

Beyond Memoryless Distributions: Model Checking Semi-Markov Chains

Gabriel G. Infante López, Holger Hermanns, and Joost-Pieter Katoen

Formal Methods and Tools Group, Faculty of Computer Science
University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands

Abstract. Recent investigations have shown that the automated verification of continuous-time Markov chains (CTMCs) against CSL (Continuous Stochastic Logic) can be performed in a rather efficient manner. The state holding time distributions in CTMCs are restricted to negative exponential distributions. This paper investigates model checking of semi-Markov chains (SMCs), a model in which state holding times are governed by general distributions. We report on the semantical issues of adopting CSL for specifying properties of SMCs and present model checking algorithms for this logic.

1 Introduction

Model checking is a technique that is more and more used to ascertain properties of computer software, hardware circuits, communication protocols, and so forth. In this approach, properties are specified via an appropriate temporal logic, such as CTL or LTL, while systems are represented as (usually finite) state-transition diagrams. More recently, model checking techniques have been extended to stochastic processes such as continuous-time Markov chains (CTMCs, for short). In particular, efficient verification algorithms have been developed for CSL (Continuous Stochastic Logic [3,4,5]), a stochastic variant of CTL. CSL supports the specification of sophisticated steady-state and time-dependent properties. CTMCs are widely used in practice, mainly because they combine a reasonable modelling flexibility with well-established efficient analysis techniques for transient and steady-state probabilities that form the basis for determining performance measures such as throughput, utilisation and latencies. The stochastic processes described by CTMCs are characterised by the fact that the state holding times, indicating the amount of time the system stays in a state, are restricted to negative exponential distributions. As a result of their so-called memoryless property, the probability of moving from one state to another is independent of the amount of time the system has spent in the current state.

Although exponential distributions appropriately model a significant number of phenomena – related to mass effects – of random nature, in many occasions they are inadequate to faithfully model the stochastic behaviour of the system under consideration. For example, file sizes of documents transferred via the Internet, cycle times in hardware circuits, timeouts in communication protocols,

human behaviour, hardware failures, and jitter in multi-media communication systems cannot be appropriately modelled. In order to model these phenomena in an adequate manner, *general* distributions such as heavy-tail [10] (for file sizes), deterministic (for cycle times and timeouts), log-normal (for human response behaviour [21]), Weibull (for hardware failures [20]), and normal distributions (for jitter [13]) are used. To adopt the model checking approach to these distributions, the simplest solution is to approximate general distributions by the mean times to absorption of a CTMCs with an absorbing state, representing a so-called phase-type distribution. Although the resulting CTMC can be analysed using the existing verification algorithms and prototype tools for CSL, such approximations (i) easily give rise to a state-space explosion – the number of states increases significantly with the accuracy of the approximation and the degree of determinism of the desired distribution – are (ii) not easy to handle in case of a choice between stochastic delays – a race condition between the entire approximated distributions is decisive – and (iii) require the ability to fit the desired distribution by an appropriate phase-type distribution – a non-trivial problem in general, see e.g., [2].

Therefore as an alternative approach, this paper investigates direct model checking of semi-Markov chains (SMCs, for short) [8,18], a natural extension of CTMCs in which state holding times are determined by *general* continuous distributions. First, the semantics of CSL on SMCs is studied. In particular, the formal characterisation of the CSL steady-state operator is adapted as limit state probabilities are not guaranteed to exist for finite-state SMCs, in contrast to finite-state CTMCs. Instead, the behaviour of SMCs on the long run is characterised using the average fraction of time the system resides in a state. For instance, the formula $\mathcal{S}_{<0.01}(\text{error})$ is valid in state s iff on the long run for at most 1% of the time on average the system is in an error state when starting in state s . For finite CTMCs this interpretation is equal to the characterisation using the limit state probabilities. Secondly, model checking algorithms are proposed to verify CSL over finite-state SMCs. Although long-run properties are semantically characterised in a slightly different way, they can be checked as for CTMCs: a graph analysis to determine the bottom strongly connected components and solving a linear system of equations for each such component suffice. (In the literature, strongly connected SMCs are also known as irreducible SMCs.) Time-bounded until formulas can be checked, like for CTMCs, by a reduction to transient analysis of SMCs. These include probabilistic timed reachability properties such as: can the system reach a goal-state within a certain time-bound with some minimal (or maximal) probability? Whereas such transient analysis for CTMCs can be solved via stable and efficient numerical techniques such as uniformisation, for SMCs it requires solving a set of non-trivial Volterra equations whose solution algorithms have a worst case time complexity of $\mathcal{O}(N^4)$, where N is the number of states of the SMC under consideration.

In the context of logical specification formalisms and automated verification, stochastic processes with general distributions have received scant attention in the literature so far. Three related works are known to the authors. Van Hung

and Chaochen [17] have defined a probabilistic variant of the duration calculus to express properties over SMCs, but did not report on any verification algorithms. De Alfaro [12] discusses model checking of long-run average properties and expected reachability times on semi-Markov decision processes. These models can be considered as SMCs extended with non-determinism. Time-bounded formulas are not considered. Kwiatkowska *et al.* [19] have recently considered the verification of a stochastic variant of timed automata, with clocks that are governed by general distributions, against properties in probabilistic timed CTL. They show that a finite-state semantics of such timed automata can be obtained using the region-based technique [1] where regions are partitioned to cater for the stochastic behaviour. Due to the intrinsic complexity of the model checking algorithm, it seems practically infeasible.

Organisation of the paper. Section 2 introduces the basic concepts of SMCs. Section 3 recalls the logic CSL and defines the semantics of CSL over SMCs. Model checking algorithms for long-run properties and time-bounded until formulas are described in Section 4. Section 5 concludes the paper.

2 Semi-Markov Chains

A semi-Markov chain (SMC) can be considered as a Kripke structure in which the transitions are labelled by information about the speed at which the chain evolves from one state to another. In a SMC, the delay between two successive state changes can be *generally* distributed. This property has to be contrasted with continuous-time Markov chains (CTMCs) where these delays need to be governed by negative exponential distributions. In this section, we introduce the basic concepts of SMCs. A more thorough treatment of SMCs can be found in [8,18].

Semi-Markov chains. Let AP be a fixed, finite set of atomic propositions. A (labelled) SMC \mathcal{M} is a tuple $(S, \mathbf{P}, \mathbf{Q}, L)$ where S is a finite set of states, $\mathbf{P} : S \times S \rightarrow [0, 1]$ is the *transition probability matrix* (satisfying $\sum_{s' \in S} \mathbf{P}(s, s') = 1$ for each s), $\mathbf{Q} : S \times S \times (\mathbb{R}_{\geq 0} \rightarrow [0, 1])$ is a matrix of continuous probability distribution functions (such that $\mathbf{P}(s, s') = 0$ implies $\mathbf{Q}(s, s', t) = 1$), and $L : S \rightarrow 2^{AP}$ is the labelling function. Function L assigns to each state $s \in S$ the set $L(s)$ of atomic propositions $a \in AP$ that are valid in s .

The intuitive interpretation of a SMC is as follows. There exists a transition from state s to s' (which possibly equals s) if and only if $\mathbf{P}(s, s') > 0$. Matrix \mathbf{P} determines the (discrete) probabilistic behaviour when changing from one state to another, i.e., $\mathbf{P}(s, s')$ is the probability to move from state s to state s' . Note that this is identical to the probabilistic branching of a discrete-time Markov chain (DTMC); (S, \mathbf{P}, L) is often called the *embedded* DTMC of SMC \mathcal{M} . Once a next state s' of state s has been selected, the state holding time of state s is determined according to the probability distribution function $\mathbf{Q}(s, s', t)$. Thus, $\mathbf{Q}(s, s', t)$ denotes the probability to move from state s to s' within at most

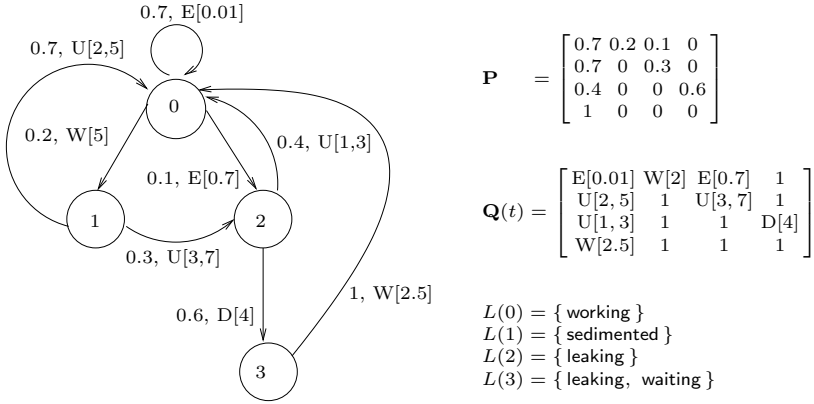


Fig. 1. A SMC describing a boiler.

t time-units, given that a transition from s to s' will be taken. A state s is absorbing if $\mathbf{P}(s, s) = 1$ and $\mathbf{Q}(s, s)$ is some arbitrary nontrivial distribution. The distribution function \mathbf{H} of state s , defined by

$$\mathbf{H}(s, t) = \sum_{s' \in S} \mathbf{P}(s, s') \cdot \mathbf{Q}(s, s', t)$$

denotes the total holding time distribution in s regardless of which successor is selected.

We assume that the system will stay in a state with at least some non-zero probability, or more formally we demand for arbitrary s that there is some $t' > 0$ and some $\varepsilon > 0$ such that $\mathbf{H}(s, t') < 1 - \varepsilon$. We further require the mean of each state holding time distribution to be finite, i.e., $E[\mathbf{H}(s)] \neq \infty$.

Example 1. As a simple example of a SMC we model a boiler. The system can be in four different states, state 0 where the boiler is working properly, state 1, where the boiler has too much sediment that needs to be removed, state 2 where a pipe is leaking that either needs to be fixed or needs to be replaced, and finally state 3 where the system is waiting for a new pipe to arrive for replacement. The model is schematically depicted in Fig. 1, together with the matrices \mathbf{P} and \mathbf{Q} ('E' denotes an exponential distribution, 'U' a uniform distribution, 'D' a deterministic distribution, and 'W' a Weibull distribution with appropriate parameters) and labelling L . The total holding time distributions can be computed from the matrices. For instance

$$\mathbf{H}(1, t) = \begin{cases} 0 & \text{if } t \leq 2, \\ 0.23 t - 0.46 & \text{if } 2 < t \leq 3, \\ 0.31 t - 0.7 & \text{if } 3 < t \leq 5, \\ 0.08 t - 0.24 & \text{if } 5 < t \leq 7, \\ 1 & \text{otherwise.} \end{cases}$$

To describe how the system evolves from state to state, suppose that the boilers starts in state 0. Matrix \mathbf{P} immediately determines the probability to move to a next state. State 2 is chosen, for instance with probability $\mathbf{P}(0, 2) = 0.1$. In this case a sample is immediately drawn from distribution $\mathbf{Q}(0, 2, t) = 1 - e^{-0.7 t}$, say 5.3. The system thus holds state 0 for 5.3 time units before moving to state 2. In state 2, again matrix \mathbf{P} is used to determine the next successor, say state 0, whence a random sample is drawn from the distribution $\mathbf{Q}(2, 0)$ to determine the holding time in state 2 before moving back to state 0. \square

Paths. Let $\mathcal{M} = (S, \mathbf{P}, \mathbf{Q}, L)$ be a SMC. A sequence $s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \xrightarrow{t_2} \dots$, with $s_i \in S$ and $t_i \in \mathbb{R}_{\geq 0}$ such that $\mathbf{P}(s_i, s_{i+1}) > 0$ for all i , is called a *path* through \mathcal{M} . For path σ and $i \in \mathbb{N}$, let $\sigma[i] = s_i$, the $(i+1)$ -st 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 with $t \leq \sum_{j=0}^i t_j$ let $\sigma@t = \sigma[i]$, the state in σ occupied at time t .

Let $Path^{\mathcal{M}}$ denote the set of paths in the SMC \mathcal{M} , and $Path^{\mathcal{M}}(s)$ the set of paths in \mathcal{M} that start in s . The superscript \mathcal{M} is omitted unless needed for distinction purposes.

Borel space. A probability measure \Pr on sets of paths through a SMC is defined using the standard cylinder construction as follows. Let $s_0, \dots, s_k \in S$ with $\mathbf{P}(s_i, s_{i+1}) > 0$, ($0 \leq i < k$), 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 $\sigma \in Path(s_0)$ such that $\sigma[i] = s_i$ ($i \leq k$), and $\delta(\sigma, i) \in I_i$ ($i < k$). Let $\mathcal{F}(Path)$ 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 $s = s_0$, $\mathbf{P}(s_i, s_{i+1}) > 0$ ($0 \leq i < k$), and I_0, \dots, I_{k-1} ranges over all sequences of non-empty intervals in $\mathbb{R}_{\geq 0}$. The probability measure \Pr on $\mathcal{F}(Path(s))$ is the unique measure defined by induction on k by $\Pr(C(s_0)) = 1$ and for $k \geq 0$:

$$\Pr(C(s_0, I_0, \dots, s_k, I', s')) = \Pr(C(s_0, I_0, \dots, s_k)) \cdot \mathbf{P}(s_k, s') \cdot (\mathbf{Q}(s_k, s', b) - \mathbf{Q}(s_k, s', a))$$

where $a = \inf I'$ and $b = \sup I'$. With this definition, a path $s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \dots$ corresponds to a sequence $(s_0, 0), (s_1, t_0), (s_2, t_0 + t_1), \dots$ of bivariate random variables satisfying the properties of Markov renewal sequences [18]. This observation links our definition of SMCs to the standard definition found in the literature.

On the basis of the probability measure \Pr , we can define various measures determining the behaviour of a SMC as time passes. For instance,

$$\pi(s, s', t) = \Pr\{\sigma \in Path(s) \mid \sigma@t = s'\}$$

defines the probability distribution on S (ranged over by s') at time t if starting in state s at time 0. We are particularly interested in two specific measures discussed below.




Fig. 2. A SMC without steady-state.

First passage time analysis. We are interested in a measure that describes the probability

$$F(s, s', t) = \Pr\{\sigma \in \text{Path}(s) \mid \exists t' \in [\delta(\sigma, 0), t] . \sigma @ t' = s'\}$$

of reaching state s' for the first time within t time units when starting in state s . Note that even if $s = s'$ only paths are considered that leave the state s , since t' has to be at least $\delta(\sigma, 0)$ which is the time needed to leave s . From [18] we have that $F(s, s', t)$ (with $s, s' \in S$) satisfies the following system of equations:

$$F(s, s', t) = \mathbf{P}(s, s') \mathbf{Q}(s, s', t) + \sum_{s'' \neq s'} \int_0^t \mathbf{P}(s, s'') \frac{d\mathbf{Q}(s, s'', x)}{dx} F(s'', s', t-x) dx$$

Intuitively, the probability to reach state s' from state s for the first time within t time units equals the sum of the probability of taking a direct transition from s to s' (within t time units) and the probability of moving via some intermediate state s'' at time x , yet reaching state s' in the remaining time interval $t - x$. It can be proven that this equation system has a unique solution if the state holding time for any state in the SMC is positive with nonzero probability (as we have assumed) [18].

Long-run average analysis. The long-run average behaviour of a SMC is not as homogeneous as it is for CTMCs. In particular the steady-state behaviour (usually defined as the limit of $\pi(s, s', t)$ for $t \rightarrow \infty$) may not exist.

Example 2. Consider for instance, the SMC depicted in Figure 2. For any $t \geq 0$ the probability $\pi(s, s', t)$ does not equal $\pi(s, s', t + 1)$, because the probability mass alternates between the two states. Thus, a limit for $t \rightarrow \infty$ of $\pi(s, s', t)$ does not exist. \square

However, we can define a related measure based on the average amount of time spent in some state, similar to [12]. For this purpose, we fix a state s , and let σ_s be a path taken randomly from the set $\text{Path}(s)$. Then, the quantity $\mathbf{1}_{s'}(\sigma_s @ t)$ is a random variable, indicating whether the state s' is occupied at time t when starting in s . Here we use the characteristic function $\mathbf{1}_{s'}(s'') = 1$ if $s' = s''$ and 0 otherwise.

On the basis of this, we can define a random variable that cumulates the time spent in some state s' up to time t (starting in s) by $\int_0^t \mathbf{1}_{s'}(\sigma_s @ x) dx$, and

normalise it by the time t in order to obtain a measure of the fraction of time spent in state s' up to time t . Since this is still a random variable, we can derive its expected value. This value corresponds to the average fraction of time spent in state s' in the time frame up to t . For the long-run average fraction of time, we consider the limit $t \rightarrow \infty$, as in [23].

Definition 1. *The average fraction of time $T(s, s')$ spent in state s' on the long run when starting in state s is given by:*

$$T(s, s') = \lim_{t \rightarrow \infty} E \left[\frac{1}{t} \int_0^t \mathbf{1}_{s'}(\sigma_s @ x) dx \right]$$

where σ_s ranges randomly over $\text{Path}(s)$.

This measure exists for SMCs whenever the expected values of all the distributions $\mathbf{Q}(s, s')$ are finite (as we have assumed). Note that for finite CTMCs the measure $T(s, s')$ agrees with the usual steady-state limit $\lim_{t \rightarrow \infty} \pi(s, s', t)$. In this sense, T conservatively extends the steady-state measure of CTMCs.

3 CSL on Semi-Markov Chains

This section recalls the syntax of the continuous stochastic logic CSL, and defines its semantics in terms of semi-Markov chains.

Syntax. CSL is a branching-time temporal logic à la CTL [9] with state- and path-formulas based on [5,4].

Definition 2. *Let $p \in [0, 1]$, $\trianglelefteq \in \{\leq, \geq\}$, $t \in \mathbb{R}_{\geq 0}$, and $a \in AP$. The syntax of CSL state-formulas is defined by the following grammar:*

$$\Phi ::= \mathbf{true} \mid a \mid \Phi \wedge \Phi \mid \neg \Phi \mid \mathcal{S}_{\trianglelefteq p}(\Phi) \mid \mathcal{P}_{\trianglelefteq p}(\varphi)$$

where for $t \in \mathbb{R}_{\geq 0}$ path-formulas are defined by

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

Other boolean connectives are derived in the usual way, i.e. $\mathbf{false} = \neg \mathbf{true}$, $\Phi_1 \vee \Phi_2 = \neg(\neg \Phi_1 \wedge \neg \Phi_2)$, and $\Phi_1 \rightarrow \Phi_2 = \neg \Phi_1 \vee \Phi_2$.

The intended meaning of the temporal operators \mathcal{U} (“until”) and X (“next step”) is standard. We recall from [5] the intuitive meaning of $\mathcal{U}^{\leq t}$, \mathcal{P} and \mathcal{S} : The path-formula $\Phi_1 \mathcal{U}^{\leq t} \Phi_2$ is satisfied iff there is some $x \in [0, t]$ such that Φ_1 continuously holds during the interval $[0, x[$ and Φ_2 becomes true at time instant x . $\mathcal{P}_{\trianglelefteq p}(\varphi)$ asserts that the probability measure of the paths satisfying φ falls in the interval $I_{\trianglelefteq p}$. The state formula $\mathcal{S}_{\trianglelefteq p}(\Phi)$ asserts that the long-run average fraction of time for a Φ -state falls in the interval $I_{\trianglelefteq p} = \{q \in [0, 1] \mid q \trianglelefteq p\}$. Temporal operators like \diamond , \square and their real-time variants $\diamond^{\leq t}$ or $\square^{\leq t}$ can be derived, e.g. $\mathcal{P}_{\trianglelefteq p}(\diamond^{\leq t} \Phi) = \mathcal{P}_{\trianglelefteq p}(\mathbf{true} \mathcal{U}^{\leq t} \Phi)$ and $\mathcal{P}_{\geq p}(\square \Phi) = \mathcal{P}_{\leq 1-p}(\diamond \neg \Phi)$.

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

$$\begin{array}{ll} s \models \mathbf{true} & \text{for all } s \in S, \\ s \models a & \text{iff } a \in L(s), \\ s \models \neg\Phi & \text{iff } s \not\models \Phi, \\ s \models \Phi_1 \wedge \Phi_2 & \text{iff } s \models \Phi_i, i \in \{1, 2\}, \\ s \models \mathcal{S}_{\leq p}(\Phi) & \text{iff } T_{\text{Sat}(\Phi)}(s) \in I_{\leq p}, \\ s \models \mathcal{P}_{\leq p}(\varphi) & \text{iff } \text{Prob}(s, \varphi) \in I_{\leq p}. \end{array}$$

Here, $T_{S'}(s)$ denotes the average fraction of time spent in $S' \subseteq S$ with respect to state s , i.e.

$$T_{S'}(s) = \sum_{s' \in S'} T(s, s').$$

Recall that $T(s, s')$ conservatively extends the definition of a steady-state distribution for CTMCs. $\text{Prob}(s, \varphi)$ denotes the probability measure of all paths $\sigma \in \text{Path}(s)$ satisfying φ , i.e.

$$\text{Prob}(s, \varphi) = \Pr\{\sigma \in \text{Path}(s) \mid \sigma \models \varphi\}.$$

The fact that, for each state s , the set $\{\sigma \in \text{Path}(s) \mid \sigma \models \varphi\}$ is measurable, follows by easy verification. The satisfaction relation (also denoted \models) for the path-formulas is defined as usual:

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

4 Model Checking SMCs against CSL

Model checking SMCs against CSL follows the usual strategy: Given a model $\mathcal{M} = (S, \mathbf{P}, \mathbf{Q}, L)$ and a state-formula Φ , the set $\text{Sat}(\Psi)$ is recursively computed for the sub-formulas of Φ . This can proceed via well studied means [16,5] (on the embedded DTMC (S, \mathbf{P}, L)) except for the time-bounded until operator $\mathcal{U}^{\leq t}$, and for the long-run operator \mathcal{S} . These two operators require specific care.

Time-bounded until. For computing the probability of satisfying a time-bounded until formula, we closely follow the strategy of [6], and reduce the problem to a well studied transient measure. More precisely, it will turn out that we can compute the time-bounded until probabilities via a first passage time analysis in a derived SMC, where certain subsets of states are made absorbing. To this end, we let $\mathcal{M}[\Phi]$ (for SMC \mathcal{M} and state formula Φ) denote the SMC obtained from \mathcal{M} by making all Φ -states absorbing. We have:

Theorem 1. *Let $\mathcal{M} = (S, \mathbf{P}, \mathbf{Q}, L)$ be a SMC, and Φ_1 and Φ_2 be CSL state-formulas. Then*

$$\begin{aligned} \text{Prob}^{\mathcal{M}}(s, \Phi_1 \mathcal{U}^{\leq t} \Phi_2) &= \text{Prob}^{\mathcal{M}[\neg\Phi_1 \vee \Phi_2]}(s, \diamond^{\leq t} \Phi_2) \\ &= \begin{cases} 1 & \text{if } s \models \Phi_2, \\ \sum_{s' \models \Phi_2} F^{\mathcal{M}[\neg\Phi_1 \vee \Phi_2]}(s, s', t) & \text{otherwise.} \end{cases} \end{aligned}$$




Fig. 3. Satisfaction of $\mathcal{P}_{\geq p}(\text{leaking } \mathcal{U}^{\leq t} \text{working})$ for state 2 and 3 of the boiler example.

Proof: The proof of the first equality is based on a bijection between the paths in \mathcal{M} satisfying $\Phi_1 \mathcal{U}^{\leq t} \Phi_2$ and the paths in $\mathcal{M}[-\Phi_1 \vee \Phi_2]$ satisfying $\diamond^{\leq t} \Phi_2$, up to the state where Φ_2 becomes satisfied, and hence over the whole path-prefixes contributing to the two probability measures $\Pr\{\sigma \in \text{Path}^{\mathcal{M}} \mid \sigma \models \Phi_1 \mathcal{U}^{\leq t} \Phi_2\}$ and $\Pr\{\sigma \in \text{Path}^{\mathcal{M}[-\Phi_1 \vee \Phi_2]} \mid \sigma \models \diamond^{\leq t} \Phi_2\}$. With respect to the second equality we only consider the case $s \neq \Phi_2$. In this case $\sigma \models \diamond^{\leq t} \Phi_2$ can be shown to hold if and only if $\exists t' \in [\delta(\sigma, 0), t] . \sigma @ t' \models \Phi_2$, since $\sigma[0] \not\models \Phi_2$. The proof follows from the definition of F and the fact that Φ_2 -states are absorbing, justifying the summation over all Φ_2 -states. \square

Example 3. Returning to the boiler example of Fig. 1, let us check the time-bounded until formula $\mathcal{P}_{\geq p}(\text{leaking } \mathcal{U}^{\leq t} \text{working})$. First, we observe that state 1 does neither satisfy *leaking* nor *working*, and hence state 1 does not satisfy the path-formula $\text{leaking } \mathcal{U}^{\leq t} \text{working}$ with positive probability. In contrast, according to Theorem 1, state 0 satisfies the path formula with probability 1, because $0 \models \text{working}$.

The remaining states 2 and 3 are more interesting. Following Theorem 1 we need to investigate a SMC where state 0 and 1 are made absorbing, and compute the probability of satisfying $\text{leaking } \mathcal{U}^{\leq t} \text{working}$ via the values of $F(2, 0, t)$, respectively $F(3, 0, t)$ in this SMC. The values of these functions are plotted in Fig. 3. One can see that for pairs (p, t) above the plot the formula is invalid, while it is valid for pairs below the plot (and for the plot itself). \square

While in the above example the values of F can be calculated directly, the situation is more involved in general. Recall that $F(s, s', t)$ is the unique solution of the system of equations

$$F(s, s', t) = \mathbf{P}(s, s') \mathbf{Q}(s, s', t) + \sum_{s'' \neq s'} \int_0^t \mathbf{P}(s, s') \frac{d\mathbf{Q}(s, s'', x)}{dx} F(s'', s', t-x) dx.$$

This system of equations can be classified as a system of Volterra equations of the second type. In principle it is possible to solve them by appropriate numerical

methods, such as Volterra-Runge-Kutta methods. A complete guide for these methods can be found in [7]. A solution of the equations can also be obtained in the Laplace domain. This approach works good for small systems and sometimes even allows a closed-form solution to be found by hand. For larger systems one is faced with two problems. One has to invert a matrix of functions in a complex variable, and to reverse the transform to the time domain.

As described in [14,15], the asymptotic space complexity of the latter method is $\mathcal{O}(N^2)$ and the asymptotic time complexity is $\mathcal{O}(N^4)$ where $N = |S|$ is the number of states. It is therefore not applicable to larger systems. Moreover, the numerical Laplace transform inversion can encounter numerical problems under some conditions. It is also possible to solve this Volterra system by transforming it to a system of partial differential equation, a system of ordinary differential equation, initial and boundary conditions, and a system of integral equations. [14] contains a comparison of these two approaches together with numerical considerations.

Long-run average. For model checking the operator $\mathcal{S}_{\leq p}(\Phi)$ one needs to accumulate the average fraction of time quantities $T(s, s')$ for each state s' satisfying Φ . If \mathcal{M} is a strongly connected¹ SMC, $T(s, s')$ can be obtained via the equilibrium probability vector π of the embedded DTMC (S, \mathbf{P}, L) , which in turn is given as the unique solution of the linear equation system

$$\pi(s) = \sum_{s' \in S} \mathbf{P}(s', s) \cdot \pi(s') \text{ such that } \sum_{s \in S} \pi(s) = 1.$$

Theorem 2. [8] *Let $\mathcal{M} = (S, \mathbf{P}, \mathbf{Q}, L)$ be a strongly connected SMC, and π be as above. Then*

$$T(s, s') = \frac{\pi(s')\mu(s')}{\sum_{s'' \in S} \mu(s'')\pi(s'')}$$

where $\mu(s'')$ is the expected holding time in state, i.e., $\mu(s'') = E[\mathbf{H}(s'')]$.

Notice that $T(s, s')$ is independent of the starting state s in this case. If otherwise \mathcal{M} is not strongly connected, we proceed as in [5], and isolate the bottom strongly connected subsets of S via a graph algorithm [22]. Whenever state s' is not a member of any bottom strongly connected subset of S , we have $T(s, s') = 0$. The following result allows model checking the \mathcal{S} operator in the other cases. We let $\text{Prob}(s, \diamond B)$ denote the probability of eventually reaching the set $B \subseteq S$ from state s . This quantity can be computed via the embedded DTMC (S, \mathbf{P}, L) [16].

Theorem 3. *Let $\mathcal{M} = (S, \mathbf{P}, \mathbf{Q}, L)$ be a SMC, B a bottom strongly connected subset of S , and $s' \in B$. Then:*

$$T(s, s') = \text{Prob}(s, \diamond B) \cdot T^B(s', s')$$

where the superscript B refers to the strongly connected SMC \mathcal{M}^B spanned by B .

¹ A SMC is strongly connected if there is some k such that $\mathbf{P}^k(s, s') > 0$ for each s, s' .

Proof: We only consider the case where $s \notin B$ can reach B with positive probability. The idea of the proof is to count the average time that the SMC \mathcal{M} spends in class B . Once we have isolated this quantity we are able to compute the fraction of time \mathcal{M} spends in a particular state of this class. Let $\mathbf{1}_B(s') = 1$ if $s' \in B$ and 0 otherwise. We shall calculate the exact value for

$$E \left[\frac{1}{t} \int_0^t \mathbf{1}_B(\sigma_s @ x) dx \right]$$

where σ_s ranges over $Path(s)$. Let \tilde{t} be the time of absorption in B (if σ_s touches B otherwise $\tilde{t} = \infty$), \tilde{t} is a random variable and depends on the path σ_s drawn from $Path(s)$. The distribution of \tilde{t} is given by $\Pr\{\tilde{t} \leq t'\} = F(s, B, t') = \sum_{s' \in B} F(s, s', t')$ where the latter is the first passage time distribution mentioned earlier.

Since B is bottom strongly connected, the function $\mathbf{1}_B(\sigma_s @ x)$ will be constant 1 from \tilde{t} on. So, for $t \geq \tilde{t}$ we have that

$$\int_0^t \mathbf{1}_B(\sigma_s @ x) dx = \frac{t - \tilde{t}}{t}$$

and otherwise (i.e., $t < \tilde{t}$) the integral equals 0. So, for fixed t the above integral describes a random variable R_t as follows:

$$R_t(\tilde{t}) = \begin{cases} \frac{t - \tilde{t}}{t} & \text{if } t \geq \tilde{t}, \\ 0 & \text{otherwise.} \end{cases}$$

The distribution of R_t is

$$\Pr\{R_t \leq x\} = \Pr\{(t - \tilde{t})/t \leq x\} + \Pr\{R_t = 0\}$$

which can be rewritten, using that $F(s, B, x)$ is the distribution of \tilde{t} , to

$$\Pr\{R_t \leq x\} = 1 - F(s, B, t - xt) + \Pr\{R_t = 0\}.$$

Now, the expected value $E[R_t]$ is obviously

$$\int_0^1 u \frac{d(1 - F(s, B, t - ut))}{du} du + 0 \frac{d(\Pr\{R_t = 0\})}{du} = \int_0^1 u t \frac{dF(s, B, t - ut)}{du} du.$$

Substituting $u = \frac{t-y}{t}$ we get

$$E[R_t] = \frac{1}{t} \int_0^t (t-y) \frac{dF(s, B, y)}{dy} dy.$$

What we are looking for is the limit of this quantity as $t \rightarrow \infty$ given by

$$\lim_{t \rightarrow \infty} F(s, B, t) - \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t y \frac{dF(s, B, y)}{dy} dy = \text{Prob}(s, \diamond B) - \lim_{t \rightarrow \infty} \frac{E[\tilde{t}]}{t}$$

where $E[\tilde{t}]$ is the expected value of \tilde{t} . Recall that \tilde{t} is distributed according to $F(s, B, t)$. Also note that $\lim_{t \rightarrow \infty} F(s, B, t) = \text{Prob}(s, \diamond B)$. Since we have assumed that state s can reach B with positive probability (and all distributions have finite means) $E[\tilde{t}]$ needs to be finite and hence

$$\lim_{t \rightarrow \infty} E \left[\frac{1}{t} \int_0^t \mathbf{1}_B(\sigma_s @ x) dx \right] = \text{Prob}(s, \diamond B).$$

The proof of the theorem follows from this result by two observations. First, the time of entering B is a renewal point, i.e., a time instant where the future behaviour of the stochastic process does only depend on the currently occupied state. Second, the fraction of time spent in a particular state inside B is independent of the starting state – due to strong connectedness – if assuming to start inside B . \square

Example 4. Let us check a long-run average property for the example boiler system, such as $\mathcal{S}_{\leq p}(\text{working})$. We first observe that the SMC in Fig. 1 is strongly connected. Theorem 2 requires the computation of the expected holding times for each state of the SMC, resulting from weighted sums of the involved distributions. We get $\mu(0) = 70.319$, $\mu(1) = 3.95$, $\mu(2) = 3.2$, and $\mu(3) = 0.887$. Next, we solve the embedded DTMC, and obtain a vector $\pi = [0.686, 0.1373, 0.109, 0.065]$. Finally we compute $T_{\text{Sat}(\text{working})}(s) = 0.981$. Since the SMC is strongly connected, this value is independent of the state s chosen, and hence $\mathcal{S}_{\leq p}(\text{working})$ is satisfied (for all states) whenever $0.981 \leq p$. \square

Apart from the need to derive expected values of general distributed random variables, the numerical algorithms needed for model checking the long-run average operator are the same as the ones needed for checking CTMCs [5].

5 Concluding Remarks

In this paper, we investigated adapting CSL model checking to semi-Markov chains, an extension of CTMCs in which state holding times are governed by general distributions. To achieve a smooth extension of the theory we developed an enhanced definition of long-run properties and proved novel results required for model checking not strongly connected SMCs. On the practical side, the conclusion we draw from our investigation is partially negative: verifying a CSL-formula can become numerically very complex when dropping the memoryless property. This is caused by the involved procedure needed for checking time-bounded formulas such as timed probabilistic reachability properties. We proved that long-run properties and (untimed) eventualities can be checked without an increase in complexity compared to the CTMC case, though.

The SMC model considered in this paper incorporates general distributions, but is known to be of limited use to model concurrent delays. Compositional extensions of SMCs – such as generalised semi-Markov chains or stochastic automata [11]– are more elegant to apply in this context. It is worth to highlight that our practically negative result concerning the model checking of time-bounded formulas carries over to these models. Further research is needed to

investigate whether abstraction techniques or weaker temporal properties – like expected time properties – yield a practical solution for such models.

References

1. R. Alur, C. Courcoubetis and D. Dill. Model-checking in dense real-time. *Inf. and Comp.*, **104**: 2–34, 1993.
2. S. Asmussen, O. Nerman and M. Olsson. Fitting phase-type distributions via the EM algorithm. *Scand. J. Statist.*, **23**: 420–441, 1996.
3. A. Aziz, K. Sanwal, V. Singhal and R. Brayton. Verifying continuous time Markov chains. In R. Alur and T.A. Henzinger (eds), *Computer-Aided Verification*, LNCS 1102, pp. 269–276, Springer, 1996.
4. A. Aziz, K. Sanwal, V. Singhal and R. Brayton. Model checking continuous time Markov chains. *ACM Trans. on Computational Logic*, **1**(1): 162–170, 2000.
5. C. Baier, J.-P. Katoen and H. Hermanns. Approximate symbolic model checking of continuous-time Markov chains. In J.C.M. Baeten and S. Mauw (eds), *Concurrency Theory*, LNCS 1664, pp. 146–162, Springer, 1999.
6. C. Baier, B. R. Haverkort, H. Hermanns and J.-P. Katoen. Model checking continuous-time Markov chains by transient analysis. In E.A. Emerson and A.P. Sistla (eds), *Computer Aided Verification*, LNCS 1855, pp. 358–372, Springer, 2000.
7. H. Brunner and P. van der Houwen, *The Numerical Solution of Volterra Equations*. North Holland, 1986.
8. E. Cinlar. *Introduction to Stochastic Processes*. Prentice-Hall Inc., 1975.
9. E. Clarke, E. Emerson and A. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems*, **8**: 244–263, 1986.
10. M.E. Crovella. Performance evaluation with heavy tailed distributions (extended abstract). In B. Haverkort, H. Bohnenkamp and C. Smith (eds), *Computer Performance Evaluation*, LNCS 1786, pp. 1–9, Springer, 2000.
11. P.R. D’Argenio, J.-P. Katoen and E. Brinksma. An algebraic approach to the specification of stochastic systems (extended abstract). In D. Gries and W.-P. de Roever (eds), *Programming Concepts and Methods*. Chapman & Hall, pp. 126–147, 1998.
12. L. de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, 1997.
13. A. Feyzi Ates, M. Bilgic, S. Saito, and B. Sarikaya. Using timed CSP for specification, verification and simulation of multimedia synchronization. *IEEE J. on Sel. Areas in Comm.*, **14**:126–137, 1996.
14. R. German *Performance Analysis of Communication Systems: Modeling with Non-Markovian Stochastic Petri Nets*. John Wiley & Sons, 2000.
15. R. German, D. Logothe, and K.S. Trivedi. Transient analysis of Markov regenerative stochastic Petri nets: A comparison of approaches. *Proc. 6th Int. Workshop on Petri Nets and Performance Models*, pages 103–112, IEEE CS Press, 1995.
16. H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing* **6**: 512–535, 1994.
17. D. Van Hung and Z. Chaochen. Probabilistic duration calculus for continuous time. *Formal Aspects of Computing*, **11**: 21–44, 1999.
18. V. Kulkarni. *Modeling and Analysis of Stochastic Systems*. Chapman & Hall, 1995.

19. M.Z. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. Verifying quantitative properties of continuous probabilistic timed automata. In C. Palamadessi (ed), *Concurrency Theory*, LNCS 1877, pp. 123-137, Springer, 2000.
20. W. Nelson. Weibull analysis of reliability data with few or no failures. *Journal of Quality Technology* **17** (3), 140-146, 1985.
21. A.D. Swain and H.E. Guttmann. Handbook of human reliability analysis with emphasis on nuclear power plant applications - final report. Technical Report NRC FIN A 1188 NUREG/CR-1278 SAND80-0200, US Nuclear Regulatory Commission, 1983.
22. R.E. Tarjan. Depth-first search and linear graph algorithms. *SIAM Journal of Computing*, **1**: 146–160, 1972.
23. H. Taylor and S. Karlin. *An Introduction To Stochastic Modeling*. Academic Press, 1998.