

Model checking continuous-time Markov chains by transient analysis

CHRISTEL BAIER^a, BOUDEWIJN HAVERKORT^b,
HOLGER HERMANN^c AND JOOST-PIETER KATOEN^{c,*}

^a*Institut für Informatik I, University of Bonn
Römerstraße 164, D-53117 Bonn, Germany*

^b*Dept. of Computer Science, RWTH Aachen
Ahornstraße 55, D-52056 Aachen, Germany*

^c*Dept. of Computer Science, University of Twente
P.O. Box 217, 7500 AE Enschede, The Netherlands*

Abstract. The verification of continuous-time Markov chains (CTMCs) against continuous stochastic logic (CSL) [3, 6], a stochastic branching-time temporal logic, is considered. CSL facilitates among others the specification of steady-state properties and the specification of probabilistic timing properties of the form $\mathcal{P}_{\bowtie p}(\Phi_1 U^I \Phi_2)$, for state formulas Φ_1 and Φ_2 , comparison operator \bowtie , probability p , and real interval I . The main result of this paper is that model checking probabilistic timing properties can be reduced to the problem of computing transient state probabilities for CTMCs. This allows us to verify such properties by using efficient techniques for transient analysis of CTMCs such as uniformisation. A second result is that a variant of ordinary lumping equivalence (i.e., bisimulation), a well-known notion for aggregating CTMCs, preserves the validity of all CSL-formulas.

1 Introduction

Continuous-time Markov chains (CTMCs) have been widely used to determine important system performance and reliability characteristics. To mention just a few applications, these models have been used to quantify the throughput of production lines, to determine the mean time between failure in safety-critical systems, or to identify bottlenecks in communication networks. Due to the rapidly increasing size and complexity of systems, obtaining such models in a direct way becomes more and more cumbersome and error-prone. An effective solution to this problem is to generate CTMCs from higher-level specifications, like queueing networks, stochastic Petri nets [1], or stochastic process algebras [20, 24].

Although these approaches have shown to be rather valuable — several (industrial) case studies have been carried out and mature tool-support is available [19, 21] — the specification of the measure of interest is mostly done informally.

* Contact author. Tel.: +31-53-4895675, fax: +31-53-4893247, e-mail: katoen@cs.utwente.nl.

The analysis of the CTMC most often boils down to the determination of *steady-state* and *transient state probabilities*. Steady-state probabilities refer to the system behaviour on the “long run” while the transient probabilities consider the system at a fixed time instant t .

In [3] measures of interest of CTMCs are specified in a branching-time logic **CSL** (*continuous stochastic logic*) that includes a **TCTL**-like time-bounded until operator \mathcal{U}^I , where I is a time-interval, and a probabilistic operator $\mathcal{P}_{\bowtie p}(\cdot)$ to reason about the probabilities of timing properties. As in the logic **PCTL** [17], a probabilistic variant of **CTL** interpreted over DTMCs, the operator $\mathcal{P}_{\bowtie p}(\varphi)$ replaces the usual **CTL** path quantifiers \forall and \exists and refers to the probability for the event specified by the path formula φ . The subscript $\bowtie p$ (where \bowtie is a comparison operator and $p \in [0, 1]$) specifies a lower or upper bound for the “allowed” probabilities. The combination of the probabilistic operator with the temporal operator $\diamond^{[t, t]}$ (which can be derived from the time-bounded until operator) analyses the quantitative behaviour at time instant t and can be used to reason about transient probabilities. For instance, $\mathcal{P}_{\leq 0.001}(\diamond^{[4, 4]} \text{error})$ asserts that the probability for a system error at time instant 4 is at most 10^{-3} .

In [6], **CSL** was extended by the usual next step and until operator and by a novel steady-state operator, e.g. the formula $\mathcal{S}_{\geq 0.98}(up)$ asserts that the steady-state probability for the system “being up” is at least 0.98. Moreover, [6] presented a model checking algorithm for the extended version of **CSL** that uses a variant of multi-terminal BDDs [12, 4]; thus, obtaining a single framework that combines the traditional approach of steady-state and transient analysis of CTMCs with the symbolic BDD-based model checking approach for temporal logics.

While [6] focuses on model checking with techniques that have been proven to be very efficient for non-stochastic systems, viz. BDDs, in this paper we investigate the complementary question and present a **CSL** model checking algorithm that operates with well-understood efficient techniques for analyzing CTMCs, namely *transient analysis* of CTMCs represented by sparse matrices. The main difficulty is the treatment of $\mathcal{P}_{\bowtie p}(\varphi)$ applied to a path formula φ of the form $\Phi_1 \mathcal{U}^I \Phi_2$.¹ Our main result states that, for a given CTMC \mathcal{M} and state s in \mathcal{M} , the measure $Prob^{\mathcal{M}}(s, \varphi)$ for the event that φ holds when the system starts in state s , can be calculated by means of a transient analysis of the CTMC \mathcal{M}' , which can easily be derived from \mathcal{M} . This allows us to adopt efficient techniques for performing transient analysis of CTMCs, like *uniformisation* [15, 16, 25, 27], for model checking probabilistic timing properties. In addition, we show that (ordinary) lumping-equivalence — a notion on Markov chains to aggregate state spaces [10, 24] that can be viewed as a continuous variant of probabilistic bisimulation [26] — preserves the validity of all **CSL**-formulas. This allows us to switch from the original state space to the (possibly much smaller) quotient space under lumping equivalence. Using this property, we indicate how the state

¹ The steady-state operator and the probabilistic operator applied to next step or (unbounded) until require essentially matrix operations like multiplication and solving linear equation systems that can be treated by standard tools for sparse matrices.

space for checking probabilistic timing properties on the derived CTMC \mathcal{M}' can be obtained.

Organisation of the paper. Section 2 introduces CTMCs and **CSL**. Section 3 presents a reduction of the model checking problem for time-bounded until to a transient analysis on CTMCs. Section 4 discusses lumping equivalence and preservation of **CSL**-properties. Section 5 reports on the checking of properties on a large plain-old telephone system [21]. Section 6 concludes the paper.

2 CTMCs and CSL

In this section, we briefly recall the basic concepts of CTMCs [28] and the logic **CSL** [3, 6]. We slightly depart from the standard notations for CTMCs and consider a CTMC as an ordinary transition system (Kripke structure) where the edges are equipped with probabilistic timing information. Let AP be a fixed, finite set of atomic propositions.

CTMCs. A (labelled) CTMC \mathcal{M} is a tuple (S, \mathbf{R}, L) where S is a finite set of *states*, $\mathbf{R} : S \times S \rightarrow \mathbb{R}_{\geq 0}$ the *rate matrix*², and $L : S \rightarrow 2^{AP}$ the *labelling* function which assigns to each state $s \in S$ the set $L(s)$ of atomic propositions $a \in AP$ that are valid in s . A state s is called *absorbing* iff $\mathbf{R}(s, s') = 0$ for all states s' . We assume that for any state s , AP contains an atomic proposition a_s which is characteristic for s , i.e., $a_s \in L(s)$ and $a_s \notin L(s')$ for any $s' \neq s$.

Intuitively, $\mathbf{R}(s, s') > 0$ iff there is a transition from s to s' ; $1 - e^{-\mathbf{R}(s, s') \cdot t}$ is the probability that the transition $s \rightarrow s'$ can be triggered within t time units. Thus, the delay of transition $s \rightarrow s'$ is governed by an exponential distribution with rate $\mathbf{R}(s, s')$. If $\mathbf{R}(s, s') > 0$ for more than one state s' , a competition between the transitions originating in s exists, known as the *race condition*. The probability to move from non-absorbing state s to a particular state s' within t time units, i.e., $s \rightarrow s'$ wins the race, is given by

$$\mathbf{P}(s, s', t) = \frac{\mathbf{R}(s, s')}{\mathbf{E}(s)} \cdot \left(1 - e^{-\mathbf{E}(s) \cdot t}\right)$$


where $\mathbf{E}(s) = \sum_{s' \in S} \mathbf{R}(s, s')$ denotes the *total rate* at which any transition emanating from state s is taken. More precisely, $\mathbf{E}(s)$ specifies that the probability of leaving s within t time-units is $1 - e^{-\mathbf{E}(s) \cdot t}$, due to the fact that the minimum of exponential distributions (competing in a race) is characterised by the sum of their rates. Consequently, the probability of moving from a non-absorbing state s to s' by a single transition, denoted $\mathbf{P}(s, s')$, is determined by the probability that the delay of going from s to s' finishes before the delays of other outgoing

² We do not set $\mathbf{R}(s, s) = -\sum_{s' \neq s} \mathbf{R}(s, s')$, as is usual for CTMCs. In our setting, self-loops at a state s are possible and can be modelled by $\mathbf{R}(s, s) > 0$. The inclusion of self-loops does neither alter the transient nor the steady-state behaviour of the CTMC, but allows the usual interpretation of linear-time temporal operators like next step and unbounded or time-bounded until.

edges from s ; formally, $\mathbf{P}(s, s') = \mathbf{R}(s, s')/\mathbf{E}(s)$. For an absorbing state s , the total rate $\mathbf{E}(s)$ is 0. (In this case, we have $\mathbf{P}(s, s') = 0$ for any state s' .)

The initial state probabilities of $\mathcal{M} = (S, \mathbf{R}, L)$ are given by an *initial distribution* $\alpha : S \rightarrow [0, 1]$ with $\sum_{s \in S} \alpha(s) = 1$. In case we have a unique initial state s , the initial distribution is denoted α_s^1 , where $\alpha_s^1(s) = 1$ and $\alpha^1(s') = 0$ for any $s' \neq s$.

Example 1. As a running example we address a *triple modular redundant system* (TMR) taken from [18], a fault-tolerant computer system consisting of three processors and a single (majority) voter that we model by a CTMC where state $s_{i,j}$ models that i processors and j voters are operational. Initially all components are functioning correctly (i.e., $\alpha = \alpha_{s_{3,1}}^1$). The failure rate of a processor is λ and of the voter ν failures per hour (fph). The expected repair time of a processor is $1/\mu$ and of the voter $1/\delta$ hours. The system is operational if at least two processors and the voter are functioning correctly. If the voter fails, the entire system is assumed to have failed, and after a repair (with rate δ) the system is assumed to start “as good as new”. The details of the CTMC are:



We have e.g., $\mathbf{P}(s_{2,1}, s_{3,1}) = \mu/(\mu + 2\lambda + \nu)$ and $\mathbf{P}(s_{0,1}, s_{0,0}) = \nu/(\mu + \nu)$. ■

Paths. A *path* σ is a finite or infinite sequence $s_0, t_0, s_1, t_1, s_2, t_2, \dots$ with, for $i \in \mathbb{N}$, $s_i \in S$ and $t_i \in \mathbb{R}_{>0}$ such that $\mathbf{R}(s_i, s_{i+1}) > 0$, if σ is infinite. For an infinite path σ and $i \in \mathbb{N}$ let $\sigma[i] = s_i$, the i -th state of σ , and $\delta(\sigma, i) = t_i$, the time spent in s_i . For $t \in \mathbb{R}_{\geq 0}$ and i the smallest index i with $t \leq \sum_{j=0}^i t_j$ let $\sigma@t = \sigma[i]$, the state of σ at time t . If σ is finite and ends in s_l , we require that s_l is absorbing, and $\mathbf{R}(s_i, s_{i+1}) > 0$ for all $i < l$. For finite σ , $\sigma[i]$ and $\delta(\sigma, i)$ are defined for $i \leq l$ in the above way, whereas $\delta(\sigma, l) = \infty$, and $\sigma@t = s_l$ for $t > \sum_{j=0}^{l-1} t_j$. Let *Path* denote the set of paths in \mathcal{M} , *Path*(s) the set of paths starting in s .

Borel space. An initial distribution α yields a probability measure Pr_α on paths as follows. Let $s_0, \dots, s_k \in S$ with $\mathbf{R}(s_i, s_{i+1}) > 0$, ($0 \leq i < k$), and I_0, \dots, I_{k-1} non-empty intervals in $\mathbb{R}_{\geq 0}$. Then, $C(s_0, I_0, \dots, I_{k-1}, s_k)$ denotes the *cylinder set* consisting of all paths $\sigma \in \text{Path}(s_0)$ such that $\sigma[i] = s_i$ ($i \leq k$), and $\delta(\sigma, i) \in I_i$ ($i < k$). Let $\mathcal{F}(\text{Path})$ be the smallest σ -algebra on *Path* which contains all sets $C(s, I_0, \dots, I_{k-1}, s_k)$ where s_0, \dots, s_k ranges over all state-sequences with $s = s_0$, $\mathbf{R}(s_i, s_{i+1}) > 0$ ($0 \leq i < k$), and I_0, \dots, I_{k-1} ranges over all sequences of non-empty intervals in $\mathbb{R}_{\geq 0}$. The probability measure Pr_α on $\mathcal{F}(\text{Path})$ is the unique

measure defined by induction on k by $\Pr_\alpha(C(s_0)) = \alpha(s_0)$ and for $k \geq 0$:

$$\Pr_\alpha(C(s_0, \dots, s_k, I', s')) = \Pr_\alpha(C(s_0, \dots, s_k)) \cdot \mathbf{P}(s_k, s') \cdot \left(e^{-\mathbf{E}(s_k) \cdot a} - e^{-\mathbf{E}(s_k) \cdot b} \right)$$

where $a = \inf I'$ and $b = \sup I'$. (For $b = \infty$ and $\lambda > 0$ let $e^{-\lambda \cdot \infty} = 0$.)

Remark 1. For an infinite path $\sigma = s_0, t_0, s_1, t_1, \dots$ we do not assume *time divergence*. Although $\sum_{j \geq 0} t_j$ might converge, in which case σ represents an “unrealistic” computation where infinitely many transitions are taken in a finite amount of time, the probability measure of such non-time-divergent paths is 0 (independent of α). This allows a lazy treatment of the notation $\sigma@t$ in the description of measurable sets of paths for which we just refer to the probability measure. ■

Steady state and transient probabilities. For a CTMC two major types of state probabilities are distinguished: steady-state probabilities where the system is considered “on the long run” i.e., when an equilibrium has been reached, and transient probabilities where the system is considered at a given time instant t . Formally, the transient probability

$$\pi^{\mathcal{M}}(\alpha, s', t) = \Pr_\alpha\{\sigma \in Path \mid \sigma@t = s'\}$$

stands for the probability to be in state s' at time t given the initial distribution α . Steady-state probabilities are defined as $\pi^{\mathcal{M}}(\alpha, s') = \lim_{t \rightarrow \infty} \pi^{\mathcal{M}}(\alpha, s', t)$. This limit always exists for finite CTMCs. For $S' \subseteq S$, let $\pi^{\mathcal{M}}(\alpha, S') = \sum_{s' \in S'} \pi^{\mathcal{M}}(\alpha, s')$ denote the steady-state probability for S' given α , i.e.,

$$\pi^{\mathcal{M}}(\alpha, S') = \lim_{t \rightarrow \infty} \Pr_\alpha\{\sigma \in Path \mid \sigma@t \in S'\}.$$

We let $\pi^{\mathcal{M}}(\alpha, \emptyset) = 0$. We often omit the superscript \mathcal{M} if the CTMC \mathcal{M} is clear from the context. In case of a unique initial state s , i.e., $\alpha = \alpha_s^1$, we write \Pr_s for \Pr_α , $\pi(s, s', t)$ for $\pi(\alpha, s', t)$, and $\pi(s, s')$ for $\pi(\alpha, s')$.

Syntax of CSL. CSL is a branching-time temporal logic where the state-formulas are interpreted over states of a CTMC [3, 6]. As in [6] we consider the extension of CSL of [3] with $\mathcal{S}_{\bowtie p}(\cdot)$ to reason about steady-state probabilities. We generalise CSL as defined in [6] with a time-bounded until operator that is parametrized by an arbitrary time-interval I . Let $a \in AP$, $p \in [0, 1]$ and $\bowtie \in \{\leq, \geq\}$. The state-formulas of CSL are defined by:

$$\Phi ::= tt \mid a \mid \Phi \wedge \Phi \mid \neg \Phi \mid \mathcal{S}_{\bowtie p}(\Phi) \mid \mathcal{P}_{\bowtie p}(\varphi)$$

where, for interval $I \subseteq \mathbb{R}_{\geq 0}$, path-formulas are defined by:

$$\varphi ::= X\Phi \mid \Phi \mathcal{U} \Phi \mid \Phi \mathcal{U}^I \Phi.$$

The other boolean connectives are derived in the usual way, i.e. $\text{ff} = \neg \text{tt}$, $\Phi_1 \vee \Phi_2 = \neg(\neg\Phi_1 \wedge \neg\Phi_2)$, and $\Phi_1 \rightarrow \Phi_2 = \neg\Phi_1 \vee \Phi_2$. The meaning of \mathcal{U} (“until”)

and X (“next step”) is standard. The temporal operator U^I is the timed variant of U ; $\Phi_1 U^I \Phi_2$ asserts that Φ_2 will be satisfied at some time instant in the interval I and that at all preceding time instants Φ_1 holds. $\mathcal{S}_{\bowtie p}(\Phi)$ asserts that the steady-state probability for a Φ -state falls in the interval $I_{\bowtie p} = \{q \in [0, 1] \mid q \bowtie p\}$. $\mathcal{P}_{\bowtie p}(\varphi)$ asserts that the probability measure of the paths satisfying φ meets the bound given by $\bowtie p$.

Temporal operators like \diamond , \square and their real-time variants \diamond^I or \square^I can be derived, e.g. $\mathcal{P}_{\bowtie p}(\diamond^I \Phi) = \mathcal{P}_{\bowtie p}(\text{tt } U^I \Phi)$ and $\mathcal{P}_{\geq p}(\square \Phi) = \mathcal{P}_{\leq 1-p}(\diamond \neg \Phi)$.

Example 2. Let $AP = \{up_i \mid 0 \leq i < 4\} \cup \{\text{down}\}$ and consider the CTMC of Example 1. $\mathcal{P}_{\leq 10^{-5}}(\diamond^{[0,10]} \text{down})$ denotes that the probability of a failure of the voter within the next 10 hours is at most 10^{-5} ; the formula $\mathcal{S}_{\geq 0.99}(up_3 \vee up_2)$ asserts that with 0.99 probability the system is operational, when the system is in equilibrium. ■

Semantics of CSL. The state-formulas are interpreted over the states of a CTMC. Let $\mathcal{M} = (S, \mathbf{R}, L)$ with labels in AP . The definition of the satisfaction relation $\models \subseteq S \times \mathbf{CSL}$ is as follows. Let $Sat(\Phi) = \{s \in S \mid s \models \Phi\}$.

$$\begin{array}{ll} s \models \text{tt} & \text{for all } s \in S \\ s \models a & \text{iff } a \in L(s) \\ s \models \neg \Phi & \text{iff } s \not\models \Phi \end{array} \qquad \begin{array}{ll} s \models \Phi_1 \wedge \Phi_2 & \text{iff } s \models \Phi_i, i=1, 2 \\ s \models \mathcal{S}_{\bowtie p}(\Phi) & \text{iff } \pi(s, Sat(\Phi)) \in I_{\bowtie p} \\ s \models \mathcal{P}_{\bowtie p}(\varphi) & \text{iff } Prob^{\mathcal{M}}(s, \varphi) \in I_{\bowtie p}. \end{array}$$

Here, $Prob^{\mathcal{M}}(s, \varphi)$ denotes the probability measure of all paths $\sigma \in Path$ satisfying φ when the system starts in state s , i.e., $Prob^{\mathcal{M}}(s, \varphi) = \Pr_s\{\sigma \in Path \mid \sigma \models \varphi\}$.³ The satisfaction relation for the path-formulas is defined as:

$$\begin{array}{ll} \sigma \models X\Phi & \text{iff } \sigma[1] \text{ is defined and } \sigma[1] \models \Phi \\ \sigma \models \Phi_1 U \Phi_2 & \text{iff } \exists k \geq 0. (\sigma[k] \models \Phi_2 \wedge \forall 0 \leq i < k. \sigma[i] \models \Phi_1) \\ \sigma \models \Phi_1 U^I \Phi_2 & \text{iff } \exists t \in I. (\sigma @ t \models \Phi_2 \wedge \forall u \in [0, t]. \sigma @ u \models \Phi_1) \end{array}$$

We note that for $I = \emptyset$ the formula $\Phi_1 U^I \Phi_2$ is not satisfiable and that $\Phi_1 U \Phi_2$ can be interpreted as an abbreviation of $\Phi_1 U^{(0, \infty)} \Phi_2$.

Remark 2. Although **CSL** does not contain an explicit transient state operator, it is possible to reason about transient state probabilities as we have: $\pi(s, s', t) = Prob(s, \diamond^{[t,t]} a_{s'})$. Thus, whereas the steady-state operator $\mathcal{S}_{\bowtie p}(\Phi)$ cannot be derived from the other operators, a transient-state operator $\mathcal{T}_{\bowtie p}^{\otimes t}(\Phi) = \mathcal{P}_{\bowtie p}(\diamond^{[t,t]} \Phi)$ can be defined. It states that the probability for a Φ -state at time point t meets the bound $\bowtie p$. ■

3 Model checking U^I by transient analysis

In [6], we presented a **CSL** model checking algorithm that essentially relies on the following ideas. The steady-state operator requires the computation of steady-state probabilities which can be obtained by a graph analysis and by solving

³ The fact that the set $\{\sigma \in Path \mid \sigma \models \varphi\}$ is measurable can be easily verified.

a linear equation system. The basis for calculating the probabilities $Prob(s, \varphi)$ are the following results. For the temporal operators next X and until U (that abstract from the amount of time spent in states but just refer to the states that are passed in an execution), we have the same characterizations for the values $Prob(s, X\Phi)$ and $Prob(s, \Phi_1 U \Phi_2)$ as in the case of DTMCs [17]:

$$\begin{aligned}
Prob(s, X\Phi) &= \sum_{s' \models \Phi} \mathbf{P}(s, s') \quad \text{and} \\
Prob(s, \Phi_1 U \Phi_2) &= \begin{cases} 1 & \text{if } s \models \Phi_2 \\ \sum_{s' \in S} \mathbf{P}(s, s') \cdot Prob(s', \Phi_1 U \Phi_2) & \text{if } s \models \Phi_1 \wedge \neg \Phi_2 \\ 0 & \text{otherwise.} \end{cases}
\end{aligned}$$

This amounts to matrix/vector-multiplication for next and solving a linear equation system for until.

For the time-bounded until operator, [6] suggested an iterative method which relies on the observation that the function $(s, t, t') \mapsto Prob(s, \Phi_1 U^{[t, t']} \Phi_2)$ can be characterized as the least fixed point of a higher-order operator Ω where $\Omega(F)$ is defined by means of Volterra integrals. This fixed-point characterization then serves as a basis for an iterative method that uses numerical integration techniques. First experiments in a non-symbolic setting have shown that this approach can be rather time-consuming and that numerical stability is hard to achieve [22]. [6] suggested a symbolic approach by combining (MT)BDD-techniques [9, 12, 4] with an operator for solving integrals by quadrature formulas. Here, we propose an alternative strategy that reduces the model checking problem for the time-bounded until operator to the problem of calculating transient probabilities in CTMCs. This observation allows us to implement **CSL** model checking on the basis of well-established transient analysis techniques for CTMCs (see below).

Four correctness-preserving transformations. We first observe that it suffices to consider time bounds specified by compact intervals, since:

$$Prob(s, \Phi_1 U^I \Phi_2) = Prob(s, \Phi_1 U^{cl(I)} \Phi_2)$$

where $cl(I)$ denotes the closure of I . Secondly, unbounded time intervals $[t, \infty)$ can be treated by combining time-bounded until and unbounded until, since:

$$Prob(s, \Phi_1 U^{[t, \infty)} \Phi_2) = \sum_{s' \in S} Prob(s, \Phi_1 U^{[t, t]} a_{s'}) \cdot Prob(s', \Phi_1 U \Phi_2).$$

In the sequel, we treat 4 types of time-bounded until-formulas with a compact interval I and show how they all can be reduced to instances of two simple base cases. For CTMC $\mathcal{M} = (S, \mathbf{R}, L)$ and **CSL**-state formula Φ let CTMC $\mathcal{M}[\Phi]$ result from \mathcal{M} by making all Φ -states in \mathcal{M} absorbing; i.e., $\mathcal{M}[\Phi] = (S, \mathbf{R}', L)$ where $\mathbf{R}'(s, s') = \mathbf{R}(s, s')$ if $s \not\models \Phi$ and 0 otherwise. Note that $\mathcal{M}[\Phi_1][\Phi_2] = \mathcal{M}[\Phi_1 \vee \Phi_2]$.

Case A: Bounded until for absorbing Φ_2 -states. Let $\varphi = \Phi_1 U^{[0, t]} \Phi_2$ and assume that all Φ_2 -states are absorbing, i.e., once a Φ_2 -state is reached it will not be

left anymore. We first observe that once a $(\neg\Phi_1 \wedge \neg\Phi_2)$ -state is reached, φ will be invalid, regardless of the future evolution of the system. As a result, we may switch from \mathcal{M} to $\mathcal{M}[\neg\Phi_1 \wedge \neg\Phi_2]$ and consider the property on the obtained CTMC. The assumption that all Φ_2 -states are absorbing allows us to conclude that φ is satisfied once a Φ_2 -state is reached at time t . Thus,

Lemma 1. *If all Φ_2 -states are absorbing in \mathcal{M} (i.e., $\mathcal{M} = \mathcal{M}[\Phi_2]$) then:*

$$Prob^{\mathcal{M}}(s, \Phi_1 \mathcal{U}^{[0,t]} \Phi_2) = Prob^{\mathcal{M}'}(s, \diamond^{[t,t]} \Phi_2) = \sum_{s'' \models \Phi_2} \pi^{\mathcal{M}'}(s, s'', t)$$

for $\mathcal{M}' = \mathcal{M}[\neg\Phi_1 \wedge \neg\Phi_2]$.

Case B: Point-interval until for $\Phi_2 \rightarrow \Phi_1$. Let $\varphi = \Phi_1 \mathcal{U}^{[t,t]} \Phi_2$ and assume $\Phi_2 \rightarrow \Phi_1$. Note that such implication holds in case of \diamond -properties. With the same motivation as for the previous case, we make $(\neg\Phi_1 \wedge \neg\Phi_2)$ -states absorbing. Since $\Phi_2 \rightarrow \Phi_1$ it follows that $Prob(s, \varphi)$ equals the probability to be in a Φ_2 -state at time t in the obtained CTMC:

Lemma 2. *If $\Phi_2 \rightarrow \Phi_1$ we have for any CTMC \mathcal{M} :*

$$Prob^{\mathcal{M}}(s, \Phi_1 \mathcal{U}^{[t,t]} \Phi_2) = Prob^{\mathcal{M}'}(s, \diamond^{[t,t]} \Phi_2) = \sum_{s'' \models \Phi_2} \pi^{\mathcal{M}'}(s, s'', t).$$

for $\mathcal{M}' = \mathcal{M}[\neg\Phi_1 \wedge \neg\Phi_2]$.

Case C: Bounded until. Let $\varphi = \Phi_1 \mathcal{U}^{[0,t]} \Phi_2$ and consider an arbitrary CTMC \mathcal{M} . This property is fulfilled if a Φ_2 -state is reached before (or at) time t via some Φ_1 -path. Once such Φ_2 -state has been reached, the future behaviour of the CTMC is irrelevant for the validity of φ . Accordingly, the Φ_2 -states can be safely made absorbing without affecting the validity of φ . As a result, it suffices to consider the probability of being in a Φ_2 -state at time t for $\mathcal{M}[\Phi_2]$, thus reducing to the case in Lemma 1. As $\mathcal{M}[\Phi_2][\neg\Phi_1 \wedge \neg\Phi_2] = \mathcal{M}[\neg\Phi_1 \vee \Phi_2]$ we obtain:

Theorem 1. *For any CTMC \mathcal{M} :*

$$Prob^{\mathcal{M}}(s, \Phi_1 \mathcal{U}^{[0,t]} \Phi_2) = Prob^{\mathcal{M}[\Phi_2]}(s, \Phi_1 \mathcal{U}^{[0,t]} \Phi_2) = \sum_{s'' \models \Phi_2} \pi^{\mathcal{M}[\neg\Phi_1 \vee \Phi_2]}(s, s'', t).$$

Case D: Interval-until. Let $\varphi = \Phi_1 \mathcal{U}^{[t,t']} \Phi_2$ with $0 < t \leq t'$ and let \mathcal{M} be an arbitrary CTMC⁴. We first observe that for any path σ with $\sigma \models \varphi$: (i) Φ_1 continuously holds in the interval $[0, t]$ (i.e., $\sigma \models \square^{[0,t]} \Phi_1$), in particular, $s' = \sigma @ t$ is a Φ_1 -state, and (ii) $\sigma' \in Path(s')$, the suffix of σ that starts at time t , fulfills the path formula $\Phi_1 \mathcal{U}^{[0, t'-t]} \Phi_2$.⁵ Let the intermediate state $s' \in Sat(\Phi_1)$ and consider the set $\Sigma(s')$ of paths $\sigma \in Path(s)$ where $\sigma @ t = s'$ and $\sigma \models \varphi$. Then

⁴ Note that $Prob(s, \Phi_1 \mathcal{U}^{[t,t']} \Phi_2) \neq Prob(s, \Phi_1 \mathcal{U}^{[0,t']} \Phi_2) - Prob(s, \Phi_1 \mathcal{U}^{[0,t]} \Phi_2)$.

⁵ Formally, σ' is the unique path with $\sigma' @ x = \sigma @ (t+x)$ for any positive real x .

$\Pr_s(\Sigma(s'))$ equals $Prob(s, \Phi_1 \mathcal{U}^{[t,t]} a_{s'})$ times $Prob(s', \Phi_1 \mathcal{U}^{[0,t'-t]} \Phi_2)$. As the sets $\Sigma(s')$ for $s' \in Sat(\Phi_1)$ are pairwise disjoint we obtain:

$$Prob(s, \Phi_1 \mathcal{U}^{[t,t']} \Phi_2) = \sum_{s' \models \Phi_1} Prob(s, \Phi_1 \mathcal{U}^{[t,t]} a_{s'}) \cdot Prob(s', \Phi_1 \mathcal{U}^{[0,t'-t]} \Phi_2).$$

To compute the probabilities $Prob(s, \Phi_1 \mathcal{U}^{[t,t]} a_{s'})$ for $s' \models \Phi_1$ we use Lemma 2, i.e., we switch from \mathcal{M} to $\mathcal{M}_1 = \mathcal{M}[\neg\Phi_1]$ and compute the transient probabilities for any Φ_1 -state s' at time t in \mathcal{M}_1 . The probabilities $Prob(s', \Phi_1 \mathcal{U}^{[0,t'-t]} \Phi_2)$ can be obtained as in Theorem 1. This yields the following result:

Theorem 2. *For any CTMC \mathcal{M} and $0 < t \leq t'$:*

$$Prob^{\mathcal{M}}(s, \Phi_1 \mathcal{U}^{[t,t']} \Phi_2) = \sum_{s' \models \Phi_1} \sum_{s'' \models \Phi_2} \pi^{\mathcal{M}[\neg\Phi_1]}(s, s', t) \cdot \pi^{\mathcal{M}[\neg\Phi_1 \vee \Phi_2]}(s', s'', t'-t).$$

With Theorem 2 we can calculate the values $Prob^{\mathcal{M}}(s, \varphi)$, using one of the following two methods. Let $\mathcal{M}_2 = \mathcal{M}[\neg\Phi_1 \vee \Phi_2]$. Either we calculate the matrices

$$\begin{aligned} \mathbf{A} &= (\pi^{\mathcal{M}_1}(s, s', t))_{s \in S, s' \in Sat(\Phi_1)}, \\ \mathbf{B} &= (\pi^{\mathcal{M}_2}(s', s'', t'-t))_{s' \in Sat(\Phi_1), s'' \in Sat(\Phi_2)} \end{aligned}$$

and then take the product $\mathbf{A} \cdot \mathbf{B}$. Or, we first calculate the distributions α_s for $s \in S$, given by $\alpha_s(s') = \pi^{\mathcal{M}_1}(s, s', t)$ and then compute $Prob^{\mathcal{M}}(s, \varphi) = \sum_{s'' \models \Phi_2} \pi^{\mathcal{M}_2}(\alpha_s, s'', t'-t)$. This alternative is based on the observation

$$\sum_{s' \models \Phi_1} \alpha_s(s') \cdot \pi^{\mathcal{M}_2}(s', s'', t'-t) = \pi^{\mathcal{M}_2}(\alpha_s, s'', t'-t)$$

for any Φ_2 -state s'' .⁶

Example 3. Consider our TMR with initial distribution $\alpha = \alpha_{s_{3,1}}^1$ and let $\Phi = \mathcal{P}_{\geq 0.15}(\Phi_1 \mathcal{U}^{[3,7]} \Phi_2)$ for $\Phi_1 = up_3 \vee up_2$ and $\Phi_2 = up_2 \vee up_1$. According to Theorem 2 model checking Φ boils down to first computing the transient probabilities at time 3, i.e., $\alpha_2 = (\pi^{\mathcal{M}_1}(s_{3,1}, 3))_{s \in S}$ in CTMC \mathcal{M}_1 of Fig. 1(a) where all $\neg\Phi_1$ -states are made absorbing. We obtain $\alpha_2 = (0.968, 0.0272, 0.011, 0, 0.003)$ with a precision of $\varepsilon = 10^{-6}$ for $\lambda = 0.01$, $\nu = 0.001$, $\mu = 1.0$ and $\delta = 0.2$. In the second phase, we compute the transient probabilities at time 4 in CTMC \mathcal{M}_2 of Fig. 1(b) starting from initial distribution α_2 , i.e., computing $\sum_{s'' \models up_2} \pi^{\mathcal{M}_2}(\alpha_2, s'', 4) \approx 0.1365$. Thus, the property Φ is violated.

Uniformisation. Based on the general principle of uniformisation [25], efficient techniques to compute transient state probabilities for CTMCs have been proposed [16, 15]. With uniformisation, the transient probabilities of a CTMC are

⁶ For both alternatives, in the worst case, for any state s , we need a transient analysis in \mathcal{M}_1 (with initial state s) and \mathcal{M}_2 (with initial distribution α_s). The second alternative might be preferable if there are only a few different distributions α_s .




Fig. 1. CTMCs to be analysed for checking $\mathcal{P}_{\geq 0.15}((up_3 \vee up_2) \mathcal{U}^{[3,7]}(up_2 \vee up_1))$

computed via a so-called *uniformised DTMC* which characterises the CTMC at state transition epochs.

Denoting with $\underline{\pi}(\alpha, t)$ the vector of state probabilities at time t , i.e., $\underline{\pi}(\alpha, t) = (\pi(\alpha, s_1, t), \dots, \pi(\alpha, s_N, t))$ (with $N = |S|$ the number of states), the Chapman-Kolmogorov differential equations characterise the transient behaviour: $\underline{\pi}'(\alpha, t) = \underline{\pi}(\alpha, t) \cdot \mathbf{Q}$, where $\mathbf{Q} = \mathbf{R} - \text{diag}(\mathbf{E})$. A formal solution is then given by the Taylor-series expansion:

$$\underline{\pi}(\alpha, t) = \alpha \cdot e^{\mathbf{Q} \cdot t} = \alpha \cdot \sum_{i=0}^{\infty} \frac{(\mathbf{Q} \cdot t)^i}{i!}.$$

This solution, however, should not be used as the basis for a numerical algorithm since: (i) it suffers from numerical instability due to the fact that \mathbf{Q} contains both positive and negative entries; (ii) the matrix powers will become less and less sparse for large i , thus requiring $O(N^2)$ storage; (iii) it is difficult to find a proper truncation criterion for the infinite summation.

Instead, by choosing $q = \max_i \{\mathbf{E}(s_i)\}$, we construct the uniformised DTMC with transition probability matrix $\mathbf{P} = \mathbf{I} + \mathbf{Q}/q$. By the choice of q , \mathbf{P} is a stochastic matrix. Substituting $\mathbf{Q} = q(\mathbf{P} - \mathbf{I})$ in the above solution, we obtain

$$\underline{\pi}(\alpha, t) = \alpha \cdot \sum_{i=0}^{\infty} e^{-q \cdot t} \frac{(q \cdot t)^i}{i!} \mathbf{P}^i,$$

which can be rewritten as

$$\underline{\pi}(\alpha, t) = \sum_{i=0}^{\infty} PP(i) \cdot \underline{\pi}_i,$$

where $PP(i) = e^{-q \cdot t} \frac{(q \cdot t)^i}{i!}$ is the i -th Poisson probability with parameter qt , and $\underline{\pi}_i = \underline{\pi}_{i-1} \mathbf{P}$ and $\underline{\pi}_0 = \alpha$. The Poisson probabilities can be computed in a stable way with the algorithm of Fox and Glynn [14]. There is no need to compute explicit powers of the matrix \mathbf{P} . Furthermore, since the terms in the summation are all between 0 and 1, the number of terms to be taken given a required accuracy, can be computed a priori. For large values of qt , this number is of order $O(qt)$. Notice, however, that for large values of qt , the DTMC described by \mathbf{P} might

even have reached steady-state, so that a further reduction in computational complexity is reached. For further details, see [28, 18].

Regarding storage complexity, we note that we require $O(3N)$ storage for the probability vectors and $O(\eta N)$ for the matrix \mathbf{P} , where η denotes the (average) number of transitions originating from a single state in the DTMC (typically $\eta \ll N$). Regarding computational complexity, to compute $\underline{\pi}(\alpha, t)$ we require the sum of $O(qt)$ vectors, each of which is the result of a matrix-vector multiplication. Given a sparse implementation of the latter, we require $O(\eta N)$ scalar multiplications for that, so that we have an overall computational complexity of $O(qt \cdot \eta N)$.

4 Abstraction with bisimulation (lumping) equivalence

In this section, we discuss some techniques to reduce the state space of a CTMC. These techniques are mainly based on the observation that (a slight variant of) ordinary *lumping equivalence* (i.e., bisimulation) preserves all **CSL**-formulas. This result is in the spirit of [8] where bisimilar states of an ordinary transition system are shown to satisfy the same **CTL**-formulas. Similar results have been established for many types of transition systems and branching-time logics; e.g., in the probabilistic setting, [2] shows that probabilistic bisimulation on DTMCs preserves **PCTL** [17]. Our result below can be considered as the continuous version of that result. Let $\mathcal{M} = (S, \mathbf{R}, L)$ be a CTMC, F a set of **CSL**-formulas, and $L_F : S \rightarrow 2^F$ a labelling defined by $L_F(s) = \{ \Phi \in F \mid s \models \Phi \}$.

Definition 1. *An F -bisimulation on $\mathcal{M} = (S, \mathbf{R}, L)$ is an equivalence R on S such that whenever $(s, s') \in R$ then $L_F(s) = L_F(s')$ and $\mathbf{R}(s, C) = \mathbf{R}(s', C)$ for all $C \in S/R$. States s and s' are F -bisimilar iff there exists an F -bisimulation R that contains (s, s') .*

Here, S/R denotes the quotient space and $\mathbf{R}(s, C)$ abbreviates $\sum_{s' \in C} \mathbf{R}(s, s')$. F -bisimulation is a slight variant of Markovian bisimulation (which is defined on CTMCs with action-labelled transitions) on CTMCs with labelled states. Markovian bisimulation coincides with (ordinary) lumping equivalence [11], a well-known notion to aggregate CTMCs.

For $s \in S$, let $[s]_R$ denote the equivalence class of s under R . For $\mathcal{M} = (S, \mathbf{R}, L)$ we define the CTMC $\mathcal{M}/R = (S/R, \mathbf{R}_R, L_R)$ with $\mathbf{R}_R([s]_R, C) = \mathbf{R}(s, C)$ and $L_R([s]_R) = L_F(s)$. That is, \mathcal{M}/R results from \mathcal{M} by building the quotient space under R and labelling states with F (rather than AP). \mathcal{M}/R can be computed by a modified version of the partition refinement algorithm for ordinary bisimulation without an increase of the worst case complexity [23]. Let \mathbf{CSL}_F denote the smallest set of **CSL**-formulas that includes F and that is closed under all **CSL**-operators. In the following we write $\models_{\mathcal{M}}$ for the satisfaction relation \models (on **CSL**) on \mathcal{M} .

Theorem 3. *Let R be an F -bisimulation on \mathcal{M} and s a state in \mathcal{M} . Then:*

- (a) *For all \mathbf{CSL}_F -formulas Φ : $s \models_{\mathcal{M}} \Phi$ iff $[s]_R \models_{\mathcal{M}/R} \Phi$*

(b) For all \mathbf{CSL}_F path-formulas φ : $Prob^{\mathcal{M}}(s, \varphi) = Prob^{\mathcal{M}/R}([s]_R, \varphi)$.

In particular, F -bisimilar states satisfy the same \mathbf{CSL}_F formulas.

Proof. Straightforward by structural induction on Φ and φ .

Theorem 3 allows to verify \mathbf{CSL} -formulas on the possibly much smaller \mathcal{M}/R rather than on \mathcal{M} , for AP -bisimulation R .

In addition, we can exploit the above result to our transformations of the previous section by using the following observation. From Theorem 3(b) and Remark 2 it follows:

$$\sum_{s' \models_{\mathcal{M}} \Phi} \pi^{\mathcal{M}}(s, s', t) = \sum_{S' \models_{\mathcal{M}/R} \Phi} \pi^{\mathcal{M}/R}([s]_R, S', t) \quad (1)$$

for any \mathbf{CSL}_F formula Φ and F -bisimulation R . This observation allows us to simplify the CTMCs $\mathcal{M}[\dots]$ that occur in the cases A–D of our model checking procedure presented in the previous section in the following way. For cases C and D, we compute the transient probabilities for Φ_2 -states in the CTMC $\mathcal{M}' = \mathcal{M}[\neg\Phi_1 \vee \Phi_2]$. Let $F = \{\neg\Phi_1 \wedge \neg\Phi_2, \Phi_2\}$ and R be the smallest equivalence on the state space S of \mathcal{M}' that identifies all Φ_2 -states and all $(\neg\Phi_1 \wedge \neg\Phi_2)$ -states. Clearly, R is an F -bisimulation on \mathcal{M}' . The state space of \mathcal{M}'/R is

$$S/R = Sat(\Phi_1 \wedge \neg\Phi_2) \cup [Sat(\Phi_2)]_R \cup [Sat(\neg\Phi_1 \wedge \neg\Phi_2)]_R$$

Since Φ_2 is a \mathbf{CSL}_F -formula, equation (1) yields

$$\sum_{s' \models_{\mathcal{M}'} \Phi_2} \pi^{\mathcal{M}'}(s, s', t) = \pi^{\mathcal{M}'/R}(s, [Sat(\Phi_2)]_R, t)$$

for any state $s \in Sat(\Phi_1 \wedge \neg\Phi_2)$. Similar arguments are applicable to case A and B. As a result, the sets $[Sat(\neg\Phi_1 \wedge \neg\Phi_2)]_R$ and $[Sat(\Phi_2)]_R$ in cases A–D can be considered as single states. This may yield a substantial reduction of the state space of the CTMC under consideration. From a computational point of view, the switch from \mathcal{M} to the modified $\mathcal{M}[\dots]/R$ is quite simple as we just collapse certain states into a single absorbing state. The generator matrix \mathbf{R}_R for $\mathcal{M}[\dots]/R$ can be obtained by simple manipulations of the generator matrix \mathbf{R} for \mathcal{M} (matrix multiplication).

Example 4. According to the above observations, in the CTMC of Fig. 1(a) we may aggregate states $[Sat(\Phi_1)]_R = \{s_{0,1}, s_{0,0}, s_{1,1}\}$ into a single state. This new state is reachable from $s_{3,1}$ with rate ν and from $s_{2,1}$ with rate $2\lambda + \nu$. In the CTMC of Fig. 1(b) we may collapse $[Sat(\Phi_2)]_R = \{s_{2,1}, s_{1,1}\}$ and $[Sat(\neg\Phi_1 \wedge \neg\Phi_2)]_R = \{s_{0,0}, s_{0,1}\}$ into single states.

5 Model checking a telephone system

In this section, we report on model checking the stochastic behaviour of an instance of the plain-old telephone system (POTS), where two users concurrently

try to get connected to each other. In [21], we have shown how a formal specification of the POTS (in LOTOS) can be augmented with stochastic timing constraints, leading to a model of more than 10^7 states. We aggregated this model compositionally using appropriate stochastic extensions of (strong and weak) bisimulation [23], to come up with a lumped CTMC \mathcal{M} of 720 states. Here we model check the resulting CTMC using transient analysis. In short, the following atomic propositions are used: *conn* characterises states where both partners are connected to each other, and conversation is running. *fed_up* characterises states where either of the user is hooking the phone because he is apparently out of luck, unable to reach his conversation partner. Our basic time unit is 1 minute. The following properties are checked:

- $\mathcal{P}_{\triangleright p}(\diamond^{[0,t]}conn)$, the probability of being connected within t minutes.
- $\mathcal{P}_{\triangleright p}(\diamond^{[t,t]}conn)$, the probability of being connected after exactly t minutes.
- $\mathcal{P}_{\triangleright p}(\diamond^{[100,100+t]}conn)$, the probability of being connected at some time between 100 and $100+t$ minutes.
- $\mathcal{P}_{\triangleright p}(\neg fed_up \mathcal{U}^{[t,100+t]}conn)$, the probability of a running conversation between t and $100+t$ minutes, without failing to get connected beforehand.

Note that only the first property can be checked with the current implementation [22] of the non-symbolic model checking algorithm based on numerical integration [6]. We do not instantiate p , as the execution times and computed probabilities will be the same for all $p \in]0, 1[$. Statistics of the computation time needed to check these formulas globally, i.e., for all states, are depicted in Table 1. They have been obtained by means of a trial implementation of the uniformisation method written in *C*, running on a 300 MHz SUN Ultra 5 workstation with 256 MB memory under the Solaris 2.6 operating system. (In all cases reported below, the memory requirements are less than 20 MB.)

t	$\mathcal{P}_{\triangleright p}(\diamond^{[0,t]}conn)$		$\mathcal{P}_{\triangleright p}(\diamond^{[t,t]}conn)$		$\mathcal{P}_{\triangleright p}(\diamond^{[100,100+t]}conn)$		$\mathcal{P}_{\triangleright p}(\neg fed_up \mathcal{U}^{[t,100+t]}conn)$	
	MV-mult.	time	MV-mult.	time	MV-mult.	time	MV-mult.	time
0.1	59	7.75	138	41.91	15,325	3,549.70	5,234	351.75
1	102	10.75	267	75.04	15,368	3,552.69	5,349	378.05
10	583	38.40	1,714	416.38	15,848	3,580.27	6,795	747.21
100	5,081	303.33	15,265	3,541.84	20,347	3,845.27	20,347	3,956.01
1000	8,901	619.51	39,155	8,835.01	24,167	4,161.41	151,815	34,405.14
10000	8,901	624.68	39,155	8,902.41	24,167	4,166.52	151,815	34,567.23

Table 1. Computation time (in *sec*) and number of matrix-vector multiplications (times 10^3) needed for checking **CSL** properties by means of uniformisation

From these statistics, we draw the following conclusions. (1) We observe a roughly linear dependency between the time bound t and the run-time of uniformisation, due to the fact that the (precomputed) number of iterations needed by uniformisation is $\mathcal{O}(t)$. (2) The number of iterations needed for $t \geq 1000$ is constant, due to the fact that our algorithm has a built-in steady-state detection. In other words, the chain at time 1000 is already behaving close to equilibrium,

up to a truncation error ε (set to 10^{-6} in all experiments). For time-bounds larger than 1000, transient analysis could in fact be replaced by a (much cheaper) steady-state analysis. (3) The times needed to check $\mathcal{P}_{\bowtie p}(\diamond^{[t,t]} conn)$ are approximately one order of magnitude higher than those needed for $\mathcal{P}_{\bowtie p}(\diamond^{[0,t]} conn)$. This is a consequence of the fact that the pruning of transitions in $\mathcal{M}[conn]$ (cf. Theorem 1) leads to a CTMC containing mutually unreachable parts, and for each state we perform transient analysis only on the reachable (lumped) CTMC. On average, this chain has only 62 states, explaining the order of magnitude difference. Note that according to Lemma 2, transient analysis can take the original chain (with 720 states) *unchanged* to check $\mathcal{P}_{\bowtie p}(\diamond^{[t,t]} conn)$, since $conn \rightarrow tt$. The two rightmost formulas involve more than one (iterative) transient solution, on different lumped CTMCs (cf. Theorem 2). Their time consumption is mainly determined by the size of the lower time bound (an observation that does not hold in general).

6 Concluding remarks

The main result of this paper is that the verification problem for probabilistic timing properties, i.e., **CSL**-formulas of the form $\mathcal{P}_{\bowtie p}(\Phi_1 U^I \Phi_2)$, is reducible to a transient analysis of CTMCs. Thus, efficient techniques for transient analysis, such as uniformisation, can be adopted for model checking these formulas. In addition, we showed that a slight variant of (ordinary) lumpability on CTMCs preserves all **CSL**-formulas. We illustrated these results by analysing a plain-old telephone system. Future work includes the adaption of partial uniformisation [27] to our setting (thus allowing a partial search of the state space) and considering a symbolic variant of our presented approach using multi-terminal BDDs [4, 12] in order to compare this approach with the symbolic (numerical integration) approach in [6]. The extension of our approach towards Markov reward models is reported in [7].

References

1. M. Ajmone Marsan, G. Balbo, G. Conte, S. Donatelli, and G. Franceschinis. *Modelling with Generalized Stochastic Petri Nets*. John Wiley & Sons, 1995.
2. A. Aziz, V. Singhal, F. Balarin, R. Brayton and A. Sangiovanni-Vincentelli. It usually works: the temporal logic of stochastic systems. In *CAV*, LNCS 939, pp. 155–165, 1995.
3. A. Aziz, K. Sanwal, V. Singhal and R. Brayton. Verifying continuous time Markov chains. In *CAV*, LNCS 1102, pp. 269–276, 1996.
4. I. Bahar, E. Frohm, C. Gaona, G. Hachtel, E. Macii, A. Padro and F. Somenzi. Algebraic decision diagrams and their applications. *Form. Meth. in Syst. Design*, **10**(2/3): 171–206, 1997.
5. C. Baier. On algorithmic verification methods for probabilistic systems. Habilitation thesis, Univ. of Mannheim, 1999.
6. C. Baier, J.-P. Katoen and H. Hermanns. Approximate symbolic model checking of continuous-time Markov chains. In *CONCUR*, LNCS 1664, pp. 146–162, 1999.

7. C. Baier, B. Haverkort, H. Hermanns and J.-P. Katoen. On the logical characterisation of performability properties. In *ICALP*, LNCS, 2000 (to appear).
8. M. Brown, E. Clarke, O. Grumberg. Characterizing finite Kripke structures in propositional temporal logic. *Th. Comp. Sc.*, **59**: 115–131, 1988.
9. R. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Trans. on Comp.*, **C-35**(8): 677–691, 1986.
10. P. Buchholz. Exact and ordinary lumpability in finite Markov chains. *J. of Appl. Prob.*, **31**: 59–75, 1994.
11. P. Buchholz. Markovian process algebra. Tech. Rep. 500, Fachbereich Informatik, Univ. of Dortmund, 1994.
12. E. Clarke, M. Fujita, P. McGeer, J. Yang and X. Zhao. Multi-terminal binary decision diagrams: an efficient data structure for matrix representation. In *Proc. IEEE Int. Workshop on Logic Synthesis*, pp. 1–15, 1993.
13. C. Courcoubetis and M. Yannakakis. Verifying temporal properties of finite-state probabilistic programs. In *Proc. IEEE Symp. on Found. of Comp. Sc.*, pp. 338–345, 1988.
14. B.L. Fox and P.W. Glynn. Computing Poisson probabilities. *Comm. of the ACM* **31**(4): 440–445, 1988.
15. W.K. Grassmann. Finding transient solutions in Markovian event systems through randomization. In W.J. Stewart, ed, *Num. Sol. of Markov Chains*, pp. 357–371, Marcel Dekker, 1991.
16. D. Gross and D.R. Miller. The randomization technique as a modeling tool and solution procedure for transient Markov chains. *Oper. Res.* **32**(2): 343–361, 1984.
17. H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Form. Asp. of Comp.* **6**: 512–535, 1994.
18. B.R. Haverkort. *Performance of Computer Communication Systems: A Model-Based Approach*. John Wiley & Sons, 1998.
19. B.R. Haverkort and I. Niemegeers. Performability modelling tools and techniques. *Perf. Ev.*, **25**: 17–40, 1996.
20. H. Hermanns, U. Herzog and J.-P. Katoen. Process algebra for performance evaluation. *Th. Comp. Sc.*, 2000 (to appear).
21. H. Hermanns and J.-P. Katoen. Automated compositional Markov chain generation for a plain-old telephone system. *Sci. of Comp. Programming*, **36**(1): 97–127, 2000.
22. H. Hermanns, J.-P. Katoen, J. Meyer-Kayser and M. Siegle. A Markov chain model checker. In *TACAS*, LNCS 1785, pp. 347–362, 2000.
23. H. Hermanns and M. Siegle. Bisimulation algorithms for stochastic process algebras and their BDD-based implementation. In *ARTS*, LNCS 1601, pp. 244–265, 1999.
24. J. Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press, 1996.
25. A. Jensen. Markov chains as an aid in the study of Markov processes. *Skand. Aktuarietidskrift* **3**: 87–91, 1953.
26. K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Inf. and Comp.*, **94**(1): 1–28, 1992.
27. A.P.A. van Moorsel and B.R. Haverkort. Probabilistic evaluation for the analytical solution of large Markov models. *Microelectron. and Reliab.* **36**(6): 733–755, 1996.
28. W.J. Stewart. *Introduction to the Numerical Solution of Markov Chains*. Princeton Univ. Press, 1994.