

Lumpability for Uncertain Continuous-Time Markov Chains

Luca Cardelli¹, Radu Grosu², Kim G. Larsen³
Mirco Tribastone⁴, Max Tschaikowski³, and Andrea Vandin⁵

¹ University of Oxford, UK

² TU Wien, Austria

³ Aalborg University, Denmark

⁴ IMT School for Advanced Studies Lucca, Italy

⁵ Sant’Anna School of Advanced Studies, Pisa, Italy

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. In this paper we consider uncertain CTMCs, where rates are assumed to vary non-deterministically with time from bounded continuous intervals. This leads to a semantics which associates each state with the reachable set of its probability under all possible choices of the uncertain rates. We develop a notion of lumpability which identifies a partition of states where each block preserves the reachable set of the sum of its probabilities, essentially lifting the well-known CTMC ordinary lumpability to the uncertain setting. We proceed with this analogy with two further contributions: a logical characterization of uncertain CTMC lumping in terms of continuous stochastic logic; and a polynomial time and space algorithm for the minimization of uncertain CTMCs by partition refinement, using the CTMC lumping algorithm as an inner step. As a case study, we show that the minimizations in a substantial number of CTMC models reported in the literature are *robust* with respect to uncertainties around their original, fixed, rate values.

1 Introduction

Motivation. Continuous-time Markov chains (CTMCs) are a fundamental tool for describing a wide range of natural and engineered systems and serve as the underlying semantics for several formalisms such as stochastic Petri Nets [16], stochastic process algebra (e.g., [33, 34]), and chemical reaction networks [26]. A CTMC is typically characterized by a number of parameters such as arrival and service rates in a queuing network [51], transmission and infection rates of epidemic processes [48], and the kinetic rates of a chemical reaction. 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 to be estimated 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 [14]. In addition, sometimes the modeler wishes to be deliberately imprecise about the value of certain parameters in order to explicitly account for the disagreement between the real system and its model.

These motivations have stimulated a vigorous line of research into quantitative modeling frameworks where *uncertainty* is a first-class citizen, with the basic idea to replace known constants with *sets* of values which can be nondeterministically assigned to parameters. A prominent instance is Jonnson and Larsen’s interval specification systems [35] (equivalent to interval-valued finite Markov chains [39]), where the probability of making a transition between two states of a discrete-time Markov chain is assumed to be taken from a continuous interval of possible values, later generalized to polynomial constraints [10].

Contributions. In this paper we consider uncertain CTMCs (UCTMCs). They allow time-varying nondeterministic uncertainty in the values of the rate parameters within given bounded intervals. This is essentially the continuous-time analogue of the model of nondeterminism in [24, 47], and can be seen as an over-approximation for a time-invariant interpretation of uncertainty which underlies a family of CTMCs, one for each possible choice of rate parameter values [35].


Here we study minimization of UCTMCs, motivated by the appeal to work with models of smaller size that still preserve quantities of interest for analysis and verification purposes. We proceed by means of analogies with the well-known CTMCs counterpart of ordinary lumpability [38, 6] (reviewed in Section 2):

- CTMC ordinary lumpability identifies a partition of the state space which induces a lumped chain where each macro-state represents a partition block; the probability of being in each macro-state at all time points is equal to the sum of the probabilities of the states of the original CTMC belonging to that block [6]. The semantics of UCTMCs associates each state with the reachable set of the probabilities of that state under all possible values of the uncertain transition rates at any time point. *Mutatis mutandis*, our notion of lumpability is such that the lumped UCTMC preserves reachable sets of sums of the states in each block. In fact, UCTMC lumpability turns out to be a conservative extension of CTMC lumpability.
- We study the logical characterization of UCTMC lumpability. Similarly to the characterization of continuous stochastic logic (CSL) [1] by F-bisimulation [1], a notion closely related to ordinary lumpability, we prove that UCTMC lumpability preserves a conservative extension of CSL to UCTMCs, where a CSL formula is satisfied by a UCTMC if it is true for all possible rate values.
- CTMCs enjoy an efficient minimization algorithm based on partition refinement which computes the coarsest ordinarily lumpable partition that refines a given initial partition of states [20, 52]. Here we develop an analogous algorithm for UCTMCs where the CTMC lumping algorithm is used as an inner step: the coarsest UCTMC lumpable partition is the coarsest one that refines both of the two time-homogeneous CTMCs derived by choosing the lower and upper bounds for all uncertainty intervals, respectively. Thus, the minimization algorithm takes $\mathcal{O}(rs \log s)$ steps in the worst case, where r is the number of transitions and s is the number of states of the UCTMC.

As an application, we consider the problem of analyzing the “robustness” of CTMC lumping, i.e., to what extent the minimization depends on the specific choice of rate parameters of a model. Using a prototype implementation, we study how adding uncertainty intervals around the constant values of the rates of a CTMC model preserves the original CTMC lumping.

Further related work. A UCTMC can be seen as a continuous-time Markov decision process (MDP). Indeed, we formally show in Section 3.3 that the UCTMC can be alternatively given as a time-inhomogeneous continuous-time MDP with an uncountable action space, which represents the values within the uncertainty intervals, see [47] and [27, Section 2.2]. 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 [44, 7, 29, 9]. Another related model is that of parametric Markov chains and parametric MDPs [40, 30, 17], 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.

Most notions of lumpability and bisimulation for these models of uncertainty impose constraints that must hold for all actions (in the case of MDPs [44, 30, 49]) or, analogously, for all parameter evaluations (for parametric Markov chains [30]). Instead, our notion of lumping 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 lumping requires these two expressions to be equal for all possible assignments of the parameters [30]—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 lumping is the alternating probabilistic bisimulation considered in [31] 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 [31] it is not proved that the bisimulation is indeed necessary for the preservation of such properties. We relate UCTMC lumping to alternating probabilistic bisimulation by defining an approximation for the continuous-time MDP interpretation of the UCTMC that discretizes both time and the action space using an MDP with a probabilistic scheduler. On this discretized model, we show that a UCTMC lumping does correspond to a probabilistic alternating bisimulation, see Section 3.4.

Paper organization. Section 2 provides the background, while Section 3 introduces UCTMCs and discusses techniques for their analysis. Section 4, instead, introduces UCTMC lumpability, its quantitative and logical characterization, and an algorithm for the computation of the coarsest UCTMC lumpability. Section 5 continues with an evaluation of UCTMC lumpability on a set of benchmarks from the literature, while Section 6 concludes the paper.

2 Preliminaries

In this section we fix the notation and briefly recall the definitions of CTMCs and lumpability that will be used throughout the paper.

Notation. We use ∂_t to denote derivative with respect to time t , while x^T is the transpose of a vector x . 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{V} , 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. To facilitate later results, throughout this paper we assume that transition rates vary with time according to *uniformly piecewise analytic* functions, i.e., functions which are analytic and bounded on all intervals $[kh; (k+1)h)$, where $k \geq 0$ is an integer and $h > 0$ is a given fixed time step.

Definition 1 (CTMC). A time-varying CTMC is a tuple (\mathcal{V}, Q) where \mathcal{V} is a set of states $\mathcal{V} = \{1, \dots, n\}$, while $Q = (q_{i,j})_{i,j}$ is a time-varying transition rate matrix such that $q_{i,j} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is a uniformly piecewise analytic transition rate function from i into j . \square

The following result relates the (transient) probability distributions of (\mathcal{V}, Q) to the Kolmogorov equations for time-varying transition rates [27, Section 2.2].

Theorem 1. Given a CTMC (\mathcal{V}, 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 equation⁶

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

⁶ Proofs are given in the extended version available at doi.org/10.5281/zenodo.4699211

Thanks to Theorem 1, ordinary lumpability for time-varying CTMCs is a straightforward generalization of ordinary lumpability for time-homogeneous CTMCs (e.g., [6]). The next well-known result provides a quantitative characterization of ordinary lumpability.

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

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

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

- States $\hat{\mathcal{V}} := \{i_H \mid H \in \mathcal{H}\}$, 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 $(\hat{\mathcal{V}}, \hat{Q})$ 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 $(\hat{\mathcal{V}}, \hat{Q})$ are denoted by $\hat{\pi}$, the following holds.

- If \mathcal{H} is an ordinary lumping, 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 lumping.

3 Uncertain Continuous-time Markov Chains

UCTMCs allow transition rates to vary non-deterministically with time within bounded continuous intervals. After the formal introduction of the model (Section 3.1), we provide the semantics of UCTMCs both in terms of reachable sets of their probability distributions using the Kolmogorov equations (Section 3.2) and by means of an encoding into a time-inhomogeneous continuous-time MDP (Section 3.3). The time and action-space discrete approximation of the latter semantics is presented in Section 3.4.

3.1 Model Definition

Definition 2 (Uncertain CTMC). *An uncertain CTMC (\mathcal{V}, m, M) is a set of states $\mathcal{V} = \{1, \dots, n\}$ and non-negative matrices $m = (m_{i,j})_{i,j}$ and $M = (M_{i,j})_{i,j}$, with $m \leq M$, describing the lower and upper bounds of the transition rates, respectively. \square*

According to the above definition, a UCTMC (\mathcal{V}, m, M) induces two *extremal* (time-homogeneous) CTMCs (\mathcal{V}, m) and (\mathcal{V}, M) by fixing all lower and upper bounds, respectively, for each transition rate.

Example. Throughout of this paper, we will use the UCTMC depicted in

Figure 1 as a running example. To favor intuition, it can be interpreted as a symmetric model of two components (e.g., two virtual machines) with a binary state (e.g., down/0 and up/1). Assuming independent events, each UCTMC state tracks a possible configuration of the two machines. Each transition is labeled with the interval within which the rates can vary; we use distinct symbols α , β , γ to indicate different activities of an hypothetical system under study, such as start-up, shut-down or machine migration, respectively. When all parameters are precisely known, i.e., $\underline{\alpha} = \bar{\alpha}$, $\underline{\beta} = \bar{\beta}$, and $\underline{\gamma} = \bar{\gamma}$, there is the ordinary lumping consisting of blocks $\{11\}$, $\{01, 10\}$, and $\{00\}$.

In this paper we will develop the theory to capture such symmetry for UCTMCs. For this, here we also observe that the aforementioned ordinary lumping carries over to the two extremal time-homogeneous CTMCs. Indeed, it turns out that these are the only two CTMCs needed to consider for UCTMC lumpability, although the UCTMC admits time-varying behaviors that do not satisfy the conditions of CTMC ordinary lumpability stated in Theorem 2 and also mentioned in Section 1.




Fig. 1. Running example.

3.2 Reachable-set Semantics

Analogously to the probability distribution of a CTMC obeying the Kolmogorov equations, the semantics of a UCTMC is given by the set of reachable probability distributions under all possible time-varying values of the transition rate matrix.

Definition 3 (UCTMC reachable-set semantics). *The semantics of a given UCTMC $\mathcal{U} = (\mathcal{V}, m, M)$ is provided by the reachable sets*

$$\mathcal{R}_{\mathcal{U}}(H, \tau, \pi[0]) = \left\{ \sum_{i \in H} \pi_i(\tau) \mid \partial_t \pi(t)^T = \pi(t)^T Q(t) \right. \\ \left. \text{such that } \pi(0) = \pi[0] \text{ and } Q \text{ is admissible} \right\},$$

where $\tau \geq 0$ and $H \subseteq \mathcal{V}$, while $Q = (q_{i,j})_{i,j}$ is admissible if, for all $t \geq 0$ and $i \neq j$, $q_{i,j}(t) \in [m_{i,j}; M_{i,j}]$ and $q_{i,j}$ is uniformly piecewise analytic function of time. \square

Remark 1. The common notion of reachable sets is recovered by restricting H to singleton blocks only, i.e., $\{\{i\} \mid i \in \mathcal{V}\}$. We allow for general blocks because our ultimate goal is to relate sums of reachable probability distributions of a UCTMC to the reachable probability distributions of a lumped UCTMC.

The reachable-set semantics gives a concrete operational view of the model. Indeed, reachable sets can be analyzed in two ways. The first is by computing the transient probabilities by sampling time-varying transition rate matrices $Q(t)$ satisfying $m_{i,j} \leq q_{i,j}(t) \leq M_{i,j}$ and solving the ODE in Eq. 1. The second way is to compute reachable intervals by formal under- and over-approximation, using well-established techniques for uncertain dynamical systems, of which Eq. 1 are an instance, such as those implemented in C2E2, Flow* or SpaceEx, see [2, 15, 23] and references therein. We remark, however, that these are methods applicable to nonlinear dynamical systems in general, of which UCTMCs are a specific instance. It is also possible to approximate UCTMCs by carefully chosen discrete-time MDPs (DTMDP), as discussed further below. In both cases, the presence of time-varying uncertainty may result in computationally challenging problems, thus further motivating the development of efficient reduction techniques.

3.3 CTMDP Semantics

As introduced in Section 1, a UCTMC can also be seen as an instance of a time-inhomogeneous continuous-time MDP (CTMDP). To see this, we consider a CTMDP with the scheduler model as in [27, Section 2.2], which can be intuitively described as follows. For a sufficiently small time step $h > 0$, a CTMDP that is in state $i \in \mathcal{V}$ 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 DTMC under action a_i at step k . Indeed, in the special case when the transition rates are time-invariant, the *time-homogeneous* CTMDP admit a characterization in terms of sojourn times and an embedded discrete time Markov chain, according to which the choice of action $a_i \in \mathcal{A}(i)$ upon entering state $i \in \mathcal{V}$ 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 [45]).

Under this model, the discussion in [27, 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 3. *For a given UCTMC (\mathcal{V}, m, M) , consider the CTMDP $(\mathcal{V}, \mathcal{A}, \mathcal{M})$ 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 $i \in \mathcal{V}$ 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 uniformly piecewise analytic functions $a : [0; \infty) \rightarrow \prod_{i \in \mathcal{V}} \mathcal{A}(i)$.

For such a CTMDP, the maximization (respectively, minimization) of the probability of reaching a state in block H at the time τ corresponds to the computation of the maximal (respectively, minimal) value of the reachable set from Definition 3.

3.4 Discrete-time Approximation of the CTMDP Semantics

In this appendix we present a discrete-time approximation of the CTMDP semantics which is of interest for a two-fold purpose. First, we show that the resulting DTMDP can be analyzed to obtain approximations of the maximal and minimal reachable probabilities for each state using dynamic programming. Second, in the proof of Theorem 9, this approximate DTMDP is used to relate the notion UCTMC lumping with the alternating probabilistic bisimulation of [31].

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 [27, Section 2.2], the foregoing example can be generalized as follows.

Proposition 1. For a given UCTMC (\mathcal{V}, m, M) , consider the CTMDP $(\mathcal{V}, \mathcal{A}', \mathcal{M}')$ 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 $i \in \mathcal{V}$ 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 $i \in \mathcal{V}$ 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 $[0; \infty) \rightarrow \prod_{i \in \mathcal{V}} \mathcal{D}(\mathcal{A}'(i))$.

Then, the policy sets \mathcal{M} and \mathcal{M}' , where \mathcal{M} refers to the policy set given in Theorem 3, induce the same set of time-inhomogeneous CTMCs.

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 [25], 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 make the key observation that the Euler method defines a time-inhomogeneous DTMC. Indeed, similarly to the discussion in Section 3.3, $I + hQ(kh, \mu(kh))$ describes the transition probability matrix of the embedded time-inhomogeneous DTMC.

Together with Theorem 3 and Proposition 1, the next result allows us to formally relate UCTMCs to time-inhomogeneous DTMDPs.

Theorem 4. *Given UCTMC (\mathcal{V}, m, M) , set*

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

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 $(\mathcal{V}, \mathcal{A}', \mathcal{M}'_h)$ given as:

- *The states are \mathcal{V} , while the actions in state $i \in \mathcal{V}$ 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 $\mathcal{M}'_h = \{\nu \mid \nu : \mathbb{N}_0 \rightarrow \prod_{i \in \mathcal{V}} \mathcal{D}(\mathcal{A}'(i))\}$. 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 $\tau > 0$ 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

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

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 $|kh - \tau|$.

Theorem 4 ensures that any $\tau \geq 0$ and any admissible transition rate matrix of the UCTMC can be matched by an approximate DTMDP such that the transition probabilities of both, the so-induced DTMC and the so-induced CTMC, are matching up to an ε at $\tau \geq 0$.

We state our first major result which relates reachability- and MDP-semantics.

Theorem 5. *For $\tau > 0$, a UCTMC $\mathcal{U} = (\mathcal{V}, m, M)$ and $H \subseteq \mathcal{V}$, let k be such that $\tau = kh$. Then, the maximal (minimal) probability for reaching a block H at τ coincides, by Theorem 3, with the maximum (minimum) of $\mathcal{R}_{\mathcal{U}}(H, \tau, \pi[0])$ from Definition 3 and can be computed in*

$$\mathcal{O}\left(k\left(\sum_{i \in \mathcal{V}} \deg_o^{all}(i)\right)\left(\sum_{i \in \mathcal{V}} 2^{\deg_o(i)}\right)\right),$$

where $\deg_o^{all}(i) = |\{j \neq i \mid 0 < M_{i,j}\}|$ is the number of outgoing transitions from state i , while $\deg_o(i) = |\{j \neq i \mid m_{i,j} < M_{i,j}\}|$ is the number of outgoing non-deterministic transitions from state i .

The complexity bound from Theorem 5 is polynomial in the number of states and exponential in $\max_i \deg_o(i)$, i.e., the maximal number of outgoing non-deterministic transitions of the approximate DTMDP.

4 UCTMC Lumpability

In Section 4.1 we prove that UCTMC lumpability characterizes the preservation of sums of reachable probability distributions. The logical characterization of UCTMC lumpability with respect to continuous stochastic logic is presented in Section 4.2. The UCTMC lumping algorithm is discussed in Section 4.3.

4.1 UCTMC Lumpability

Definition 4 (UCTMC Lumpability). *A partition \mathcal{H} of \mathcal{V} is a UCTMC lumping of UCTMC (\mathcal{V}, m, M) if it is an ordinary lumping of both CTMCs (\mathcal{V}, m) and (\mathcal{V}, M) . \square*

For instance, $\mathcal{H} = \{\{11\}, \{01, 10\}, \{11\}\}$ is a UCTMC lumping of the UCTMC from Figure 1. The lumped UCTMC is obtained in a similar way as for ordinary lumpability.

Definition 5 (Lumped UCTMC). *Assume that \mathcal{H} is a UCTMC lumping of (\mathcal{V}, m, M) and fix, for each $H \in \mathcal{H}$, some representative $i_H \in H$. The lumped UCTMC has states $\hat{\mathcal{V}} := \{i_H \mid H \in \mathcal{H}\}$ and bounds $\hat{m}_{i_H, i_{H'}} := \sum_{j \in H'} m_{i_H, j}$ and $\hat{M}_{i_H, i_{H'}} := \sum_{j \in H'} M_{i_H, j}$. \square*

Example. In the case of the UCTMC from Figure 1, the UCTMC lumping

$\mathcal{H} = \{\{11\}, \{01, 10\}, \{11\}\}$ induces the lumped UCTMC in Figure 2. Each state is labeled with a representative of the corresponding partition block. It is interesting to note that the transitions between states 01 and 10 in the original UCTMC correspond to self-loops in the lumped UCTMC. However, since self-loops induce self canceling terms at the level of forward Kolmogorov equations, they do not have an impact on system's dynamics and can be ignored.

Sums of reachable probability distributions of a UCTMC coincide with the reachable probability distributions of the corresponding lumped UCTMC.




Fig. 2. Lumped UCTMC.

Theorem 6 (Preservation of Reachability). *Assume that \mathcal{H} is a UCTMC lumping of $\mathcal{U} = (\mathcal{V}, m, M)$. Then, for any time $\tau \geq 0$, block $H \in \mathcal{H}$ and initial probability distribution $\pi[0]$, it holds that*

$$\mathcal{R}_{\mathcal{U}}(H, \tau, \pi[0]) = \mathcal{R}_{\hat{\mathcal{U}}}(\{i_H\}, \tau, \hat{\pi}[0]),$$

where $\hat{\mathcal{U}}$ refers to the lumped UCTMC induced by \mathcal{H} and $\hat{\pi}[0]_{i_H} = \sum_{i \in H} \pi[0]_i$ for all $H \in \mathcal{H}$.

Example. In the case of the running example, Theorem 6 ensures, for instance, that $\mathcal{R}_{\mathcal{U}}(\{10, 01\}, t, \pi[0]) = \mathcal{R}_{\hat{\mathcal{U}}}(\{10\}, t, \hat{\pi}[0])$ for all $t \geq 0$ and $\pi[0]$.

We next present a modification of Theorem 6 that allows one to over-approximate sums of reachable probability distributions when \mathcal{H} is not a UCTMC. It resembles [37] which provides over-approximations of uniformized CTMCs.

Theorem 7 (Over-Approximation). *For a given UCTMC $\mathcal{U} = (\mathcal{V}, m, M)$ and partition \mathcal{H} of \mathcal{V} , assume that*

- $m' \leq m$ such that \mathcal{H} is an ordinary lumping of the CTMC (\mathcal{V}, m') ;
- $M \leq M'$ such that \mathcal{H} is an ordinary lumping of the CTMC (\mathcal{V}, M') .

Then, \mathcal{H} is a UCTMC lumping of $\mathcal{U}' = (\mathcal{V}, m', M')$ and for any initial probability distribution $\pi[0]$, the lumped UCTMC $\hat{\mathcal{U}}'$ induced by \mathcal{U}' and \mathcal{H} satisfies

$$\mathcal{R}_{\mathcal{U}}(H, \tau, \pi[0]) \subseteq \mathcal{R}_{\hat{\mathcal{U}}'}(\{i_H\}, \tau, \hat{\pi}'[0])$$

for all $\tau \geq 0$ and $H \in \mathcal{H}$, provided that $\hat{\pi}'[0]_{i_H} = \sum_{i \in H} \pi[0]_i$ for all $H \in \mathcal{H}$.

Our next result is the converse of Theorem 6. Together with Theorem 6, it provides a quantitative characterization of UCTMC lumpability.

Theorem 8 (Quantitative Characterization). *Let $\mathcal{U} = (\mathcal{V}, m, M)$ be some UCTMC and $\hat{\mathcal{U}} = (\hat{\mathcal{V}}, \hat{m}, \hat{M})$ a UCTMC with $\hat{\mathcal{V}} = \{i_H \mid H \in \mathcal{H}\}$ where \mathcal{H} is a partition of \mathcal{V} and for any time $\tau \geq 0$, block $H \in \mathcal{H}$ and initial probability distribution $\pi[0]$, it holds that*

$$\mathcal{R}_{\mathcal{U}}(H, \tau, \pi[0]) = \mathcal{R}_{\hat{\mathcal{U}}}(\{i_H\}, \tau, \hat{\pi}[0])$$

whenever $\hat{\pi}[0]_{i_H} = \sum_{i \in H} \pi[0]_i$ for all $H \in \mathcal{H}$. Then, \mathcal{H} is a UCTMC lumping and $\hat{\mathcal{U}}$ the underlying lumped UCTMC.

Remark 2. By Theorem 8, the subset relation of Theorem 7 becomes an identity only if \mathcal{H} is a UCTMC lumping of (\mathcal{V}, m, M) and $m = m', M = M'$. In particular, over-approximations due to Theorem 7 are proper in general.

We end this section by relating UCTMC lumpability to other notions. First, we observe that UCTMC lumpability is a conservative generalization of ordinary lumpability.

Lemma 1 (Generalization). *Assume that \mathcal{H} is a UCTMC lumping of a UCTMC (\mathcal{V}, m, M) which is deterministic, i.e., $m = M$. Then, \mathcal{H} is an ordinary lumping.*

Second, we prove that any UCTMC admits a DTMDP approximation that discretizes time and action spaces and for which the notions of UCTMC lumpability and alternating probabilistic bisimulation (cf. [31]) coincide.

Theorem 9. *Fix a UCTMC (\mathcal{V}, m, M) , an equivalence relation $\mathfrak{R} \subseteq \mathcal{V} \times \mathcal{V}$ and let $\mathcal{H} = \mathcal{V}/\mathfrak{R}$. Then \mathcal{H} is a UCTMC lumpability of (\mathcal{V}, m, M) if and only if \mathfrak{R} is an alternating probabilistic bisimulation of the DTMDP from Theorem 4.*

As mentioned earlier, alternating probabilistic bisimulation only preserves logical and quantitative properties [31] on the domain of DTMDPs, while UCTMC lumping characterizes these on the domain of UCTMCs.

4.2 Logical Characterization

We extend CSL to UCTMCs by defining a formula to be true when it is satisfied by all admissible $Q = (q_{i,j})_{(i,j)}$. This allows one to study safety properties in presence of uncertainty, aligning with [44], which considers CSL for CTMDPs with finite action spaces.

Definition 6 (CSL for UCTMCs). *Given a UCTMC (\mathcal{V}, m, M) , 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:

- $i, t \models a$ 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 admissible q ;

– $i, t \models \mathcal{P}_{\bowtie p}^{\forall}(\phi_1 U^{[t_0; t_1]} \phi_2)$ iff $i, t \models \mathcal{P}_{\bowtie p}(\phi_1 U^{[t_0; t_1]} \phi_2)$ for all admissible q . \square

Similarly to [44], existential quantification is given by $\mathcal{P}_{\bowtie p}^{\exists}(\Phi) := \neg \mathcal{P}_{\neg \bowtie p}^{\forall}(\Phi)$, where $\neg \bowtie$ is defined in the obvious manner (e.g., $\neg \leq$ is $>$). Likewise, \vee, \rightarrow are defined using \wedge, \neg .

Theorem 10 (Preservation of CSL). *Let \mathcal{H} be a UCTMC lumping 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 $\pi[0]$.

The next result is a converse of Theorem 10 and establishes a logical characterization of UCTMC lumpability.

Theorem 11 (Logical Characterization). *Fix a UCTMC (\mathcal{V}, m, M) , a partition \mathcal{H} of \mathcal{V} and let $\mathcal{L}, \hat{\mathcal{A}}$ and $\hat{\mathcal{L}}$ be as in Theorem 10. Assume further that there exists a UCTMC $(\hat{\mathcal{V}}, \hat{m}, \hat{M})$ such that $\hat{\mathcal{V}} = \{i_H \mid H \in \mathcal{H}\}$ and*

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

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 lumping and $(\hat{\mathcal{V}}, \hat{m}, \hat{M})$ the underlying lumped UCTMC.

4.3 UCTMC Lumping Algorithm

We next present an algorithm for the efficient computation of the coarsest UCTMC lumping that refines a given partition \mathcal{H} . Its steps are as follows.

- A1 With \mathcal{H} being the current partition, compute the coarsest ordinary lumping \mathcal{H}' of the CTMC (\mathcal{V}, m) that refines \mathcal{H} ;
- A2 Compute the coarsest ordinary lumping \mathcal{H}'' of the CTMC (\mathcal{V}, M) that refines \mathcal{H}' ;
- A3 If $\mathcal{H}'' = \mathcal{H}$, return \mathcal{H}'' ; Otherwise, set $\mathcal{H} := \mathcal{H}''$ and go to A1.

Obviously, if A1 does not refine \mathcal{H} and A2 does not refine \mathcal{H}' , then \mathcal{H} is a UCTMC lumping of (\mathcal{V}, m, M) . The algorithm terminates because \mathcal{V} 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 lumping.

The next result summarizes the above discussion. The complexity statement follows thanks to the fact that A1 and A2 can be processed via efficient CTMC lumping algorithms such as [20, 52].

Theorem 12. *Given a UCTMC (\mathcal{V}, m, M) , let \mathcal{H} be a partition of \mathcal{V} . Then, the following can be shown.*

Algorithm 1 Partition refinement algorithm for the computation of the coarsest UCTMC lumping \mathcal{H} from the proof of Theorem 12.

Require: Uncertain CTMC (\mathcal{V}, m, M) and initial partition \mathcal{H}

```

1: while true do
2:    $\mathcal{H}' \leftarrow$  coarsest ordinary lumping of CTMC  $(\mathcal{V}, m)$  that refines  $\mathcal{H}$ 
3:    $\mathcal{H}'' \leftarrow$  coarsest ordinary lumping of CTMC  $(\mathcal{V}, M)$  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

```

- 1) *Algorithm 1 computes the coarsest UCTMC lumping refining \mathcal{H} .*
- 2) *The time and space complexity required for one while loop iteration does not exceed $\mathcal{O}(r \log(s))$, where $r := |\{(i, j) \mid m_{i,j} > 0 \text{ or } M_{i,j} > 0\}|$ and $s := |\mathcal{V}|$. The number of while loop iterations, instead, is at most s .*

We conclude this section with two remarks regarding the lumping algorithm. First, we note that it simplifies to the CTMC lumping algorithm if applied to a deterministic UCTMC (\mathcal{V}, m, M) , i.e., a UCTMC that satisfies $m = M$. Second, using the correspondence between UCTMC lumpability and probabilistic alternating bisimulation from Theorem 9, it would be possible to apply the algorithm for the largest alternating probabilistic bisimulation [31] of the approximate DTMDP from Theorem 9. However, in contrast to the UCTMC algorithm, such an approach would incur an exponential dependence on the maximal number of outgoing non-deterministic transitions of the approximate DTMDP.

5 Evaluation

Here we assess UCTMC lumpability in terms of both its computational tractability and reduction power with respect to ordinary lumpability. To this end, we consider uncertain variants of CTMCs of increasing size generated from benchmark models in PRISM [40].


Tool-support and replicability. In our experiments we used a prototype implementation of our algorithm based on the tool ERODE [12]. ERODE supports CTMC minimization as a special case of lumping algorithms for nonlinear ordinary differential equations [13]. Given that CTMC ordinary lumpability is the most important inner step of our algorithm, other tools implementing CTMC lumping could have been used, such as MRMC [36], STORM [18], and CoPaR [22, 19]. All runtimes reported refer to the execution of ERODE on a common desktop machine with 8 GB RAM. All the material to replicate the experiments is available at <https://www.erode.eu/examples.html>.

Set-up. For this evaluation we used CTMCs in the MRMC format [36], generated from PRISM models. We considered CTMCs which describe: a dependable

Original model (CTMC)				CTMC Lumpability		UCTMC Lumpability	
N	r	s	$ \mathcal{H}_0 $	Red.(s)	$ \mathcal{H} $	Red.(s)	$ \mathcal{H} $
WORKSTATION CLUSTER							
128	2 908 192	597 012	4	2.21E+1	298 893	2.64E+1	identical
192	6 524 960	1 337 876	4	6.78E+1	669 517	8.04E+1	identical
256	11 583 520	2 373 652	4	1.55E+2	1 187 597	1.85E+2	identical
320	18 083 872	3 704 340	4	2.81E+2	1 853 133	3.58E+2	identical
384	26 026 016	5 329 940	4	out of memory		out of memory	
WIRELESS GROUP COMMUNICATION PROTOCOL							
16	686 153	103 173	2	2.26E+0	4 846	2.97E+0	identical
24	3 183 849	453 125	2	1.34E+1	20 476	1.61E+1	identical
32	10 954 382	1 329 669	2	4.49E+1	58 906	5.50E+1	identical
40	22 871 849	3 101 445	2	1.35E+2	135 752	1.61E+2	identical
48	46 574 793	6 235 397	2	out of memory		out of memory	
CELL CYCLE CONTROL IN EUKARYOTES							
2	18 342	4 666	3	1.76E-1	3 514	1.97E-1	4 000
3	305 502	57 667	3	8.21E-1	40 667	9.81E-1	48 147
4	2 742 012	431 101	3	6.45E+0	282 956	7.80E+0	339 368
5	16 778 785	2 326 666	3	4.58E+1	1 424 935	9.15E+1	1 712 322
6	78 768 799	9 960 861	3	out of memory		out of memory	

Table 1. Quantitative comparison of CTMC and UCTMC lumpability. Entries *identical* denote cases with identical CTMC and UCTMC lumpable partitions.

cluster of workstations [32]; a protocol for wireless group communication [43, 3]; a model of the cell cycle control in eukaryotes [41, 46]. in [31, 28], we considered uncertain relaxations of such CTMCs by adding uncertainty to the transition rates. In particular, in each model we replaced every transition rate value with an interval of fixed length (arbitrarily fixed equal to 20% of the smallest transition rate in the model) centered at the original rate value itself. This relaxation is exemplified in the right inset, where the grayed-out transition rates are those of some given CTMC. This example demonstrates that an ordinary lumping of the original CTMC is not necessarily a UCTMC lumping of the so-constructed UCTMC, whereas the converse follows from Lemma 1. Indeed, we note that $\{\{A, B\}, \{C, D\}\}$ is an ordinary lumping of the original CTMC, while it is not a UCTMC lumping because $4.8 \neq 1.8 + 2.8$ and $5.2 \neq 2.2 + 3.2$.



Results. The results are provided in Table 1. We report the number of transitions and states of the obtained CTMCs in the second and third column, respectively,

<i>CTMC lumpability</i>		<i>UCTMC lumpability</i>
$\Downarrow [38, 11, 6]$	Dynamics	\Downarrow Thm. 6, 8
$\sum_{i \in H} \pi_i(t) = \hat{\pi}_{i_H}(t)$		$\mathcal{R}(\sum_{i \in H} \pi_i, t, \pi[0]) = \mathcal{R}(\hat{\pi}_{i_H}, t, \hat{\pi}[0])$
$\Downarrow [1]$	Logics	\Downarrow Thm. 10, 11
$i, t \models_V \phi \iff i_H, t \models_{\hat{V}} \phi$		$i, t \models_{V, m, M} \phi \iff i_H, t \models_{\hat{V}, \hat{m}, \hat{M}} \phi$
$\mathcal{O}(r \log(s))$ [20, 52]	Complexity	$\mathcal{O}(sr \log(s))$ Thm. 12

Table 2. Summary of results. UCTMC lumpability generalizes the well-known dynamical, logical and algorithmic properties of ordinary lumpability for CTMCs (in statements concerning complexity, s refers to the numbers of states, while r denotes the number of transitions).

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. The comparison of the runtimes of the minimization 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 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 lumping can be evaluated by comparing the size of the coarsest UCTMC lumpings with their corresponding CTMC counterparts. Notably, CTMC and UCTMC lumpability coincide on the first two families of models, while in the third family UCTMC lumpability leads to finer (at most 18% more blocks) partitions than the CTMC counterpart.

6 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. UCTMC lumpability enjoys a polynomial time and space algorithm for the computation of the largest UCTMC lumping. Similarly to CTMC lumping that characterizes the preservation of sums of probability distributions, UCTMC lumping characterizes the preservations of reachable sets of sums of probability distributions. We have provided a logical characterization of UCTMC lumpability to uncertain time-varying parameters. Overall, the results in this paper can be put in analogy with the corresponding CTMC counterparts, as summarized in Table 2. The applicability of UCTMC lumpability has been established by presenting substantial reductions in benchmark models. The discretization of a UCTMC as a DTMDP has offered the means to relating UCTMC lumpability to bisimulations for DTMDPs. Future work will consider model-checking algorithms for UCTMCs.

References

1. Christel Baier, Boudewijn R. Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Model-Checking Algorithms for Continuous-Time Markov Chains. *IEEE Trans. Software Eng.*, 29(6):524–541, 2003.
2. Sergiy Bogomolov, Goran Frehse, Radu Grosu, Hamed Ladan, Andreas Podelski, and Martin Wehrle. A Box-Based Distance between Regions for Guiding the Reachability Analysis of SpaceEx. In *CAV*, pages 479–494, 2012.
3. Andrea Bondavalli, Andrea Coccoli, and Felicita Di Giandomenico. *QoS Analysis of Group Communication Protocols in Wireless Environment*, pages 169–188. Springer US, Boston, MA, 2002.
4. Luca Bortolussi and Jane Hillston. Fluid Model Checking. In *CONCUR*, pages 333–347, 2012.
5. Luca Bortolussi and Jane Hillston. Model checking single agent behaviours by fluid approximation. *Inf. Comput.*, 242:183–226, 2015.
6. Peter Buchholz. Exact and ordinary lumpability in finite Markov chains. *Journal of Applied Probability*, 31(1):59–75, 1994.
7. Peter Buchholz, Ernst Moritz Hahn, Holger Hermanns, and Lijun Zhang. Model Checking Algorithms for CTMDPs. In *CAV*, pages 225–242, 2011.
8. Peter Buchholz and Ingo Schulz. Numerical analysis of continuous time markov decision processes over finite horizons. *Comput. Oper. Res.*, 38(3):651–659, 2011.
9. Yuliya Butkova, Hassan Hatefi, Holger Hermanns, and Jan Krcál. Optimal continuous time Markov decisions. In *ATVA*, pages 166–182, 2015.
10. Benoît Caillaud, Benoît Delahaye, Kim G. Larsen, Axel Legay, Mikkel L. Pedersen, and Andrzej Wasowski. Constraint Markov chains. *Theoretical Computer Science*, 412(34):4373 – 4404, 2011.
11. Luca Cardelli, Mirco Tribastone, Max Tschaikowski, and Andrea Vandin. Symbolic computation of differential equivalences. In *POPL*, pages 137–150, 2016.
12. Luca Cardelli, Mirco Tribastone, Max Tschaikowski, and Andrea Vandin. ERODE: A tool for the evaluation and reduction of ordinary differential equations. In *TACAS*, 2017.
13. Luca Cardelli, Mirco Tribastone, Max Tschaikowski, and Andrea Vandin. Maximal aggregation of polynomial dynamical systems. *Proceedings of the National Academy of Sciences*, 114(38):10029–10034, 2017.
14. Milan Česka, Frits Dannenberg, Nicola Paoletti, Marta Kwiatkowska, and Luboš Brim. Precise parameter synthesis for stochastic biochemical systems. *Acta Informatica*, 54(6):589–623, 2017.
15. Xin Chen, Erika Ábrahám, and Sriram Sankaranarayanan. Flow*: An Analyzer for Non-linear Hybrid Systems. In *CAV*, pages 258–263, 2013.
16. René David and Hassane Alla. *Discrete, Continuous, and Hybrid Petri Nets*. Springer, 2005.
17. Christian Dehnert, Sebastian Junges, Nils Jansen, Florian Corzilius, Matthias Volk, Harold Brountjes, Joost-Pieter Katoen, and Erika Ábrahám. PROPhESY: A PRObabilistic ParamETER SYnthesis Tool. In *CAV*, pages 214–231, 2015.
18. Christian Dehnert, Sebastian Junges, Joost-Pieter Katoen, and Matthias Volk. A storm is coming: A modern probabilistic model checker. In *Computer Aided Verification - 29th International Conference, CAV 2017, Heidelberg, Germany, July 24-28, 2017, Proceedings, Part II*, pages 592–600, 2017.
19. Hans-Peter Deifel, Stefan Milius, Lutz Schröder, and Thorsten Wißmann. Generic partition refinement and weighted tree automata. In *FM*, 2019. To Appear.

20. Salem Derisavi, Holger Hermanns, and William H. Sanders. Optimal state-space lumping in markov chains. *Inf. Process. Lett.*, 87(6):309–315, 2003.
21. Josu Doncel, Nicolas Gast, Mirco Tribastone, Max Tschaikowski, and Andrea Vandin. UTOPIA: under-approximation through optimal control. In David Parker and Verena Wolf, editors, *QEST*, volume 11785, pages 277–291. Springer, 2019.
22. Ulrich Dorsch, Stefan Milius, Lutz Schröder, and Thorsten Wißmann. Efficient coalgebraic partition refinement. In *CONCUR*, pages 32:1–32:16, 2017.
23. Chuchu Fan, Bolun Qi, Sayan Mitra, Mahesh Viswanathan, and Parasara Sridhar Duggirala. Automatic Reachability Analysis for Nonlinear Hybrid Models with C2E2. In *CAV*, pages 531–538, 2016.
24. Harald Fecher, Martin Leucker, and Verena Wolf. *Don't Know* in probabilistic systems. In *SPIN*, pages 71–88, 2006.
25. C. William Gear. *Numerical Initial Value Problems in Ordinary Differential Equations*. Prentice Hall PTR, 1971.
26. D.T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, 81(25):2340–2361, December 1977.
27. Xianping Guo and Onesimo Hernandez-Lerma. *Continuous-Time Markov Decision Processes*. Springer Verlag, 2009.
28. Ernst Moritz Hahn, Vahid Hashemi, Holger Hermanns, and Andrea Turrini. Exploiting robust optimization for interval probabilistic bisimulation. In Gul Agha and Benny Van Houdt, editors, *Quantitative Evaluation of Systems*, pages 55–71, Cham, 2016. Springer International Publishing.
29. Ernst Moritz Hahn, Holger Hermanns, Ralf Wimmer, and Bernd Becker. Transient Reward Approximation for Continuous-Time Markov Chains. *IEEE Trans. Reliability*, 64(4):1254–1275, 2015.
30. Ernst Moritz Hahn, Holger Hermanns, and Lijun Zhang. Probabilistic reachability for parametric Markov models. *STTT*, 13(1):3–19, 2011.
31. Vahid Hashemi, Andrea Turrini, Ernst Moritz Hahn, Holger Hermanns, and Khaled M. Elbassioni. Polynomial-Time Alternating Probabilistic Bisimulation for Interval MDPs. In *SETTA*, pages 25–41, 2017.
32. B.R. Haverkort, H. Hermanns, and J.-P. Katoen. On the use of model checking techniques for dependability evaluation. In *SRDS*, pages 228–237, 2000.
33. Holger Hermanns and Michael Rettelsbach. Syntax, semantics, equivalences, and axioms for MTIPP. In *Proceedings of Process Algebra and Probabilistic Methods*, pages 71–87, Erlangen, 1994.
34. Jane Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press, 1996.
35. Bengt Jonsson and Kim G. Larsen. Specification and refinement of probabilistic processes. In *LICS*, pages 266–277, 1991.
36. J.-P. Katoen, M. Khattri, and I.S. Zapreev. A Markov Reward Model Checker. In *QEST*, pages 243–244, 2005.
37. Joost-Pieter Katoen, Daniel Klink, Martin Leucker, and Verena Wolf. Three-valued abstraction for continuous-time Markov chains. In *CAV*, pages 311–324, 2007.
38. J.G. Kemeny and J.L. Snell. *Finite Markov Chains*. Springer New York, Heidelberg, Berlin, 1976.
39. Igor O. Kozine and Lev V. Utkin. Interval-valued finite Markov chains. *Reliable Computing*, 8(2):97–113, Apr 2002.
40. M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: Verification of Probabilistic Real-time Systems. In *CAV*, pages 585–591, 2011.
41. Paola Lecca and Corrado Priami. Cell cycle control in eukaryotes: A biospi model. *Electr. Notes Theor. Comput. Sci.*, 180(3):51–63, 2007.

42. Daniel Liberzon. *Calculus of Variations and Optimal Control Theory: A Concise Introduction*. Princeton University Press, Princeton, NJ, USA, 2011.
43. M. Massink, J.-P. Katoen, and D. Latella. Model checking dependability attributes of wireless group communication. In *DSN*, pages 711–720, 2004.
44. Martin R. Neuhäuser and Joost-Pieter Katoen. Bisimulation and Logical Preservation for Continuous-Time Markov Decision Processes. In *CONCUR*, pages 412–427, 2007.
45. James R. Norris. *Markov chains*. Cambridge University Press, 1998.
46. 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*, 353(1378):2063–2076, 1998.
47. Koushik Sen, Mahesh Viswanathan, and Gul Agha. Model-checking Markov chains in the presence of uncertainties. In *TACAS*, pages 394–410, 2006.
48. Péter L. Simon, Michael Taylor, and Istvan Z. Kiss. Exact epidemic models on graphs using graph-automorphism driven lumping. *Journal of Mathematical Biology*, 62(4):479–508, 2010.
49. Lei Song, Lijun Zhang, and Jens Chr. Godskesen. Bisimulations and Logical Characterizations on Continuous-Time Markov Decision Processes. In *VMCAI*, pages 98–117, 2014.
50. F. M. Spieksma. Markov decision processes. 2015. <http://www.math.leidenuniv.nl/spieksma/colleges/besliskunde/MDP.pdf>.
51. William J. Stewart. *Probability, Markov Chains, Queues, and Simulation*. Princeton University Press, 2009.
52. Antti Valmari and Giuliana Franceschinis. Simple $O(m \log n)$ time Markov chain lumping. In *TACAS*, pages 38–52, 2010.