On the logical characterisation of performability properties

Christel Baier[a], Boudewijn Haverkort[b], Holger Hermanns[c], 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. Markov-reward models, as extensions of continuous-time Markov chains, have received increased attention for the specification and evaluation of performance and dependability properties of systems. Until now, however, the specification of reward-based performance and dependability measures has been done manually and informally. In this paper, we change this undesirable situation by the introduction of a continuous-time, reward-based stochastic logic. We argue that this logic is adequate for expressing performability measures of a large variety. We isolate two important sub-logics, the logic **CSL** [1, 3], and the novel logic **CRL** that allows one to express reward-based properties. These logics turn out to be complementary, which is formally established in our main duality theorem. This result implies that reward-based properties expressed in **CRL** for a particular Markov reward model can be interpreted as **CSL** properties over a derived continuous-time Markov chain, so that model checking procedures for **CSL** [3, 2] can be employed.

1 Introduction

With the advent of fault-tolerant and distributed computer and communication systems, the classical separation between performance evaluation and dependability (i.e., reliability, availability and timeliness) evaluation does not make sense anymore. Instead, the combined performance and dependability of a system is of critical importance. This observation led to development of the performability evaluation framework [12, 13]. This framework allows one to specify models that include both performance-related and dependability-related events in a natural way. Furthermore, the choice of Markov-reward models (MRMs) [11] as mathematical basis allows one to specify a wide variety of measures of interest, albeit at times in a slightly cumbersome way. An MRM is a continuous-time Markov chain (CTMC) augmented with a *reward structure* assigning a real-valued reward to each state in the model. Such reward can be interpreted as bonus, gain, or dually, as cost. Typical measures of interest express the amount of gain accumulated by the system, over a finite or infinite time-horizon.

Given the fact that the model is stochastic, the measures of interest are stochastic variables. MRMs have shown to pair a reasonable modelling flexibility and expressiveness with manageable computational expenses for the model evaluation. To increase the modelling flexibility, a number of application-oriented model specification techniques and supporting tools have been developed [8].

The specification of the measure-of-interest for a given MRM can not always be done conveniently, nor can all possible measures-of-interest be expressed conveniently. In particular, until recently it has not been possible to directly express measures where state *sequences* or paths matter, nor to accumulate rewards only in certain subsets of states, if the rewards outside these subsets are non-zero. Such measures are then either "specified" informally, with all its negative implications, or require a manual tailoring of the model so as to address the right subsets of states. An example of a measure that is very difficult to specify directly is the expected amount of gain obtained from the system until a particular state is reached, provided that all paths to that state obey certain properties.

Recently, Obal and Sanders have proposed a technique to specify so-called path-based reward variables [14] by which the specification of measures over state sequences becomes more convenient, because it avoids the manual tailoring of the model. In the context of the stochastic process algebra PEPA, Clark *et al.* recently proposed the use of a probabilistic modal logic to ease the specification of reward *structures* of MRM [5], as opposed to the specification of reward-based *measures*, as we do.

In [3] we proposed to specify measures of interest for CTMCs in the logic **CSL** (Continuous Stochastic Logic), a superset of the (equally named) logic introduced by Aziz *et al.* [1]. **CSL** includes a timed **CTL**-like time-bounded until operator, and a steady-state operator. Using this logic, very complex measures can be expressed easily; model-checking algorithms for **CSL** have been proposed [3, 2] (and implemented [10]). Notice however, that **CSL** is interpreted over CTMCs only, and is hence not able to address reward-based measures. The current paper extends this work, in that Markov-*reward* models are evaluated, i.e., CTMCs augmented with a reward structure.

In this paper, we introduce a novel continuous-time, stochastic reward-based logic **CSRL**, that is adequate for expressing performability measures of a large variety. It includes next and until operators, that are equipped with time-interval- as well as reward-interval-bounds. We present syntax and formal semantics of the logic, and isolate two important sub-logics: the logic **CSL**, and the logic **CRL** (Continuous Reward Logic) that allows one to express time-independent reward-based properties. These logics turn out to be complementary, which is formally established in a main duality theorem, showing that time- and reward-intervals are interchangeable. More precisely, we show that for each MRM $\mathcal{M}$ and formula Φ the set of states satisfying Φ equals the set of states of a derived MRM $\mathcal{M}^{-1}$ satisfying formula Φ^{-1}, where the latter is obtained from Φ by simply swapping time- and reward-intervals. The transformation of $\mathcal{M}$ is inspired by [4]. The fixpoint characterisations for the **CRL** path operators (interpreted over an MRM) reduce to the characterisations that are used for model checking **CSL** (over a CTMC). As a consequence of the duality result, the model checking problem for **CRL** is reducible to the model checking problem for **CSL** and hence solvable with existing techniques for **CSL**.

The paper is organised as follows. Section 2 introduces MRMs and typical measures of interest for them. In Section 3 the logic **CSRL** and its sub-logics are defined, whereas Section 4 presents the main duality theorem. Section 5 discusses its consequences for model checking and highlights that most reward-based performability measures having appeared in the literature can be expressed as simple formulas of (a minor extension of) the logic. Section 6 concludes the paper.

2 Markov reward models

In this section we introduce the basic concepts of MRMs [11]. We slightly depart from the standard notation for MRMs (and CTMCs) and consider an MRM as an ordinary transition system, i.e., a Kripke structure, where the edges are equipped with probabilistic timing information and the states are augmented with a real number that indicates the earned reward per unit of time for staying in a state. This then allows the usual interpretation of linear-time temporal operators like next step and unbounded or time-bounded until.

MRMs. Let AP be a fixed, finite set of atomic propositions.

Definition 1. *A (labelled) CTMC $\mathcal{C}$ is a tuple $(S, \mathbf{R}, L)$ where S is a finite set of* states, $\mathbf{R} : S \times S \to \mathbb{R}_{\geqslant 0}$ *the* rate matrix, *and $L : S \to 2^{AP}$ 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* terminal *(or* absorbing*) iff $\mathbf{R}(s, s') = 0$ for all states 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 $s \to s'$ can be triggered within t time units. Thus the delay of transition $s \to s'$ 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 exists, known as the *race condition*. The probability to move from non-absorbing s to s' within t time units, i.e., $s \to s'$ to win the race, is given by

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

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 $1 - e^{-\mathbf{E}(s)\cdot t}$, because 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 moving from s to s' finishes before the delays of other outgoing edges from s; formally, $\mathbf{P}(s, s') = \mathbf{R}(s, s')/\mathbf{E}(s)$. For absorbing states, the total rate $\mathbf{E}(s) = 0$; we then have $\mathbf{P}(s, s') = 0$ for any state s'.

Definition 2. *A (labelled) MRM $\mathcal{M}$ is a pair $(\mathcal{C}, \rho)$ where $\mathcal{C}$ is a (labelled) CTMC, and $\rho : S \to I\!R_{\geqslant 0}$ is a* reward structure *that assigns to each state $s \in S$ a reward $\rho(s)$, also called gain or bonus or dually, cost.*

Example 1. As a running example we consider a fault-tolerant multiprocessor system inspired by [15]. The system consists of three processors, three memories, and a single interconnection network that allows a processor to access any memory. We model this system by a CTMC, depicted below, where state $(i, j, 1)$ models that i processors and j memories $(1 \leqslant i, j < 4)$ are operational and are connected by a single network. Initially all components are functioning correctly, i.e., the initial state is $(3, 3, 1)$. The minimal operational configuration of the system is $(1, 1, 1)$. The failure rate of a processor is λ, of a memory μ, and of the network γ failures per hour (fph). It is assumed that a single repair unit is present to repair all types of components. The expected repair time of a processor is $1/\nu$ and of a memory $1/\eta$ hours. In case all memories, all processors, or the network has failed the system moves to state F. After a repair in state F, we assume the system to restart in state $(3, 3, 1)$ with rate δ.

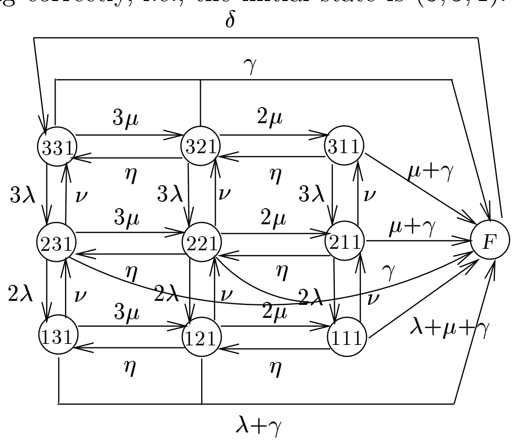

The reward structure can be instantiated in different ways so as to specify a variety of performability measures. The following reward structures are taken from [15]. The simplest reward structure (leading to an availability model) divides the states into operational and non-operational states: $\rho_1(F) = 0$ and $\rho_1(i, j, k) = 1$. A reward structure in which varying levels of performance of the system are represented is for instance based on the capacity of the system: $\rho_2(F) = 0$ and $\rho_2(i, j, k) = min(i, j)$. The third reward structure does consider processors contending for the memories, by taking as reward for operational states the expected available memory bandwidth: $\rho_3(F) = 0$ and $\rho_3(i, j, k) = m \cdot \big(1 - (1 - 1/m)^l\big)$ where $l = min(i, j)$ and $m = max(i, j)$. ■

Let $\mathcal{M} = (\mathcal{C}, \rho)$ be an MRM with underlying CTMC $\mathcal{C} = (S, \mathbf{R}, L)$.

Paths. An infinite *path* σ is a sequence $s_0, t_0, s_1, t_1, s_2, t_2, \ldots$ with for $i \in I\!N$, $s_i \in S$ and $t_i \in \mathbb{R}_{>0}$ such that $\mathbf{R}(s_i, s_{i+1}) > 0$. For $i \in I\!N$ 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 $t \in \mathbb{R}_{\geqslant 0}$ and i the smallest index with $t \leqslant \sum_{j=0}^{i} t_j$ let $\sigma@t = \sigma[i]$, the state in σ at time t. For $t = \sum_{j=0}^{k-1} t_j + t'$ with $t' \leqslant t_k$ we define $y(\sigma, t) = \sum_{j=0}^{k-1} t_j \cdot \rho(s_j) + t' \cdot \rho(s_k)$, the cumulative reward along σ up to time t.

A finite path σ is a sequence $s_0, t_0, s_1, t_1, s_2, t_2, \ldots, t_{l-1}, s_l$ where 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 only defined for $i \leqslant l$; they are defined as above for $i < l$, and $\delta(\sigma, l) = \infty$. For $t > \sum_{j=0}^{l-1} t_j$ we let $\sigma@t = s_l$ and let the cumulative reward $y(\sigma, t) = \sum_{j=0}^{l-1} t_j \cdot \rho(s_j) + (t - \sum_{j=0}^{l-1} t_j) \cdot \rho(s_l)$; for the other cases, $\sigma@t$ and $y(\sigma, t)$ are defined as above.

Let $Path^{\mathcal{M}}(s)$ denote the set of (finite and infinite) paths starting in s.

Borel space. Any state $s = s_0$ yields a probability measure Pr on $Path^{\mathcal{M}}(s)$ as follows. Let $s_0, \ldots, s_k \in S$ with $\mathbf{R}(s_i, s_{i+1}) > 0$, $(0 \leqslant i < k)$, and $I_0, \ldots, I_{k-1}$ non-empty intervals in $\mathbb{R}_{\geqslant 0}$. Then, $C(s_0, I_0, \ldots, I_{k-1}, s_k)$ denotes the *cylinder set* consisting of all paths $\sigma \in Path^{\mathcal{M}}(s)$ such that $\sigma[i] = s_i$ $(i \leqslant k)$, and $\delta(\sigma, i) \in I_i$ $(i < k)$. Let $\mathcal{F}(Path^{\mathcal{M}}(s))$ be the smallest σ-algebra on $Path^{\mathcal{M}}(s)$ which contains all sets $C(s, I_0, \ldots, I_{k-1}, s_k)$ where $s_0, \ldots, s_k$ ranges over all state-sequences with $s = s_0$, $\mathbf{R}(s_i, s_{i+1}) > 0$ $(0 \leqslant i < k)$, and $I_0, \ldots, I_{k-1}$ ranges over all sequences of non-empty intervals in $\mathbb{R}_{\geqslant 0}$. The probability measure Pr on $\mathcal{F}(Path^{\mathcal{M}}(s))$ is the unique measure defined by induction on k: $\Pr(C(s_0)) = 1$, and for $k \geqslant 0$,

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

where $a = \inf I'$ and $b = \sup I'$. (For $b = \infty$ and $\lambda > 0$ let $e^{-\lambda \cdot \infty} = 0$.) Note that $e^{-\mathbf{E}(s_k) \cdot a} - e^{-\mathbf{E}(s_k) \cdot b}$ is the probability of leaving state s_k in the interval I'.

Remark. For infinite paths we do not assume *time divergence*. Although such paths represent "unrealistic" computations where infinitely many transitions are taken in a finite amount of time, the probability measure of such Zeno paths is 0. This justifies a lazy treatment of the notations $\sigma@t$ and $y(\sigma, t)$ when we refer to the probability of a measurable ste of paths. ∎

Steady-state and transient probabilities. For a CTMC $\mathcal{C}$ 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*

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

stands for the probability to be in state s' at time t given the initial state s. Note that this set is measurable. *Steady-state probabilities* are defined as

$$\pi^{\mathcal{C}}(s, s') = \lim_{t \to \infty} \pi^{\mathcal{C}}(s, s', t).$$

This limit always exists for finite CTMCs. For $S' \subseteq S$, $\pi^{\mathcal{C}}(s, S') = \sum_{s' \in S'} \pi^{\mathcal{C}}(s, s')$ denotes the steady-state probability for set S'. In the sequel, we will often use $\mathcal{M}$ rather than $\mathcal{C}$ (the underlying CTMC of $\mathcal{M}$) as superscript.

3 Stochastic CTL with time and rewards

This section introduces a stochastic logic to reason about reward-based as well as time-based constraints, and identifies two important sub-logics of it. For explanatory purposes, we first introduce a simple branching time logic without any support for real time or reward constraints.

Basic logic. The base stochastic logic **SL**, a stochastic variant of **CTL** (Computational Tree Logic), is a continuous-time variant of **PCTL** [7].

Syntax. For $a \in AP$, $p \in [0,1]$ and $\bowtie \in \{ \leqslant, <, \geqslant, > \}$, the state-formulas of **SL** are defined by the grammar

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

where path-formulas are defined by $\varphi ::= X\Phi \;\Big|\; \Phi\, \mathcal{U}\, \Phi.$

Other boolean connectives such as $\vee$ and $\rightarrow$ are derived in the obvious way. As usual $\Diamond\Phi = \mathrm{tt}\,\mathcal{U}\,\Phi$ and the $\Box$-operator can be obtained by, for example, $\mathcal{P}_{\geqslant p}(\Box\Phi) = \neg\mathcal{P}_{\geqslant 1-p}(\Diamond\neg\Phi)$. The state-formula $\mathcal{S}_{\bowtie p}(\Phi)$ asserts that the steady-state probability for the set of Φ-states meets the bound $\bowtie p$. For the running example, the formula $\mathcal{S}_{\geqslant 0.8}(\mathit{2pup})$ expresses that the steady-state probability to be in a state with two operational processors is at least 0.8 where $\mathit{2pup}$ holds in state $(2, j, 1)$, $1 \leqslant j < 4$. The operator $\mathcal{P}_{\bowtie p}(.)$ replaces the usual **CTL** path quantifiers $\exists$ and $\forall$. $\mathcal{P}_{\bowtie p}(\varphi)$ asserts that the probability measure of the paths satisfying φ meets the bound $\bowtie p$. For example, $\mathcal{P}_{\geqslant 0.3}(\Diamond F)$ denotes that the probability to eventually reach the failure state of the multi-processor system is at least 0.3.

Semantics. The **SL** state-formulas are interpreted over the states of a CTMC $\mathcal{C} = (S, \mathbf{R}, L)$ (or an MRM $\mathcal{M}$ with underlying CTMC $\mathcal{C}$) with proposition labels in AP. Let $Sat^{\mathcal{C}}(\Phi) = \{ s \in S \mid s \models \Phi \}$.

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

Here, $Prob^{\mathcal{C}}(s,\varphi)$ denotes the probability measure of all paths satisfying φ given that the system starts in state s, i.e.,

$$Prob^{\mathcal{C}}(s,\varphi) \;=\; \Pr\{ \sigma \in Path^{\mathcal{C}}(s) \mid \sigma \models \varphi \}.$$

The fact that the set $\{ \sigma \in Path^{\mathcal{C}}(s) \mid \sigma \models \varphi \}$ is measurable can be easily verified. The intended meaning of the temporal operators $\mathcal{U}$ and X is standard:

$$\begin{array}{ll} \sigma \models X\Phi & \text{iff } \sigma[1] \text{ is defined and } \sigma[1] \models \Phi \\ \sigma \models \Phi_1\,\mathcal{U}\,\Phi_2 & \text{iff } \exists k \geqslant 0.\, (\sigma[k] \models \Phi_2 \wedge \forall 0 \leqslant i < k.\, \sigma[i] \models \Phi_1). \end{array}$$

Alternative characterisations. For next-formulas we have, as for DTMCs [7]:

$$Prob^{\mathcal{C}}(s, X\Phi) = \mathbf{P}(s,\Phi) \tag{1}$$

where $\mathbf{P}(s,\Phi) = \sum_{s' \in Sat^{\mathcal{C}}(\Phi)} \mathbf{P}(s,s')$, the probability to reach a Φ-state in one step from s. For until-formulas we have that the probability $Prob^{\mathcal{C}}(s, \Phi_1\,\mathcal{U}\,\Phi_2)$

is the least solution[1] of the following set of equations: $Prob^{\mathcal{C}}(s, \Phi_1 \,\mathcal{U}\, \Phi_2)$ equals 1 if $s \models \Phi_2$, equals

$$\sum_{s' \in S} \mathbf{P}(s, s') \cdot Prob^{\mathcal{C}}(s', \Phi_1 \,\mathcal{U}\, \Phi_2) \tag{2}$$

if $s \models \Phi_1 \wedge \neg \Phi_2$, and 0 otherwise. This probability can be computed as the solution of a regular system of linear equations by standard means such as Gaussian elimination [6] or can be approximated by an iterative approach.

The full logic. We now extend **SL** by providing means to reason about both time constraints and cumulative reward constraints. We refer to this logic as **CSRL**. Later we will identify fragments of **CSRL** that refer to only time, respectively only reward constraints.

Syntax. The syntax (and semantics) of the state formulas of **CSRL** are defined as for the basic logic. Path-formulas φ are defined for intervals $I, J \subseteq \mathbb{R}_{\geqslant 0}$ by:

$$\varphi \ ::= \ X_J^I \, \Phi \ \Big| \ \Phi \, \mathcal{U}_J^I \, \Phi.$$

In a similar way as before, we define $\Diamond_J^I \Phi = \mathrm{tt} \, \mathcal{U}_J^I \, \Phi$ and $\mathcal{P}_{\bowtie p}(\Box_J^I \Phi) = \neg \mathcal{P}_{\bowtie p}(\Diamond_J^I \neg \Phi)$. Interval I can be considered as a timing constraint whereas J represents a bound for the cumulative reward. The path-formula $X_J^I \, \Phi$ asserts that a transition is made to a Φ-state at time point $t \in I$ such that the earned cumulative reward r until time t meets the bounds specified by J, i.e., $r \in J$. The semantics of $\Phi_1 \, \mathcal{U}_J^I \, \Phi_2$ is as for $\Phi_1 \, \mathcal{U} \, \Phi_2$ with the additional constraints that the Φ_2-state is reached at some time point t in I and the earned cumulative reward up to t lies in J. As an example property for the multi-processor system, $\mathcal{P}_{\geqslant 0.95}(\Diamond_{[0,2]}^{[60,60]} \mathrm{tt})$ denotes that with probability at least 0.95 the cumulative reward (e.g., the expected capacity of the system for reward structure ρ_2) at time instant 60 is at most 2. Given that the reward of a state indicates the number of jobs processed per time-unit, property $\mathcal{P}_{\geqslant 0.98}(\mathit{3mup} \, \mathcal{U}_{[7,\infty)}^{[0,30]} \, \mathit{mdown})$ expresses that with probability at least 0.98 at least 7 jobs have been processed (starting from the initial state) before the first memory unit fails within 30 time units, where *3mup* is valid in states $(i, 3, 1)$, $1 \leqslant i < 4$ and *mdown* is valid in states $(i, 2, 1)$, $0 \leqslant i < 4$.

Semantics. The semantics of the **CSRL** path-formulas is defined as follows:

$$\begin{array}{ll} \sigma \models X_J^I \, \Phi & \text{iff } \sigma[1] \text{ is defined and } \sigma[1] \models \Phi \wedge \delta(\sigma, 0) \in I \wedge y(\sigma, \delta(\sigma, 0)) \in J \\ \sigma \models \Phi_1 \, \mathcal{U}_J^I \, \Phi_2 & \text{iff } \exists t \in I. \, (\sigma@t \models \Phi_2 \wedge (\forall t' \in [0, t). \, \sigma@t' \models \Phi_1) \wedge y(\sigma, t) \in J). \end{array}$$

Special cases occur for $I = [0, \infty)$ and $J = [0, \infty)$:

$$X \Phi = X_{[0,\infty)}^{[0,\infty)} \, \Phi \ \text{ and } \Phi_1 \, \mathcal{U} \, \Phi_2 = \Phi_1 \, \mathcal{U}_{[0,\infty)}^{[0,\infty)} \, \Phi_2.$$

[1] Strictly speaking, the function $s \mapsto Prob^{\mathcal{C}}(s, \Phi_1 \,\mathcal{U}\, \Phi_2)$ is the least fixpoint of a higher-order function on $(S \to [0,1]) \to (S \to [0,1])$ where the underlying partial order on $S \to [0,1]$ is defined for $F_1, F_2 : S \to [0,1]$ by $F_1 \leqslant F_2$ iff $F_1(s) \leqslant F_2(s)$ for all $s \in S$.

Thus, **SL** is a proper subset of this logic. The logic **CSL** [1,3] (or, timed stochastic CTL) is obtained in case $J = [0, \infty)$ for all sub-formulas. Similarly, we obtain the new logic **CRL** (reward-based stochastic CTL) in case $I = [0, \infty)$ for all sub-formulas. In the sequel, intervals of the form $[0, \infty)$ are often omitted from the operators.

We recall that $y(\sigma, t)$ denotes the cumulative reward along the prefix of σ up to time t. The intuition behind $y(\sigma, t)$ depends on the formula under consideration and the interpretation of the rewards in the MRM $\mathcal{M}$ under consideration. For instance, for $\varphi = \Diamond good$ and path σ that satisfies φ, the cumulative reward $y(\sigma, t)$ can be interpreted as the cost to reach a *good* state within t time units. For $\varphi = \Diamond bad$, it may be interpreted as the gain earned before reaching a *bad* state within t time units.

Alternative characterisations. We first observe that it suffices to consider time and reward bounds specified by closed intervals. Let $K = \{ x \in I \mid \rho(s) \cdot x \in J \}$ for closed intervals I and J. The probability of leaving state s at some time point x within the interval I such that the earned reward $\rho(s) \cdot x$ lies in J is can be expressed by

$$\mathbf{P}^I_J(s) = \int_K \mathbf{E}(s) \cdot e^{-\mathbf{E}(s) \cdot x} \; dx.$$

For instance, $\mathbf{P}^{[0,t]}_{[0,\infty)}(s) = 1 - e^{-\mathbf{E}(s) \cdot t}$, the probability to leave state s within t time units where the reward earned is irrelevant. If $\rho(s) = 2$, $I = [1, 3]$ and $J = [9, 11]$ then $K = \varnothing$ and $\mathbf{P}^I_J(s) = 0$. For $X^I_J \Phi$ we obtain:

$$Prob^{\mathcal{M}}(s, X^I_J \Phi) = \mathbf{P}^I_J(s) \cdot \mathbf{P}(s, \Phi).$$

For the case $I = J = [0, \infty)$ this reduces to equation (1).

Let $I \ominus x$ denote $\{ t - x \mid t \in I, t \geqslant x \}$. For $\varphi = \Phi_1 \, \mathcal{U}^I_J \, \Phi_2$ we have that $Prob^{\mathcal{M}}(s, \varphi)$ is the least solution of the following set of equations: $Prob^{\mathcal{M}}(s, \varphi) =$ 1 if $s \models \neg \Phi_1 \wedge \Phi_2$, $\inf I = 0$ and $\inf J = 0$,

$$\int_0^{\sup K} \sum_{s' \in S} \mathbf{P}(s, s', x) \cdot Prob^{\mathcal{M}}(s', \Phi_1 \, \mathcal{U}^{I \ominus x}_{J \ominus \rho(s) \cdot x} \, \Phi_2) \; dx \qquad (3)$$

if $s \models \Phi_1 \wedge \neg \Phi_2$, and

$$e^{-\mathbf{E}(s) \cdot \inf K} + \int_0^{\inf K} \sum_{s' \in S} \mathbf{P}(s, s', x) \cdot Prob^{\mathcal{M}}(s', \Phi_1 \, \mathcal{U}^{I \ominus x}_{J \ominus \rho(s) \cdot x} \, \Phi_2) \; dx$$

if $s \models \Phi_1 \wedge \Phi_2$, and 0 otherwise, where $\mathbf{P}(s, s', x) = \mathbf{R}(s, s') \cdot e^{-\mathbf{E}(s) \cdot x}$ denotes the probability of moving from state s to s' within x time units. The above characterisation is justified as follows. If s satisfies Φ_1 and $\neg \Phi_2$, the probability of reaching a Φ_2-state from s within the interval I by earning a reward $r \in J$ equals the probability of reaching some direct successor s' of s within x time units ($x \leqslant \sup I$ and $\rho(s) \cdot x \leqslant \sup J$, that is, $x \leqslant \sup K$), multiplied by the probability of reaching a Φ_2-state from s' in the remaining time interval $I \ominus x$

while earning a reward of $r-\rho(s)\cdot x$. If s satisfies $\Phi_1 \wedge \Phi_2$, the path-formula φ is satisfied if no transition outgoing from s is taken for at least $\inf K$ time units (first summand).[2] Alternatively, state s should be left before $\inf K$ in which case the probability is defined in a similar way as for the case $s \models \Phi_1 \wedge \neg\Phi_2$ (second summand). Note that $\inf K = 0$ is possible (if e.g., $\inf J = \inf I = 0$). In this case, $s \models \Phi_1 \wedge \Phi_2$ yields that any path starting in s satisfies $\Phi_1\, \mathcal{U}_J^I\, \Phi_2$ and $Prob^{\mathcal{M}}(s, \Phi_1\, \mathcal{U}_J^I\, \Phi_2) = 1$.

If the reward constraint is trivial, i.e., $J = [0,\infty)$, and I is of the form $[0,t]$ for $t \in \mathbb{R}_{\geqslant 0}$, then the characterisation for $\mathcal{U}^I$ reduces to the least solution of the following set of equations: $Prob^{\mathcal{M}}(s, \Phi_1\, \mathcal{U}^{[0,t]}\, \Phi_2)$ equals 1 if $s \models \Phi_2$, equals

$$\int_0^t \sum_{s' \in S} \mathbf{P}(s,s',x) \cdot Prob^{\mathcal{M}}(s', \Phi_1\, \mathcal{U}^{[0,t-x]}\, \Phi_2)\; dx \tag{4}$$

if $s \models \Phi_1 \wedge \neg\Phi_2$, and 0 otherwise. This coincides with the characterisation for time-bounded until in [3]. For the special case $I = J = [0,\infty)$ we obtain $K = [0,\infty)$ and hence the characterisation for $\mathcal{U}^I$ reduces to (2).

4 Duality

In this section we present the main result of the paper, a duality theorem that has important consequences for model checking sub-logics of **CSRL**. The basic idea behind this duality, inspired by [4], is that the progress of time can be regarded as the earning of reward and vice versa. First we obtain a duality result for MRMs where all states have a positive reward. After that we consider the (restricted) applicability of the duality result to MRMs with zero rewards.

Transformation of MRMs. Let $\mathcal{M} = (S, \mathbf{R}, L, \rho)$ be an MRM that satisfies $\rho(s) > 0$ for any state s. Define MRM $\mathcal{M}^{-1} = (S, \mathbf{R}', L, \rho')$ that results from $\mathcal{M}$ by: (i) rescaling the transition rates by the reward of their originating state (as originally proposed in [4]), i.e., $\mathbf{R}'(s,s') = \mathbf{R}(s,s')/\rho(s)$ and, (ii) inverting the reward structure, i.e., $\rho'(s) = 1/\rho(s)$. Intuitively, the transformation of $\mathcal{M}$ into $\mathcal{M}^{-1}$ stretches the residence time in state s with a factor that is proportional to the reciprocal of its reward $\rho(s)$ if $\rho(s) > 1$, and it compresses the residence time by the same factor if $0 < \rho(s) < 1$. The reward structure is changed similarly. Note that $\mathcal{M} = (\mathcal{M}^{-1})^{-1}$.

One might interpret the residence of t time units in $\mathcal{M}^{-1}$ as the earning of t reward in state s in $\mathcal{M}$, or (reversely) an earning of a reward r in state s in $\mathcal{M}$ corresponds to a residence of r in $\mathcal{M}^{-1}$. Thus, the notions of time and reward in $\mathcal{M}$ are reversed in $\mathcal{M}^{-1}$. Accordingly:

Lemma 1. *For MRM $\mathcal{M} = (S, \mathbf{R}, L, \rho)$ with $\rho(s) > 0$ for all $s \in S$ and* **CSRL** *state-formulas Φ, Φ_1 and Φ_2:*

1. $Prob^{\mathcal{M}}(s, X_J^I\, \Phi) = Prob^{\mathcal{M}^{-1}}(s, X_I^J\, \Phi)$

[2] By convention, $\inf \varnothing = \infty$.

2. $Prob^{\mathcal{M}}(s, \Phi_1\, \mathcal{U}^I_J\, \Phi_2) = Prob^{\mathcal{M}^{-1}}(s, \Phi_1\, \mathcal{U}^J_I\, \Phi_2)$.

We informally justify *2.* for $I = [0, t]$ and $J = [0, r]$ with $r, t \in \mathbb{R}_{\geqslant 0}$. Let MRM $\mathcal{M} = (S, \mathbf{R}, L, \rho)$ with $\rho(s) > 0$ for all $s \in S$. Let $s \in S$ be such that $s \models \Phi_1 \wedge \neg\Phi_2$. From equation (3) we have that $Prob^{\mathcal{M}^{-1}}(s, \Phi_1\, \mathcal{U}^J_I\, \Phi_2)$ equals

$$\int_{K'} \sum_{s' \in S} \mathbf{P}(s, s', x) \cdot Prob^{\mathcal{M}^{-1}}(s', \Phi_1\, \mathcal{U}^{J \ominus x}_{I \ominus \rho'(s) \cdot x}\, \Phi_2)\; dx.$$

for $K' = \{\, x \in [0, t] | \rho'(s) \cdot x \in [0, r]\,\}$, i.e., $K' = [0, min(t, \frac{r}{\rho'(s)})]$. By the definition of $\mathcal{M}^{-1}$ this equals

$$\int_{K'} \sum_{s' \in S} \frac{\mathbf{R}(s, s')}{\rho(s)} \cdot e^{-\frac{\mathbf{E}(s)}{\rho(s)} \cdot x} \cdot Prob^{\mathcal{M}^{-1}}(s', \Phi_1\, \mathcal{U}^{J \ominus x}_{I \ominus \frac{x}{\rho(s)}}\, \Phi_2)\; dx.$$

By substitution $y = \frac{x}{\rho(s)}$ this integral reduces to:

$$\int_{K} \sum_{s' \in S} \mathbf{R}(s, s') \cdot e^{-\mathbf{E}(s) \cdot y} \cdot Prob^{\mathcal{M}^{-1}}(s', \Phi_1\, \mathcal{U}^{J \ominus \rho(s) \cdot y}_{I \ominus y}\, \Phi_2)\; dy$$

where $K = [0, min(\frac{t}{\rho(s)}, r)]$. Thus, the function that maps (s, I, J) onto $Prob^{\mathcal{M}^{-1}}(s, \Phi_1\, \mathcal{U}^J_I\, \Phi_2)$ meets the fixed point equation for $Prob^{\mathcal{M}}(s, \Phi_1\, \mathcal{U}^I_J\, \Phi_2)$. Using arguments of fixed point theory, i.e., Tarski's theorem for least fixed points of monotonic functions on lattices, it can be shown that these fixed points agree (as they both are the least fixed point of the same operator). Thus, we obtain

$$\int_{K} \sum_{s' \in S} \mathbf{P}(s, s', y) \cdot Prob^{\mathcal{M}}(s', \Phi_1\, \mathcal{U}^{I \ominus y}_{J \ominus \rho(s) \cdot y}\, \Phi_2)\; dy$$

and this equals $Prob^{\mathcal{M}}(s, \Phi_1\, \mathcal{U}^I_J\, \Phi_2)$ for $s \models \Phi_1 \wedge \neg\Phi_2$, cf. (3).

For **CSRL** state-formula Φ let Φ^{-1} be defined as Φ where for each sub-formula in Φ of the form X^I_J or $\mathcal{U}^I_J$ the intervals I and J are swapped. This notion can be easily defined by structural induction on Φ and its definition is omitted here. For instance, for $\Phi = \mathcal{P}_{\geqslant 0.9}(\neg F \mathcal{U}^{[50,50]}_{[10,\infty)} F)$ we have $\Phi^{-1} = \mathcal{P}_{\geqslant 0.9}(\neg F \mathcal{U}^{[10,\infty)}_{[50,50]} F)$. We now have:

Theorem 1. *For MRM* $\mathcal{M} = (S, \mathbf{R}, L, \rho)$ *with* $\rho(s) > 0$ *for all* $s \in S$ *and* **CSRL** *state-formula* Φ*:*

$$Sat^{\mathcal{M}}(\Phi) \quad = \quad Sat^{\mathcal{M}^{-1}}(\Phi^{-1}).$$

If $\mathcal{M}$ contains states equipped with a zero reward, this duality result does not hold, as the reverse of earning a zero reward in $\mathcal{M}$ when considering Φ should correspond to a residence of 0 time units in $\mathcal{M}^{-1}$ for Φ^{-1}, which — as the advance of time in a state cannot be halted — is in general not possible. However, the result of Theorem 1 applies to some restricted, though still practical, cases, viz.

if (i) for each sub-formula of Φ of the form $X_J^I \Phi'$ we have $J = [0, \infty)$, and (ii) for each sub-formula of the form $\Phi_1 \,\mathcal{U}_J^I\, \Phi_2$ we either have $J = [0, \infty)$ or $Sat^{\mathcal{M}}(\Phi_1) \subseteq \{ s \in S \mid \rho(s) > 0 \}$, i.e., all Φ_1-states are positively rewarded. The intuition is that either the reward constraint (i.e., time constraint) is trivial in Φ (in Φ^{-1}), or that zero-rewarded states are not involved in checking the reward constraint. Here, we define $\mathcal{M}^{-1}$ by setting $\mathbf{R}'(s, s') = \mathbf{R}(s, s')$ and $\rho'(s) = 0$ in case $\rho(s) = 0$ and as defined above otherwise. For instance, Theorem 1 applies to the property $\mathcal{P}_{\geqslant 0.9}(\neg F \,\mathcal{U}_{[10,\infty)}^{[50,50]}\, F)$ for the multi-processor example, since all $\neg F$-states have a positive reward.

5 Application of the logic

In this section, we discuss model checking of **CSRL**. We furthermore illustrate that **CSRL** and its fragments **CSL** and **CRL** provide ample means for the specification of performability measures.

Model checking. CSL model checking can be carried out in the following way. $\mathcal{S}_{\bowtie p}(\Phi)$ gives rise to a system of linear equations for each bottom strongly connected component of the graph underlying the CTMC [3]. The probability to satisfy $\mathcal{U}$- and X-path formulas can be obtained as the solution of a system of linear equations, resp. a single matrix-vector multiplication [7], based on (1) and (2). Finally, the probability to satisfy a $\mathcal{U}^I$-formula can be obtained as the solution of a system of Volterra integral equations (4), that can be computed by either numerical integration [3], or transient analysis of the CTMC [2]. From Theorem 1, we can conclude that model checking an MRM against a **CRL**-formula can be performed using the algorithms established for model checking CTMCs against **CSL**:

Corollary 1. *For an MRM without any zero rewards, model checking* **CRL** *is reducible to model checking* **CSL**.

In a number of interesting, albeit restricted cases (cf. Sec 4), the corollary carries over to MRMs with zero rewards. The duality theorem does not provide an algorithmic recipe for **CSRL**, but a direct solution using numerical integration can be constructed based on the fixpoint characterisation for $\mathcal{U}_J^I$. An investigation of the feasibility of applying known efficient performability evaluation algorithms to model checking **CSRL** is ongoing.

Typical performability measures. Performability measures that frequently appear in the literature, e.g., [15], can be specified by simple **CSRL**-formulas. This is illustrated by Table 1 where we listed a (non-exhaustive) variety of typical performability measures for the multi-processor system together with the corresponding **CSRL** formulas. Measure (a) expresses a bound on the steady-state availability of the system and (b) expresses (a bound on) the probability to be not in a failed state at time t, i.e., the instantaneous availability at time t. Measure (c) expresses the time until a failure, starting from a non-failed state. Evaluating this measure for varying t, gives us the distribution of the time to

	performability measure	formula	logic
(a)	steady-state availability	$\mathcal{S}_{\bowtie p}(\neg F)$	**SL**
(b)	instantaneous availability at time t	$\mathcal{P}_{\bowtie p}(\Diamond^{[t,t]}\neg F)$	**CSL**
(c)	distribution of time to failure	$\mathcal{P}_{\bowtie p}(\neg F\, \mathcal{U}^{[0,t]}\, F)$	**CSL**
(d)	distribution of reward until failure	$\mathcal{P}_{\bowtie p}(\neg F\, \mathcal{U}_{[0,r]}\, F)$	**CRL**
(e)	distribution of cumulative reward until t	$\mathcal{P}_{\bowtie p}(\Diamond^{[t,t]}_{[0,r]}\mathrm{tt})$	**CSRL**

Table 1. Performability measures and their logical specification

failure. Measure (d) complements this by expressing the distribution of the reward accumulated until failure. Measure (e) generalises (c) and (d) by expressing the simultaneous distribution of the accumulated reward against time, i.e., it expresses the probability for the reward accumulated at t to be at most r. This measure coincides with the performability distribution as proposed in the seminal paper [12]. Note that for the computation of all these measures efficient algorithms do exist [9]. We emphasize that, in its full generality, **CSRL** allows to specify much more complex performability measures than previous ad hoc methods.

A possible extension of CSRL. Consider state s in MRM $\mathcal{M}$. For time t and set of states S', the *instantaneous reward* $\rho^{\mathcal{M}}(s, S', t)$ equals $\sum_{s' \in S'} \pi^{\mathcal{M}}(s, s', t) \cdot \rho(s')$ and denotes the rate at which reward is earned in some state in S' at time t. The *expected (or long run) reward rate* $\rho^{\mathcal{M}}(s, S')$ equals $\sum_{s' \in S'} \pi^{\mathcal{M}}(s, s') \cdot \rho(s')$. We can now add the following operators to our framework:

$$\begin{array}{lll} s \models \mathcal{E}_J(\Phi) & \text{iff} & \rho^{\mathcal{M}}(s, Sat^{\mathcal{M}}(\Phi)) \in J \\ s \models \mathcal{E}^t_J(\Phi) & \text{iff} & \rho^{\mathcal{M}}(s, Sat^{\mathcal{M}}(\Phi), t) \in J \\ s \models \mathcal{C}^I_J(\Phi) & \text{iff} & \int_I \rho^{\mathcal{M}}(s, Sat^{\mathcal{M}}(\Phi), u)\; du \in J \end{array}$$

Although the duality principle is not applicable to the new operators, their model checking is rather straightforward. The first two formulas require the summation of the Φ-conforming steady-state or transient state probabilities (as computed for measure (a) and (b)) multiplied with the corresponding rewards. The operator $\mathcal{C}^I_J(\Phi)$ states that the expected amount of reward accumulated in Φ-states during the interval I lies in J. It can be evaluated using a variant of uniformisation [9, 16]. Some example properties are now: $\mathcal{E}_J(\neg F)$, which expresses the expected reward rate (e.g., the system's capacity) for an operational system, $\mathcal{E}^t_J(\mathrm{tt})$ expresses the expected instantaneous reward rate at time t and $\mathcal{C}^{[0,t]}_J(\mathrm{tt})$ expresses the amount of cumulated reward up to time t.

6 Concluding remarks

We introduced a continuous-time, reward-based stochastic logic which is adequate for expressing performability measures of a large variety. Two important sub-logics were identified, viz. **CSL** [1, 3], and the novel logic **CRL** that allows one to express reward-based properties. The main result of the paper is that **CSL** and **CRL** are complementary, implying that **CRL**-properties for a Markov reward model can be interpreted as **CSL**-properties over a derived CTMC, so that existing model checking procedures for **CSL** can still be employed. The model checking of the full logic **CSRL**, in particular properties in which time- and reward-bounds are combined, is left for future work.

Acknowledgement. We thank the reviewers for their helpful comments.

References

1. A. Aziz, K. Sanwal, V. Singhal and R. Brayton. Verifying continuous time Markov chains. In *CAV*, LNCS 1102, pp. 269–276, 1996.
2. C. Baier, B.R. Haverkort, H. Hermanns and J.-P. Katoen. Model checking continuous-time Markov chains by transient analysis. In *CAV*, LNCS, 2000.
3. 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.
4. M.D. Beaudry. Performance-related reliability measures for computing systems. *IEEE Trans. on Comp. Sys.*, **27**(6): 540–547, 1978.
5. G. Clark, S. Gilmore, and J. Hillston. Specifying performance measures for PEPA. In *Form. Meth. for Real-Time and Prob. Sys. (ARTS)*, LNCS 1601, pp. 211–227, 1999.
6. C. Courcoubetis and M. Yannakakis. Verifying temporal properties of finite-state probabilistic programs. In *Found. of Comp. Sc. (FOCS)*, pp. 338–345, 1988.
7. H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Form. Asp. of Comp.*, **6**: 512–535, 1994.
8. B.R. Haverkort and I.G. Niemegeers. Performability modelling tools and techniques. *Perf. Ev.*, **25**: 17–40, 1996.
9. B.R. Haverkort. *Performance of Computer Communication Systems: A Model-Based Approach.* John Wiley & Sons, 1998.
10. H. Hermanns, J.-P. Katoen, J. Meyer-Kayser and M. Siegle. A Markov chain model checker. In *TACAS*, LNCS 1785, pp. 347–362, 2000.
11. R.A. Howard. *Dynamic Probabilistic Systems; Vol. I, II.* John Wiley & Sons, 1971.
12. J.F. Meyer. On evaluating the performability of degradable computing systems. *IEEE Trans. on Comp.*, **29**(8): 720–731, 1980.
13. J.F. Meyer. Performability evaluation: where it is and what lies ahead. In *1st IEEE Int. Comp. Perf. and Dependability Symp. (IPDS)*, pp. 334–343, 1995.
14. W.D. Obal and W.H. Sanders. State-space support for path-based reward variables. In *3rd IEEE Int. Comp. Perf. and Dependability Symp. (IPDS)*, pp. 228–237, 1998.
15. R.M. Smith, K.S. Trivedi and A.V. Ramesh. Performability analysis: measures, an algorithm and a case study. *IEEE Trans. on Comp.*, **37**(4): 406–417, 1988.
16. E. de Souza e Silva and H.R. Gail. Performability analysis of computer systems: from model specification to solution. *Perf. Ev.*, **14**: 157–196, 1992.
17. W.J. Stewart. *Introduction to the Numerical Solution of Markov Chains.* Princeton Univ. Press, 1994.
18. K.S. Trivedi, J.K. Muppala, S.P. Woolet, and B.R. Haverkort. Composite performance and dependability analysis. *Perf. Ev.*, **14**: 197–215, 1992.
19. M.Y. Vardi. Automatic verification of probabilistic concurrent finite state programs. *Found. of Comp. Sc. (FOCS)*, pp 327–338, 1985.