

Model Checking Action- and State-Labelled Markov Chains

Christel Baier
Universität Bonn
Institut für Informatik
baier@cs.uni-bonn.de

Lucia Cloth, Boudewijn Haverkort
University of Twente
Department of Computer Science
[lucia,brh]@cs.utwente.nl

Matthias Kuntz, Markus Siegle
University of the Federal Armed Forces Munich
Department of Computer Science
[kuntz,siegle]@informatik.unibw-muenchen.de

Abstract

In this paper we introduce the logic asCSL, an extension of continuous stochastic logic (CSL), which provides powerful means to characterise execution paths of action- and state-labelled Markov chains. In asCSL, path properties are characterised by regular expressions over actions and state-formulas. Thus, the executability of a path not only depends on the available actions but also on the validity of certain state formulas in intermediate states. Our main result is that the model checking problem for asCSL can be reduced to CSL model checking on a modified Markov chain, which is obtained through a product automaton construction. We provide a case study of a scalable cellular phone system which shows how the logic asCSL and the model checking procedure can be applied in practice.

1. Introduction

It becomes increasingly important to assert that distributed hardware and software systems are working correctly and that they meet high performance and dependability constraints. To reason about stochastic phenomena, several high-level models such as stochastic Petri nets, stochastic process algebras, queueing networks, etc., have been established. Typically, the verification of quantitative properties relies on a transformation of these high-level models into a Markov chain, on which the actual analysis is carried out. For the model-based verification of functional properties, temporal logics provide powerful means to specify complex requirements that a system has to satisfy, see e.g. [1]. Over the past 10 years, several researchers have adapted the temporal-logical approach to reason about

probabilistic properties. One result of these efforts is the logic CSL (continuous stochastic logic), introduced in [2] and extended in [3], which is a continuous-time variant of PCTL (probabilistic computation tree logic) [4] and can be used as specification formalism for performance and dependability properties. For instance, the CSL-formula $\mathcal{P}_{\geq 0.99}(\text{legal} \mathcal{U}^{\leq 5} \text{goal})$ specifies the state-property asserting that “there is at least a 99% chance to reach a goal state within the next 5 time units while passing only legal states before”. The goal states and legal states can be formalized, e.g., by atomic propositions that are attached to the states or by complex CSL-formulas. A so-called steady-state operator allows to reason about stationary probabilities. Formula $\mathcal{S}_{\geq 0.75}(\text{green})$ states that the probability mass accumulated in green states is at least 75%. An extension of CSL to reason about rewards has been introduced in [5]. However, the specification of these measures is completely state-oriented. For action-oriented modelling formalisms, such as stochastic process algebra, CSL is not expressive enough, since it is not possible to characterise sequences of actions. In [6] an action-based variant of CSL, aCSL, was proposed, and in [7] it was shown how to employ this logic for performability modelling. A further step is the logic aCSL+ [8], which combines state-oriented and action-oriented features, and employs regular expressions for characterising more general path-based properties.

In this paper we propose a new logic, called asCSL (for action- and state-oriented CSL), which subsumes all the above mentioned logics. (Preliminary work on logics similar to asCSL has been published in [9] and [10].) asCSL can be seen as being motivated either by the method of path-based reward variables as described in [11], or by the propositional dynamic logic [12] and extended linear temporal logic [13]. With asCSL, paths are characterised by regular expressions, also called programs, but in addition it is possible to express that a program is executable only if

the current state satisfies a given state property. This makes it possible to combine in an elegant way state- and action-oriented behavior. Unlike extended linear temporal logic [13], we do not allow for ω -regular expressions (representing infinite behaviors). Instead, in **asCSL** the regular expressions are used in combination with lower and upper time bounds. Thus, the switch from CSL to **asCSL** is orthogonal to the extension **PCTL*** [14] of **PCTL** which is concerned with specifying “complex” properties of infinite computations, whereas **asCSL** focusses on “complex” properties of finite computations with real-time constraints (e.g. deadlines).

This paper is further organised as follows: In Section 2 we define action- and state-labelled Markov chains. Section 3 presents syntax and semantics of the new logic **asCSL**. Section 4 is dedicated to the model checking procedure for **asCSL**. In Section 5 we apply the new technique to derive properties of the handover procedure in a cellular radio network. The paper ends with a short summary and conclusions.

2. Action and state labelled Markov chains

In this section we explain the notation for Markov chains with action and state labels. The reader is supposed to be familiar with Markov chains, c.f. [15]. Action names are used to label the transitions. The state labels are formalized by a set **AP** of atomic propositions, which e.g. can assert that the system (or one of its subprocesses) is at a certain control point or that a program variable has a certain value, and a function L that assigns to any state the set of atomic propositions that are assumed to hold in the given state.

Definition 1 (Action- and state-labelled Markov chain)

An action- and state-labelled continuous-time Markov chain (ASMC) is a tuple $\mathcal{M} = (S, \text{Act}, \text{AP}, L, \mathbf{R})$, where S is a finite set of states, **Act** is a finite set of action labels, **AP** is a finite set of atomic propositions, $L : S \rightarrow 2^{\text{AP}}$ a state labelling function, and $\mathbf{R} : S \times \text{Act} \times S \rightarrow \mathbb{R}_{\geq 0}$ a rate matrix. We require that for any state s there exists a pair $(a, s') \in \text{Act} \times S$ with $\mathbf{R}(s, a, s') > 0$. \square

Intuitively, if $\mathbf{R}(s, a, s') = \lambda > 0$ then there is an a -labelled transition from state s to state s' whose delay is specified by an exponential distribution with rate λ . For $S' \subseteq S$ we write $\mathbf{R}(s, a, S') = \sum_{s' \in S'} \mathbf{R}(s, a, s')$ to denote the total rate to move from state s via action a to state-set S' .

Definition 2 (Paths in \mathcal{M})

A finite path in \mathcal{M} is a finite word $\sigma = (s_0, a_0, t_0), (s_1, a_1, t_1), \dots, (s_{n-1}, a_{n-1}, t_{n-1}), s_n \in (S \times \text{Act} \times \mathbb{R}_{>0})^* \times S$, where $\mathbf{R}(s_i, a_i, s_{i+1}) > 0$ for $i = 0, \dots, n-1$. An infinite path in \mathcal{M} is an infinite word $\zeta = (s_0, a_0, t_0), (s_1, a_1, t_1), \dots \in (S \times \text{Act} \times \mathbb{R}_{>0})^\omega$ with $\mathbf{R}(s_i, a_i, s_{i+1}) > 0$ for all $i \geq 0$ and such that the infinite series $\sum_i t_i$ is divergent (i.e., $t_0 + t_1 + t_2 + \dots = \infty$). \square

We write paths as sequences of transitions, e.g., for the finite path σ above, we use the notation

$$\sigma = s_0 \xrightarrow{a_0, t_0} s_1 \xrightarrow{a_1, t_1} s_2 \xrightarrow{a_2, t_2} \dots \xrightarrow{a_{n-2}, t_{n-2}} s_{n-1} \xrightarrow{a_{n-1}, t_{n-1}} s_n.$$

Let σ be a finite path as before, then $|\sigma| = n$ denotes the length of σ , i.e., number of transitions in σ , $\sigma[i] = s_i$ the $(i+1)$ st state on σ . We refer to $\tau(\sigma) = \sum_{j=0}^{n-1} t_j$ as the execution time of σ . For $t \leq \tau(\sigma)$, $\sigma@t$ denotes the state that is occupied at time t on path σ , that is, $\sigma@t = \sigma[k]$, where k is the smallest index for which $t < \sum_{j=0}^k t_j$. We write $\sigma(i, j)$ to denote the fragment of path σ starting at the $(i+1)$ -st state s_i and ending at the $(j+1)$ -st state s_j ($i \leq j$).

If σ_1, σ_2 are finite paths such that the first state of σ_2 agrees with the last state of σ_1 then the concatenation $\sigma_1\sigma_2$ is a path of length $|\sigma_1| + |\sigma_2|$ which is defined in the obvious way. Similar notation is used for infinite paths. $\text{Path}_{\text{fin}}^{\mathcal{M}}$ denotes the set of all finite paths in \mathcal{M} , whereas $\text{Path}_{\text{fin}}^{\omega, \mathcal{M}}$ stands for the set of infinite paths in \mathcal{M} . By $\text{Path}_{\text{fin}}^{\mathcal{M}}(s)$, resp. $\text{Path}_{\text{fin}}^{\omega, \mathcal{M}}(s)$, we denote the set of all finite, resp. infinite, paths in \mathcal{M} with initial state s .

In the sequel, we shall deal with the standard probability measure $\Pr_s^{\mathcal{M}}$ on $\text{Path}_{\text{fin}}^{\omega, \mathcal{M}}(s)$ (where the underlying σ -field can be defined with the help of basic cylinders as in [16]). For measurable $X \subseteq \text{Path}_{\text{fin}}^{\omega, \mathcal{M}}(s)$, we often skip the parameter \mathcal{M} and/or s and simply write $\Pr(X)$ or $\Pr^{\mathcal{M}}(X)$. The transient state probabilities $\pi^{\mathcal{M}}(s, s', t)$ are given by

$$\pi^{\mathcal{M}}(s, s', t) = \Pr^{\mathcal{M}} \{ \zeta \in \text{Path}_{\text{fin}}^{\omega, \mathcal{M}}(s) : \zeta@t = s' \},$$

and the steady state probability of being in state s' on the long run, provided that the system started in state s , is $\pi^{\mathcal{M}}(s, s') = \lim_{t \rightarrow \infty} \pi^{\mathcal{M}}(s, s', t)$. For $S' \subseteq S$, we put $\pi^{\mathcal{M}}(s, S') = \sum_{s' \in S'} \pi^{\mathcal{M}}(s, s')$.

3. Syntax and semantics of **asCSL**

We now extend CSL (continuous stochastic logic) as introduced in [2, 3] by path formulas that specify constraints for action- and state-sequences in combination with lower and/or upper time bounds.

3.1. Syntax of **asCSL**

The syntax of **asCSL** is defined according to Def. 3 (state formulas) and Def. 4 (programs or path formulas). Here, we assume that the sets **Act** of actions and **AP** of atomic propositions are fixed.

Definition 3 (State formulas of **asCSL)** State formulas of **asCSL** are given by the following grammar:

$$\phi ::= \mathbf{q} \mid \neg\phi \mid \phi \vee \phi \mid \mathcal{L}_{\triangleright \Delta p}(\phi) \mid \mathcal{P}_{\triangleright \Delta p}(\alpha^I),$$

where $q \in \text{AP}$ is an atomic proposition, $p \in [0, 1]$ denotes a probability value, $\bowtie \in \{<, \leq, >, \geq\}$ a comparison operator, $I = [t, t'] \subseteq \mathbb{R}_{\geq 0}$ a time interval and α a program as defined in Def. 4. We refer to α^I as a path-formula and use Φ to denote the set of state formulas of asCSL. \square

The logical connectives \neg and \vee have their usual meaning. Using negation \neg and disjunction \vee , the constants true, false and all other boolean connectives such as conjunction \wedge , implication \rightarrow , etc. can be derived. The so-called steady-state operator $\mathcal{S}_{\bowtie p}(\phi)$ asserts that the probability of being in a ϕ -state on the long run is within the interval $\bowtie p$. The operator $\mathcal{P}_{\bowtie p}(\alpha^I)$ asserts that the probability measure of all infinite paths which have a prefix that satisfies α^I , where α is a program and I is a real interval specifying the time bound, is within the interval $\bowtie p$.

The program α specifies a property for finite paths via a regular set of finite words whose atomic symbols are pairs (ϕ, b) (in the sequel we simply write ϕb) consisting of an asCSL-state formula ϕ (which is viewed as a test for the current state of a path) and an action $b \in \text{Act}$ or $b = \surd$ (where $\surd \notin \text{Act}$). The symbol \surd can be viewed as a pseudo-action which is always immediately executable and does not change the current state in the ASMC. Formally, the programs are regular expressions over the alphabet

$$\Sigma = \Phi \times (\text{Act} \cup \{\surd\}) = \{\phi b \mid \phi \in \Phi \wedge b \in (\text{Act} \cup \{\surd\})\}.$$

Definition 4 (Programs) asCSL-programs are defined by the following grammar:

$$\alpha ::= \varepsilon \mid \phi b \mid \alpha; \alpha \mid \alpha \cup \alpha \mid \alpha^*,$$

where $\phi b \in \Sigma$ (i.e., $\phi \in \Phi$, $b \in \text{Act} \cup \{\surd\}$). The language $\mathcal{L}(\alpha) \subseteq \Sigma^*$ is defined in the standard way. \square

The symbol ε stands for the empty word (as an element of Σ^*). The intuitive meaning of ϕb is that the current state s fulfills the test ϕ (that is, the formula ϕ holds in s) and, if $b \in \text{Act}$, state s has an outgoing b -transition. If $b = \surd$, there is no statement about outgoing transitions. Note that the test can be empty, i.e., $\phi = \text{true}$. The operator $;$ denotes sequential composition (concatenation), \cup denotes alternative choice (union), and $*$ denotes the n -fold sequential composition for arbitrary $n \geq 0$ (Kleene star). As examples for the corresponding languages, $\mathcal{L}(\phi_1 a; (\phi_2 \surd \cup \phi_3 b))$ consists of the two words $\phi_1 a, \phi_2 \surd$ and $\phi_1 a, \phi_3 b$ whereas $\mathcal{L}((\phi_1 a; \phi_2 b; \phi_3 \surd)^*; \phi_4 c)$ is an infinite language that contains the words $\phi_4 c$ and $\phi_1 a, \phi_2 b, \phi_3 \surd, \phi_4 c$ and $\phi_1 a, \phi_2 b, \phi_3 \surd, \phi_1 a, \phi_2 b, \phi_3 \surd, \phi_4 c$, and so on.

In the context of asCSL, the meaning of a program (which will be formally defined in the next subsection) is a set of finite paths in the underlying ASMC.

3.2. Semantics of asCSL

The semantics of asCSL is provided by means of a satisfaction relation \models for the state- and path formulas. In the sequel, we assume a fixed ASMC $\mathcal{M} = (S, \text{Act}, \text{AP}, L, \mathbf{R})$. For state s in \mathcal{M} and state formula ϕ , $\mathcal{M}, s \models \phi$ means that the state-property specified by ϕ holds for s . Similarly, for an infinite path ζ , $\mathcal{M}, \zeta \models \alpha^I$ denotes that the behaviour specified by the path formula α^I is fulfilled by ζ . The formal definition of the satisfaction relation \models for the state and path formulas is by structural induction on the syntax of the formulas.

Definition 5 (Semantics of asCSL) The satisfaction relation \models for the state formulas is defined as follows:

$$\begin{aligned} \mathcal{M}, s \models q & \iff q \in L(s) \\ \mathcal{M}, s \models \neg \phi & \iff \mathcal{M}, s \not\models \phi \\ \mathcal{M}, s \models \phi_1 \vee \phi_2 & \iff \mathcal{M}, s \models \phi_1 \text{ or } \mathcal{M}, s \models \phi_2 \\ \mathcal{M}, s \models \mathcal{S}_{\bowtie p}(\phi) & \iff \pi^{\mathcal{M}}(s, \text{Sat}^{\mathcal{M}}(\phi)) \bowtie p \\ \mathcal{M}, s \models \mathcal{P}_{\bowtie p}(\alpha^I) & \iff \text{Prob}^{\mathcal{M}}(s, \alpha^I) \bowtie p \end{aligned}$$

where $\text{Sat}^{\mathcal{M}}(\phi) = \{s \in S : \mathcal{M}, s \models \phi\}$ denotes the satisfaction set of ϕ in \mathcal{M} and

$$\text{Prob}^{\mathcal{M}}(s, \alpha^I) = \text{Pr}^{\mathcal{M}}\{\zeta \in \text{Path}_{\omega}^{\mathcal{M}}(s) \mid \mathcal{M}, \zeta \models \alpha^I\}.$$

The meaning of the path formulas is formalized as follows. If ζ is an infinite path in \mathcal{M} then $\mathcal{M}, \zeta \models \alpha^I$ iff there exists a finite prefix σ of ζ with $\sigma \in \text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha)$ and $\tau(\sigma) \in I$.¹ Here, the set $\text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha)$ consists of all finite paths σ in \mathcal{M} that can be viewed as an instance of α . Formally:

$$\begin{aligned} \sigma \in \text{Path}_{\text{fin}}^{\mathcal{M}}(\varepsilon) & \iff |\sigma| = 0 \\ \sigma \in \text{Path}_{\text{fin}}^{\mathcal{M}}(\phi a) & \iff \exists t > 0 \text{ s.t. } \sigma = s \xrightarrow{a,t} s' \\ & \text{and } \mathcal{M}, s \models \phi \\ \sigma \in \text{Path}_{\text{fin}}^{\mathcal{M}}(\phi \surd) & \iff \sigma = s \text{ and } \mathcal{M}, s \models \phi \\ \sigma \in \text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha_1; \alpha_2) & \iff \exists i \in \{0, 1, \dots, |\sigma|\} \text{ s.t.} \\ & \sigma(0, i) \in \text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha_1) \text{ and } \sigma(i, |\sigma|) \in \text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha_2) \\ \sigma \in \text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha_1 \cup \alpha_2) & \iff \sigma \in \text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha_1) \cup \text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha_2) \\ \sigma \in \text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha^*) & \iff \exists i \geq 0 \text{ s.t. } \sigma \in \text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha^i) \end{aligned}$$

where $\alpha^{i+1} = \alpha; \alpha^i$ and $\alpha^0 = \varepsilon$ (the empty word in Σ^*).² In the sequel, we often write $s \models \phi$ and $\zeta \models \alpha^I$ rather than $\mathcal{M}, s \models \phi$ and $\mathcal{M}, \zeta \models \alpha^I$ respectively. \square

We simply write ϕA for $\bigcup_{a \in A} \phi a$. The reader should notice the difference between the programs $\alpha_1 = (\text{true Act})^* \phi \surd$ and $\alpha_2 = (\text{true Act})^* \phi \text{Act}$. Program α_1 stands for all finite paths σ that end in a ϕ -state, whereas α_2 stands for all finite paths whose prefinal state satisfies ϕ .

A time-bounded until operator can be derived from the syntax of asCSL by $\phi_1 \mathcal{U}^I \phi_2 = ((\phi_1 \text{Act})^* \phi_2 \surd)^I$. From

¹Recall that $\tau(\sigma)$ denotes the execution time of σ .

²Note that all paths σ with $|\sigma| = 0$ belong to $\text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha^*)$.

this, we may derive the time-bounded eventually-operator $\diamond^I \phi = \text{true} \mathcal{W}^I \phi$. For instance, $\mathcal{P}_{\leq 0.05}(\diamond^{\leq 5} \text{error})$ states that the probability to reach an error state within 5 time units is bounded above by 0.05. Its dual, the time-bounded always-operator, is obtained (as in CSL) through $\mathcal{P}_{\geq p}(\square^I \phi) = \mathcal{P}_{\leq 1-p}(\diamond^I \neg \phi)$, stating that ϕ continuously holds in the time interval I with probability at least p .

3.3. asCSL-equivalence and bisimulation

The correctness proof of the model checking algorithm presented in Section 4 will make use of the observation that asCSL-equivalence (which identifies those states that cannot be distinguished by asCSL-formulas) agrees with bisimulation equivalence for ASMCs. The latter can be viewed as a variant of lumpability [17, 18] and essentially agrees with Markovian bisimulation as introduced in [19, 20] for action-labelled Markov chains. Formally, bisimulation equivalence \sim for an ASMC $\mathcal{M} = (S, \text{Act}, \text{AP}, L, \mathbf{R})$ is the coarsest equivalence on S such that for all $s_1 \sim s_2$: (i) $L(s_1) = L(s_2)$ and (ii) $\mathbf{R}(s_1, a, C) = \mathbf{R}(s_2, a, C)$ for all actions $a \in \text{Act}$ and all equivalence classes $C \in S/\sim$.

It is well-known that bisimulation equivalence \sim agrees with CSL-equivalence [16, 21] (which means the equivalence that identifies exactly those states that satisfy the same CSL state formulas). In particular, bisimulation equivalent states have the same transient and steady state probabilities. Using structural induction on the syntax of state- and path-formulas of asCSL the preservation property for asCSL and bisimulation equivalence can be established in the following sense: If $s_1 \sim s_2$ then we have $s_1 \models \phi$ iff $s_2 \models \phi$ for all state-formulas ϕ of asCSL: and $\text{Prob}^{\mathcal{M}}(s_1, \alpha^I) = \text{Prob}^{\mathcal{M}}(s_2, \alpha^I)$ for all path-formulas α^I of asCSL. As asCSL subsumes the logic CSL we obtain by the completeness result established in [21]:

Proposition 1 *Bisimulation equivalence agrees with asCSL-equivalence.*

As Markovian testing equivalence [22] is weaker than bisimulation equivalence, states that fulfill the same asCSL-formulas are Markovian testing equivalent.

4. Model Checking asCSL

The model checking procedure for asCSL is similar to that for CTL [23]. Given the asCSL state formula ϕ and an ASMC \mathcal{M} , we successively consider the subformulas ψ of ϕ and calculate the satisfaction sets $\text{Sat}^{\mathcal{M}}(\psi) = \{s \in S : \mathcal{M}, s \models \psi\}$. This technique allows us to treat subformulas as atomic propositions. The treatment of subformulas whose top-level operator is a boolean connective (negation or disjunction) is obvious. Subformulas of the form $\mathcal{S}_{\bowtie p}(\phi)$ can be handled with the same procedure as for

CSL, see [24]. The new and challenging case is the treatment of formulas of type $\phi = \mathcal{P}_{\bowtie p}(\alpha^I)$. For each state s we have to compute the probability $\text{Prob}^{\mathcal{M}}(s, \alpha^I) = \text{Pr}^{\mathcal{M}}\{\zeta \in \text{Path}_{\omega}^{\mathcal{M}}(s) \mid \mathcal{M}, \zeta \models \alpha^I\}$ and check whether it lies within the specified bound $\bowtie p$. The approach to calculate the values $\text{Prob}^{\mathcal{M}}(s, \alpha^I)$ is to build the product of \mathcal{M} and a finite automaton \mathcal{A}_{α} (representing the program α), which yields a new Markov chain, denoted $\mathcal{M} \times \mathcal{A}_{\alpha}$, and then to apply the CSL model checking procedure to calculate the probabilities in $\mathcal{M} \times \mathcal{A}_{\alpha}$ to reach a state $\langle s', q \rangle$, with s' a state in \mathcal{M} and q a final state in \mathcal{A}_{α} , within the time interval I .

Section 4.1 is devoted to the construction of the automaton \mathcal{A}_{α} , whereas Section 4.2 presents the construction of the product Markov chain $\mathcal{M} \times \mathcal{A}_{\alpha}$.

4.1. Program automata

Since programs are regular expressions, we can apply standard techniques to construct a finite automaton for a given program. We call this a nondeterministic program automaton (NPA).

Definition 6 (NPA) *An NPA is a quintuple $\mathcal{A} = (Z, \Sigma', \delta, Z_0, F)$, where Z is a finite set of states, Σ' a finite subset of Σ (the input alphabet)³, $\delta : Z \times \Sigma' \rightarrow 2^Z$ is the transition function, $Z_0 \subseteq Z$ the set of initial states and $F \subseteq Z$ the set of accepting (final) states. $\mathcal{L}(\mathcal{A}) \subseteq (\Sigma')^*$ denotes the accepted language of \mathcal{A} which is defined in the standard way. \square*

We now describe how a program automaton \mathcal{A} can be used to describe the path-set $\text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha)$ for a program α . Thus, we consider NPAs as *acceptors for finite paths* in \mathcal{M} (rather than as acceptors for finite words over the alphabet Σ'). The intuitive behaviour of an NPA \mathcal{A} for the input path $\sigma = s \xrightarrow{a^I} \sigma'$ is as follows. The automaton starts in one of its initial states $z_0 \in Z_0$. If the current automata-state is z , then \mathcal{A} chooses nondeterministically between one of the outgoing transitions $z \xrightarrow{\phi b} z'$, where $s \models \phi$ and either $b = \surd$ or $b = a$, and then moves to state z' . In the latter case, i.e., if $a = b$, \mathcal{A} proceeds in the same way for state z' and the input path σ' . In the former case, i.e., if $b = \surd$, no input symbol is consumed, i.e., the procedure is repeated with state z' and the input path σ . If there is no outgoing transition from z which can be taken for the input path σ then \mathcal{A} rejects. As soon as \mathcal{A} reaches a final state (a state in F) and the whole input path has been consumed, the automaton accepts. The formal definition of acceptance for paths is provided by means of runs which are sequences of automaton-states that can be generated by the operational behaviour of \mathcal{A} as sketched above.

³ Σ is the alphabet of programs as defined in Def. 4.

Definition 7 (Runs in NPAs, accepted paths) Let \mathcal{A} be a NPA and \mathcal{M} an ASMC as before, $z \in Z$ and σ a finite path in \mathcal{M} . Then, we define $\text{Runs}(z, \sigma)$ as the greatest set of sequences $z, z_1, \dots, z_n \in Z^+$ such that the following two conditions are fulfilled:

- (i) $z \in \text{Runs}(z, \sigma)$ iff $|\sigma| = 0$
- (ii) If $z, z_1, \dots, z_n \in \text{Runs}(z, \sigma)$ and $n \geq 1$ then there exists $\phi b \in \Sigma'$ such that $z_1 \in \delta(z, \phi b)$, $\sigma[0] \models \phi$ and
 - if $b \in \text{Act}$ then $\sigma = s \xrightarrow{b, t} \sigma'$ with $z_1, \dots, z_n \in \text{Runs}(z_1, \sigma')$
 - if $b = \surd$ then $z_1, \dots, z_n \in \text{Runs}(z_1, \sigma)$.

Let $Z' \subseteq Z$. The elements of $\text{Runs}(Z', \sigma) = \bigcup_{z \in Z'} \text{Runs}(z, \sigma)$ are called runs for σ in \mathcal{A} with starting state in Z' . A run z_0, z_1, \dots, z_n for σ is called accepting iff it is initial (i.e., $z_0 \in Z_0$) and $z_n \in F$. The set of accepted paths, $\text{Path}_{\text{fin}}^{\mathcal{M}}(\mathcal{A})$, denotes the set of finite paths in \mathcal{M} that have an accepting run in \mathcal{A} . \square

Acceptance of \mathcal{A} as an ordinary finite automaton (acceptor for finite words over Σ') and acceptance of \mathcal{A} as an NPA (acceptor for finite paths) are related in the same way as the language $\mathcal{L}(\alpha)$ of a program α and the induced path-set $\text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha)$. Hence, it is easy to verify the following proposition:

Proposition 2 If $\mathcal{L}(\alpha) = \mathcal{L}(\mathcal{A})$ then $\text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha) = \text{Path}_{\text{fin}}^{\mathcal{M}}(\mathcal{A})$.

We now extend the transition function δ of \mathcal{A} to a transition relation $\hat{\delta}^{\mathcal{M}}$ which associates with any pair (Z', σ) consisting of a set Z' of automaton-states and a finite path σ in the ASMC \mathcal{M} , the set of automaton states z such that z is the last state of a run for σ that starts in a Z' -state. The idea behind the definition of $\hat{\delta}^{\mathcal{M}}$ is similar to the definition of the transition relation of the deterministic finite automaton obtained from \mathcal{A} (viewed as an acceptor for finite words) via the standard powerset construction. However, the following remark shows that the determinisation process of an NPA as an acceptor for finite paths in \mathcal{M} has to be done “in conjunction” with \mathcal{M} .

Remark. An NPA $\mathcal{A} = (Z, \Sigma', \delta, Z_0, F)$ is called deterministic if Z_0 is a singleton set and $|\delta(z, \phi b)| \leq 1$ for all states $z \in Z$ and input symbols $\phi b \in \Sigma'$. Given a deterministic NPA \mathcal{A} , the “behaviour” of \mathcal{A} for an input word over Σ' is deterministic, i.e., there is at most one run, whereas the “behaviour” of a deterministic NPA \mathcal{A} for an input path σ can be *nondeterministic*, even if \mathcal{A} does not contain \surd -transitions (i.e., transitions that are labelled with an input symbol $\phi \surd \in \Sigma'$). The reason is that the current automaton-state z might have two transitions $z \xrightarrow{\phi a} z'$ and $z \xrightarrow{\psi a} z''$, where a is the first action of the input path σ and where the first state of σ satisfies both ϕ and ψ . \square

We now return to the formal definition of the extended transition function $\hat{\delta}^{\mathcal{M}}$. If σ is a path of length 0, i.e., $\sigma = s$ for some state s , then $\hat{\delta}^{\mathcal{M}}(Z', \sigma) = \hat{\delta}^{\mathcal{M}}(Z', s)$ consists of all automaton-states \tilde{z} that are reachable in \mathcal{A} from a Z' -state via \surd -transitions $z_1 \xrightarrow{\phi, \surd} z_2$ where state s fulfills the state formulas ϕ . This corresponds to the so-called \surd -closure of Z' for state s which is defined as follows.

Definition 8 (\surd -closure) $\surd\text{Closure}(Z', s)$ denotes the least subset of Z such that

$$Z' \cup \bigcup_{z \in \surd\text{Closure}(Z', s)} \bigcup_{\substack{\phi \in \Phi \text{ s.t.} \\ s \models \phi}} \delta(z, \phi \surd) \subseteq \surd\text{Closure}(Z', s). \quad \square$$

We now have all ingredients to define $\hat{\delta}^{\mathcal{M}}(Z', \sigma)$ by induction on the length of σ :

Definition 9 (Extended transition function) The function $\hat{\delta}^{\mathcal{M}} : 2^Z \times \text{Path}_{\text{fin}}^{\mathcal{M}} \rightarrow 2^Z$ is given by:

$$\hat{\delta}^{\mathcal{M}}(Z', s) = \surd\text{Closure}(Z', s)$$

and $\hat{\delta}^{\mathcal{M}}(Z', s \xrightarrow{a, t} \sigma') = \hat{\delta}^{\mathcal{M}}(Y, \sigma')$ where

$$Y = \bigcup_{z \in \surd\text{Closure}(Z', s)} \bigcup_{\substack{\phi \in \Phi \text{ s.t.} \\ s \models \phi}} \delta(z, \phi a). \quad \square$$

Note that Y stands for the set of all automaton states y that are reachable in \mathcal{A} from a state $z' \in Z'$ via transitions labelled with elements $\psi \surd \in \Sigma$ such that $s \models \psi$ followed by a transition with a label $\phi a \in \Sigma$ such that $s \models \phi$.

It can be shown by induction on $|\sigma|$ that $\hat{\delta}^{\mathcal{M}}(Z', \sigma)$ consists of all states that are reachable in \mathcal{A} via a run starting in Z' for σ , i.e.,

$$\hat{\delta}^{\mathcal{M}}(Z', \sigma) = \{z \in Z \mid \exists z_0, z_1, \dots, z_n \in \text{Runs}(Z', \sigma) : z_n = z\}.$$

For $Z' = Z_0$, we obtain that $\hat{\delta}^{\mathcal{M}}(Z_0, \sigma)$ consists of all automaton-states that can be reached via an initial run for σ . Thus, $\text{Path}_{\text{fin}}^{\mathcal{M}}(\mathcal{A}) = \{\sigma \in \text{Path}_{\text{fin}}^{\mathcal{M}} : \hat{\delta}^{\mathcal{M}}(Z_0, \sigma) \cap F \neq \emptyset\}$. From this observation we obtain:

Proposition 3 If α is a program and \mathcal{A} an NPA with $\mathcal{L}(\alpha) = \mathcal{L}(\mathcal{A})$ then we have:

$$\text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha) = \{\sigma \in \text{Path}_{\text{fin}}^{\mathcal{M}} : \hat{\delta}^{\mathcal{M}}(Z_0, \sigma) \cap F \neq \emptyset\}.$$

Example 1 We consider the ASMC shown in Fig. 1(a) and the program

$$\alpha = (\phi a)^*; (\psi a); ((\text{true } c) \cup (\xi b))$$




Figure 1. ASMC \mathcal{M} (a) and NPA $\mathcal{A} = \mathcal{A}_\alpha$ (b)

with corresponding program automaton $\mathcal{A} = \mathcal{A}_\alpha$ as shown in Fig. 1(b). For the relevant state formulas we assume $\text{Sat}^{\mathcal{M}}(\phi) = \{s_1, s_2\}$, $\text{Sat}^{\mathcal{M}}(\psi) = \{s_4\}$ and $\text{Sat}^{\mathcal{M}}(\xi) = \{s_5\}$. The finite paths that start in state s_1 and end in state s_6 have the form

$$\begin{aligned} \sigma_1 &= s_1 \xrightarrow{a, \cdot} s_2 \xrightarrow{a, \cdot} s_4 \xrightarrow{a, \cdot} s_5 \xrightarrow{b, \cdot} s_6 \\ \text{or } \sigma_2 &= s_1 \xrightarrow{c, \cdot} s_3 \xrightarrow{d, \cdot} s_4 \xrightarrow{a, \cdot} s_5 \xrightarrow{b, \cdot} s_6, \end{aligned}$$

with arbitrary sojourn times. We then have $\hat{\delta}^{\mathcal{M}}(\{z_1, z_2\}, \sigma_1) = \{z_4\}$ and $\hat{\delta}^{\mathcal{M}}(\{z_1, z_2\}, \sigma_2) = \emptyset$. Therefore, all paths σ_1 belong to $\text{Path}_{\text{fin}}^{\mathcal{M}}(\mathcal{A})$ as we have $\hat{\delta}^{\mathcal{M}}(Z_0, \sigma_1) \cap F \neq \emptyset$, whereas none of the paths σ_2 belong to $\text{Path}_{\text{fin}}^{\mathcal{M}}(\mathcal{A})$ because $\hat{\delta}^{\mathcal{M}}(Z_0, \sigma_2) \cap F = \emptyset$. \square

4.2. The product Markov chain

We now return to the question of how to calculate the satisfaction set $\text{Sat}^{\mathcal{M}}(\phi)$ where $\phi = \mathcal{P}_{\times p}(\alpha^I)$. We first apply recursively an aCSL-model checking algorithm to the state formulas that occur in the program α . As soon as the satisfaction sets $\text{Sat}^{\mathcal{M}}(\psi)$ are known for all state formulas ψ in α we can treat them as atomic propositions. Then, we apply standard algorithms to construct a (nondeterministic) finite automaton \mathcal{A} for α (viewed as an ordinary regular expression over the alphabet Σ). We then consider \mathcal{A} as an NPA and build the product of the ASMC \mathcal{M} and \mathcal{A} (which is defined below) and finally apply a CSL model checking algorithm to $\mathcal{M} \times \mathcal{A}$ to calculate the probability to reach a final automaton state within the given time interval I .

Definition 10 (Product Markov chain \mathcal{M}^\times) Let $\mathcal{M} = (S, \text{Act}, \text{AP}, L, \mathbf{R})$ be an ASMC and $\mathcal{A} = (Z, \Sigma, \delta, Z_0, F)$ an NPA. The product ASMC is defined as $\mathcal{M} \times \mathcal{A} = (S^\times, \text{Act}^\times, \text{AP}^\times, L^\times, \mathbf{R}^\times)$ where

$$S^\times = \{\langle s, Z' \rangle \mid s \in S \wedge Z' \in 2^Z\},$$

$\text{Act}^\times = \text{Act}$, $\text{AP}^\times = \text{AP} \cup \{\text{accept}\}$ (where $\text{accept} \notin \text{AP}$) and $L^\times(\langle s, Z' \rangle) = L(s) \cup \{\text{accept}\}$ if $Z' \cap F \neq \emptyset$ and

$L^\times(\langle s, Z' \rangle) = L(s)$ otherwise. The rate matrix \mathbf{R}^\times is defined by $\mathbf{R}^\times(\langle s_1, Z_1 \rangle, a, \langle s_2, Z_2 \rangle) = \mathbf{R}(s_1, a, s_2)$, if $Z_2 = \hat{\delta}^{\mathcal{M}}(Z_1, s_1 \xrightarrow{a, \cdot} s_2)$, and $\mathbf{R}^\times(\cdot) = 0$ otherwise. \square

The idea behind the definition of \mathbf{R}^\times is to copy the transitions from \mathcal{M} , provided that the corresponding transition is possible in the current set of states of \mathcal{A} .

Our goal is to show that the values $\text{Prob}^{\mathcal{M}}(s, \alpha^I)$ can be calculated using a CSL model checking procedure for $\mathcal{M} \times \mathcal{A}$ and the CSL-path formula $\diamond^I \text{accept}$ (which states that a state labelled with the atomic proposition accept will be reached at some point in the time interval I). To establish this result, we first observe that \mathcal{M} and $\mathcal{M} \times \mathcal{A}$ are state-wise bisimulation equivalent when the set of atomic propositions in $\mathcal{M} \times \mathcal{A}$ is restricted to AP, i.e., we deal with the labeling function L_{AP}^\times which is given by $L_{\text{AP}}^\times(\langle s, Z \rangle) = L(s)$ rather than L^\times . This follows by the fact that the coarsest equivalence \mathcal{R} on $S \uplus (S \times 2^Z)$ which identifies any state s with any of its copies $\langle s, Z' \rangle$ where $Z' \subseteq Z$ is a bisimulation. Hence, $s \sim \langle s, Z' \rangle$ for all states s in \mathcal{M} and all subsets Z' of Z . Using Prop. 1, we obtain:

Proposition 4 For any state s of \mathcal{M} we have:

$$\text{Prob}^{\mathcal{M}}(s, \alpha^I) = \text{Prob}^{\mathcal{M} \times \mathcal{A}}(\langle s, Z_0 \rangle, \alpha^I)$$

Next we observe the one-to-one-correspondence between paths in \mathcal{M} and paths in $\mathcal{M} \times \mathcal{A}$ (when we fix the states $\langle s, Z_0 \rangle$ as starting states). Clearly, by removing the automaton-component of any state in a path in $\mathcal{M} \times \mathcal{A}$ one obtains a path in \mathcal{M} . Vice versa, each finite path

$$\sigma = s_0 \xrightarrow{a_0, t_0} s_1 \xrightarrow{a_1, t_1} \dots \xrightarrow{a_{n-1}, t_{n-1}} s_n \text{ in } \mathcal{M}$$

can be lifted to a path σ^\times in $\mathcal{M} \times \mathcal{A}$ by extending the states by sets of automaton-states with the help of $\hat{\delta}^{\mathcal{M}}$:

$$\sigma^\times = \langle s_0, Z_0 \rangle \xrightarrow{a_0, t_0} \langle s_1, Z_1 \rangle \xrightarrow{a_1, t_1} \dots \xrightarrow{a_{n-1}, t_{n-1}} \langle s_n, Z_n \rangle$$

where $Z_k = \hat{\delta}^{\mathcal{M}}(Z_{k-1}, s_{k-1} \xrightarrow{a_{k-1}, \cdot} s_k)$, $k = 0, 1, \dots, n$. Hence, if $\mathcal{L}(\alpha) = \mathcal{L}(\mathcal{A})$ then (by Prop. 3):

$$\begin{aligned} \langle s_n, Z_n \rangle \models \text{accept} &\text{ iff } \text{accept} \in L^\times(\langle s_n, Z_n \rangle) \\ &\text{ iff } Z_n \cap F \neq \emptyset \\ &\text{ iff } \hat{\delta}^{\mathcal{M}}(Z_0, \sigma) \cap F \neq \emptyset \\ &\text{ iff } \sigma \in \text{Path}_{\text{fin}}^{\mathcal{M}}(\mathcal{A}) = \text{Path}_{\text{fin}}^{\mathcal{M}}(\alpha). \end{aligned}$$

Thus, for all infinite paths $\zeta^\times \in \text{Path}_{\omega}^{\mathcal{M} \times \mathcal{A}}(\langle s, Z_0 \rangle)$ we have: $\mathcal{M} \times \mathcal{A}, \zeta^\times \models \diamond^I \text{accept}$ iff $\mathcal{M} \times \mathcal{A}, \zeta^\times \models \alpha^I$. Hence, for all states s in \mathcal{M} we have

$$\text{Prob}^{\mathcal{M} \times \mathcal{A}}(\langle s, Z_0 \rangle, \alpha^I) = \text{Prob}^{\mathcal{M} \times \mathcal{A}}(\langle s, Z_0 \rangle, \diamond^I \text{accept}).$$

Using this observation and Prop. 4 we obtain the following theorem.

Theorem 1 If α is an asCSL-program, \mathcal{A} an NPA with $\mathcal{L}(\alpha) = \mathcal{L}(\mathcal{A})$ and s a state in \mathcal{M} then

$$\text{Prob}^{\mathcal{M}}(s, \alpha^I) = \text{Prob}^{\mathcal{M} \times \mathcal{A}}(\langle s, Z_0 \rangle, \diamond^I \text{accept}).$$

Theorem 1 shows that the problem of computing the satisfaction set $\text{Sat}^{\mathcal{M}}(\phi)$ for the asCSL-formula $\phi = \mathcal{P}_{\triangleright p}(\alpha^I)$ is reducible to the problem of calculating the satisfaction set $\text{Sat}^{\mathcal{M} \times \mathcal{A}}(\phi_{\text{CSL}})$ for the CSL-state formula $\phi_{\text{CSL}} = \mathcal{P}_{\triangleright p}(\diamond^I \text{accept})$. In summary, to calculate $\text{Sat}^{\mathcal{M}}(\phi)$ where ϕ is as above we

- apply standard techniques to generate a nondeterministic finite automaton \mathcal{A} for α (viewed as an ordinary regular expression over the alphabet Σ);
- calculate the product ASMC $\mathcal{M} \times \mathcal{A}$, where it suffices to calculate the reachable part of $\mathcal{M} \times \mathcal{A}$ with an on-the-fly construction that starts with the states $\langle s, Z_0 \rangle$, and to ignore the action-labels in the sense that the rates of “parallel” transitions are cumulated;
- apply a CSL-model checker to calculate the values $p_s = \text{Prob}^{\mathcal{M} \times \mathcal{A}}(\langle s, Z_0 \rangle, \diamond^I \text{accept})$ for all states s in \mathcal{M} , e.g. with the help of a transient analysis of the Markov chain which is obtained from $\mathcal{M} \times \mathcal{A}$ when all states labelled by accept and all states from which one cannot reach a state labelled by accept (especially those that have an empty automaton part) are made absorbing [16, 24];
- return the set $\{s \in \mathcal{S} : p_s \triangleright p\}$.

In [24] it was shown that the time complexity of the uniformisation-based model checking algorithm for CSL-formulas of type $\mathcal{P}_{\triangleright p}(\phi_1 \mathcal{W}^{[t, t']} \phi_2)$ is $\mathcal{O}(M \cdot q \cdot t')$, where M is the number of transitions of the model and q is the uniformisation rate (which is given by the largest exit-rate of a state of the model). In our approach, an asCSL-formula of type $\mathcal{P}_{\triangleright p}(\alpha^I)$ is checked by first constructing an NPA \mathcal{A}_α which has $|Z| = \mathcal{O}(|\alpha|)$ states, and then constructing the product Markov chain, which has at most $M \cdot 2^{|Z|}$ transitions. The uniformisation rate and the time bound t' are not affected by the product automaton construction. Therefore, the overall time complexity of our algorithm to calculate the satisfaction set for an asCSL-formula of type $\mathcal{P}_{\triangleright p}(\alpha^{[t, t']})$ is bounded by $\mathcal{O}(M \cdot 2^{|\alpha|} \cdot q \cdot t')$.

5. Handover in a cellular radio network

In this section we present a case study in order to illustrate the techniques we have developed. We consider a scalable cellular phone setting, where base stations and switching center operate at different load levels. We track a single distinguished mobile radio station (MS) moving from one cell to another, thereby possibly triggering a so-called handover procedure. We are especially interested in the behaviour of the system concerning this distinguished user. The model

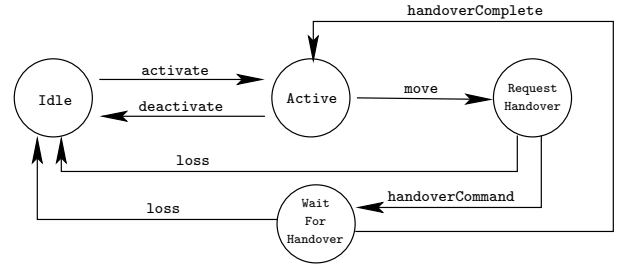


Figure 2. State machine for the MS behaviour

is inspired by the description of the GSM handover procedure in [25] and [26]. We describe the system as a set of synchronising processes, namely the cells, the switching center and the user. Each of these processes will be represented by a stochastic Petri net. The properties of interest are expressed with asCSL-formulas involving programs. We show the corresponding NPAs and relate the size of the resulting product Markov chains to the size of the original model.

5.1. The model

In the model, each cell corresponds to one base station subsystem (BSS). A BSS is modelled with only two states: either it has still free capacity or it is full and does not accept further connections. The number N of cells is a parameter of the model. The mobile services switching center (MSC) is modelled similarly: it has either low, medium or high load. The time needed for the handover command procedure depends on the current load. Under high load, the MSC does not process any request for handover.

Our distinguished MS has predetermined possibilities of moving between cells. For the time being, we assume that cells are positioned in a ring order, that is, if the MS is located in cell i it can only move to cells $(i+1) \bmod N$ and $(i-1) \bmod N$. We could easily change this to other topological cell orderings.

Finally, we have a model of the MS behaviour (beside its spatial position). When not being active with a connection, the MS is idle. At any time, the MS can become active, meaning that it has established a (radio) connection. After a while, the connection is terminated and the MS becomes idle again. If it moves from one cell to another while being active, the corresponding BSS commands a handover to the new cell from the MSC. If the new cell has free capacity, the handover is eventually completed and the MS returns to state active (note that the connection is continued during the entire handover procedure). If the handover procedure is not completed in time, the connection might also be lost. The connection is then terminated (assume, that the distance to the former cell has become too large) and the MS is back in idle state.

process	action	rate	description
BSS	block	0.002	cell will not accept further connections
	free	0.008	cell will accept further connections
MSC	lowtoMedium	0.5	from low load to medium load
	mediumToHigh	1.0	from medium load to high (blocking) load
	highToMedium	3.0	from high (blocking) load to medium load
	mediumToLow	1.0	from medium to low load
MS position	move	0.02	from cell i to $(i-1) \bmod N$ or $(i+1) \bmod N$
MS behaviour	activate	0.000625	average time between connections is 1600 seconds
	deactivate	0.008	connections last on average 125 seconds
	handoverCommand	1.0/0.5	for low/medium load of MSC, not available if MSC is blocking
	handoverComplete	1.0	only activated if new cell is not blocking
	loss	0.1	might happen during handover procedure

Table 1. Action labels and rates of the transitions of the cellular network model

Figure 2 shows a state machine for the user call behaviour. Transitions are labelled with action names. The move transition synchronises with the spatial movement of the user whenever he is active. Table 1 states the rates for transitions labelled with the given actions. Note that all numbers are educated guesses made on the basis of [27].

5.2. asCSL-properties

Out-dated handover. When the MS moves from one cell to the next, the BSS requests a handover to the new cell. However, the model does not prevent the MS from moving on to yet another cell. This behaviour is not explicitly visible in the model: here a handover is simply made to the cell the MS is in, no matter where it has been in between. In reality this type of movement could cause a problem. So, we would like to know whether the probability of such an outdated handover is lower than 2%. As asCSL-formula, this becomes: $\Phi_1 = \mathcal{P}_{\leq 0.02}(\alpha_1^{[0, \infty]})$, with

$$\alpha_1 = (\text{Active}, \text{move}); \quad (1)$$

$$(\text{RequestHandover} \vee \text{WaitForHandover}, \quad (2)$$

$$\text{Act}\{\text{handoverComplete}, \text{move}\}^*; \quad (3)$$

$$(\text{RequestHandover} \vee \text{WaitForHandover}, \text{move})^* \quad (4)$$

A move while the MS is active triggers a handover. Lines (2/3) describe the system inside the handover procedure. A move (4) leads to an outdated handover. An NPA for the program α_1 is given in Fig. 3(a).

Return without interruption. Assume that the MS initiates a connection while in cell 1. It is free to move between cells. We would like it to leave cell 1 and to return within 10 minutes (600 seconds) without terminating or losing the connection. Is the probability for this scenario at least 50%? Coded into an asCSL-formula this

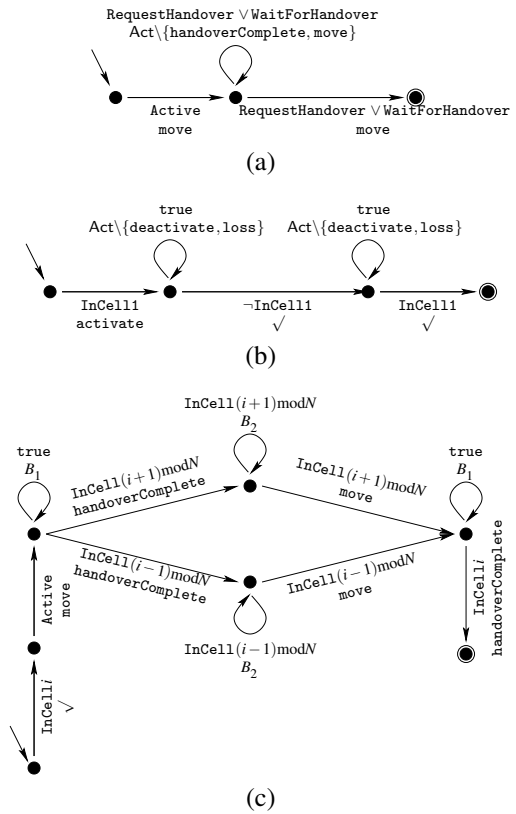


Figure 3. NPAs for the programs α_1 , α_2 and β_i

reads $\Phi_2 = \mathcal{P}_{> 0.5}(\alpha_2^{[0, 600]})$, with

$$\alpha_2 = (\text{InCell1}, \text{activate}); \quad (5)$$

$$(\text{true}, \text{Act}\{\text{deactivate}, \text{loss}\})^*; \quad (6)$$

$$(\neg \text{InCell1}, \checkmark); \quad (7)$$

$$(\text{true}, \text{Act}\{\text{deactivate}, \text{loss}\})^*; \quad (8)$$

$$(\text{InCell1}, \checkmark) \quad (9)$$

The regular expression first ensures that the user activates a connection while being in cell 1 (5). Then the user can behave arbitrarily, as long as the connection is not ended via a deactivate or loss event (6). At some time, the user must have left cell 1 (7) and can again behave arbitrarily, as long the connection remains established (8). Finally, he should return to cell 1 (9). Figure 3(b) shows an NPA for the program α_2 .

Ping-pong. Sometimes there are handovers from a cell i to a neighbouring cell and back to cell i within a short time interval. From a performance point of view this is not desirable since presumably the call could have remained in cell i . If having an active connection, is the probability of such a ping-pong handover to occur within 10 seconds at most 10%? As asCSL-formula: $\Phi_3 = \mathcal{P}_{\leq 0.1}(\alpha_3^{[0, 10]})$. A

ping-pong originating from cell i is described by the program

$$\beta_i = (\text{InCell}i, \sqrt{ }); (\text{Active}, \text{move}); (\text{true}, B_1)^*; \quad (10)$$

$$((\text{InCell}(i+1) \bmod N, \text{handoverComplete})) \quad (11)$$

$$(\text{InCell}(i+1) \bmod N, B_2)^*; \quad (12)$$

$$(\text{InCell}(i+1) \bmod N, \text{move}) \cup \quad (13)$$

$$(\text{InCell}(i-1) \bmod N, \text{handoverComplete}); \quad (14)$$

$$(\text{InCell}(i-1) \bmod N, B_2)^*; \quad (15)$$

$$(\text{InCell}(i-1) \bmod N, \text{move}); \quad (16)$$

$$(\text{true}, B_1)^*; (\text{InCell}i, \text{handoverComplete}) \quad (17)$$

with $B_1 = \text{Act} \setminus \{\text{move}, \text{loss}, \text{handoverComplete}\}$ and $B_2 = \text{Act} \setminus \{\text{deactivate}, \text{move}\}$. An NPA for this program is given in Fig. 3(c). If there are N cells, all possible ping-pong situations are described by $\alpha_3 = \beta_0 \cup \beta_1 \cup \dots \cup \beta_{N-1}$. The NPA for α_3 consists of N replicas of the automaton in Fig. 3(c), instantiated with $i = 0, \dots, N-1$. It has N initial and N final states.

5.3. Tool support


A prototype implementation that performs the construction of the product Markov chain given an ASMC and an NPA was done in C++.

For the evaluation we used a stochastic Petri net (SPN) model of the cellular phone system. All components of the systems are described by simple state machines, we therefore omit their SPN representation. The SPN is described in a variant of CSPL, which also allows the specification of marking-dependent properties, which can be seen as atomic propositions in the underlying Markov chain. The state space generation code of [28] has been extended in order to record these properties and the transition names (as action labels) and generates an ASMC.


Programs have been described directly via their corresponding NPA. The prototype implementation takes the ASMC and the NPA as input and computes the reduced product Markov chain where only reachable states are generated and accept-states and states $\langle s, \emptyset \rangle$ with empty automaton part are merged into two special absorbing states. The final computation of the satisfaction relation of the corresponding CSL formula can then be done using an existing CSL model checker, for example ETMCC [7].

5.4. Results

Figure 4 shows the number of states and transitions of the original ASMC model of a cellular radio network and of the product Markov chains needed for the model checking procedure of the three given asCSL-formulas as a function of the number N of cells that ranges from 2 to 11.



(a) $|S^x|$



(b) $|R^x|$

Figure 4. Size of state space (a) and number of transitions (b) in the original ASMC and the product Markov chains

The original model has about 270000 states and 3.6 million transitions for $N = 11$. For all three presented programs, the number of states in the product Markov chain is larger than in the original ASMC. This could be expected, since the state space is a subset of the Cartesian product of S and $2^{|Z|}$. The increase in size varies between a factor 1.5 and 4. The largest state space is the one of the ping-pong property, it has more than one million states for $N = 11$. If we keep only those states from which the accept-state is reachable and merge the others into one absorbing state, the number of states can also become smaller than the original state space.

For the properties “outdated handover” and “return without interruption”, the number of transitions in the product Markov chain is smaller than in the original ASMC. The corresponding program automata are very restrictive, in the sense that in each state of the original ASMC only a subset of all outgoing transitions is allowed by the NPA. The NPA for “ping-pong” allows a broader range of different combinations of states and transitions. Consequently, it shows the largest growth in state space, and the number of transitions

is larger than in the original model.

6. Conclusions

We introduced the logic asCSL as a new specification formalism to reason about performability measures for Markov chains with both action- and state-labels. It subsumes CSL (with time intervals $[0, t]$) and several variants that have been suggested in the literature, such as aCSL, aCSL+, pathCSL, SPDL [6], [8], [9],[10]. Our model checking algorithm for formulas of type $\mathcal{P}_{\triangleright p}(\alpha^I)$ relies on a reduction to the CSL model checking problem via a product construction of the Markov chain \mathcal{M} and an automaton for the path formula α^I . The case study in section 5 demonstrates how asCSL formulas can specify complex properties that refer to both action and state labels in a rather simple way.

References

- [1] E. Clarke, O. Grumberg, and D. Peled, *Model Checking*. MIT Press, 1999.
- [2] A. Aziz, K. Sanwal, V. Singhal, and R. Brayton, “Verifying continuous time Markov chains,” in *CAV’96, LNCS 1102*, 1996, pp. 269–276.
- [3] C. Baier, J.-P. Katoen, and H. Hermanns, “Approximate symbolic model checking of continuous-time Markov chains,” in *CONCUR’99, LNCS 1664*, 1999, pp. 146–161.
- [4] H. Hansson and B. Jonsson, “A logic for reasoning about time and reliability,” *Formal Aspects of Computing*, vol. 6, no. 5, pp. 512–535, 1994.
- [5] C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen, “On the logical characterisation of performability properties,” in *ICALP’00, LNCS 1853*, 2000, pp. 780–792.
- [6] H. Hermanns, J. Katoen, J. Meyer-Kayser, and M. Siegle, “Towards model checking stochastic process algebra,” in *IFM’00, LNCS 1945*, 2000, pp. 420–439.
- [7] H. Hermanns, J.-P. Katoen, J. Meyer-Kayser, and M. Siegle, “Implementing a model checker for performability behaviour,” in *Proc. 5th Int’l. Workshop on Performability Modelling of Computer and Communication Systems*, Erlangen, Germany, 2001, pp. 110–115.
- [8] J. Meyer-Kayser, “Automatische Verifikation Stochastischer Systeme,” Ph.D. dissertation, Universität Erlangen-Nürnberg, Institut für Informatik, 2004.
- [9] C. Baier, L. Cloth, B. Haverkort, H. Hermanns, and J. Katoen, “Model checking pathCSL,” in *Proc. 6th Int’l. Workshop on Performability Modeling of Computer and Communication Systems*, Monticello, Illinois, 2003, pp. 19–22.
- [10] M. Kuntz and M. Siegle, “A stochastic extension of the logic PDL,” in *Proc. 6th Int’l. Workshop on Performability Modeling of Computer and Communication Systems*, Monticello, Illinois, 2003, pp. 58–61.
- [11] W. Obal and W. Sanders, “State-space support for path-based reward variables,” *Perform. Eval.*, vol. 35, no. 3-4, pp. 233–251, 1999.
- [12] M. Fischer and R. Ladner, “Propositional dynamic logic of regular programs,” *J. Comput. Syst. Sci.*, vol. 8, pp. 194–211, 1979.
- [13] P. Wolper, “Specification and synthesis of communicating processes using an extended temporal logic,” in *Proc. 9th Symposium on Principles of Programming Languages*, 1982, pp. 20–33.
- [14] A. Aziz, V. Singhal, F. Balarin, and R. K. Brayton, “It usually works: The temporal logic of stochastic systems,” in *CAV’95, LNCS 939*, 1995, pp. 155–165.
- [15] W. Stewart, *Introduction to the Numerical Solution of Markov Chains*. Princeton University Press, 1994.
- [16] C. Baier, B. Haverkort, J.-P. Katoen, and H. Hermanns, “Model checking continuous-time Markov chains by transient analysis,” in *CAV’00, LNCS 1855*, 2000, pp. 358–372.
- [17] J. Kemeny and J. Snell, *Finite Markov Chains*. Springer, 1976.
- [18] P. Buchholz, “Exact and ordinary lumpability in finite Markov chains,” *J. Appl. Prob.*, no. 31, pp. 59–75, 1994.
- [19] J. Hillston, “A Compositional Approach to Performance Modelling,” Ph.D. dissertation, University of Edinburgh, 1994.
- [20] H. Hermanns and M. Rettelbach, “Syntax, semantics, equivalences, and axioms for MTIPP,” in *PAPM’94*. Universität Erlangen-Nürnberg, 1994, pp. 71–88.
- [21] J. Desharnais and P. Panangaden, “Continuous stochastic logic characterizes bisimulation of continuous-time Markov processes,” *J. Logic Algebr. Progr.*, vol. 56, no. 1-2, pp. 99–115, 2003.
- [22] M. Bernardo and R. Cleaveland, “A theory of testing for Markovian processes,” in *CONCUR 2000, LNCS 1877*, 2000, pp. 305–319.
- [23] E. Clarke, E. Emerson, and A. Sistla, “Automatic verification of finite-state concurrent systems using temporal logic specifications,” *ACM Trans. Program. Lang. Syst.*, vol. 8, no. 2, pp. 244–263, 1986.
- [24] C. Baier, B. Haverkort, H. Hermanns, and J. Katoen, “Model-Checking Algorithms for Continuous-Time Markov Chains,” *IEEE Trans. Softw. Eng.*, vol. 29, no. 7, pp. 1–18, 2003.
- [25] B. H. Walke, *Mobile Radio Networks*. Wiley, 1999.
- [26] J. M. Thomsen and R. Manggaard, “Analysis of GSM handover using coloured Petri nets,” Master’s thesis, University of Aarhus, 2003.
- [27] J. Ventura Agustina, P. Zhang, and R. Kantola, “Performance evaluation of GSM handover traffic in a GPRS/GSM network,” in *Proc. 8th IEEE Int’l. Symp. on Computers and Communications*. IEEE Press, 2003, pp. 137–142.
- [28] B. Haverkort, H. Bohnenkamp, and A. Bell, “On the efficient sequential and distributed evaluation of very large stochastic Petri nets,” in *PNPM’99*. IEEE Press, 1999, pp. 12–21.