

Model Checking Expected Time and Expected Reward Formulae with Random Time Bounds

MARTA KWIATKOWSKA¹, GETHIN NORMAN¹ AND ANTÓNIO PACHECO²

¹School of Computer Science, University of Birmingham, Edgbaston,
Birmingham B15 2TT, United Kingdom

²Instituto Superior Tecnico, Department of Mathematics,
CEMAT and CLC, Av. Rovisco Pais, 1049-001 Lisboa, Portugal

Abstract—In this paper we extend CSL (Continuous Stochastic Logic) with an *expected time* and an *expected reward* operator, both of which are parameterized by a random terminal time. With the help of such operators we can state for example that the expected sojourn time in a set of goal states within some generally distributed delay is at most (at least) some time threshold. In addition, certain performance measures of systems which contain general distributions can be calculated with the aid of this extended logic. We extend the efficient model checking of CTMCs against the logic CSL developed by Katoen *et al.* [1] to cater for the new operator. Our method involves precomputing a family of mixed Poisson expected sojourn time coefficients for a range of random variables which includes Pareto, uniform and gamma distributions, but otherwise carries the same computational cost as calculating CSL until formulae. © 1900 Elsevier Science Ltd. All rights reserved.

Keywords—Continuous time Markov chains, Continuous Stochastic Logic, Expected time operator, Expected reward operator, Markov reward processes, Mixed Poisson expected sojourn times, Uniformisation.

1. INTRODUCTION

Continuous time Markov chains (CTMCs) form an important class of models widely used in performance and dependability analysis. The class is characterised by allowing only exponential distributions – the time that the system remains in a state is given by an exponential distribution. This restriction allows one to employ well established efficient analysis techniques for both transient and steady-state probabilities, and hence also for determining standard performance measures such as throughput, mean waiting time and average cost. Recently extensions of temporal logic have been proposed which can express such properties. The temporal logic CSL (Continuous Stochastic Logic) introduced by Aziz *et al.* [2,3] and since extended by Baier *et al.* [4] is based on the temporal logics CTL [5] and PCTL [6] and provides a powerful means to specify both path-based and traditional state-based performance measures on CTMCs. CSL contains a time-bounded until operator that allows one to express properties such as “the probability of 3 servers becoming faulty within 7.01 seconds is at most 0.1”. Model checking of CTMCs against

This research was partly supported by EPSRC grant GR/N22960, FCT grant SFRH/ BSAB/251/01, the projects POSI/42069/CPS/2001, POSI/EEA-CPS/58154/2004, and POCTI/MAT/55796/2004, and the Anglo-Portuguese Joint Research Programme - Treaty of Windsor 2003 - Action No. - 20-A/03.

CSL has been improved in [7,1] through the use of uniformisation [8,9] and transient analysis, and implemented in the tool PRISM [10,11]. The usefulness of this approach has been demonstrated by a number of case studies ranging from a wireless cell to a flexible manufacturing system.

However, in practice it is often the case that exponential distributions are not an adequate modelling tool for capturing the behaviour of stochastic systems. Examples of such situations include modelling file transfer over the Internet, timeouts in communication protocols and the residence time in a wireless cell. For these cases the modelling framework must be capable of handling *general distributions*, such as Pareto, Erlang, gamma or phase-type. An unfortunate consequence of including general distributions within the modelling framework, as has been demonstrated recently, for example in [12] and [13], is a considerable increase in the complexity of performance analysis, or, if using phase-type distributions, a substantial increase in the size of the state space.

In [14] we made an alternative proposal, namely, to remain in the CTMC framework and instead extend the logic CSL with a variant of the time-bounded until operator which allows generally distributed *random time bounds*. In this paper we further extend CSL with *expected time* and *expected reward* operators which are parameterized by a generally distributed random time bound. We extend the efficient model checking of CTMCs against the logic CSL developed in [1] to cater for the new operator. Our method involves precomputing a family of coefficients for a range of random variables which includes Pareto, uniform and gamma distributions, but otherwise carries the same computational cost as calculating CSL until formulae.

With the help of these operators we can state that the expected sojourn time in a set of goal states or the expected reward within some random delay is at most (at least) some time threshold. Although general distributions cannot be added explicitly to the model in a completely arbitrary way, with the aid of this extended logic, one can establish certain performance measures of systems which include generally distributed delays. As an example application consider a queue where the customers arrive with some generally distributed delay, then by letting the random time bound have *the same* distribution we can express (and verify) properties such as: “if the queue is full then, when the next customer arrives, the expected number of customers in the queue is at most k ”. The applicability of this approach is demonstrated in [11,15] and Section 4 where this extended logic has been used to analyse power management systems when the inter-arrival rate of jobs has a general distribution.

Outline of paper. We begin by recalling the definition of CTMCs and the logic CSL. Next we introduce the new expected time and expected cost operators both parameterised by a random time bound, give their semantics and a model checking algorithm which uses a family of coefficients called *mixed Poisson expected sojourn times*. Next we show that these coefficients can be calculated by means of the algorithms developed in [14] for generating *mixed Poisson probability* coefficients. In the remainder of the paper we describe some experimental results for a power management system example using these operators.

2. PRELIMINARIES

In this section we briefly recall basic concepts we require in the remainder of the paper.

Random variables. Let T denote a *nonnegative random variable*. We let F (\bar{F}) denote the distribution (survival) function of T , i.e. $F(t) = P(T \leq t)$ ($\bar{F}(t) = P(T > t)$), and $E[T]$ its expected value, that is $E[T] = \int_0^\infty \bar{F}(t) dt$.

Continuous time Markov Chains. Let AP be a finite set of atomic propositions. A (labelled) CTMC \mathcal{M} is a tuple (S, Λ, L) where S is a finite set of *states*, $\Lambda : S \times S \rightarrow \mathbb{R}_+$ is the *rate matrix* and $L : S \rightarrow 2^{\text{AP}}$ is a *labelling* function which assigns to each state s the set $L(s)$ of atomic propositions valid in s . For any state $s \in S$, the probability of leaving state s within time t is

given by $1 - e^{-E(s) \cdot t}$ where $E(s) = \sum_{s' \in S} \Lambda(s, s')$.

A path through a CTMC is an alternating sequence $\sigma = s_0 t_0 s_1 t_1 s_2 \dots$ such that $\Lambda(s_i, s_{i+1})$ and t_i are positive for all i . The time stamps t_i denote the amount of time spent in state s_i . Let $Path^{\mathcal{M}}(s)$ denote the set of paths of \mathcal{M} which start in state s (i.e. $s_0 = s$); $\sigma@t$ denote the state of σ occupied at time t , i.e. $\sigma@t = \sigma[i]$ where i is the largest index such that $\sum_{j=0}^{i-1} t_j \leq t$; and P_s denote the unique probability measure on sets of paths that start in s [4]. Let $\pi^{\mathcal{M}}(s, t)(s')$ denote the probability of being in state s' at time t given that the system starts in state s , i.e. $\pi^{\mathcal{M}}(s, t)(s') = P_s\{\sigma \in Path^{\mathcal{M}}(s) : \sigma@t = s'\}$. Moreover, let $E_s[\cdot] = E[\cdot | \mathcal{M}(0) = s]$ for $s \in S$, that is E_s is the expected value operator conditional to the CTMC \mathcal{M} starting in state s .

Reward structure. Often a CTMC $\mathcal{M} = (S, \Lambda, L)$ is extended with a reward or cost structure [16]. It takes the form of a tuple $(\underline{r}, \mathbf{R})$, where for any $s, s' \in S$: \underline{r}_s denotes the rate at which the reward (cost) is incurred *continuously* in s , and $\mathbf{R}_{s,s'}$ denotes the *instantaneous* reward (cost) associated with the transition from s to s' , where $\mathbf{R}_{s,s} = 0$ for all s . The reward Markov process associated to \mathcal{M} is then $(\mathcal{M}, \mathcal{R}^{\mathcal{M}})$, where the reward accumulated in the interval of time $(0, t]$ is

$$\mathcal{R}^{\mathcal{M}}(t) = \int_0^t \underline{r}_{\mathcal{M}(u)} du + \int_0^t \mathbf{R}_{\mathcal{M}(u-), \mathcal{M}(u)} dN^{\mathcal{M}}(u) \quad (1)$$

where $N^{\mathcal{M}}$ is the counting process of state transitions in \mathcal{M} , i.e., $N^{\mathcal{M}}(u)$ is the number of state transitions of \mathcal{M} in the time interval $(0, u]$.

Uniformisation. For CTMC $\mathcal{M} = (S, \Lambda, L)$ the *embedded uniformised* discrete time Markov chain (DTMC) (with uniformisation rate $q \geq \max\{E(s) : s \in S\}$) is $u(\mathcal{M}) = (S, \mathbf{P}^{u(\mathcal{M})}, L)$, where $\mathbf{P}^{u(\mathcal{M})} = \mathbf{I} + [\Lambda - \text{diag}(\underline{E})]/q$ is its transition probability matrix. We let $N_q = \{N_q(t), t \geq 0\}$ denote the (uniformising) Poisson process with rate q , independent of $u(\mathcal{M})$. Then (see, e.g., [17]), the original CTMC $\{\mathcal{M}(t), t \geq 0\}$ has the same distribution as the uniformised CTMC $\{\mathcal{M}_q(t) = u(\mathcal{M})_{N_q(t)}, t \geq 0\}$. Hence, the distribution of the CTMC \mathcal{M} can be characterised completely through the distribution of the embedded uniformised DTMC $u(\mathcal{M})$. In particular, the probabilities $\pi^{\mathcal{M}}(s, t)(s')$ can be computed as follows:

$$\underline{\pi}^{\mathcal{M}}(s, t) = \underline{\pi}^{\mathcal{M}}(s, 0) \cdot \sum_{k=0}^{\infty} e^{-q \cdot t} \frac{(q \cdot t)^k}{k!} [\mathbf{P}^{u(\mathcal{M})}]^k = \sum_{k=0}^{\infty} \gamma(k, q \cdot t) \cdot \underline{\pi}^{u(\mathcal{M})}(s, k) \quad (2)$$

where $\gamma(k, q \cdot t) = e^{-q \cdot t} \cdot (q \cdot t)^k / k!$ is the k th Poisson probability with parameter $q \cdot t$, and the vector $\underline{\pi}^{u(\mathcal{M})}(s, k)$ denotes the probability distribution in $u(\mathcal{M})$ after k epochs when starting in s , i.e. $\underline{\pi}^{u(\mathcal{M})}(s, k) = \underline{\pi}^{\mathcal{M}}(s, 0) \cdot [\mathbf{P}^{u(\mathcal{M})}]^k$, where $\pi^{\mathcal{M}}(s, 0)(s) = 1$ and $\pi^{\mathcal{M}}(s, 0)(s') = 0$ if $s \neq s'$.

The logic CSL. Let $a \in \text{AP}$, $p \in [0, 1]$, $\bowtie \in \{\leq, \geq\}$ and $t \in \mathbb{R}_+$ (or ∞). The syntax of CSL is:

$$\Phi ::= \text{true} \mid a \mid \Phi \wedge \Phi \mid \neg \Phi \mid \mathcal{S}_{\bowtie p}[\Phi] \mid \mathcal{P}_{\bowtie p}[\Phi \mathcal{U}^{\leq t} \Phi].$$

The semantics of CSL for the boolean operators is identical to that for CTL [5]. $\mathcal{S}_{\bowtie p}[\Phi]$ asserts that the steady-state probability for a Φ -state meets the bound $\bowtie p$, whereas $\mathcal{P}_{\bowtie p}[\Phi \mathcal{U}^{\leq t} \Phi]$ asserts that with probability $\bowtie p$, by the time t a state satisfying Ψ will be reached such that all preceding states satisfy Φ . CSL model checking algorithms can be found in [4,7,1].

In [14] the logic CSL was extended to include *random* time-bounded until formulae of the form $\mathcal{P}_{\bowtie p}[\Phi \mathcal{U}^{\leq T} \Psi]$, where T is a nonnegative random variable. The formula asserts that, with probability $\bowtie p$, by the random time T a state satisfying Ψ will be reached such that all preceding states satisfy Φ .

Mixed Poisson probabilities. For a nonnegative random variable T , we let $\alpha_T(k, q) = \int_0^\infty \gamma(k, q \cdot t) dF(t)$ denote its associated k th *mixed Poisson probability*, which is equal to the probability that exactly k renewals take place in the uniformising Poisson process with rate q until the random time T . As reported in [14], for any $k \in \mathbb{N}$:

$$\alpha_{U_{[0,t]}}(k, q) = \int_0^t \frac{1}{t} \cdot \gamma(k, q \cdot u) du = \frac{1}{q \cdot t} \sum_{j=k+1}^{\infty} \gamma(k, q \cdot t) \quad (3)$$

where $U_{[0,t]}$ denotes a random variable with uniform distribution on $[0, t]$. Efficient algorithms for computing the mixed Poisson probabilities when the distribution of the random time T is a finite discrete, uniform, gamma or Pareto distribution, or is a finite mixture of distributions of these types are provided in [14].

3. EXPECTED TIME AND REWARD FORMULAE WITH RANDOM TIME-BOUNDS

We now extend the logic CSL to include *expected time* and *expected reward* operators which are parameterized by a generally distributed nonnegative random variable T and consider model checking algorithms for such formulae. The formulae we introduce are of the form $\mathcal{E}_{\bowtie V}^{\leq T}[\Psi]$ and $\mathcal{ER}_{\bowtie V}^{\leq T}$, where T is independent of the CTMC under study, Ψ is a CSL formula and $V \in \mathbb{R}_+$. The formula $\mathcal{E}_{\bowtie V}^{\leq T}[\Psi]$ asserts that the expected amount of time on the interval $(0, T]$ that Ψ is satisfied is $\bowtie V$. Similarly, the formulae $\mathcal{ER}_{\bowtie V}^{\leq T}$ is true in a state if the the expected reward (cost) before the random terminal time T is $\bowtie V$. To introduce the semantics of these operators, we let $E^{\mathcal{M}}(s, \mathcal{E}^{\leq t} \Psi)$ denote the expected time until t that Ψ is satisfied when starting from state s , and $E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(t))$ denote the expected reward until time t starting from state s . Hence,

$$s \models \mathcal{E}_{\bowtie V}^{\leq T}[\Psi] \Leftrightarrow E^{\mathcal{M}}(s, \mathcal{E}^{\leq T} \Psi) \bowtie V \quad \& \quad s \models \mathcal{ER}_{\bowtie V}^{\leq T} \Leftrightarrow E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(T)) \bowtie V$$

where

$$E^{\mathcal{M}}(s, \mathcal{E}^{\leq T} \Psi) = \int_0^\infty E^{\mathcal{M}}(s, \mathcal{E}^{\leq t} \Psi) dF(t) \quad \text{and} \quad E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(T)) = \int_0^\infty E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(t)) dF(t).$$

Here $\mathcal{R}^{\mathcal{M}}(t)$ is given by (1) and $\mathcal{E}^{\leq t} \Psi = \int_0^t 1(\mathcal{M}(u) \models \Psi) du$, where $1(A)$ is the indicator function of A , i.e., $1(A)$ is one if A is true and is zero otherwise. Note that, in view of (1), $\mathcal{E}_{\bowtie V}^{\leq T}[\Psi]$ reduces to $\mathcal{ER}_{\bowtie V}^{\leq T}$ using the reward tuple $(\underline{\mathcal{L}}_\Psi, \mathbf{0})$, where: $\underline{\mathcal{L}}_\Psi$ characterises $Sat(\Psi)$, i.e. $\underline{\mathcal{L}}_\Psi(s) = 1$ if $s \models \Psi$, and 0 otherwise, and $\mathbf{0}$ is a matrix with all entries null.

In the following sections we develop model checking algorithms for formulae of the type $\mathcal{ER}_{\bowtie V}^{\leq T}$ based on the approach used for verifying time-bounded until formulae [1]. These algorithms may then be specialized to model check formulae of the type $E^{\mathcal{M}}(s, \mathcal{E}^{\leq T} \Psi) \bowtie V$ by setting $(\underline{r}, \mathbf{R}) = (\underline{\mathcal{L}}_\Psi, \mathbf{0})$. We begin by considering deterministic terminal times and then proceed to random terminal times. Below, we presume \mathcal{M} is a CTMC with state space S and N_q denotes the uniformising Poisson process with rate q , independent of the embedded uniformised DTMC $u(\mathcal{M})$.

3.1. Deterministic Terminal Times

The following result summarizes how $E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(T))$ may be computed for a general reward structure $(\underline{r}, \mathbf{R})$ and a deterministic terminal time.

THEOREM 1. *For any $s \in S$ and $t \in \mathbb{R}_+$:*

$$E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(t)) = \sum_{k=0}^{\infty} \bar{\gamma}_t(k, q) \cdot \underline{\pi}^{u(\mathcal{M})}(s, k) \cdot \underline{f}_q(\underline{r}, \mathbf{R}) \quad (4)$$

where

$$\bar{\gamma}_t(k, q) = E \left[\int_0^t 1(N_q(u) = k) du \right] = \int_0^t \gamma(k, q \cdot u) du = \frac{1}{q} \sum_{j=k+1}^{\infty} \gamma(j, q \cdot t) \quad (5)$$

and

$$\underline{f}_q(\underline{r}, \mathbf{R}) = \underline{r} + q \cdot [\mathbf{P}^{u(\mathcal{M})} \bullet \mathbf{R}] \cdot \underline{1} \quad (6)$$

with \bullet denoting the Schur or entrywise multiplication of matrices and $\underline{1}$ a vector with unitary entries.

Proof. We start by proving (5), which follows since, using dominated convergence [18] and the fact that $\gamma(k, q \cdot u) = P(N_q(u) = k) = E[1(N_q(u) = k)]$:

$$\bar{\gamma}_t(k, q) = E \left[\int_0^t 1(N_q(u) = k) du \right] = \int_0^t \gamma(k, q \cdot u) du = \frac{1}{q} \sum_{j=k+1}^{\infty} \gamma(j, q \cdot t)$$

where the last equality follows from (3). Moreover, since \mathcal{M} has the same distribution as the uniformised CTMC $\{\mathcal{M}_q(u) = u(\mathcal{M})_{N_q(u)}, u \geq 0\}$, it follows from (1) that

$$E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(t)) = E_s \left[\int_0^t \underline{r}_{\mathcal{M}(u)} du \right] + E_s \left[\sum_{k=1}^{N_q(t)} \mathbf{R}_{u(\mathcal{M})_{k-1}, u(\mathcal{M})_k} \right]. \quad (7)$$

Note that $\int_0^t \underline{r}_{\mathcal{M}(u)} du = \int_0^t \sum_{s' \in S} \underline{r}_{s'} \cdot 1(\mathcal{M}(u) = s') du = \sum_{s' \in S} \underline{r}_{s'} \int_0^t 1(\mathcal{M}(u) = s') du$ and, using dominated convergence,

$$E_s \left[\int_0^t 1(\mathcal{M}(u) = s') du \right] = \int_0^t E_s [1(\mathcal{M}(u) = s')] du = \int_0^t \pi^{\mathcal{M}}(s, t)(s') du.$$

These facts and (2) lead, after rearranging terms, to

$$E_s \left[\int_0^t \underline{r}_{\mathcal{M}(u)} du \right] = \int_0^t \sum_{k=0}^{\infty} \gamma(k, q \cdot u) \cdot \underline{\pi}^{u(\mathcal{M})}(s, k) \cdot \underline{r} du = \sum_{k=0}^{\infty} \bar{\gamma}_t(k, q) \cdot \underline{\pi}^{u(\mathcal{M})}(s, k) \cdot \underline{r}. \quad (8)$$

Since $E_s[\mathbf{R}_{u(\mathcal{M})_{k-1}, u(\mathcal{M})_k}] = ([\mathbf{P}^{u(\mathcal{M})}]^{k-1} \cdot [\mathbf{P}^{u(\mathcal{M})} \bullet \mathbf{R}] \cdot \underline{1})_s$ and $E[X] = E[E[X|Y]]$ for any random variables X and Y , for any $s \in S$ and $t \in \mathbb{R}_+$ we have

$$\begin{aligned} E_s \left[\sum_{k=1}^{N_q(t)} \mathbf{R}_{u(\mathcal{M})_{k-1}, u(\mathcal{M})_k} \right] &= \sum_{n=1}^{\infty} \gamma(n, q \cdot t) \cdot E_s \left[\sum_{k=1}^n \mathbf{R}_{u(\mathcal{M})_{k-1}, u(\mathcal{M})_k} \right] \\ &= \sum_{n=1}^{\infty} \gamma(n, q \cdot t) \cdot \sum_{k=1}^n \underline{\pi}^{u(\mathcal{M})}(s, k-1) \cdot [\mathbf{P}^{u(\mathcal{M})} \bullet \mathbf{R}] \cdot \underline{1} \\ &= \sum_{k=0}^{\infty} \bar{\gamma}_t(k, q) \cdot \underline{\pi}^{u(\mathcal{M})}(s, k) \cdot q \cdot [\mathbf{P}^{u(\mathcal{M})} \bullet \mathbf{R}] \cdot \underline{1}. \end{aligned} \quad (9)$$

where for the last equality we have interchanged summations and used (5). Finally, (4) follows from (7)-(9). ■

From the definition (5), $\bar{\gamma}_t(k, q)$ is the expected sojourn time in state k of the uniformising Poisson process. Hence, we call $\bar{\gamma}_t(k, q)$ the k -th *Poisson expected sojourn time* coefficient (on the interval $[0, t]$ with associated rate q). Note that, from (5), it follows that the Poisson expected sojourn time coefficients are non-increasing, take values on the interval $[0, 1/q]$, and $\sum_{k=0}^{\infty} \bar{\gamma}_t(k, q) = t$. Thus, the expected sojourn times on $[0, t]$ of the uniformising Poisson process on each of its states,

$k = 0, 1, 2, \dots$, tend to be smaller than regular expected sojourn times in states, which have expected value $1/q$, and decrease with k . Moreover, the Poisson expected sojourn time coefficients are related to the mixed Poisson probability coefficients associated with the uniform distribution on $[0, t]$, $\bar{\gamma}_t(k, q) = t \cdot \alpha_{U_{[0,t]}}(k, q)$, where $\alpha_{U_{[0,t]}}(k, q)$ is the k th mixed Poisson probability associated with the uniform distribution on $[0, t]$, as given in (3).

The vector $\underline{f}_q(\underline{r}, \mathbf{R})$ associated with the reward tuple $(\underline{r}, \mathbf{R})$ denotes the vector of expected reward per unit of time in the states of the CTMC \mathcal{M} and is the sum of the vectors \underline{r} and $q \cdot [\mathbf{P}^{u(\mathcal{M})} \bullet \mathbf{R}] \cdot \underline{1}$. The vector \underline{r} corresponds to the expected rewards per unit of time accumulated continuously in the states of \mathcal{M} , whereas $q \cdot [\mathbf{P}^{u(\mathcal{M})} \bullet \mathbf{R}] \cdot \underline{1}$ contains the expected rewards per unit of time associated with state transitions, which take place at rate q in the uniformised CTMC \mathcal{M}_q . The entry corresponding to state s of $[\mathbf{P}^{u(\mathcal{M})} \bullet \mathbf{R}] \cdot \underline{1}$ is the expected reward associated to a random transition out from state s .

Following the approach taken in [1], the computation of $E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(t))$ for all states reduces to computing the following sum over vectors of coefficients:

$$\underline{E}^{\mathcal{M}}(\mathcal{R}^{\mathcal{M}}(t)) = \sum_{k=0}^{\infty} \bar{\gamma}_t(k, q) \cdot [\mathbf{P}^{u(\mathcal{M})}]^k \cdot \underline{f}_q(\underline{r}, \mathbf{R}). \quad (10)$$

3.2. Random Terminal Times

In this section we consider formulae of the form $\mathcal{E}\mathcal{R}_{\times V}^{\leq T}$, where T is a nonnegative random variable. Below is the main observation with regards to model checking such formulae.

THEOREM 2. *For any $s \in S$ and nonnegative random time T with finite expected value:*

$$E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(T)) = \sum_{k=0}^{\infty} \bar{\alpha}_T(k, q) \cdot \underline{\pi}^{u(\mathcal{M})}(s, k) \cdot \underline{f}_q(\underline{r}, \mathbf{R})$$

where, as given in (6), $\underline{f}_q(\underline{r}, \mathbf{R}) = \underline{r} + q \cdot [\mathbf{P}^{u(\mathcal{M})} \bullet \mathbf{R}] \cdot \underline{1}$, and

$$\bar{\alpha}_T(k, q) = E \left[\int_0^T 1(N_q(t) = k) dt \right] = \int_0^{\infty} \bar{\gamma}_t(k, q) dF(t) = \frac{1}{q} \sum_{j=k+1}^{\infty} \alpha_T(j, q) \quad (11)$$

and this last relation also holds when $E[T] = \infty$.

Proof. We start by proving (11). In view of (5) and since $E[X] = E[E[X|Y]]$, for any random variables X and Y ,

$$\begin{aligned} \bar{\alpha}_T(k, q) &= E \left[\int_0^T 1(N_q(u) = k) du \right] \\ &= \int_0^{\infty} E \left[\int_0^t 1(N_q(u) = k) du \right] dF(t) \\ &= \int_0^{\infty} \bar{\gamma}_t(k, q) dF(t). \end{aligned}$$

Thus, (11) follows since using (5) and monotone convergence [18]:

$$\bar{\alpha}_T(k, q) = \int_0^{\infty} \bar{\gamma}_t(k, q) dF(t) = \int_0^{\infty} \frac{1}{q} \sum_{j=k+1}^{\infty} \gamma(j, q \cdot t) dF(t) = \frac{1}{q} \sum_{j=k+1}^{\infty} \alpha_T(j, q)$$

by definition of $\alpha_T(j, q)$. From Theorem 1 and definition of $E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(T))$, for any $s \in S$ and nonnegative random time T with finite expected value, $E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(T))$ equals:

$$\begin{aligned} \int_0^\infty E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(t)) dF(t) &= \int_0^\infty \sum_{k=0}^\infty \bar{\gamma}_t(k, q) \cdot \underline{\pi}^{u(\mathcal{M})}(s, k) \cdot \underline{f}_q(\underline{r}, \mathbf{R}) dF(t) \\ &= \sum_{k=0}^\infty \int_0^\infty \bar{\gamma}_t(k, q) dF(t) \cdot \underline{\pi}^{u(\mathcal{M})}(s, k) \cdot \underline{f}_q(\underline{r}, \mathbf{R}) \\ &= \sum_{k=0}^\infty \bar{\alpha}_T(k, q) \cdot \underline{\pi}^{u(\mathcal{M})}(s, k) \cdot \underline{f}_q(\underline{r}, \mathbf{R}) \end{aligned}$$

where the second equality follows by dominated convergence, as required. ■

Note that, from the definition, $\bar{\alpha}_T(k, q)$ is the expected sojourn time of the uniformising Poisson process in state k until the random time T . Accordingly, we call $\bar{\alpha}_T(k, q)$ the *k-th mixed Poisson expected sojourn time coefficient* (associated with the random time T and the uniformisation rate q). The mixed Poisson expected sojourn time coefficients enjoy properties similar to those of Poisson expected sojourn time coefficients. In particular, the coefficients $\bar{\alpha}_T(k, q)$ are non-increasing and take values on the interval $[0, 1/q]$, and $\sum_{k=0}^\infty \bar{\alpha}_T(k, q) = E[T]$, independently of T having finite or infinite expected value. Moreover, using the last part of (11), it follows that $\lim_{k \rightarrow \infty} \bar{\alpha}_T(k, q) = 0$ since, for any $t \in \mathbb{R}_+$: $\bar{\alpha}_T(k, q) = \frac{1}{q} \cdot P(N_q(T) > k) \leq \frac{1}{q} \cdot [P(N_q(t) > k) + P(T \geq t)]$.

The mixed Poisson expected sojourn times are equivalent to the $\alpha_{\bar{F}}$ -factors introduced in [12] which are used in the calculation of steady-state probabilities for non-Markovian stochastic Petri nets. Alternative proofs of some of the stated properties for the mixed Poisson expected sojourn times are presented in [12] under more restrictive conditions.

Note that computing $E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(T))$ for all states is similar to the case where T is deterministic and may be done as in (10) by replacing the Poisson expected sojourn time coefficients by the $\bar{\alpha}_T(k, q)$ coefficients. As $\sum_{k=0}^\infty \bar{\alpha}_T(k, q) = E[T]$, provided $E[T] < \infty$ we have $\lim_{k \rightarrow \infty} \sum_{n > k} \bar{\alpha}_T(n, q) = 0$. This result is in the basis of the pseudo-code given in Figure 1 for a generic algorithm for computing the values of $E^{\mathcal{M}}(s, \mathcal{R}^{\mathcal{M}}(T))$ with an error of at most ε , for arbitrary positive ε . The algorithm is based on the fact that if we choose $K(\varepsilon)$ such that $\sum_{n=0}^{K(\varepsilon)} \bar{\alpha}_T(n, q) \geq E[T] - \varepsilon/\alpha$, with $\alpha = \max_{s \in S} |\underline{f}_q(\underline{r}, \mathbf{R})_s|$, then

$$\left| \frac{E^{\mathcal{M}}(\mathcal{R}^{\mathcal{M}}(T))}{\alpha} - \sum_{k=0}^{K(\varepsilon)} \left(\bar{\alpha}_T(k, q) \cdot [\mathbf{P}^{u(\mathcal{M})}]^k \cdot \underline{f}_q(\underline{r}, \mathbf{R}) \right) \right| < \varepsilon.$$

Note that the DTMC $u(\mathcal{M})$ may reach steady state before $K(\varepsilon)$ and, in this case, the summation can be truncated at this earlier point [19].

3.3. Computing Mixed Poisson Expected Sojourn Times

From (11), it follows that the mixed Poisson expected sojourn times may be computed recursively using the mixed Poisson probabilities through

$$\bar{\alpha}_T(k, q) = \bar{\alpha}_T(k-1, q) - \frac{1}{q} \alpha_T(k, q)$$

for $k \in \mathbb{N}$, with $\bar{\alpha}_T(-1, q) = 1/q$. Thus, algorithms for the computation of mixed Poisson expected sojourn times when the distribution of the random time T has a finite discrete, uniform, gamma or Pareto distribution, or is a finite mixture of distributions of these types, may be obtained directly from the algorithms given in [14] for the corresponding mixed Poisson probabilities.

```

input :  $\bar{\alpha}_T(0, q), \dots, \bar{\alpha}_T(K(\varepsilon), q)$  such that  $\sum_{k=0}^{K(\varepsilon)} \bar{\alpha}_T(k, q) \geq \mathbb{E}[T] - \varepsilon/\alpha$ 

 $\underline{b} := f_q(\underline{r}, \mathbf{R})$ 
 $\underline{sol} := \underline{0}$ 
for  $k = 0$  to  $K(\varepsilon)$  do
     $\underline{sol} := \underline{sol} + \bar{\alpha}_T(k, q) \cdot \underline{b}$ 
     $\underline{b} := \mathbf{P}^{u(\mathcal{M})} \cdot \underline{b}$ 
endfor
output :  $\underline{\mathbb{E}^{\mathcal{M}}(\mathcal{R}^{\mathcal{M}}(T))} := \underline{sol}$ 

```

Figure 1. Generic algorithm for computing $\underline{\mathbb{E}^{\mathcal{M}}(\mathcal{R}^{\mathcal{M}}(T))}$

4. EXAMPLE

We consider a simple *power management system* taken from [20]. The model consists of four components: a service requester (SR) which generates requests to be served; a (finite) service request queue (SQ) which stores the requests before service; a service provider (SP) which processes requests; and a power manager (PM) which monitors the states of the other components and issues state transition commands to the SP. We suppose that the requests arrive according to a renewal process with inter-renewal time distributed as T .

The SP has three states: *sleep*, *idle* and *busy*. In *sleep* the SP is inactive, and hence no requests can be served; in *idle* the SP is active but is not working on any requests (i.e. the SQ is empty) and in *busy* requests are being served. The transitions between *sleep* and *idle* are controlled by the PM, while transitions between *idle* and *busy* correspond to the arrival of requests in the queue and the service of requests. In each state of the SP power is consumed at a certain rate and there is a switching energy associated with each pair of states – the energy needed for the SP to switch between these states. Note that requests can only be served when the SP is in state *busy* and we suppose that the service time and transition times between the states of the SP are exponentially distributed.

We consider the simple PM which switches the SP on (from *sleep* to *idle*) as soon as a request arrives (the SQ becomes nonempty) and switches the SP off (from *idle* to *sleep*) as soon as there are no longer any requests to be served (the SQ becomes empty).

For this system we compute the *average number of waiting requests in the SQ* as the performance metric and the *average power consumption of the SP* as the power metric. These measures are calculated by through the following procedure.

- Construct a restricted model of the system in which transitions corresponding to new requests are removed.
- In the restricted model, calculate the expected reward until the random terminal time T for the cases when the reward structure corresponds to the power consumption and to the size of the queue.
- Construct the embedded DTMC model of the full system taking the time of the next service request as one unit of time¹ and calculate the steady state probabilities of this DTMC.
- Combine the expected reward values and steady state probabilities using the theory of Markov regenerative processes [17] to give the performance and power metrics.

¹This can be achieved following the methodology of [12] or by calculating the probability of satisfying random time-bounded until formulae on the restricted model.

Table 1. Performance results for power management model as inter-arrival time distribution varies

performance measure	inter-arrival time distribution				
	deterministic	Erlang 10	uniform[0,b]	exponential	Pareto
power	3.0198	2.9793	2.8883	2.6714	0.3790
performance	1.6880	1.6939	1.7145	1.7343	0.7700

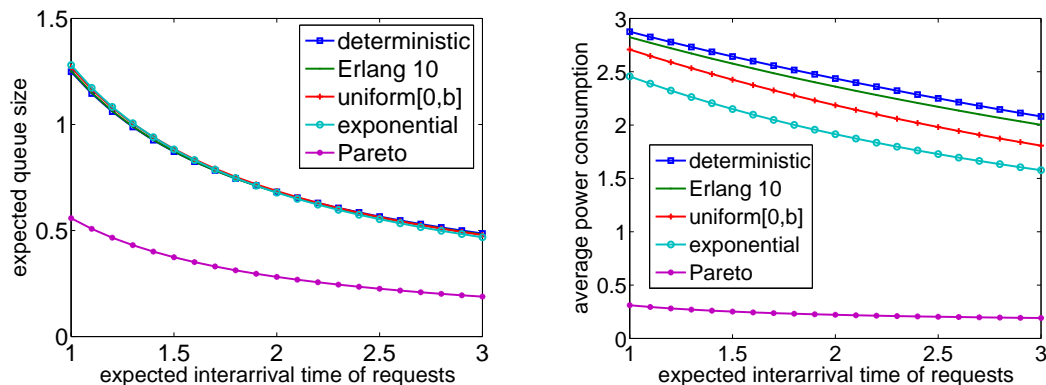


Figure 2. Performance and power results as the expected inter-arrival time varies

Note that, in the restricted model, all transitions have an exponential delay, that is, it is a CTMC, and since the service time distribution is independent of the arrival time distribution, the inter-arrival time distribution is independent of this model.

We constructed the restricted model using the probabilistic model checker PRISM [10,11] and then exported the generator matrix of this model to a prototype implementation in MATLAB to calculate the measures of interest. We consider five different distributions for the arrivals of requests: deterministic, exponential, Erlang (10 phases), uniform and Pareto. In the Table 1 we give the results in the case when the parameters of the system are those given in [20]. Moreover, in Figure 2 we have plotted the performance metrics as the expected time inter-arrival time of requests varies for each of the considered distributions.

As can be seen in both Table 1 and Figure 2 the expected queue size and the power consumption when requests arrive with a Pareto distribution are much smaller than when requests arrive with the other distributions considered. This is a result of the Pareto distribution's *heavy tail*, which means that, in the long run, many requests will not arrive for a very long time, and hence in these cases the service provider (SP) will serve all pending requests, and then the system will spend a long time with the queue empty and the SP in its *sleep* state consuming very little power. Moreover, more requests are blocked for the Pareto distribution than with the other distributions.

The performance metric (average queue size) for the remaining distributions is very similar for all the other distributions considered, which is not true of the power metric. The relation between the power consumption for the remaining distributions corresponds to the difference between the “tails” of the distributions – the larger the tail the higher the chance of the SP spending time off preserving power. For example, the deterministic distribution is zero for all values greater than the expected value, whereas the uniform distribution is zero for any value greater than two times the expected value, and the exponential distribution has a heavier tail than an Erlang (with more than one phase).

For all the distributions considered, Figure 2 shows that, as the expected inter-arrival time increases, both the average queue size and power consumption decreases. This is to be expected, as increasing the expected inter-arrival time means that, on average, there will be more time for the SP to serve requests between the arrival of successive requests, and hence, in the long run, the queue will be smaller and more time will be spent with the SP off (i.e. consuming less power).

Finally, we note the similarity between the cases for requests arriving with a deterministic or Erlang distribution; this is to be expected as the Erlang distribution is often used as a continuous approximation of a (discrete) deterministic distribution.

5. CONCLUSIONS

This paper presents an extension of CSL with expected time and expected reward formulae with random time bounds where the time bound takes the form of a general nonnegative random variable. As the example demonstrates, in certain cases, using such formulae enables us to calculate performance measures of systems which include general distributions. It is shown that model checking for such formulae can be efficiently carried out by first precomputing mixed Poisson expected sojourn time coefficients.

So far, we have only considered a prototype implementation in MATLAB. In future we aim to implement these algorithms in the probabilistic symbolic model checker PRISM in order to tackle the verification of more complex models. Additionally, we would like to work on generalising this approach to other important families of distributions, apply analytic methods to finding upper bounds for $K(\varepsilon)$, and extend our approach to express random time intervals rather than simply the time bound T .

REFERENCES

1. J.-P. Katoen, M. Kwiatkowska, G. Norman and D. Parker, Faster and Symbolic CTMC Model Checking, In *Proc. 1st Joint International Workshop on Process Algebra and Probabilistic Methods, Performance Modeling and Verification (PAPM/PROBMIV'01)*, (eds. L. de Alfaro and S. Gilmore), Lecture Notes in Computer Science, Vol. 2165, pp. 23–38, Springer-Verlag, (2001).
2. A. Aziz, K. Sanwal, V. Singhal and R. Brayton, Verifying Continuous Time Markov Chains, In *Proc. 8th International Conference on Computer Aided Verification (CAV'96)*, (eds. R. Alur and T. Henzinger), Lecture Notes in Computer Science, Vol. 1102, pp. 269–276, Springer-Verlag, (1996).
3. A. Aziz, K. Sanwal, V. Singhal and R. Brayton, Model Checking Continuous Time Markov Chains, *ACM Transactions on Computational Logic* **1** **1**, 162–170, (2000).
4. C. Baier, J.-P. Katoen and H. Hermanns, Approximate Symbolic Model Checking of Continuous-Time Markov Chains, In *Proc. 10th International Conference on Concurrency Theory (CONCUR'99)*, (eds. J. Baeten and S. Mauw), Lecture Notes in Computer Science, Vol. 1664, pp. 146–161, Springer-Verlag, (1999).
5. 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** **2**, 244–263, (1986).
6. H. Hansson and B. Jonsson, A Logic for Reasoning about Time and Probability, *Formal Aspects of Computing* **6**, 512–535, (1994).
7. C. Baier, B. Haverkort, H. Hermanns and J.-P. Katoen, Model Checking Continuous-Time Markov Chains by Transient Analysis, In *Proc. 12th International Conference on Computer Aided Verification (CAV'00)*, (eds. A. Emerson and A. Sistla), Lecture Notes in Computer Science, Vol. 1855, pp. 358–372, Springer-Verlag, (2000).
8. A. Jensen, Markov Chains as an Aid in the Study of Markov Processes, *Skandinavisk Aktuarietidskrift, Marts*, 87–91, (1953).
9. D. Gross and D. Miller, The randomization technique as a modeling tool and solution procedure for transient Markov processes, *Operations Research* **32** **2**, 343–361, (1984).
10. M. Kwiatkowska, G. Norman and D. Parker, PRISM: Probabilistic Symbolic Model Checker, In *Proc. 12th International Conference on Modelling Techniques and Tools for Computer Performance Evaluation (TOOLS'02)*, (eds. T. Field, P. Harrison, J. Bradley and U. Harder), Lecture Notes in Computer Science, Vol. 2324, pp. 200–204, Springer-Verlag, (2002).
11. PRISM Web Page: <http://www.cs.bham.ac.uk/~dxp/prism>.
12. R. German, *Performance Analysis of Communication Systems: Modeling with Non-Markovian Stochastic Petri Nets*, John Wiley and Sons, (2000).
13. G. Infante Lópes, H. Hermanns and J.-P. Katoen, Beyond Memoryless Distributions, In *Proc. 1st Joint International Workshop on Process Algebra and Probabilistic Methods, Performance Modeling and Verification (PAPM/PROBMIV'01)*, (eds. L. de Alfaro and S. Gilmore), Lecture Notes in Computer Science, Vol. 2165, pp. 57–70, Springer-Verlag, (2001).
14. M. Kwiatkowska, G. Norman and A. Pacheco, Model Checking CSL Until Formulae with Random Time Bounds, In *Proc. 2nd Joint International Workshop on Process Algebra and Probabilistic Methods, Performance Modeling and Verification (PAPM/PROBMIV'02)*, (eds. H. Hermanns and R. Segala), Lecture Notes in Computer Science, Vol. 2399, pp. 152–168, Springer-Verlag, (2002).

15. G. Norman, D. Parker, M. Kwiatkowska, S. Shukla and R. Gupta, Formal Analysis and Validation of Continuous Time Markov Chain Based System Level Power Management Strategies, In *Proc. 7th Annual IEEE International Workshop on High Level Design Validation and Test (HLDVT'02)*, (eds. W. Rosenstiel), pp. 45–50, IEEE Computer Society Press, (2002).
16. M. Telek, A. Pfening and G. Fodor, An effective numerical method to compute the moments of the completion time of Markov reward models, *Computers & Mathematics with Applications* **36** **8**, 59–65, (1998).
17. V.G. Kulkarni, *Modeling and Analysis of Stochastic Systems*, Chapman and Hall, (1995).
18. Resnick, S. I., *A Probability Path*, Birkhäuser, Boston, MA, (1999).
19. J. Muppala and K. Trivedi, Numerical Transient Solution of Finite Markovian Queueing Systems, *Queueing Systems, Queueing and Related Models*, Oxford University Press, 262–284, (1992).
20. Q. Qiu, Q. Wu and M. Pedram, Stochastic modeling of a power-managed system: construction and optimization, *IEEE Transactions on Computer Aided Design* **20** **9**, 1200–1217, (2001).