UCTMC lumpability is a conservative generalization of ordinary lumpability.

Lemma 1 (Generalization). *Assume that \mathcal{H} is a UCTMC lumpability of a deterministic UCTMC $(\mathcal{S}, m, M, r, \phi)$. Then, \mathcal{H} is an ordinary lumpability.*

Proof. Trivial. \square

The next two corollaries extend our main results. This first accounts for uncertain reward functions.

Corollary 1 (Uncertain reward functions). *Fix a UCTMC $\mathcal{U} = (\mathcal{S}, m, M, r, \phi)$, a partition \mathcal{H} of \mathcal{S} and running rewards \underline{r}, \bar{r} and final rewards $\underline{\phi}, \bar{\phi}$ that are constant on \mathcal{H} and that satisfy $\underline{r} \leq r \leq \bar{r}$ and $\underline{\phi} \leq \phi \leq \bar{\phi}$ pointwise. Then, if*

- \mathcal{H} is a UCTMC lumpability of $\bar{\mathcal{U}} = (\mathcal{S}, m, M, \bar{r}, \bar{\phi})$ and;
- \mathcal{H} is a UCTMC lumpability of $\underline{\mathcal{U}} = (\mathcal{S}, m, M, \underline{r}, \underline{\phi})$,

we obtain for all $T \geq 0$, $H \in \mathcal{H}$ and $i \in H$:

$$\begin{aligned} V_{\bar{\mathcal{U}}}^{\sup}(i, T) &\leq V_{\mathcal{U}}^{\sup}(i, T) = V_{\underline{\mathcal{U}}}^{\sup}(i_H, T) \\ V_{\underline{\mathcal{U}}}^{\inf}(i_H, T) &= V_{\mathcal{U}}^{\inf}(i, T) \leq V_{\bar{\mathcal{U}}}^{\inf}(i, T), \end{aligned} \quad (3)$$

where $\bar{\mathcal{U}}$ and $\underline{\mathcal{U}}$ is the lumped UCTMC of \mathcal{U} and $\bar{\mathcal{U}}$, respectively. Conversely, if (3) holds true for all $T \geq 0$, $H \in \mathcal{H}$, $i \in H$ and UCTMCs \mathcal{U} as described above, then we obtain a)-b).

Proof. Follows readily from Theorem 5, Theorem 6 and the fact that $\underline{r} \leq r \leq \bar{r}$ and $\underline{\phi} \leq \phi \leq \bar{\phi}$. \square

The second result describes how the optimization over general CTMC families can be over-approximated.

Corollary 2 (Over-approximation). *Fix some family of time-varying CTMCs $\mathcal{U} = (\mathcal{S}, \mathcal{F}, r, \phi)$, where*

- \mathcal{S} is a set of states;
- \mathcal{F} is some family of piecewise continuous transition rate matrices with range $[m; M]$ for some $m \leq M$ in $\mathbb{R}_{\geq 0}^S$;
- r and ϕ are running and final rewards, respectively.

Assume that $\underline{r} \leq r \leq \bar{r}$ and $\underline{\phi} \leq \phi \leq \bar{\phi}$ are such that $\underline{r}, \bar{r}, \underline{\phi}, \bar{\phi}$ are constant on some partition \mathcal{H} of \mathcal{S} and that conditions a)-b) from Corollary 1 are true. Then

$$\begin{aligned} V_{\underline{\mathcal{U}}}^{\inf}(i_H, T) &\leq \inf_{Q \in \mathcal{F}} V^Q(i, T) \\ \sup_{Q \in \mathcal{F}} V^Q(i, T) &\leq V_{\bar{\mathcal{U}}}^{\sup}(i, T), \end{aligned}$$

for all $T \geq 0$, $H \in \mathcal{H}$ and $i \in H$.

Proof. Follows from Corollary 1 and the fact that any element of \mathcal{F} is a piecewise continuous function with range $[m; M]$. \square

We end the section by noting the following.

Remark 2. One may wonder whether the aforementioned lumping results can be established by lumping ODE system (2) in Theorem 3. Specifically, one may apply forward differential equivalence [53] to (2), hoping that sums $\sum_{i \in H} \pi_i$ and $\sum_{i \in H} p_i$ are preserved and describe a closed-form expression of a hypothetical lumped UCTMC. Noting that the optimal transition rates are defined through differences $p_i - p_j$, however, we note that this is not possible.

B. Characterization via Temporal Logics

We next relate UCTMC lumpability to continuous stochastic logic for time-varying CTMCs [16], [17] (CSL).

Definition 6 (CSL). *The CSL syntax is given by*

$$\phi ::= a \mid \phi \wedge \phi \mid \neg \phi \mid \mathcal{P}_{\bowtie p}(X^{[t_0; t_1]}\phi) \mid \mathcal{P}_{\bowtie p}(\phi U^{[t_0; t_1]}\phi),$$

where $a \in \mathcal{A}$ and \mathcal{A} is the nonempty finite set of atomic propositions, $p \in [0; 1]$ is a probability, $\bowtie \in \{<, \leq, \geq, >\}$ and $0 \leq t_0 \leq t_1 < \infty$. For a time-varying CTMC with piecewise analytic transition rates $(q_{i,j})_{i,j}$, the satisfiability relation is given by structural induction over ϕ .

- $i, t \models a$ if and only if $a \in \mathcal{L}(i)$;
- $i, t \models \phi_1 \wedge \phi_2$ if and only if $i, t \models \phi_1$ and $i, t \models \phi_2$;
- $i, t \models \neg \phi$ if and only if not $i, t \models \phi$;
- $i, t \models \mathcal{P}_{\bowtie p}(X^{[t_0; t_1]}\phi)$ if and only if $\mathbb{P}\{\sigma \mid \sigma, t \models X^{[t_0; t_1]}\phi\} \bowtie p$ with $\pi(t) = e_i$;
- $i, t \models \mathcal{P}_{\bowtie p}(\phi_1 U^{[t_0; t_1]}\phi_2)$ if and only if $\mathbb{P}\{\sigma \mid \sigma, t \models \phi_1 U^{[t_0; t_1]}\phi_2\} \bowtie p$ with $\pi(t) = e_i$;
- $\sigma, t \models X^{[t_0; t_1]}\phi$ if and only if $t_\sigma[1] \in [t + t_0; t + t_1]$ and $\sigma[1], t_\sigma[1] \models \phi$;
- $\sigma, t \models \phi_1 U^{[t_0; t_1]}\phi_2$ if and only if there exists a $t' \in [t + t_0; t + t_1]$ such that $\sigma @ t', t' \models \phi_2$ and $\sigma @ t'', t'' \models \phi_1$ for all $t'' \in [t + t_0; t + t']$,

where $\mathcal{L}(i)$ is the set of atomic propositions valid in state i , \mathbb{P} is the probability measure, $\pi(t)$ the probability distribution of the CTMC at time t , σ a path of the CTMC, $\sigma @ t$ the state of the CTMC at time point t , $\sigma[1]$ the state at the time of the first jump and $t_\sigma[1]$ the corresponding time point.

The Boolean operators \wedge and \rightarrow are defined as usual; likewise, $\text{tt} := a \vee \neg a$ and $\text{ff} := a \wedge \neg a$ for some $a \in \mathcal{A}$. As in the case of classic model checking [54], \mathbf{X} refers to the next operator, while \mathbf{U} corresponds to the until operator. The assumption of piecewise analyticity is needed to ensure the well-definedness, please refer to [17], [55] for details. This motivates the following.

Definition 7. Given a UCTMC $(\mathcal{V}, m, M, r, \phi)$, any piecewise analytic transition rate matrix $Q(\cdot) \in [m; M]$ is called admissible.

We next extend CSL to UCTMCs by defining a formula to be true when it is satisfied by all admissible Q . This allows one to study safety properties in presence of uncertainty.

Definition 8 (CSL for UCTMCs). Given a UCTMC $(\mathcal{V}, m, M, r, \phi)$, the CSL syntax is

$$\phi ::= a \mid \phi \wedge \phi \mid \neg \phi \mid \mathcal{P}_{\bowtie p}^{\forall}(\mathbf{X}^{[t_0:t_1]}\phi) \mid \mathcal{P}_{\bowtie p}^{\forall}(\phi \mathbf{U}^{[t_0:t_1]}\phi)$$

For an arbitrary small but fixed time step $h > 0$, let \underline{t} denote the smallest grid point in $\{0, h, 2h, \dots\}$ that minimizes the distance to $t \geq 0$, i.e., $\underline{t} = h \cdot \lfloor t/h \rfloor$, where $\lfloor \cdot \rfloor$ is the floor function. For a given labeling function $\mathcal{L} : \mathcal{V} \rightarrow 2^{\mathcal{V}}$ and initial probability distribution $\pi[0]$, the satisfiability operator is defined by induction for any $i \in \mathcal{S}$ and $t \geq 0$ as:

- $i, t \models a$ iff $a \in \mathcal{L}(i)$;
- $i, t \models \phi_1 \wedge \phi_2$ iff $i, t \models \phi_1$ and $i, t \models \phi_2$;
- $i, t \models \neg \phi$ iff not $i, t \models \phi$;
- $i, t \models \mathcal{P}_{\bowtie p}^{\forall}(\mathbf{X}^{[t_0:t_1]}\phi_1)$ iff $i, \underline{t} \models \mathcal{P}_{\bowtie p}(\mathbf{X}^{[t_0:t_1]}\phi)$ for all $Q \in [m; M]$;
- $i, t \models \mathcal{P}_{\bowtie p}^{\forall}(\phi_1 \mathbf{U}^{[t_0:t_1]}\phi_2)$ iff $i, \underline{t} \models \mathcal{P}_{\bowtie p}(\phi_1 \mathbf{U}^{[t_0:t_1]}\phi_2)$ for all $Q \in [m; M]$.

As usual, existential quantification is given via negation, i.e., $\mathcal{P}_{\bowtie p}^{\exists}(\Phi) := \neg \mathcal{P}_{\bowtie p}^{\forall}(\neg \Phi)$, where $\neg \bowtie$ is defined in the obvious manner (e.g., $\neg \leq$ is $>$), while \vee, \rightarrow are defined using \wedge, \neg . Together with the assumption of piecewise analytic transition rate functions, the usage of \underline{t} in Definition 8 ensures that the function $t \mapsto i, t \models \phi$ has finitely many discontinuity points on any bounded time interval. Amongst other, this technical ingredient allows one to prove the following.

Theorem 7 (Preservation of CSL). Let \mathcal{H} be a UCTMC lumpability of UCTMC \mathcal{U} and let $\hat{\mathcal{U}}$ be the underlying lumped UCTMC. Further, assume that $\mathcal{L}(i) = \mathcal{L}(j)$ for all $H \in \mathcal{H}$ and $i, j \in H$. With this, define $\hat{\mathcal{A}} := \mathcal{A}$ and $\hat{\mathcal{L}}(i_H) := \mathcal{L}(i_H)$ for all $H \in \mathcal{H}$. Then

$$i, t \models_{\mathcal{U}} \phi \iff i_H, t \models_{\hat{\mathcal{U}}} \phi$$

for any $t \geq 0$, $h > 0$, block $H \in \mathcal{H}$, state $i \in H$, \mathbf{X} -operator free CSL formula ϕ and initial probability distribution $\pi[0]$.

Proof. See Appendix D. \square

The next result is a converse of Theorem 7 and establishes a logical characterization of UCTMC lumpability.

Theorem 8 (Logical Characterization). Fix a UCTMC $(\mathcal{V}, m, M, r, \phi)$, a partition \mathcal{H} of \mathcal{V} and let $\mathcal{L}, \hat{\mathcal{A}}$ and $\hat{\mathcal{L}}$ be

Algorithm 1 Computation of coarsest UCTMC lumpability.

Require: Uncertain CTMC $(\mathcal{S}, m, M, r, \phi)$, and initial partition \mathcal{H} respecting r and ϕ , i.e., r and ϕ are constant on each $H \in \mathcal{H}$.

- 1: **while true do**
- 2: $\mathcal{H}' \leftarrow$ coarsest ordinary lumpability of CTMC (\mathcal{S}, m) that refines \mathcal{H}
- 3: $\mathcal{H}'' \leftarrow$ coarsest ordinary lumpability of CTMC (\mathcal{S}, M) that refines \mathcal{H}'
- 4: **if** $\mathcal{H}'' = \mathcal{H}$ **then**
- 5: **return** \mathcal{H}''
- 6: **else**
- 7: $\mathcal{H} \leftarrow \mathcal{H}''$
- 8: **end if**
- 9: **end while**

as in Theorem 7. Assume further that there exists a UCTMC $(\hat{\mathcal{V}}, \hat{m}, \hat{M}, \hat{r}, \hat{\phi})$ such that $\hat{\mathcal{V}} = \{i_H \mid H \in \mathcal{H}\}$ and

$$i, t \models_{\mathcal{V}, m, M} \phi \iff i_H, t \models_{\hat{\mathcal{V}}, \hat{m}, \hat{M}} \hat{\phi}$$

for any $t \geq 0$, $h > 0$, $H \in \mathcal{H}$, $i \in H$ and \mathbf{X} -operator free CSL formula ϕ . Then, \mathcal{H} is a UCTMC lumpability and $(\hat{\mathcal{V}}, \hat{m}, \hat{M}, \hat{r}, \hat{\phi})$ the underlying lumped UCTMC.

Proof. See Appendix D. \square

C. Computation via Partition Refinement

We now present an algorithm for computing the coarsest UCTMC lumpability that refines a partition of states \mathcal{H} . Its steps are as follows.

- 1) Compute the coarsest ordinary lumpability \mathcal{H}' of the CTMC (\mathcal{S}, m) that refines the current partition \mathcal{H} .
- 2) Compute the coarsest ordinary lumpability \mathcal{H}'' of the CTMC (\mathcal{S}, M) that refines \mathcal{H}' .
- 3) If $\mathcal{H}'' = \mathcal{H}$, return \mathcal{H}'' ; else, set $\mathcal{H} := \mathcal{H}''$ and go to 1).

Obviously, if 1) does not refine \mathcal{H} and 2) does not refine \mathcal{H}' , then \mathcal{H} is a UCTMC lumpability of (\mathcal{S}, m, M) . It terminates because \mathcal{S} is finite. Moreover, it can be shown that the algorithm indeed computes the coarsest UCTMC partition because each refinement produces a partition which, itself, is still refined by the coarsest UCTMC lumpability.

The next result summarizes the above discussion.

Theorem 9. Given a UCTMC $(\mathcal{S}, m, M, r, \phi)$, let \mathcal{H} be a partition of \mathcal{S} . Then, the following can be shown.

- 1) The coarsest UCTMC lumpability refining \mathcal{H} exists and is computed by Algorithm 1.
- 2) The time and space complexity required for one while loop iteration does not exceed $\mathcal{O}(p \log(n))$, where $p := |\{(i, j) \mid m_{i,j} > 0 \text{ or } M_{i,j} > 0\}|$ and $n := |\mathcal{S}|$. The number of while loop iterations, instead, is at most n .

Proof. Since the trivial partition $\{\{i\} \mid i \in \mathcal{S}\}$ is an ordinary lumpability of both, (\mathcal{S}, m) and (\mathcal{S}, M) , we observe that the set of UCTMC lumpable partitions refining some partition \mathcal{H} of \mathcal{S} is not empty. Moreover, assume that \mathcal{H}' and \mathcal{H}'' refine \mathcal{H} and are ordinary lumpable partitions of both (\mathcal{S}, m) and

(\mathcal{S}, M) . Then, with the asterisk denoting the transitive closure of relations and with equivalence relations \sim' and \sim'' given by $\mathcal{H}' = \mathcal{S}/\sim'$ and $\mathcal{H}'' = \mathcal{S}/\sim''$, respectively, it holds that $\mathcal{S}/(\sim' \cup \sim'')^*$ refines \mathcal{H} and is an ordinary lumpable partition of both (\mathcal{S}, m) and (\mathcal{S}, M) . This foregoing discussion allows us to conclude that there exists a coarsest UCTMC lumpability refining \mathcal{H} . The complexity estimation of CTMC lumpability algorithms like [18], [19] implies the complexity estimation of Algorithm 1. As for the correctness, let \mathcal{H}^ω denote the coarsest UCTMC lumpability refining the given initial partition \mathcal{H} . Then, for any partition \mathcal{H}' of \mathcal{S} , the following observations hold true:

- If \mathcal{H}' refines \mathcal{H}^ω , then \mathcal{H}^ω refines the coarsest ordinary lumpability of (\mathcal{S}, m) that refines \mathcal{H}' .
- If \mathcal{H}' refines \mathcal{H}^ω , then \mathcal{H}^ω refines the coarsest ordinary lumpability of (\mathcal{S}, M) that refines \mathcal{H}' .

Since \mathcal{S} is finite, the sequence of partitions computed in Algorithm 1 will eventually reach a fixed point that cannot be refined. Noting that such a fixed point is an ordinary lumpability of (\mathcal{S}, m) and (\mathcal{S}, M) , the proof is complete. \square

The complexity statement follows from the fact that 1) and 2) can be processed by CTMC lumpability algorithms [18], [19]. We note that Algorithm 1 simplifies to the CTMC lumpability algorithm if applied to a deterministic UCTMC.

IV. EVALUATION

Using a prototype implemented in the tool ERODE [56], here we assess UCTMC lumpability in terms of its computational tractability and reduction power with respect to ordinary lumpability. After this, we study the advantages offered by UCTMC lumpability in terms of analysis speedup. In the latter case, UCTMCs were solved using MATLAB scripts automatically generated by our prototype. All material is available at <https://www.erode.eu/examples.html> for purposes of replicability. Reported experiments refer to a common laptop (a Mac with a 2.4 GHz Intel Core i5), where for our reduction we used 5 GB RAM and did not exploit parallelization.

Set-up. We considered the following benchmark CTMC models from the model repository of the PRISM probabilistic model checker [27]: a dependable cluster of workstations [57]; a protocol for wireless group communication [58], [59]; a model of the cell cycle control in eukaryotes [60], [61]; and a model of a cyclic server polling system from [62]. Notably, the first and the last models belong also to the QComp benchmark suite <https://qcomp.org/benchmarks/>. In addition, we created a CTMC PRISM model from our running example by setting $\lambda = \bar{\lambda} = 2$ and $\mu = \bar{\mu} = 4$.

We introduced uncertain relaxations of such CTMCs by replacing every transition rate with an interval of fixed length (arbitrarily fixed equal to 20% of the smallest transition rate in the model) centered at the original value. This relaxation is exemplified in the right inset, where the grayed-out transition rates are those of some given CTMC. The example given in the inset demonstrates that an ordinary lumpability of the original CTMC is not necessarily a UCTMC lumpability of the so-constructed UCTMC, whereas the converse follows from Lemma 1. Indeed, $\{\{A, B\}, \{C, D\}\}$

| Original model (CTMC) | | | | Lumpability | | | |
|---------------------------------------|------------|------------|-------------------|---------------|-----------------|-------|------|
| | | | | CTMC | | UCTMC | |
| N | p | n | $ \mathcal{H}_0 $ | Red.(s) | $ \mathcal{H} $ | RR | SR |
| WORKSTATION CLUSTER | | | | | | | |
| 128 | 2 908 192 | 597 012 | 4 | 2.21E+1 | 298 893 | 1.19 | 1.00 |
| 192 | 6 524 960 | 1 337 876 | 4 | 6.78E+1 | 669 517 | 1.19 | 1.00 |
| 256 | 11 583 520 | 2 373 652 | 4 | 1.55E+2 | 1 187 597 | 1.19 | 1.00 |
| 320 | 18 083 872 | 3 704 340 | 4 | 2.81E+2 | 1 853 133 | 1.27 | 1.00 |
| 384 | 26 026 016 | 5 329 940 | 4 | out of memory | | | |
| WIRELESS GROUP COMMUNICATION PROTOCOL | | | | | | | |
| 16 | 686 153 | 103 173 | 2 | 2.26E+0 | 4 846 | 1.19 | 1.00 |
| 24 | 3 183 849 | 453 125 | 2 | 1.34E+1 | 20 476 | 1.20 | 1.00 |
| 32 | 10 954 382 | 1 329 669 | 2 | 4.49E+1 | 58 906 | 1.22 | 1.00 |
| 40 | 22 871 849 | 3 101 445 | 2 | 1.35E+2 | 135 752 | 1.22 | 1.00 |
| 48 | 46 574 793 | 6 235 397 | 2 | out of memory | | | |
| CELL CYCLE CONTROL IN EUKARYOTES | | | | | | | |
| 2 | 18 342 | 4 666 | 3 | 1.76E-1 | 3 514 | 1.12 | 1.14 |
| 3 | 305 502 | 57 667 | 3 | 8.21E-1 | 40 667 | 1.19 | 1.18 |
| 4 | 2 742 012 | 431 101 | 3 | 6.45E+0 | 282 956 | 1.21 | 1.20 |
| 5 | 16 778 785 | 2 326 666 | 3 | 4.58E+1 | 1 424 935 | 2.00 | 1.20 |
| 6 | 78 768 799 | 9 960 861 | 3 | out of memory | | | |
| POLLING | | | | | | | |
| 14 | 2 695 168 | 344 064 | 2 | 3.64E+0 | 344 064 | 1.71 | 1.00 |
| 15 | 6 144 000 | 737 280 | 2 | 8.64E+0 | 737 280 | 1.77 | 1.00 |
| 16 | 13 893 632 | 1 572 864 | 2 | 1.69E+1 | 1 572 864 | 1.80 | 1.00 |
| 17 | 31 195 136 | 3 342 336 | 2 | 7.16E+1 | 3 342 336 | 1.91 | 1.00 |
| 18 | 69 599 232 | 7 077 888 | 2 | out of memory | | | |
| M/M/1/N QUEUE WITH 2 JOB CLASSES | | | | | | | |
| 15 | 131 068 | 65 535 | 2 | 7.48E-1 | 16 | 1.41 | 1.00 |
| 17 | 524 284 | 262 143 | 2 | 1.82E+0 | 18 | 1.75 | 1.00 |
| 19 | 2 097 148 | 1 048 575 | 2 | 6.86E+0 | 20 | 1.98 | 1.00 |
| 21 | 8 388 604 | 4 194 303 | 2 | 3.21E+1 | 22 | 2.15 | 1.00 |
| 23 | 33 554 428 | 16 777 215 | 2 | out of memory | | | |

TABLE I: Robustness of lumpability to uncertainty in the CTMC parameters. Columns *RR* and *SR* give the ratios of the runtimes and of the state-space sizes, respectively, between the UCTMC and the CTMC lumpability.

is an ordinary lumpability of the original CTMC, while it is not a UCTMC lumpability because $4.8 \neq 1.8 + 2.8$ and $5.2 \neq 2.2 + 3.2$. As a result, this setup indicates how much lumpability partitions of CTMCs are robust to the presence of uncertainties in transition rates.

Every benchmark can be associated with a parameter N which allows us to increase the state space size of the UCTMC. For instance, in our running example we let N denote the buffer size, i.e., we considered $M/M/1/N$ queues.

UCTMC lumpability vs CTMC lumpability. The results are provided in Table I. We report the number of transitions and states of the obtained CTMCs and UCTMCs in the second and third column, respectively, as a function of the scaling parameter N . The initial input partition of states, denoted by \mathcal{H}_0 , was induced by the original model specification by creating blocks of states characterized by the same atomic propositions used for the original model checking purposes. For the $M/M/1/N$ system, the initial partition isolated the empty state to be able to recover the utilization in the reduction. For CTMC lumpability we provide the sizes of the computed partitions ($|\mathcal{H}|$), which required between 0.1 to 281 seconds across all tests. For UCTMC lumpability we provide the ratios of the runtimes (*RR*) and of the obtained state-space sizes (*SR*) over the CTMC case.

The comparison of the runtimes of the lumpability algorithms provides an indication of the increased overhead for the reduction (which is proportional to the number of states in the worst case). In all our tests, UCTMC lumpability had, up to about a factor of two, the same runtime as the CTMC version. This is because in all models at most two iterations of our algorithm were necessary. The effectiveness of UCTMC lumpability can be evaluated by comparing the state-space size ratio. Notably, the CTMC and UCTMC reductions coincide in four families of models (workstation cluster, wireless group communication protocol, polling, and $M/M/1/N$ queue), while in the cell-cycle model UCTMC lumpability leads to finer (at most 20% more blocks) partitions than the CTMC counterpart. In particular, all families of models that admit CTMC reductions, also admit UCTMC reductions, while the polling family does not admit any reduction.

Analysis speedup due to UCTMC lumpability. Here we study the speedups that UCTMC lumpability bring to UCTMC analysis. We consider the same models as in Table II, but we omit the polling family because it does not admit reduction, and thus is not of interest here. The results are provided in Table II. The table has a structure similar to Table I with the difference that: column $|\mathcal{H}|$ gives the number of states in the reduced UCTMC, while $Red.(s)$ the time in seconds to obtain it. Finally, columns $Opt.(s)$ provide the time to solve the original and reduced UCTMCs according to Section II-A, while column $Opt.$ contains the computed optimal maximal value (which is always the same in the original and reduced models). In all cases we fix an arbitrary time horizon of 1.

PRISM models can be equipped with so-called *reward structures*.¹ In particular, our notion of running reward coincides with PRISM's *state rewards*. Therefore, when solving the UCTMCs according to Section II-A, we use (one of) the actual rewards specified in the models. For example, for the workstation cluster models we use a reward named *percent_op* related to the “percentage of operational workstations stations” in a state, while for the cell cycle control in eukaryotes we use one named *cdc14* related to the “abound of molecules CDC14”. In the case of queues, we considered instances of our queue example $M/M/1/N$ with capacity N , arrival rate $\lambda = \bar{\lambda} = 2$, and service rate bounds $\underline{\mu} = 3$ and $\bar{\mu} = 5$. In other words, we have no uncertainties on the arrival rates, and uncertainty on the service rates. Furthermore, we assume that the modeler is interested in achieving as utilization $k = \lceil N/2 \rceil$. The latter can be realized by the running reward $r_s = |k - (|s|_u + |s|_b)|$, where $|s|_b$ and $|s|_u$ denote, respectively, the number of background and user jobs in state s .

The initial partition of states, denoted in the table by \mathcal{H}_0 , was induced by the rewards assigned to each state, guaranteeing that lumped states have same reward value. The comparison of the analysis runtimes for the original and reduced models provides an indication of the analysis speedup offered by UCTMC lumpability. For each family of models we provide instances where the analysis succeeded in both the original and reduced models, showing speedups of up to 3 orders of magnitude, and

| Original model (as UCTMC) | | | | UCTMC lumpability | | | | |
|---------------------------------------------|--------|--------|-----------|-------------------|-----------------|-----------|-----------|---------|
| N | p | n | $Opt.(s)$ | $ \mathcal{H}_0 $ | $ \mathcal{H} $ | $Red.(s)$ | $Opt.(s)$ | $Opt.$ |
| WORKSTATION CLUSTER (max) | | | | | | | | |
| 8 | 12832 | 2772 | 2.41E+1 | 17 | 1017 | 2.80E-1 | 3.75E+0 | 9.99E+1 |
| 16 | 48160 | 10132 | 4.12E+2 | 33 | 3621 | 4.52E-1 | 4.62E+1 | 9.99E+1 |
| 32 | 186400 | 38676 | timeout | 65 | 13629 | 9.00E-1 | 1.01E+3 | 9.99E+1 |
| WIRELESS GROUP COMMUNICATION PROTOCOL (max) | | | | | | | | |
| 4 | 5369 | 1125 | 1.62E+0 | 18 | 83 | 4.10E-2 | 5.10E-2 | 3.80E-2 |
| 8 | 54953 | 9477 | 1.32E+2 | 34 | 520 | 1.98E-1 | 4.60E-1 | 1.90E-2 |
| 16 | 686153 | 103173 | timeout | 66 | 4894 | 2.01E+0 | 2.72E+1 | 9.00E-3 |
| CELL CYCLE CONTROL IN EUKARYOTES (max) | | | | | | | | |
| 1 | 300 | 142 | 1.01E-1 | 3 | 90 | 2.90E-2 | 4.90E-2 | 2.00E+0 |
| 2 | 18342 | 4666 | 2.60E+1 | 5 | 3714 | 1.10E-1 | 1.54E+1 | 4.00E+0 |
| 3 | 305502 | 57667 | timeout | 7 | 46460 | 1.49E+0 | timeout | timeout |
| M/M/1/N QUEUE WITH 2 JOB CLASSES (min) | | | | | | | | |
| 11 | 8188 | 4095 | 2.82E+1 | 7 | 12 | 9.70E-2 | 4.00E-2 | 4.78E+0 |
| 13 | 32764 | 16383 | 5.62E+2 | 8 | 14 | 2.12E-1 | 4.00E-2 | 5.77E+0 |
| 15 | 131068 | 65535 | timeout | 9 | 16 | 8.15E-1 | 4.00E-2 | 6.76E+0 |

TABLE II: Analysis speedup offered by UCTMC lumpability. Columns $Opt.(s)$ provide the time to solve the original and reduced UCTMCs according to Section II-A, while column $Opt.$ provides the obtained optimal value (maximal for the first three families, minimal for the queue family), which always coincides in the original and reduced models.

one instance where the analysis failed for the original model due to an arbitrarily chosen timeout of 90 minutes.

V. CONCLUSION

Uncertain continuous-time Markov chains (UCTMCs) generalize continuous-time Markov chains (CTMCs) by allowing transition rates to non-deterministically take values within given bounded intervals. We presented UCTMC lumpability as a conservative generalization of ordinary lumpability to UCTMCs. It enjoys an efficient algorithm for the computation of the largest UCTMC lumpability that refines a given initial partition of UCTMC states. Similarly to CTMC lumpability that characterizes the preservation of sums of probability distributions, UCTMC lumpability characterizes the preservation of value functions and temporal logical formulae. The applicability of UCTMC lumpability has been established by presenting substantial reductions in large benchmark models.

APPENDIX

A. Proof of Theorem 3

Theorem 3. Pontryagin's principle [46] ensures that any optimal transition rate matrix Q^* of the UCTMC satisfies the equations

$$\begin{aligned}
 \partial_t \pi(t) &= (\partial_p H)(\pi(t), Q^*(t), p(t)) = \pi(t)^T Q^*(t) \\
 \partial_t p(t) &= -(\partial_\pi H)(\pi(t), Q^*(t), p(t)) = -Q^*(t)p(t) + r(t) \\
 Q^*(t) &= \arg \max_{Q \in [m; M]} H(\pi(t), Q, p(t)),
 \end{aligned}$$

where $H(\pi, Q, p) = \pi^T(Qp - r)$ is the so-called Hamiltonian for the running cost $\sum_i r_i \pi_i$. To facilitate the proof, we considered a perturbed Hamiltonian given by $H_\eta(\pi, Q, p) = H(\pi, Q, p) - \eta \sum_{j \neq i} q_{i,j}^2$, where $\eta \geq 0$ is an arbitrarily small regularization constant. Standard continuity results of ODE

¹ www.prismmodelchecker.org/manual/ThePRISMLanguage/CostsAndRewards

systems [63, Chapter 3] ensure that $\lim_{\eta \rightarrow 0} V_\eta^* = V^*$. A direct computation yields then

$$\begin{aligned} \partial_t p_i - r_i &= - \sum_j q_{i,j} p_j = - \sum_{j \neq i} q_{i,j} p_j - q_{i,i} p_i \\ &= - \sum_{j \neq i} q_{i,j} p_j + \sum_{j \neq i} q_{i,j} p_i \\ &= \sum_{j \neq i} q_{i,j} (p_i - p_j) = - \sum_{j \neq i} (p_j - p_i) q_{i,j} \end{aligned}$$

Moreover, if $\eta > 0$, the $\arg \max$ has a unique solution. Indeed, using the above expression of Qp and the definition of H_η , we first infer

$$Q^* = \max_{Q \in [m; M]} \left[\sum_i \pi_i \left(\sum_{j \neq i} (p_j - p_i) q_{i,j} - r_i \right) - \eta \sum_{j \neq i} q_{i,j}^2 \right]$$

For fixed $i \neq j$, this in turn implies

$$q_{i,j}^* = \max_{q_{i,j} \in [m_{i,j}; M_{i,j}]} \left[\pi_i (p_j - p_i) q_{i,j} - \eta q_{i,j}^2 \right]$$

Since the convex function $q_{i,j} \mapsto \pi_i (p_j - p_i) q_{i,j} - \eta q_{i,j}^2$ attains, over the entire \mathbb{R} , its maximum at $q_{i,j}^* = \frac{\pi_i}{2\eta} (p_j - p_i)$, we infer

$$q_{i,j}^*(t) = \begin{cases} m_{i,j} & , q_{i,j}^* < m_{i,j} \\ M_{i,j} & , q_{i,j}^* > M_{i,j} \\ q_{i,j}^* & , q_{i,j}^* \in [m_{i,j}; M_{i,j}] \end{cases}$$

Taking the limit $\eta \rightarrow 0$ yields then the claim. \square

B. Discrete-time Approximation of UCTMCs

In this appendix we present a DTMDP-approximation of UCTMCs. Apart from relating UCTMCs to DTMDPs, the discrete-time approximation can be used to relate UCTMC lumpability to alternating probabilistic bisimulation [32].

Instrumental to the DTMDP approximation is an alternative CTMDP encoding which uses finite action spaces, at the expense of probabilistic (instead of deterministic) policies. Before giving this encoding, we convey the main underlying idea on an illustrative example. Let us assume that we are given a CTMDP that can move from state i only into state j and that the corresponding time-dependent deterministic transition rate function is $q_{i,j}(t, a(t)) = a_{i,j}(t)$, where $m_{i,j} = 1$, $M_{i,j} = 2$ and $a_{i,j}(t) = 2 - e^{-t}$. With this, we first replace the continuous interval $[1; 2]$ with the discrete action set $\{\mathbf{m}_{i,j}, \mathbf{M}_{i,j}\}$, where the symbols $\mathbf{m}_{i,j}$ and $\mathbf{M}_{i,j}$ represent the boundary values $m_{i,j} = 1$ and $M_{i,j} = 2$, respectively. Then, the idea is to choose suitable probability functions $\mu_{\mathbf{m}_{i,j}}(t)$ and $\mu_{\mathbf{M}_{i,j}}(t)$ such that the average transition rate from state i into state j at time t , given by $1\mu_{\mathbf{m}_{i,j}}(t) + 2\mu_{\mathbf{M}_{i,j}}(t)$, is identical to $a_{i,j}(t)$. It can be easily verified that $\mu_{\mathbf{m}_{i,j}}(t) = e^{-t}$ and $\mu_{\mathbf{M}_{i,j}}(t) = 1 - e^{-t}$ induce $a_{i,j}$. Following [8, Section 2.2], the foregoing discussion generalizes as follows.

Proposition 1. *For a given UCTMC (S, m, M, r, ϕ) , consider the CTMDP $(S, \mathcal{A}', \mathcal{M}', r, \phi)$ where an action in state i at time t is taken randomly, is denoted by $a_i(t)$, and is a row vector such that each row entry $a_{i,j}(t) \in \{\mathbf{m}_{i,j}, \mathbf{M}_{i,j}\}$ determines the transition rate from i into j at time t accordingly. Formally, we have the following.*

- The set of actions in state $i \in S$ is given by $\mathcal{A}'(i) = \prod_{j \neq i} \{\mathbf{m}_{i,j}, \mathbf{M}_{i,j}\}$.
- The transition rate of from i into j at time t under action $a_i \in \mathcal{A}'(i)$ is $q_{i,j}(t, a_i) = v(a_{i,j})$, where $v(a_{i,j}) = m_{i,j}$ if $a_{i,j} = \mathbf{m}_{i,j}$ and $v(a_{i,j}) = M_{i,j}$ when $a_{i,j} = \mathbf{M}_{i,j}$.
- The set of policies, \mathcal{M}' , constitutes non-negative uniformly piecewise analytic functions μ satisfying $\sum_{a_i \in \mathcal{A}'(i)} \mu_{a_i}(t) = 1$ for all $i \in S$ and $t \geq 0$. In particular, with $\mathcal{D}(X)$ denoting the set of probability measures on a set X , it holds that \mathcal{M}' is a proper subset of $[0; \infty) \rightarrow \prod_{i \in S} \mathcal{D}(\mathcal{A}'(i))$.

Then, the policy sets \mathcal{M} and \mathcal{M}' , where \mathcal{M} refers to the policy set given in Proposition 4, induce the same set of time-inhomogeneous CTMCs.

Proof. Given some $i \in S$, we first note that $\mathcal{A}(i)$ describes a hypercube. Noting that $\mathcal{A}'(i)$ uniquely identifies the edges of $\mathcal{A}(i)$ via the function v , it thus suffices to show that any $z_i \in \mathcal{A}(i)$ can be expressed as a convex combination of the edges of $\mathcal{A}(i)$. To see this, assume that we are given some $z_i \in \mathcal{A}(i)$ and consider the system of linear equations

$$\begin{aligned} \sum_{\substack{a_i \in \mathcal{A}'(i), \\ a_{i,j} = \mathbf{M}_{i,j}}} \nu(a_i) &= \frac{z_{i,j} - m_{i,j}}{M_{i,j} - m_{i,j}}, \quad j \neq i \\ \sum_{\substack{a_i \in \mathcal{A}'(i), \\ a_{i,j} = \mathbf{m}_{i,j}}} \nu(a_i) &= 1 - \frac{z_{i,j} - m_{i,j}}{M_{i,j} - m_{i,j}}, \quad j \neq i. \end{aligned}$$

Noting that

$$z_{i,j} = M_{i,j} \frac{z_{i,j} - m_{i,j}}{M_{i,j} - m_{i,j}} + m_{i,j} \left(1 - \frac{z_{i,j} - m_{i,j}}{M_{i,j} - m_{i,j}} \right),$$

we observe that a solution of the linear system of equations yields the claim. This, in turn, can be identified as

$$\nu(a_i) = \prod_{j \in J_M} \frac{z_{i,j} - m_{i,j}}{M_{i,j} - m_{i,j}} \cdot \prod_{j \in J_m} \left(1 - \frac{z_{i,j} - m_{i,j}}{M_{i,j} - m_{i,j}} \right),$$

where $j \in J_M$ if $a_{i,j} = \mathbf{M}_{i,j}$ and $j \in J_m$ if $a_{i,j} = \mathbf{m}_{i,j}$. (Following standard notation, the product over an empty set is defined to be one). \square

For a policy $\mu \in \mathcal{M}'$, the Kolmogorov equations $\partial_t \pi(t)^T = \pi(t)^T Q(t, \mu(t))$ describing the transient probabilities of the time-inhomogeneous CTMC can be solved numerically by invoking the Euler method [64], a classic approach for the numeric solution of systems of differential equations. Specifically, by discretizing time into the set $\{0, h, 2h, \dots\}$, the probability distribution at time kh , denoted by $\pi(kh)$, is approximated by $\pi[k]$, where

$$\pi[k+1]^T := \pi[k]^T (I + hQ(kh, \mu(kh))),$$

$\pi[0] := \pi(0)$ and I is the identity matrix. Additionally to the known fact that the approximation error is $\mathcal{O}(h)$, we note that the Euler method defines a time-inhomogeneous DTMC. Indeed, $I + hQ(kh, \mu(kh))$ describes the transition probability matrix of the embedded time-inhomogeneous DTMC.

The next result relates UCTMCs to DTMDPs.

Proposition 2. Given UCTMC $(\mathcal{S}, m, Mr, \phi)$, set

$$\Lambda = \max_{i \in \mathcal{S}} \left(\sum_{j \neq i} M_{i,j} + \sum_{j \neq i} M_{j,i} \right)$$

and fix $h \leq 1/\Lambda$. Then, $I + hQ(kh, \mu(kh))$ is a stochastic matrix for all $\mu \in \mathcal{M}'$ and $k \geq 0$. With this, consider the DTMDP $(\mathcal{S}, \mathcal{A}', \mathcal{M}'_h, r, \phi)$ given as:

- The states are \mathcal{S} , while the actions in state $i \in \mathcal{S}$ are given by $\mathcal{A}'(i) = \prod_{j \neq i} \{\mathbf{m}_{i,j}; \mathbf{M}_{i,j}\}$.
- The transition probability from state i into state j at step $k \geq 0$ for $a_i \in \mathcal{A}'(i)$ is

$$p_{i,j}(k, a_i) = \begin{cases} hv(a_{i,j}) & , j \neq i \\ 1 - h \sum_{j \neq i} v(a_{i,j}) & , j = i \end{cases}$$

- The set of policies is $\mathcal{M}'_h = \{\nu \mid \nu : \mathbb{N}_0 \rightarrow \prod_{i \in \mathcal{S}} \mathcal{D}(\mathcal{A}'(i))\}$. In particular, for a given policy ν , the transition probability from state i into state j at step $k \geq 0$ is given by $p_{i,j}(k, \nu(k)) = \sum_{a_i \in \mathcal{A}'(i)} \nu_{a_i}(k) p_{i,j}(k, a_i)$.

Then, for any time $T > 0$ and policy $a \in \mathcal{M}$ such that the modulus of the derivative of each $a_{i,j}$ is bounded by $\lambda \geq 0$ almost everywhere, there exists a policy $\nu \in \mathcal{M}'_h$ such that

$$\max_{i \in \mathcal{S}} |\pi_i[k] - \pi_i(T)| \leq h \left[\frac{3\Lambda}{2} + \frac{\lambda}{\Lambda} \max_{i \in \mathcal{S}} \deg(i) \right] (e^{\Lambda T} - 1),$$

where $\deg(i) = |\{j \neq i \mid m_{i,j} < M_{i,j}\}| + |\{j \neq i \mid m_{j,i} < M_{j,i}\}|$ are the incoming and outgoing non-deterministic transitions of i , while $k \geq 0$ minimizes $|kh - T|$.

Proof. We first note that $\max_{i \in \mathcal{S}} |\pi_i[k] - \pi_i(T)| = \|\pi[k] - \pi(T)\|_\infty$, where $\|\cdot\|_\infty$ denotes the maximum norm. Second, we recall that the Lipschitz constant with respect to the maximum norm of a linear ODE system $\partial_t x^T = x^T B$ is given by the matrix maximum norm $\|B^T\|_\infty = \max_{i \in \mathcal{S}} \sum_{j \in \mathcal{S}} |b_{j,i}|$. With this, the error estimation follows from $\|\pi(t)\|_\infty \leq 1$ and by evaluating the constants L, K and Z in the error term provided in [64, Eq. 1.14]. More specifically, it holds that $L \leq \Lambda$, $Z \leq \Lambda$ and $K = \lambda \cdot \max_{i \in \mathcal{S}} \deg(i)$. Additionally to that, we note that [64, Eq. 1.13] carries over to the multi-dimensional case. Indeed, while different coordinates i may require to use different values $\theta_i \in [0; 1]$ in [64, Eq. 1.13], thanks to the fact that we use the maximum norm, it suffices to consider the coordinate with the largest error $d_{n,i}$. \square

While theoretically appealing, estimating the approximation error via Proposition 2 is computationally challenging. This is because useful estimations require h to be at most of order $\frac{1}{|\mathcal{S}|}/e^{\Lambda T}$. We next relate UCTMC lumpability to alternating probabilistic bisimulation of DTMDPs [32].

Proposition 3. Fix a UCTMC $(\mathcal{S}, m, M, r, \phi)$, an equivalence relation $\mathfrak{R} \subseteq \mathcal{S} \times \mathcal{S}$ and let $\mathcal{H} = \mathcal{S}/\mathfrak{R}$. Then \mathcal{H} is a UCTMC lumpability of $(\mathcal{S}, m, M, r, \phi)$ if and only if \mathfrak{R} is an alternating probabilistic bisimulation of the DTMDP from Proposition 2.

Proof. By definition, \mathcal{H} is a UCTMC lumpability if and only if for any $H \in \mathcal{H}$ and $i_1, i_2 \in H$, it holds that $\sum_{j \in H'} m_{i_1,j} = \sum_{j \in H'} m_{i_2,j}$ and $\sum_{j \in H'} M_{i_1,j} = \sum_{j \in H'} M_{i_2,j}$ for all $H' \in \mathcal{H}$. This holds true if and only if for any action $a_{i_1} \in \mathcal{A}(i_1)$ there exists an action $a_{i_2} \in \mathcal{A}(i_2)$ (and vice versa) such

that, for all $H' \in \mathcal{H}$, it holds that $\sum_{j \in H'} q_{i_1,j}(t, a_{i_1}) = \sum_{j \in H'} q_{i_2,j}(t, a_{i_2})$. Thanks to the proof of Proposition 1, this holds true if and only if for any distribution $\mu_{i_1} \in \mathcal{D}(\mathcal{A}'(i_1))$ there exists a distribution $\mu_{i_2} \in \mathcal{D}(\mathcal{A}'(i_2))$ (and vice versa) such that, for all $H' \in \mathcal{H}$, one has

$$\sum_{j \in H'} \sum_{a \in \mathcal{A}'(i_1)} \mu_{i_1}(a) \cdot q_{i_1,j}(t, a) = \sum_{j \in H'} \sum_{a \in \mathcal{A}'(i_2)} \mu_{i_2}(a) \cdot q_{i_2,j}(t, a)$$

Taking into account the definition of the DTMDP in Proposition 2, the foregoing statement holds true if and only if for any distribution $\nu_{i_1} \in \mathcal{D}(\mathcal{A}'(i_1))$ there exists a distribution $\nu_{i_2} \in \mathcal{D}(\mathcal{A}'(i_2))$ (and vice versa) such that, for all $H' \in \mathcal{H}$, it holds that

$$\sum_{j \in H'} \sum_{a \in \mathcal{A}'(i_1)} \nu_{i_1}(a) \cdot p_{i_1,j}(k, a) = \sum_{j \in H'} \sum_{a \in \mathcal{A}'(i_2)} \nu_{i_2}(a) \cdot p_{i_2,j}(k, a)$$

Observing that the existence of ν_{i_1} and ν_{i_2} ensures that \mathfrak{R} is an alternating probabilistic bisimulation of the DTMDP yields the claim (indeed, i_1, i_2, ν_{i_1} and ν_{i_2} are playing the role of s, t, ρ_s and ρ_t in [32, Definition 3], respectively). \square

C. Proofs of Theorems 5 and 6

For the proofs of Theorem 5-6, the following auxiliary notions will be needed.

Definition 9. Let $\mathcal{U} = (\mathcal{S}, m, M, r, \phi)$ be some UCTMC and \mathcal{H} a partition of \mathcal{S} . Moreover, let $\hat{\mathcal{U}} = (\hat{\mathcal{S}}, \hat{m}, \hat{M}, \hat{r}, \hat{\phi})$ be a UCTMC with $\hat{\mathcal{S}} = \{i_H \mid H \in \mathcal{H}\}$.

- We write $\mathcal{U} \subseteq \hat{\mathcal{U}}$ if for every Q of \mathcal{U} there exists a \hat{Q} of $\hat{\mathcal{U}}$ such that

$$\hat{\pi}_{i_H}^{\hat{Q}}(t) = \sum_{i \in H} \pi_i^Q(t), \quad \forall H \in \mathcal{H}, t > 0 \quad (4)$$

provided that (4) is valid at $t = 0$.

- We write $\hat{\mathcal{U}} \subseteq \mathcal{U}$ if for every \hat{Q} of $\hat{\mathcal{U}}$ there exists a Q of \mathcal{U} such that (4) holds.

Further, $\mathcal{R}_{\mathcal{U}}(H, T, \pi[0]) = \{\sum_{i \in H} \pi_i^Q(T) \mid Q \in [m; M]\}$ denotes for $H \in \mathcal{H}$ the H -reachability probabilities at $T \geq 0$.

Proof of Theorem 5. It suffices to consider H -reachability probabilities. Specifically, note that $\mathcal{U} \subseteq \hat{\mathcal{U}}$ implies $V_{\hat{\mathcal{U}}}^{\inf}(i_H, T) \leq V_{\mathcal{U}}^{\inf}(i, T)$ and $V_{\hat{\mathcal{U}}}^{\sup}(i_H, T) \geq V_{\mathcal{U}}^{\sup}(i, T)$; likewise, $\hat{\mathcal{U}} \subseteq \mathcal{U}$ implies $V_{\hat{\mathcal{U}}}^{\inf}(i_H, T) \geq V_{\mathcal{U}}^{\inf}(i, T)$ and $V_{\hat{\mathcal{U}}}^{\sup}(i_H, T) \leq V_{\mathcal{U}}^{\sup}(i, T)$ for all $H \in \mathcal{H}, i \in H$ and $T \geq 0$. With this, the claim follows from Proposition 4 given next. \square

Proposition 4. Fix a UCTMC $\mathcal{U} = (\mathcal{S}, m, M, r, \phi)$ and assume that \mathcal{H} is an ordinary lumpability of CTMCs (\mathcal{S}, m) and (\mathcal{S}, M) . Then, $\mathcal{U} \subseteq \hat{\mathcal{U}}$ and $\hat{\mathcal{U}} \subseteq \mathcal{U}$.

Proof. Note that \mathcal{H} is an ordinary lumpability of CTMCs (\mathcal{S}, m) and (\mathcal{S}, M) . Let \mathcal{G} be a partition of $\{(i, j) \in \mathcal{S}^2 \mid i \neq j\}$ such that $(i_k, j_k), (i_l, j_l) \in G$ for some $G \in \mathcal{G}$ if and only if there is $H, H' \in \mathcal{H}$ such that $i_k, i_l \in H$ and $j_k, j_l \in H'$.

- $G_{H \rightarrow H'} \in \mathcal{G}$ is the unique block of edges originating in H and ending in H' .
- $G_{H \rightarrow H'} \in \mathcal{G}$ is called invariant if $H = H'$; the set of invariant blocks is denoted by \mathcal{G}_i .

For arbitrary $G \in \mathcal{G}$ and $(i_k, j_k) \in G$, let $f_i^{i_k, j_k}$ denote the change in π_i due to q_{i_k, j_k} . More formally, if $f(\pi) := \pi^T Q$ for all $\pi \in \mathbb{R}^S$, then $f_i^{i_k, j_k} := \partial_{q_{i_k, j_k}} f_i$. One can note that

$$f_i^{i_k, j_k} = \begin{cases} -\pi_{i_k} & , i = i_k \\ \pi_{i_k} & , i = j_k \end{cases}$$

For an arbitrary $H \in \mathcal{H}$, we note that

$$\begin{aligned} \partial_t \left(\sum_{i \in H} \pi_i(t) \right) &= \sum_{i \in H} \sum_{G \in \mathcal{G}} \sum_{(i_k, j_k) \in G} q_{i_k, j_k}(t) f_i^{i_k, j_k}(\pi(t)) \\ &= \sum_{G \in \mathcal{G}} \sum_{(i_k, j_k) \in G} q_{i_k, j_k}(t) \sum_{i \in H} f_i^{i_k, j_k}(\pi(t)) \\ &= \sum_{G \in \mathcal{G}} q_{G, H}(t) \sum_{(i_l, j_l) \in G} \sum_{i \in H} f_i^{i_l, j_l}(\pi(t)), \end{aligned}$$

provided that $q_{G, H}$ satisfies

$$\begin{aligned} q_{G, H}(t) \sum_{(i_l, j_l) \in G} \sum_{i \in H} f_i^{i_l, j_l}(\pi(t)) \\ = \sum_{(i_k, j_k) \in G} q_{i_k, j_k}(t) \sum_{i \in H} f_i^{i_k, j_k}(\pi(t)) \end{aligned}$$

When $\sum_{(i_l, j_l) \in G} \sum_{i \in H} f_i^{i_l, j_l}(\pi(t)) \neq 0$, it must obviously hold

$$q_{G, H}(t) = \sum_{(i_k, j_k) \in G} q_{i_k, j_k}(t) \frac{\sum_{i \in H} f_i^{i_k, j_k}(\pi(t))}{\sum_{(i_l, j_l) \in G} \sum_{i \in H} f_i^{i_l, j_l}(\pi(t))} \quad (5)$$

If the denominator is zero, instead, the value $q_{G, H}(t)$ can be chosen arbitrarily. We next show that, setting $q_{G, H} := q_G$, where $G = G_{H_1 \rightarrow H_2}$ for some $H_1, H_2 \in \mathcal{H}$, does the job:

$$q_G(t) = \begin{cases} \text{any value in } [\hat{m}_{i_{H_1}, i_{H_2}}; \hat{M}_{i_{H_1}, i_{H_2}}] & , G \in \mathcal{G}_i \\ \text{any value in } [\hat{m}_{i_{H_1}, i_{H_2}}; \hat{M}_{i_{H_1}, i_{H_2}}] & , \sum_{(i_l, j_l) \in G} \pi_{i_l}(t) = 0 \\ \sum_{(i_k, j_k) \in G} q_{i_k, j_k}(t) \frac{\pi_{i_k}(t)}{\sum_{(i_l, j_l) \in G} \pi_{i_l}(t)} & , \text{otherwise} \end{cases}$$

Key to this is to prove that the value of the fraction term in (5) is, whenever defined, invariant with respect to $H \in \mathcal{H}$. To see this, fix an arbitrary $G \in \mathcal{G}$, $(i_k, j_k) \in G$ and $H, H' \in \mathcal{H}$ such that $H \neq H'$. We consider the following case distinction.

- $i_k \in H \wedge j_k \in H'$: By the choice of \mathcal{G} , it holds that $i_l \in H \wedge j_l \in H'$ for all $(i_l, j_l) \in G$. Hence

$$\frac{\sum_{i \in H} f_i^{i_k, j_k}(\pi(t))}{\sum_{(i_l, j_l) \in G} \sum_{i \in H} f_i^{i_l, j_l}(\pi(t))} = \frac{-\pi_{i_k}(t)}{-\sum_{(i_l, j_l) \in G} \pi_{i_l}(t)}$$

and

$$\frac{\sum_{i \in H'} f_i^{i_k, j_k}(\pi(t))}{\sum_{(i_l, j_l) \in G} \sum_{i \in H'} f_i^{i_l, j_l}(\pi(t))} = \frac{\pi_{i_k}(t)}{\sum_{(i_l, j_l) \in G} \pi_{i_l}(t)},$$

meaning that both fraction terms are either identical or undefined. In the latter case, neither H nor H' constrains the value of q_G .

- $i_k \in H \wedge j_k \in H$: By the choice of \mathcal{G} , it holds that $i_l \in H \wedge j_l \in H$ for all $(i_l, j_l) \in G$. Hence

$$\sum_{(i_l, j_l) \in G} \sum_{i \in H} f_i^{i_l, j_l}(\pi(t)) = \sum_{(i_l, j_l) \in G} (\pi_{i_l}(t) - \pi_{i_l}(t)) = 0$$

for all $t \geq 0$, meaning that H does not constrain the value of q_G (note that in this case G is invariant).

- $i_k \notin H \wedge j_k \notin H$: Let $H_1, H_2 \in \mathcal{H}$ be such that $i_k \in H_1$ and $j_k \in H_2$. By the choice of \mathcal{G} , it holds that $i_l \in H_1 \wedge j_l \in H_2$ for all $(i_l, j_l) \in G$. Hence $\sum_{(i_l, j_l) \in G} \sum_{i \in H} f_i^{i_l, j_l}(\pi(t)) = 0$ for all $t \geq 0$, meaning that H does not constrain the value of q_G .

For an arbitrary $H \in \mathcal{H}$, the above discussion implies that

$$\begin{aligned} \partial_t \left(\sum_{i \in H} \pi_i(t) \right) &= \sum_{i \in H} \sum_{G \in \mathcal{G}} \sum_{(i_k, j_k) \in G} q_{i_k, j_k}(t) f_i^{i_k, j_k}(\pi(t)) \\ &= \sum_{i \in H} \sum_{G \in \mathcal{G}} q_G(t) \sum_{(i_k, j_k) \in G} f_i^{i_k, j_k}(\pi(t)) \\ &= \sum_{H' \neq H} q_{G_{H \rightarrow H'}}(t) \sum_{i \in H} \sum_{(i_l, j_l) \in G_{H \rightarrow H'}} f_i^{i_l, j_l}(\pi(t)) \\ &\quad + \sum_{H' \neq H} q_{G_{H' \rightarrow H}}(t) \sum_{i \in H} \sum_{(i_l, j_l) \in G_{H' \rightarrow H}} f_i^{i_l, j_l}(\pi(t)) \\ &= - \sum_{H' \neq H} q_{G_{H \rightarrow H'}}(t) \sum_{i \in H} |H'| \pi_i(t) \\ &\quad + \sum_{H' \neq H} q_{G_{H' \rightarrow H}}(t) \sum_{i \in H} \sum_{j \in H'} \pi_j(t) \\ &= - \sum_{H' \neq H} q_{G_{H \rightarrow H'}}(t) |H'| \left(\sum_{i \in H} \pi_i(t) \right) \\ &\quad + \sum_{H' \neq H} q_{G_{H' \rightarrow H}}(t) |H| \left(\sum_{j \in H'} \pi_j(t) \right) \quad (6) \end{aligned}$$

We next show that $\hat{q}_{i_H, i_{H'}}(t) := q_{G_{H \rightarrow H'}}(t) |H'| \in [\hat{m}_{i_H, i_{H'}}; \hat{M}_{i_H, i_{H'}}]$ for all $t \geq 0$ and $H, H' \in \mathcal{H}$ with $H \neq H'$. To this end, we note that

$$\begin{aligned} q_{G_{H \rightarrow H'}}(t) |H'| &= \sum_{(i_k, j_k) \in G_{H \rightarrow H'}} q_{i_k, j_k}(t) \frac{|H'| \sum_{i \in H} f_i^{i_k, j_k}(\pi(t))}{\sum_{i \in H} \sum_{(i_l, j_l) \in G_{H \rightarrow H'}} f_i^{i_l, j_l}(\pi(t))} \\ &= \sum_{(i_k, j_k) \in G_{H \rightarrow H'}} q_{i_k, j_k}(t) \frac{-|H'| \pi_{i_k}(t)}{-\sum_{i \in H} |H'| \pi_i(t)} \\ &= \sum_{(i_k, j_k) \in G_{H \rightarrow H'}} q_{i_k, j_k}(t) \frac{\pi_{i_k}(t)}{\sum_{i \in H} \pi_i(t)} \\ &= \sum_{i_k \in H} \frac{\pi_{i_k}(t)}{\sum_{i \in H} \pi_i(t)} \sum_{j_k \in H'} q_{i_k, j_k}(t) \end{aligned}$$

Since \mathcal{H} is an ordinary lumpability of CTMCs (\mathcal{S}, m) and (\mathcal{S}, M) , for all $i_k, i'_k \in H$ it holds that

$$\sum_{j_k \in H'} m_{i_k, j_k} = \sum_{j_k \in H'} m_{i'_k, j_k}, \quad \sum_{j_k \in H'} M_{i_k, j_k} = \sum_{j_k \in H'} M_{i'_k, j_k}.$$

With this, for all $t \geq 0$ it holds that $\hat{q}_{i_H, i_{H'}}(t) \in [\hat{m}_{i_H, i_{H'}}; \hat{M}_{i_H, i_{H'}}]$ which, together with (6), implies $\mathcal{U} \subseteq \hat{\mathcal{U}}$.

To show the converse relation, let us assume that we are given transition rate functions $(\hat{q}_{i_H, i_{H'}})_{H, H'}$. For $H, H' \in \mathcal{H}$ with $H \neq H'$ and $(i_k, j_k) \in G_{H \rightarrow H'}$, we set $q_{i_k, j_k}(t)$ to

$$m_{i_k, j_k} + \frac{M_{i_k, j_k} - m_{i_k, j_k}}{\sum_{j_l \in H'} (M_{i_k, j_l} - m_{i_k, j_l})} (\hat{q}_{i_H, i_{H'}}(t) - \hat{m}_{i_H, i_{H'}})$$

Since \mathcal{H} is an ordinary lumpability of CTMCs (\mathcal{S}, m) and (\mathcal{S}, M) , it holds that

$$\sum_{j_l \in H'} (M_{i_k, j_l} - m_{i_k, j_l}) = \sum_{j_k \in H'} (M_{i_k, j_k} - m_{i_k, j_k}).$$

Hence, $\sum_{j_k \in H'} q_{i_k, j_k}(t) = \hat{q}_{i_H, i_{H'}}(t)$ and $q_{i_k, j_k}(t) \in [m_{i_k, j_k}; M_{i_k, j_k}]$ for all $t \geq 0$. By choosing the so-constructed $(q_{i, j})_{i, j}$ and repeating the argumentation from the first part of the proof, we observe that $\sum_{i \in H} \pi_i(t) = \hat{\pi}_{i_H}(t)$ for all $H \in \mathcal{H}$ and $t \geq 0$. This yields the converse $\hat{\mathcal{U}} \subseteq \mathcal{U}$. \square

Proof of Theorem 6. Set $r = 0$ and $\phi_H = \mathbb{1}_H$. With this, we obtain $V_{\mathcal{U}}^{\Theta}(i', T) = \Theta\{\mathcal{R}_{\mathcal{U}}(H, T, \mathbb{1}_{i'})\}$ and $V_{\hat{\mathcal{U}}}^{\Theta}(i_{H'}, T) = \Theta\{\mathcal{R}_{\hat{\mathcal{U}}}(\{i_H\}, T, \mathbb{1}_{i_{H'}})\}$ for $H, H' \in \mathcal{H}$, $i' \in H'$ and $\Theta \in \{\inf, \sup\}$. Since the H -reachability probabilities constitute intervals, the assumption of Theorem 6 implies the assumption of Proposition 5 which, in turn, yields the claim. \square

Proposition 5. Fix the triples $\mathcal{U} = (\mathcal{S}, m, M)$ and $\hat{\mathcal{U}} = (\hat{\mathcal{S}}, \hat{m}, \hat{M})$ where $\hat{\mathcal{S}} = \{i_H \mid H \in \mathcal{H}\}$ and \mathcal{H} is a partition of \mathcal{S} . Assume that for any $T \geq 0$, $H, H' \in \mathcal{H}$ and $i' \in H'$

$$\mathcal{R}_{\mathcal{U}}(H, T, \mathbb{1}_{i'}) = \mathcal{R}_{\hat{\mathcal{U}}}(\{i_H\}, T, \mathbb{1}_{i_{H'}})$$

Then, \mathcal{H} is an ordinary lumpability of (\mathcal{S}, m) and (\mathcal{S}, M) .

Proof. Fix arbitrary $i, i' \in H'$. Picking $\mathbb{1}_i$ as the initial condition, we infer that $\sum_{j \in H} \pi_j(t) = (\sum_{j \in H} q_{i, j}(0))t + o(t)$ for small $t \geq 0$ because q is piecewise continuous function with right limits. Thus, with $\text{Int}(\mathcal{R}_{\mathcal{U}}(H, t, \mathbb{1}_i)) = (\alpha_i(t); \beta_i(t))$, we obtain²

- $\alpha_i(t) = (\sum_{j \in H} m_{i, j})t + o(t)$ for small $t \geq 0$.
- $\beta_i(t) = (\sum_{j \in H} M_{i, j})t + o(t)$ for small $t \geq 0$.

Likewise, if the initial condition is $\mathbb{1}_{i'}$, we obtain that

- $\alpha_{i'}(t) = (\sum_{j \in H} m_{i', j})t + o(t)$ for small $t \geq 0$.
- $\beta_{i'}(t) = (\sum_{j \in H} M_{i', j})t + o(t)$ for small $t \geq 0$.

Moreover, the assumption ensures $\mathcal{R}_{\mathcal{U}}(H, T, \mathbb{1}_i) = \mathcal{R}_{\mathcal{U}}(H, T, \mathbb{1}_{i'})$ for all $T \geq 0$. This implies that $\alpha_i(0) = \alpha_{i'}(0)$ and $\beta_i(0) = \beta_{i'}(0)$, thus showing that \mathcal{H} is an ordinary lumpability of (\mathcal{S}, m) and (\mathcal{S}, M) . \square

D. Proof of Theorem 7

We prove Theorem 7 by exploiting the fact that the validity of an until formula can be expressed in terms of a reachability probability [16], [17], [55]. In this section, we drop rewards for sake of readability and assume that transition rate functions are piecewise analytic, a property that has been required for the well-definedness of the satisfiability operator in [17], [55]. We begin by introducing a version of the auxiliary CTMC from [55]

²By Filippov's theorem [46, Section 4.5], measurability of rate functions ensures that the reachable sets are closed intervals $[\alpha(t); \beta(t)]$.

that is tailored to our needs. To improve readability, we omit the explicit mentioning of the running and final rewards.

Definition 10 (Auxiliary UCTMC). Assume that \mathcal{H} is a UCTMC lumpability of the UCTMC (\mathcal{S}, m, M) . Moreover, let $U, \mathcal{T} \subseteq \mathcal{S}$ be such that both U and \mathcal{T} can be written as unions of blocks from \mathcal{H} . With this, $(\tilde{\mathcal{V}}, \tilde{m}, \tilde{M})$ is given by $\tilde{\mathcal{V}} = \mathcal{S} \cup \tilde{\mathcal{S}}$, where $\tilde{\mathcal{S}} = \{\tilde{i} \mid i \in \mathcal{S}\}$, and

- $\tilde{m}_{i, j} := m_{i, j}$ and $\tilde{M}_{i, j} := M_{i, j}$ if $i \notin U \cup \mathcal{T}$ and $j \notin \mathcal{T}$;
- $\tilde{m}_{i, \tilde{j}} := m_{i, j}$ and $\tilde{M}_{i, \tilde{j}} := M_{i, j}$ if $i \notin U \cup \mathcal{T}$ and $j \in \mathcal{T}$;
- all other entries of \tilde{m} and \tilde{M} are zero.

As observed next, a UCTMC lumpability of (\mathcal{S}, m, M) induces a UCTMC lumpability of $(\tilde{\mathcal{V}}, \tilde{m}, \tilde{M})$.

Lemma 2. Assume that \mathcal{H} is a UCTMC lumpability of the UCTMC (\mathcal{S}, m, M) . Moreover, let $U, \mathcal{T} \subseteq \mathcal{S}$ be such that both U and \mathcal{T} can be written as unions of blocks from \mathcal{H} . Then, $\tilde{\mathcal{H}} = \mathcal{H} \cup \tilde{\mathcal{H}}$ is a UCTMC lumpability.

Proof. Follows by noting that the auxiliary UCTMC arises by redirecting the lower and upper bounds blockwise. \square

The model checking of until formulae is ultimately related to the probability that a time-varying target set can be reached by avoiding a time-varying set of unsafe states [17], [55]. The next definition formalizes this in our context.

Definition 11. Assume that \mathcal{H} is a UCTMC lumpability of (\mathcal{S}, m, M) . Further, let $U, \mathcal{T} : [0; \infty) \rightarrow \text{Powerset}(\mathcal{S})$ be such that

- U, \mathcal{T} have, on any bounded time interval at most finitely many discontinuity points with respect to the discrete topology;
- both $U(t)$ and $\mathcal{T}(t)$ can be written, for any $t \geq 0$, as unions of blocks from \mathcal{H} .

Then, $\mathcal{P}_{\text{reach}}(Q, t, T, \mathcal{T}, U)[i]$ is the probability of the set of paths underlying a given $Q = (q_{i, j})_{i, j}$ reaching a (target) state in $\mathcal{T}(T)$ at time $T \in [t; t + T]$ without passing through a (unsafe) state in $U(T')$ for any $T' \in [t; T]$, when starting in state $i \in \mathcal{S}$ at time t .

The following result is key for the proof of Theorem 7.

Proposition 6. Assume that \mathcal{H} is a UCTMC lumpability of (\mathcal{S}, m, M) . Let \mathcal{H} induce $(\tilde{\mathcal{V}}, \tilde{m}, \tilde{M})$ and $U, \mathcal{T} : [0; \infty) \rightarrow \text{Powerset}(\mathcal{S})$ be such that

- U, \mathcal{T} have, on any bounded time interval at most finitely many discontinuity points with respect to the discrete topology;
- $U(T)$ and $\mathcal{T}(T)$ can be written, for any $T \geq 0$, as unions of blocks from \mathcal{H} .

Set $\hat{U}(T) := \{i_H \mid H \in \mathcal{H} \wedge H \subseteq U(T)\}$ and $\hat{\mathcal{T}}(T) := \{i_H \mid H \in \mathcal{H} \wedge H \subseteq \mathcal{T}(T)\}$. Then, for given $T > 0$, $t \geq 0$, $H \in \mathcal{H}$ and $i \in H$:

- for admissible $(q_{i, j})_{(i, j)}$, we construct admissible $(\hat{q}_{i_H, i_{H'}})_{(i_H, i_{H'})}$;
- instead, for admissible $(\hat{q}_{i_H, i_{H'}})_{(i_H, i_{H'})}$, we construct admissible $(q_{i, j})_{(i, j)}$,

such that $\mathcal{P}_{\text{reach}}(Q, t, T, U, \mathcal{T})[i] = \mathcal{P}_{\text{reach}}(\hat{Q}, t, T, \hat{U}, \hat{\mathcal{T}})[i_H]$.

Proof. Let $t = T_0 < T_1 < \dots < T_{\kappa+1} = t + T$ be the time points in $[t; t + T]$ at which discontinuities of U or \mathcal{T} may arise. Following [55], we set $W(s) = \mathcal{S} \setminus (U(s) \cup \mathcal{T}(s))$ and let $\zeta_W(T_\nu)$ be the $n \times n$ matrix equal to 1 only on the diagonal elements corresponding to states ι belonging to both $W(T_\nu^-)$ and $W(T_\nu^+)$ (i.e., states that are safe and not a target both before and after T_ν), and equal to 0 elsewhere. Furthermore, let $\zeta_{\mathcal{T}}(T_\nu)$ be the $n \times n$ matrix equal to 1 in the diagonal elements corresponding to states ι belonging to $W(T_\nu^-) \cap \mathcal{T}(T_\nu^+)$ and zero elsewhere. Finally, let $\zeta(T_\nu)$ be the $2n \times 2n$ matrix

$$\zeta(T_\nu) := \begin{pmatrix} \zeta_W(T_\nu) & \zeta_{\mathcal{T}}(T_\nu) \\ 0 & I_{n \times n} \end{pmatrix}$$

Let us assume that we are given an admissible Q and write e_i for the unit vector of state i . Thanks to the fact that Q is piecewise analytic with finitely many discontinuity points on any bounded time interval, the discussion in [55] ensures that

$$\mathcal{P}_{\text{reach}}(Q, t, T, U, \mathcal{T})[i] = \sum_{\bar{j} \in \mathcal{S}} \Upsilon(t, t + T)_{i, \bar{j}} + \mathbb{1}\{i \in \mathcal{T}(t)\},$$

where $\mathbb{1}$ denotes the characteristic function, while

$$\Upsilon(t, t + T) = \tilde{\Pi}(t, T_1) \zeta(T_1) \tilde{\Pi}(T_1, T_2) \zeta(T_2) \dots \dots \cdot \zeta(T_\kappa) \tilde{\Pi}(T_\kappa, t + T)$$

is such that $\tilde{\Pi}(t_1, t_2)$ is the $2n \times 2n$ matrix where $e_{\iota'}^T \tilde{\Pi}(t_1, t_2) e_{\iota'}$ is the probability that the auxiliary CTMC is in state $\iota' \in \tilde{\mathcal{V}}$ at time t_2 , provided that it was initialized with state $\iota \in \tilde{\mathcal{V}}$ at time t_1 . The auxiliary CTMC in turn is given by Definition 10 and

- $U_\nu := U(\frac{T_{\nu-1} + T_\nu}{2})$ and $\mathcal{T}_\nu := \mathcal{T}(\frac{T_{\nu-1} + T_\nu}{2})$;
- $\tilde{\pi}[T_\nu]^T := \tilde{\pi}[T_{\nu-1}]^T \cdot \tilde{\Pi}_{|\mathcal{S} \times \mathcal{S}}(T_{\nu-1}, T_\nu)$ and $\tilde{\pi}[T_0] := e_i^T$;
- Q on $[T_{\nu-1}; T_\nu]$ is induced by U_ν , \mathcal{T}_ν , $\tilde{\pi}[T_{\nu-1}]$ and Q .

Since $\tilde{\Pi}(T_{\nu-1}, T_{\nu-1}) = I_{2n \times 2n}$, matrix $\tilde{\Pi}(T_{\nu-1}, T_\nu)$ can be obtained by solving the forward Kolmogorov equation $\partial_T \tilde{\Pi}(T_{\nu-1}, T) = \tilde{\Pi}(T_{\nu-1}, T) \cdot \tilde{Q}(T)$ on the interval $T \in [T_{\nu-1}; T_\nu]$. In particular, $e_{\iota'}^T \cdot \tilde{\Pi}(T_{\nu-1}, T_\nu)$ is given by $\tilde{\pi}(T_\nu)$ when $\tilde{\pi}(T_{\nu-1}) = e_\iota$ and $\partial_T \tilde{\pi}(T)^T = \tilde{\pi}^T(T) \cdot \tilde{Q}(T)$ for all $T \in [T_{\nu-1}; T_\nu]$. The composite term $\tilde{\Pi}(T_{\nu-1}, T_\nu) \zeta(T_\nu)$ writes as (for the benefit of presentation, we suppress the explicit time dependence in the following equation):

$$\begin{aligned} \tilde{\Pi} \cdot \zeta &= \begin{pmatrix} \tilde{\Pi}_{|\mathcal{S} \times \mathcal{S}} & \tilde{\Pi}_{|\mathcal{S} \times \mathcal{S}} \\ 0 & I_{n \times n} \end{pmatrix} \cdot \begin{pmatrix} \zeta_W & \zeta_{\mathcal{T}} \\ 0 & I_{n \times n} \end{pmatrix} \\ &= \begin{pmatrix} \tilde{\Pi}_{|\mathcal{S} \times \mathcal{S}} \cdot \zeta_W & \tilde{\Pi}_{|\mathcal{S} \times \mathcal{S}} \cdot \zeta_{\mathcal{T}} + \tilde{\Pi}_{|\mathcal{S} \times \mathcal{S}} \\ 0 & I_{n \times n} \end{pmatrix} \end{aligned}$$

Note that, for all $H \in \mathcal{H}$ and $\iota, \iota' \in H$, it holds that $e_\iota^T \zeta_W e_{\iota'} = e_{\iota'}^T \zeta_W e_\iota$ and $e_\iota^T \zeta_{\mathcal{T}} e_{\iota'} = e_{\iota'}^T \zeta_{\mathcal{T}} e_\iota$ because U_ν and \mathcal{T}_ν are unions of blocks from \mathcal{H} . Hence, ζ_W and $\zeta_{\mathcal{T}}$ are cutoff functions that are operating blockwise.

The above discussion and Lemma 2 ensure that a given probability distribution $\tilde{\pi}[T_{\nu-1}]$ induces a piecewise analytic \hat{Q} on $[T_{\nu-1}; T_\nu]$ such that

$$\sum_{\iota \in X} \tilde{\pi}_\iota(T) = \hat{\pi}_{\iota_X}(T) \text{ for all } X \in \mathcal{H} \cup \bar{\mathcal{H}} \text{ and } T \in [T_{\nu-1}; T_\nu],$$

where $\hat{\pi}$ is the transient probability of the lumped auxiliary CTMC. Hence, for all $H' \in \mathcal{H}$, it holds that

$$\begin{aligned} \tilde{\pi}[T_{\nu-1}]^T \cdot \tilde{\Pi}_{|\mathcal{S} \times \mathcal{S}}(T_{\nu-1}, T_\nu) \cdot \left(\sum_{\iota \in H'} e_\iota \right) &= \\ \hat{\pi}[T_{\nu-1}]^T \cdot \hat{\Pi}_{|\hat{\mathcal{S}} \times \hat{\mathcal{S}}}(T_{\nu-1}, T_\nu) \cdot e_{i_{H'}} & \\ \tilde{\pi}[T_{\nu-1}]^T \cdot \tilde{\Pi}_{|\mathcal{S} \times \bar{\mathcal{S}}}(T_{\nu-1}, T_\nu) \cdot \left(\sum_{\iota \in \bar{H}'} e_\iota \right) &= \\ \hat{\pi}[T_{\nu-1}]^T \cdot \hat{\Pi}_{|\hat{\mathcal{S}} \times \bar{\hat{\mathcal{S}}}}(T_{\nu-1}, T_\nu) \cdot e_{i_{\bar{H}'}} \end{aligned}$$

where $\hat{\Pi}$ is the matrix of transient probabilities of the lumped CTMC. The above discussion ensures that

$$\begin{aligned} e_i^T \Upsilon(t, t + T) &= e_i^T \tilde{\Pi}(t, T_1) \zeta(T_1) \tilde{\Pi}(T_1, T_2) \zeta(T_2) \dots \\ &\dots \zeta(T_\kappa) \tilde{\Pi}(T_\kappa, t + T) \\ &= e_{i_H}^T \hat{\Pi}(t, T_1) \hat{\zeta}(T_1) \hat{\Pi}(T_1, T_2) \hat{\zeta}(T_2) \dots \\ &\dots \hat{\zeta}(T_\kappa) \hat{\Pi}(T_\kappa, t + T) \end{aligned}$$

for all $H \in \mathcal{H}$ and $i \in H$, where $\hat{\zeta}$ is defined in the obvious manner. This implies the statement if we can find an admissible \hat{Q} such that $\hat{Q} = \tilde{Q}$. We next study \hat{Q} via the following case distinction:

- If $H \cap (U_\nu \cup \mathcal{T}_\nu) = \emptyset \wedge H' \cap \mathcal{T}_\nu = \emptyset$: Then, the proof of Theorem 5 yields

$$\hat{q}_{i_H, i_{H'}}(t) = \sum_{i_k \in H} \frac{\tilde{\pi}_{i_k}(t)}{\sum_{i \in H} \tilde{\pi}_i(t)} \sum_{j_k \in H'} q_{i_k, j_k}(t)$$

- If $H \cap (U_\nu \cup \mathcal{T}_\nu) = \emptyset \wedge H' \subseteq \mathcal{T}_\nu$: Then, the proof of Theorem 5 yields

$$\hat{q}_{i_H, \bar{i}_{H'}}(t) = \sum_{i_k \in H} \frac{\tilde{\pi}_{i_k}(t)}{\sum_{i \in H} \tilde{\pi}_i(t)} \sum_{j_k \in H'} q_{i_k, j_k}(t)$$

- Otherwise, $\hat{q}_{\iota, \mu} \equiv 0$.

The above case distinction suggests to pick

$$\hat{q}_{i_H, i_{H'}}(t) = \sum_{i_k \in H} \frac{\tilde{\pi}_{i_k}(t)}{\sum_{i \in H} \tilde{\pi}_i(t)} \sum_{j_k \in H'} q_{i_k, j_k}(t)$$

for all $H, H' \in \mathcal{H}$ with $H \neq H'$. With this, we proceed by the following case distinction.

- If $i_H \notin \hat{U}_\nu \cup \hat{\mathcal{T}}_\nu \wedge i_{H'} \notin \hat{\mathcal{T}}_\nu$: Then $\hat{q}_{i_H, i_{H'}} \equiv \hat{q}_{i_H, i_{H'}} \equiv \hat{q}_{i_H, i_{H'}}$.
- If $i_H \notin \hat{U}_\nu \cup \hat{\mathcal{T}}_\nu \wedge i_{H'} \in \hat{\mathcal{T}}_\nu$: Then $\hat{q}_{i_H, \bar{i}_{H'}} \equiv \hat{q}_{i_H, i_{H'}} \equiv \hat{q}_{i_H, \bar{i}_{H'}}$.
- Otherwise, $\hat{q}_{\iota, \mu} \equiv 0 \equiv \hat{q}_{\iota, \mu}$.

This completes the proof in the case where we are given an admissible Q and have to find an admissible \hat{Q} such that

$$\mathcal{P}_{\text{reach}}(Q, t, T, U, \mathcal{T})[i] = \mathcal{P}_{\text{reach}}(\hat{Q}, t, T, \hat{U}, \hat{\mathcal{T}})[i_H] \quad (7)$$

For the converse, let us now assume that we are given some admissible \hat{Q} and have to find an admissible Q such that (7) holds true. To this end, we construct Q from \hat{Q} as in the proof of Theorem 5, that is, we set $q_{i_k, j_k}(t)$ to

$$m_{i_k, j_k} + \frac{M_{i_k, j_k} - m_{i_k, j_k}}{\sum_{j_l \in H'} (M_{i_k, j_l} - m_{i_k, j_l})} (\hat{q}_{i_H, i_{H'}}(t) - \hat{n}_{i_H, i_{H'}})$$

To see that this Q does the job, we construct following the foregoing discussion \tilde{Q} from Q and \hat{Q}' from \tilde{Q} , respectively. This yields

$$\begin{aligned}\hat{q}'_{i_H, i_{H'}}(t) &= \sum_{i_k \in H} \frac{\tilde{\pi}_{i_k}(t)}{\sum_{i \in H} \tilde{\pi}_i(t)} \sum_{j_k \in H'} q_{i_k, j_k}(t) \\ &= \sum_{i_k \in H} \frac{\tilde{\pi}_{i_k}(t)}{\sum_{i \in H} \tilde{\pi}_i(t)} \sum_{j_k \in H'} \left[m_{i_k, j_k} \right. \\ &\quad \left. + \frac{M_{i_k, j_k} - m_{i_k, j_k}}{\sum_{j_l \in H'} (M_{i_k, j_l} - m_{i_k, j_l})} (\hat{q}_{i_H, i_{H'}}(t) - \hat{m}_{i_H, i_{H'}}) \right] \\ &= \sum_{i_k \in H} \frac{\tilde{\pi}_{i_k}(t)}{\sum_{i \in H} \tilde{\pi}_i(t)} (\hat{m}_{i_H, i_{H'}} + \hat{q}_{i_H, i_{H'}}(t) - \hat{m}_{i_H, i_{H'}}) \\ &= \hat{q}_{i_H, i_{H'}}(t)\end{aligned}$$

Since $\hat{q}' \equiv \hat{q}$, the discussion preceding (7) ensures that our choice of Q yields (7). \square

Armed with Proposition 6, we can prove Theorem 7.

Theorem 7. The proof proceeds by structural induction on ϕ .

- $\phi = \alpha$: Follows from the fact that $\hat{\mathcal{L}}(i_H) = \mathcal{L}(i)$ for all $H \in \mathcal{H}$ and $i \in H$.
- $\phi = \phi_1 \wedge \phi_2$: Follows by induction hypothesis.
- $\phi = \neg \phi_1$: Follows by induction hypothesis.
- $\phi = \mathcal{P}_{\bowtie p}^{\forall}(\phi_1 \mathbf{U}^{[t_0; t_1]} \phi_2)$: Let us define

$$U(t) := \{j \in \mathcal{S} \mid j, t \models \neg \phi_1\} \quad \mathcal{T}(t) := \{j \in \mathcal{S} \mid j, t \models \phi_2\}$$

By induction hypothesis, it holds that both $U(T)$ and $\mathcal{T}(T)$ can be written, for any $T \geq 0$, as unions of blocks from \mathcal{H} . The definition of the semantics, instead, ensures that U and \mathcal{T} have finitely many discontinuity points on any bounded time interval. Together with $\hat{U}(t) := \{i_H \mid H \in \mathcal{H} \wedge H \subseteq U(t)\}$ and $\hat{\mathcal{T}}(t) := \{i_H \mid H \in \mathcal{H} \wedge H \subseteq \mathcal{T}(t)\}$, the discussion in [55] implies that for any admissible ...

- ... q we have: $i, t \models_U \mathcal{P}_{\bowtie p}(\phi_1 \mathbf{U}^{[t_0; t_1]} \phi_2)$ iff $\mathcal{P}_{\text{reach}}(Q, t, t_1 - t_0, U, \mathcal{T})[i] \bowtie p$;
- ... \hat{q} we have: $i_H, t \models_{\hat{U}} \mathcal{P}_{\bowtie p}(\phi_1 \mathbf{U}^{[t_0; t_1]} \phi_2)$ iff $\mathcal{P}_{\text{reach}}(\hat{Q}, t, t_1 - t_0, \hat{U}, \hat{\mathcal{T}})[i_H] \bowtie p$.

With this, Proposition 6 yields the claim. \square

REFERENCES

- [1] D. Gillespie, “Exact stochastic simulation of coupled chemical reactions,” *Journal of Phys. Chemistry*, vol. 81, no. 25, pp. 2340–2361, 1977.
- [2] J. Kemeny and J. Snell, *Finite Markov Chains*. Springer, 1976.
- [3] P. D. O'Neill, “A tutorial introduction to Bayesian inference for stochastic epidemic models using Markov chain monte Carlo methods,” *Mathematical biosciences*, vol. 180, no. 1–2, pp. 103–114, 2002.
- [4] M. Hajiaghayi, B. Kirkpatrick, L. Wang, and A. Bouchard-Côté, “Efficient continuous-time markov chain estimation,” in *ICML*, 2014, pp. 638–646.
- [5] W. Wang, M. Tang, H. E. Stanley, and L. A. Braunstein, “Unification of theoretical approaches for epidemic spreading on complex networks,” *Reports on Progress in Physics*, vol. 80, no. 3, p. 036603, 2017.
- [6] D. F. Anderson and T. G. Kurtz, *Continuous Time Markov Chain Models for Chemical Reaction Networks*. Springer New York, 2011, pp. 3–42.
- [7] A. C. Babbie and M. P. H. Stumpf, “How to deal with parameters for whole-cell modelling,” *Journal of The Royal Society Interface*, vol. 14, no. 133, p. 20170237, 2017.
- [8] X. Guo and O. Hernandez-Lerma, *Continuous-Time Markov Decision Processes*. Springer Verlag, 2009.
- [9] P. Buchholz, E. M. Hahn, H. Hermanns, and L. Zhang, “Model Checking Algorithms for CTMDPs,” in *CAV*, 2011, pp. 225–242.
- [10] E. M. Hahn, H. Hermanns, R. Wimmer, and B. Becker, “Transient Reward Approximation for Continuous-Time Markov Chains,” *IEEE Trans. Reliability*, vol. 64, no. 4, pp. 1254–1275, 2015.
- [11] L. Li, T. J. Walsh, and M. L. Littman, “Towards a unified theory of state abstraction for MDPs,” in *ISAIM*, 2006.
- [12] D. Abel, D. E. Hershkowitz, and M. L. Littman, “Near optimal behavior via approximate state abstraction,” in *ICML*, 2016, pp. 2915–2923.
- [13] P. Buchholz, “Exact and ordinary lumpability in finite Markov chains,” *Journal of Applied Probability*, vol. 31, no. 1, pp. 59–75, 1994.
- [14] K. G. Larsen and A. Skou, “Bisimulation through probabilistic testing,” *Inf. Comput.*, vol. 94, no. 1, pp. 1–28, 1991.
- [15] J. Desharnais and P. Panagaden, “Continuous stochastic logic characterizes bisimulation of continuous-time Markov processes,” *J. Log. Algebraic Methods Program.*, vol. 56, no. 1–2, pp. 99–115, 2003.
- [16] C. Baier, B. R. Haverkort, H. Hermanns, and J. Katoen, “Model-Checking Algorithms for Continuous-Time Markov Chains,” *IEEE Trans. Software Eng.*, vol. 29, no. 6, pp. 524–541, 2003.
- [17] L. Bortolussi and J. Hillston, “Fluid Model Checking,” in *CONCUR*, 2012, pp. 333–347.
- [18] S. Derisavi, H. Hermanns, and W. H. Sanders, “Optimal state-space lumping in markov chains,” *Inf. Proc. Lett.*, vol. 87, pp. 309–315, 2003.
- [19] A. Valmari and G. Franceschinis, “Simple $O(m \log n)$ time Markov chain lumping,” in *TACAS*, 2010, pp. 38–52.
- [20] T. Han, J. Katoen, and A. Mereacre, “Compositional modeling and minimization of time-inhomogeneous markov chains,” in *HSCC*, M. Egerstedt and B. Mishra, Eds., vol. 4981, 2008, pp. 244–258.
- [21] M. R. Neuhäuffer and J. Katoen, “Bisimulation and Logical Preservation for Continuous-Time Markov Decision Processes,” in *CONCUR*, 2007, pp. 412–427.
- [22] M. R. Neuhäuffer, M. Stoelinga, and J. Katoen, “Delayed nondeterminism in continuous-time markov decision processes,” in *FOSSACS*, L. de Alfaro, Ed., 2009, pp. 364–379.
- [23] Y. Butkova, H. Hafezi, H. Hermanns, and J. Krcál, “Optimal continuous time Markov decisions,” in *ATVA*, 2015, pp. 166–182.
- [24] Y. Butkova and G. Fox, “Optimal time-bounded reachability analysis for concurrent systems,” in *TACAS*, vol. 11428, 2019, pp. 191–208.
- [25] T. Han, J. Katoen, and A. Mereacre, “Approximate parameter synthesis for probabilistic time-bounded reachability,” in *RTSS*. IEEE Computer Society, 2008, pp. 173–182.
- [26] E. M. Hahn, H. Hermanns, and L. Zhang, “Probabilistic reachability for parametric Markov models,” *STTT*, vol. 13, no. 1, pp. 3–19, 2011.
- [27] M. Kwiatkowska, G. Norman, and D. Parker, “PRISM 4.0: Verification of Probabilistic Real-time Systems,” in *CAV*, 2011, pp. 585–591.
- [28] M. Ceska, F. Dannenberg, N. Paoletti, M. Kwiatkowska, and L. Brim, “Precise parameter synthesis for stochastic biochemical systems,” *Acta Informatica*, vol. 54, no. 6, pp. 589–623, 2017.
- [29] C. Dehnert, S. Junges, N. Jansen, F. Corzilius, M. Volk, H. Buntjes, J. Katoen, and E. Abraham, “PROPHESY: A PRObabilistic ParamETER SYnthesis Tool,” in *CAV*, 2015, pp. 214–231.
- [30] T. S. Badings, N. Jansen, S. Junges, M. Stoelinga, and M. Volk, “Sampling-based verification of ctmc with uncertain rates,” in *CAV*, S. Shoham and Y. Vizel, Eds., vol. 13372, 2022, pp. 26–47.
- [31] L. Song, L. Zhang, and J. C. Godskesen, “Bisimulations and Logical Characterizations on Continuous-Time Markov Decision Processes,” in *VMCAI*, 2014, pp. 98–117.
- [32] V. Hashemi, A. Turrini, E. M. Hahn, H. Hermanns, and K. M. Elbassioni, “Polynomial-Time Alternating Probabilistic Bisimulation for Interval MDPs,” in *SETTA*, 2017, pp. 25–41.
- [33] G. J. Pappas and S. Simic, “Consistent abstractions of affine control systems,” *IEEE TAC*, vol. 47, no. 5, pp. 745–756, 2002.
- [34] M. Althoff, C. Le Guernic, and B. H. Krogh, “Reachable Set Computation for Uncertain Time-Varying Linear Systems,” in *HSCC*, 2011, pp. 93–102.
- [35] D. Angeli, “A Lyapunov approach to incremental stability properties,” *IEEE Trans. Automat. Contr.*, vol. 47, no. 3, pp. 410–421, 2002.
- [36] J. Lygeros, “On reachability and minimum cost optimal control,” *Automatica*, vol. 40, no. 6, pp. 917–927, 2004.
- [37] S. Prajna, “Barrier certificates for nonlinear model validation,” *Automatica*, vol. 42, no. 1, pp. 117–126, 2006.
- [38] J. K. Scott and P. I. Barton, “Bounds on the reachable sets of nonlinear control systems,” *Automatica*, vol. 49, no. 1, pp. 93 – 100, 2013.
- [39] A. J. van der Schaft, “Equivalence of dynamical systems by bisimulation,” *IEEE Transactions on Automatic Control*, vol. 49, 2004.
- [40] P. Tabuada and G. J. Pappas, “Abstractions of Hamiltonian control systems,” *Automatica*, vol. 39, no. 12, pp. 2025–2033, 2003.

- [41] S. Strubbe and A. J. van der Schaft, "Bisimulation for communicating piecewise deterministic markov processes (cpdps)," in *HSCC*, M. Morari and L. Thiele, Eds., vol. 3414, 2005, pp. 623–639.
- [42] G. J. Pappas, G. Lafferriere, and S. Sastry, "Hierarchically consistent control systems," *IEEE TAC*, vol. 45, no. 6, pp. 1144–1160, 2000.
- [43] B. C. Geiger, T. Petrov, G. Kubin, and H. Koeppl, "Optimal kullback-leibler aggregation via information bottleneck," *IEEE Transactions on Automatic Control*, vol. 60, no. 4, pp. 1010–1022, 2015.
- [44] M. Lahijanian, S. B. Andersson, and C. Belta, "Formal verification and synthesis for discrete-time stochastic systems," *IEEE Transactions on Automatic Control*, vol. 60, no. 8, pp. 2031–2045, 2015.
- [45] L. Cardelli, R. Grosu, K. G. Larsen, M. Tribastone, M. Tschaikowski, and A. Vandin, "Lumpability for uncertain continuous-time markov chains," in *QEST*, A. Abate and A. Marin, Eds., vol. 12846, 2021, pp. 391–409.
- [46] D. Liberzon, *Calculus of Variations and Optimal Control Theory: A Concise Introduction*. Princeton University Press, 2011.
- [47] L. Kleinrock, *Queueing Systems Theory*. Wiley-Intersc., 1975, vol. 1.
- [48] J. Gao, M. A. Zahran, A. Sheoran, S. Fahmy, and B. Ribeiro, "Infinity learning: Learning Markov chains from aggregate steady-state observations," in *AAAI*, vol. 34, no. 04, 2020, pp. 3922–3929.
- [49] W. A. Massey, "The analysis of queues with time-varying rates for telecommunication models," *Telecommunication Systems*, vol. 21, no. 2, pp. 173–204, 2002.
- [50] M. Tschaikowski, "Over-approximation of fluid models," *IEEE Trans. Autom. Control*, vol. 65, no. 3, pp. 999–1013, 2020.
- [51] M. I. Kamien and N. L. Schwartz, "Sufficient conditions in optimal control theory," *Journal of Econ. Theory*, vol. 3, pp. 207–214, 1971.
- [52] J. Norris, *Markov Chains*, ser. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998.
- [53] L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin, "Maximal aggregation of polynomial dynamical systems," *Proc. of the National Academy of Sciences*, vol. 114, no. 38, pp. 10 029–10 034, 2017.
- [54] E. M. Clarke, E. A. Emerson, and J. Sifakis, "Model checking: algorithmic verification and debugging," *Commun. ACM*, vol. 52, pp. 74–84, 2009.
- [55] L. Bortolussi and J. Hillston, "Model checking single agent behaviours by fluid approximation," *Inf. Comput.*, vol. 242, pp. 183–226, 2015.
- [56] L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin, "ERODE: A tool for the evaluation and reduction of ordinary differential equations," in *TACAS*, 2017.
- [57] B. Haverkort, H. Hermanns, and J. Katoen, "On the use of model checking techniques for dependability evaluation," in *SRDS'20*, pp. 228–237.
- [58] M. Massink, J. Katoen, and D. Latella, "Model checking dependability attributes of wireless group communication," in *DSN*, 2004, pp. 711–720.
- [59] A. Bondavalli, A. Coccoli, and F. Di Giandomenico, *QoS Analysis of Group Communication Protocols in Wireless Environment*. Springer, 2002, pp. 169–188.
- [60] P. Lecca and C. Priami, "Cell cycle control in eukaryotes: A BioSpi model," *ENTCS*, vol. 180, no. 3, pp. 51–63, 2007.
- [61] B. Novak, A. Csikasz-Nagy, B. Györfy, K. Nasmyth, and J. J. Tyson, "Model scenarios for evolution of the eukaryotic cell cycle," *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, vol. 353, no. 1378, pp. 2063–2076, 1998.
- [62] O. Ibe and K. Trivedi, "Stochastic Petri net models of polling systems," *IEEE Journal on Sel. Areas in Comm.*, vol. 8, pp. 1649–1657, 1990.
- [63] L. C. Piccinini, G. Stampacchia, and G. Vidossich, *Ordinary Differential Equations in Rn*. Springer, 1984.
- [64] C. W. Gear, *Numerical Initial Value Problems in Ordinary Differential Equations*. Prentice Hall PTR, 1971.



Luca Cardelli has a PhD in computer science from the University of Edinburgh. He worked at Bell Labs, Murray Hill, from 1982 to 1985, and at Digital Equipment Corporation, Systems Research Center in Palo Alto, from 1985 to 1997, before assuming a position at Microsoft Research, in Cambridge UK, where he was head of the Programming Principles and Tools and Security groups until 2012. Since 2014 he is a Royal Society Research Professor at the University of Oxford. He is a Fellow of the Royal Society, an ACM Fellow, an Elected Member of the Academia Europaea, an Elected Member of AITO, and a member of EATCS.



Radu Grosu has a PhD in computer science from the TU Munich. He is a Professor at TU Wien and a Research Professor at the Stony Brook University, US, where he has been an Associate Professor before joining TU Wien. Radu Grosu is the recipient of the NSF Career Award and an elected member of the International Federation for Information Processing, Working Group 2.2.



Kim G. Larsen is a Professor in Computer Science at Aalborg University and the director of the ICT-competence center CISS, Center for Embedded Software Systems. He is the holder of the VILLUM Investigator grant S4OS and the ERC Advanced Grant LASSO. Moreover, he is the director of the Danish Innovation Network InfinIT, as well as the Innovation Fund Denmark research center DiCyPS.



Mirco Tribastone is a Professor at IMT Lucca, Italy. Prior to joining IMT Lucca he was Associate Professor at the University of Southampton, UK, and Assistant Professor at the Ludwig-Maximilians University of Munich, Germany. He received his Ph.D. in Computer Science from the University of Edinburgh, UK, in 2010. He graduated in Computer Engineering at the University of Catania, Italy.



Max Tschaikowski is a Poul Due Jensen Associate Professor at Aalborg University, Denmark. Prior to it, he was a Lise Meitner Fellow at TU Wien, Austria, an Assistant Professor at IMT Lucca, Italy, a Research Fellow at the University of Southampton, UK, and a Research Assistant at the Ludwig-Maximilians University in Munich, Germany. He was awarded a Diplom in mathematics (equivalent to a Master) and a Ph.D. in computer science by the LMU in 2010 and 2014, respectively.



Andrea Vandin is an Associate Professor at Sant'Anna School for Advanced Studies, Pisa, Italy, and an Adjunct Associate Professor at DTU Technical University of Denmark. Prior to it he was an Associate Professor at DTU and an Assistant Professor at IMT Lucca, Italy. In 2013-2015 he was a Senior Research Assistant at the University of Southampton, UK. He received his PhD in Computer Science and Engineering from IMT Lucca, Italy. He graduated in Computer Science at the University of Pisa, Italy.