Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Information and Computation 204 (2006) 713–740

Information and Computation

www.elsevier.com/locate/ic

Quantitative analysis of probabilistic lossy channel systems☆

Alexander Rabinovich

School of Computer Science, Sackler Faculty of Exact Sciences, Tel Aviv University, Israel

Received 4 November 2003; revised 20 June 2005
Available online 4 April 2006

Abstract

Many protocols are designed to operate correctly even in the case where the underlying communication medium is faulty. To capture the behaviour of such protocols, lossy channel systems (LCS) have been proposed. In an LCS the communication channels are modelled as FIFO buffers which are unbounded, but also unreliable in the sense that they can nondeterministically lose messages. Recently, several attempts have been made to study probabilistic lossy channel systems (PLCS) in which the probability of losing messages is taken into account and the following qualitative model checking problem is investigated: to verify whether a given property holds with probability one. Here we consider a more challenging problem, namely to calculate the probability by which a certain property is satisfied. Our main result is an algorithm for the following
Quantitative model checking problem:
Instance: A PLCS, its state s, a finite state ω-automaton $\mathcal{A}$, and a rational θ>0.
*Task:*Find a rational r such that the probability of the set of computations that start at s and are accepted by $\mathcal{A}$ is between r and r+θ.

1. Introduction

Finite state machines which communicate through unbounded buffers (CFSM) have been popular in the modelling of communication protocols.A CFSM defines in a natural way an infinite

☆ This is the extended and full version of the paper presented in ICALP 2003 [18].
E-mail address: rabinoa@post.tau.ac.il.

0890-5401/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.ic.2005.11.004

state transition system. The fact that Turing machines can be simulated by CFSMs [8] implies that all the nontrivial verification problems are undecidable for CFSMs. Many protocols are designed to operate correctly even in the case where the underlying communication medium is faulty. To capture the behaviour of such protocols, lossy channel systems (LCS) [4] have been proposed as an alternative model. In an LCS the communication channels are modelled as FIFO buffers which are unbounded, but also unreliable in the sense that they can nondeterministically lose messages. Though an LCS defines in a natural way an infinite state transition system, it has been shown that the reachability problem for LCS is decidable [4], while progress properties are undecidable [3].

1.1. Probabilistic lossy channel systems

Since we are dealing with unreliable communication media, it is natural to deal with models in which the probability of losing messages is taken into account. Recently, several attempts [14,6,1,5,7] have been made to study probabilistic lossy channel systems (PLCS) which introduce randomization into the behaviour of LCS. The works in [14,6,1,5,7] define different semantics for PLCS, depending on the manner in which the messages may be lost inside the channels. All these models associate in a natural way a countable Markov Chain (M.C.) to a PLCS.

Baier and Engelen [6] consider a model which assumes that at most a single message may be lost during each step of the execution of the system. They showed decidability of the following problems under the assumption that the probability of losing messages is at least 0.5.

Qualitative Probabilistic Reachability
Instance: A PLCS M and its states s_1, s_2.
Question: Is s_2 reached from s_1 with probability one?

Qualitative Probabilistic Model-checking
Instance: A PLCS M, its state s and a finite state ω-automaton $\mathcal{A}$.
Question: Is the probability of the set of computations that start at s and are accepted by $\mathcal{A}$ equal to one?

The model in [1] assumes that messages can only be lost during send operations. Once a message is successfully sent to a channel, it continues to reside inside the channel until it is removed by a receive operation. Even the qualitative reachability problem was shown to be undecidable for this model of PLCS and losing probability $\lambda < 0.5$.

In [5,7] another semantics for PLCS was considered which is more realistic than that in [6,1]. More precisely, it was assumed that, during each step in the execution of the system, each message may be lost with a certain predefined probability. This means that the probability of losing a certain message will not decrease with the length of the channels (as it is the case with [6,1]).

For this model, the decidability of both the qualitative reachability and the qualitative model-checking problems was independently established in [5,7].

1.2. Our contribution

All the above mentioned papers consider qualitative properties of PLCS. Here we consider a more challenging problem, namely to calculate the probability by which a certain property is satisfied.

Unfortunately, we were unable to prove that the probability of reaching a state s_2 from a state s_1 in PLCS is an algebraic number, or it is explicitly expressible by standard mathematical functions.

Therefore, we will approximate the probability by which a certain property is satisfied. Our main result is that the following two problems are computable.

Quantitative Probabilistic Reachability
Instance: A PLCS $\mathcal{L}$, its states s_1, s_2 and a rational $\theta > 0$.
Task: Find a rational r such that s_2 is reached from s_1 with the probability between r and $r + \theta$.

Quantitative Probabilistic Model-checking
Instance: A PLCS $\mathcal{L}$, its state s, a finite state ω-automaton $\mathcal{A}$, and a rational $\theta > 0$.
Task: Find a rational r such that the probability of the set of computations that start at s and are accepted by $\mathcal{A}$ is between r and $r + \theta$.

To approximate the probability p of the set of computations from a state s with a property φ in a PLCS $\mathcal{L}$ one can try to compute this probability p_n for the finite sub-chain $M_n = (S_n, P_n)$ of the countable Markov chain M generated by $\mathcal{L}$, where S_n is the set of states with at most n messages. There are two problems in this approach: (a) a state which was recurrent in M might become transient in M_n; (b) how to find n which will ensure that the result p_n approximates up to θ the probability p in M.

To overcome problem (a) we analyze the structure of the recurrent classes of the Markov chain generated by a PLCS $\mathcal{L}$. For problem (b) the value for n is computed from an appropriate reduction of the Markov chains generated by PLCSs to one dimensional random walks.

1.3. Outline

In the next two sections we give basics of transition systems and countable Markov chains respectively. In addition to standard definitions about countable Markov chain, we recall the concept of *attractor sets* [5]: an attractor for a Markov chain is a set of states such that regardless of the state in which we start, we are guaranteed to reach the attractor with probability one. Markov chains with finite attractors and finite state Markov chains have many common properties. For example, the set of recurrent states is reached with the probability one; a state is recurrent iff it belongs to a bottom strongly connected component of the underlying transition system; many qualitative properties of Markov chains with finite attractors can be established on the basis of the topology of the underlying transition systems.

In Section 4, the quantitative probabilistic reachability problem over countable Markov chains is considered. We provide a naive Path Enumeration (PE) scheme for the quantitative reachability problem and prove general theorems about (countable state) Markov chains which serve as sufficient conditions for termination of this scheme. In particular the path enumeration scheme terminates over Markov chains with finite attractors. The Markov chain assigned to probabilistic lossy channel systems have finite attractors, and from the path enumeration scheme an algorithm for the quantitative reachability problem can be extracted. The path enumeration scheme is conceptually very simple, however, no information about the number of iterations before it terminates can be extracted from the theorems of Section 4. For finite state Markov chains standard algebraic methods allow to find the exact value for the probability of reaching one state from another in polynomial time; however, in this case the PE scheme finds an approximation in time exponential in the quality of the approximation. In Section 5, we consider a generalization of one dimensional random walk and from the theorems proved there we will obtain in Section 8 an algorithm for the quantitative probabilistic reachability problem over PLCS with the parametric complexity which is polynomial in the quality of the approximation.

In Sections 6 and 7, the semantics of lossy channel systems and probabilistic lossy channel systems are described. In Section 8, the algorithm for the quantitative probabilistic reachability problem over PLCS is presented and its complexity is analyzed. In Section 9, we generalize our results to the verification of the properties definable by ω-behavior of finite state automata (or equivalently formulas in the Monadic Logic of Order). Finally, we give conclusions and directions for future work in Section 10.

2. Transition systems

In this section, we recall some basic concepts of transition systems.

A *transition system* T is a pair $(S, \longrightarrow)$ where S is a (potentially infinite) set of *states*, and $\longrightarrow$ is a binary relation on S. We write $s_1 \longrightarrow s_2$ to denote that $(s_1, s_2) \in \longrightarrow$ and use $\xrightarrow{*}$ to denote the reflexive transitive closure of $\longrightarrow$. We say that s_2 is *reachable* from s_1 if $s_1 \xrightarrow{*} s_2$. For sets $Q_1, Q_2 \subseteq S$, we say that Q_2 is *reachable* from Q_1, denoted $Q_1 \xrightarrow{*} Q_2$, if there are $s_1 \in Q_1$ and $s_2 \in Q_2$ with $s_1 \xrightarrow{*} s_2$. A *path* or *computation* π (from s_0) is a (finite or infinite) sequence $s_0, s_1, \ldots, s_n, \ldots$ of states with $s_i \longrightarrow s_{i+1}$ for $i \geqslant 0$. We use $\pi(i)$ to denote s_i, and write $s \in \pi$ to denote that there is an $i \geqslant 0$ with $\pi(i) = s$. For states s and s', we say that π *leads* from s to s', written, $s \xrightarrow{\pi} s'$, if $s = s_0$ and $s' \in \pi$. We use $|\pi|$ to denote the length of π (i.e., the number of transitions in π), and $last(\pi)$ denotes the last state of the finite path π. For a finite path π_1 and a path π_2 that starts at $last(\pi_1)$, we use $\pi_1\pi_2$ to denote the concatenation of π_1 and π_2; this is the path π such that $\pi(i) = \pi_1(i)$ for $i \leqslant |\pi_1|$ and $\pi(i) = \pi_2(i + 1 - |\pi_1|)$ for $i > |\pi_1|$.

For $Q \subseteq S$, we define the *graph* of Q, denoted $Graph(Q)$, to be the transition system $(Q, \longrightarrow')$ where $s_1 \longrightarrow' s_2$ iff $s_1 \xrightarrow{*} s_2$. A *strongly connected component (SCC)* in T is a maximal set $C \subseteq S$ such that $s_1 \xrightarrow{*} s_2$ for each $s_1, s_2 \in C$. We say that C is a *bottom SCC (BSCC)* if there is no other SCC C_1 in T with $C \xrightarrow{*} C_1$. In other words, the BSCCs are the leaves in the acyclic graph of SCCs (ordered by reachability).

3. Markov chains

In this section, we recall basic properties of *Markov chains*. We also introduce attractors which play an important role in our analysis of recurrent classes of Markov chains.

A *Markov chain* M is a pair (S,P) where S is a (potentially countable) set of states and P is a mapping from $S \times S$ to the set $[0,1]$, such that $\sum_{s' \in S} P(s,s') = 1$, for each $s \in S$.

A Markov chain induces a transition system where the transition relation consists of pairs of states related by positive probabilities.

Formally, the *underlying transition system* of M is $\big(S, \longrightarrow\big)$ where $s_1 \longrightarrow s_2$ iff $P(s_1,s_2) > 0$.

In this manner, the concepts defined for transition systems can be lifted to Markov chains.

For instance, an SCC in M is a SCC in the underlying transition system.

A Markov chain (S,P) induces a natural measure on the set of computations from every state s.

Let us recall some basic notions from probability theory. A *measurable space* is a pair (Ω, Δ) consisting of a non empty set Ω and a σ-algebra Δ of its subsets that are called *measurable sets* and represent random events. A *σ-algebra* over Ω contains Ω and is closed under complementation and countable union. Adding to a measurable space a *probability measure Prob* : $\Delta \to [0,1]$ such that $Prob(\Omega) = 1$ and that is countably additive, we get a *probability space* $(\Omega, \Delta, Prob)$.

Consider a state s of a Markov chain (S,P). On the sets of computations that start at s, the probabilistic space $(\Omega, \Delta, Prob)$ is defined as follows (see [16]): $\Omega = sS^\omega$ is the set of all ω-sequences of states starting from s, Δ is the σ-algebra generated by the basic cylindric sets $D_u = uS^\omega$, for every $u \in sS^*$, and the probability measure $Prob$ is defined by $Prob(D_u) = \prod_{i=0,\dots,n-1} P(s_i, s_{i+1})$ where $u = s_0 s_1 \dots s_n$; it is well known that this measure is extended in a unique way to the elements of the σ-algebra generated by the basic cylindric sets.

For a set $G \subseteq sS^*$ of finite paths we use $Prob(G)$ to stand for $Prob(\{\bigcup_{u \in G} D_u\})$.

Consider a set $Q \subseteq S$ of states and a path π. We say that π *reaches* Q if there is an $i \geqslant 0$ with $\pi(i) \in Q$. We say that π *repeatedly reaches* Q if there are infinitely many i with $\pi(i) \in Q$. Let s be a state in S. We say that a state s is *recurrent* if $Prob\,\{\pi : \pi$ is a path from s and π repeatedly reaches $s\} = 1$. We say that a state s is *transient* if $Prob\,\{\pi : \pi$ is a path from s and π repeatedly reaches $s\} = 0$.

The next theorem summarizes standard properties of countable Markov chains [17].

Theorem 1.

(1) Every state is either transient or recurrent.

(2) If s is recurrent then all the states reachable from s are recurrent.

(3) Let C be a strongly connected component of a Markov chain. Then, either all the states in C are transient or all the states in C are recurrent.

(4) Let C be a recurrent strongly connected component of a Markov chain and $s_1 \in C$. Then $Prob\{\pi$: π starts at s_1 and repeatedly reaches every state of $C\} = 1$. For every state s and non-empty subset $B \subseteq C$ the probability to repeatedly reach every state of B from s is the same as the probability to reach B from s and is the same as the probability to reach s_1 from s.

(5) A recurrent strongly connected component is always a bottom strongly connected component.

A recurrent (transient) SCC is often called a recurrent (transient) class or component.

The next lemma is also a standard one.

Lemma 2. *Let M be a Markov chain and let s_1, s, s_2 be three of its states. Assume that s_2 is reachable from s. Then*

$$Prob(\{\pi : \pi \textit{ starts at } s_1 \textit{ and repeatedly reaches } s \textit{ and never visits } s_2\}) = 0.$$

We introduce an important concept which we use in our solution for the probabilistic reachability problem, namely that of *attractors.*

Definition 3 (*attractor*)**.** A set $A \subseteq S$ is said to be an *attractor* if for each $s \in S$, the set A is reachable from s with probability one.

In other words, regardless of the state in which we start, we are guaranteed to reach the attractor with probability one. It is clear that an attractor has a state in every recurrent class.

The following lemma is an immediate consequence of Definition 3 and holds for Markov chains with countable attractors.

Lemma 4. *If $A \subseteq S$ is an attractor then for each $s \in S$, the set A is repeatedly reachable from s with probability one.*

Markov chains with finite attractors and finite state Markov chains have many common properties. The Lemma below follows from Theorem 1 and describes properties of Markov chains with finite attractor.

Lemma 5. *Assume that a Markov chain M has a finite attractor A. Then*

(1) Each BSCC C of Graph(A) is a subset of a recurrent component in M.
(2) A state is recurrent if and only if it is reachable from a BSCC C of Graph(A).
(3) For every s in M the set of recurrent states is reached from s with probability one.
(4) The recurrent components of M are the BSCCs of M.

4. Approximating probability for countable Markov chains

Let M be a M.C. and let s_1, s_2 be its states. We use $Prob_M(s_1 \xrightarrow{*} s_2)$ for the probability with which s_2 is reached from s_1 in M. Let $\mathrm{Comp}_n(s_1)$ be the set of all the computations of length n in M from s_1. Partition $\mathrm{Comp}_n(s_1)$ into three sets:

$$\mathrm{Reach}_n(s_1, s_2) = \{\pi : \pi \in \mathrm{Comp}_n(s_1) \wedge \exists i \leqslant n.\pi(i) = s_2\}$$

$$\mathrm{Escape}_n(s_1, s_2) = \left\{\pi : \pi \in \mathrm{Comp}_n(s_1) \setminus \mathrm{Reach}_n(s_1, s_2) \wedge s_2 \textbf{ is unreachable from } \pi(n)\right\}$$

$$\mathrm{Undecided}_n(s_1, s_2) = \mathrm{Comp}_n(s_1) \setminus \mathrm{Reach}_n(s_1, s_2) \setminus \mathrm{Escape}_n(s_1, s_2).$$

All the computations in $\mathrm{Reach}_n(s_1, s_2)$ reach s_2, and no computation in $\mathrm{Escape}_n(s_1, s_2)$ extends to a computation that reaches s_2. Note that $Prob(\mathrm{Comp}_n(s_1)) = 1$. Let $p_n^+ = Prob(\mathrm{Reach}_n(s_1, s_2))$, $p_n^- =$

$Prob(\text{Escape}_n(s_1,s_2))$ and $p_n^? = Prob(\text{Undecided}_n(s_1,s_2))$. Observe that p_n^+ and p_n^- are increasing sequences, while $p_n^?$ is decreasing and

$$p_n^+ \leqslant \lim p_n^+ = Prob_M(s_1 \xrightarrow{*} s_2) \leqslant p_n^+ + p_n^?. \tag{1}$$

The Path Enumeration (PE) scheme for approximating $Prob_M(s_1 \xrightarrow{*} s_2)$ is based on (1).

Path Enumeration Scheme for Approximating Probabilistic Reachability

Instance: A M.C. M, its states s_1, s_2 and a $\theta > 0$.
Task: Find r such that s_2 is reached from s_1 with a probability between r and $r + \theta$.

```
begin
    1. $n := 0$; $\Delta := 1$;
    2. while ($\Delta > \theta$) do
            3. $n := n + 1$; Compute $r := p_n^+$; Compute $\Delta := p_n^?$
     end while
    4. return($r$)
end
```

From (1) it follows that the PE scheme is partially correct with respect to its specification, i.e., upon termination of the PE scheme it outputs r such that s_2 is reached from s_1 with the probability between r and $r + \theta$.

In the above problem, we do not assume that M is finite. Hence, these are not instances of an algorithmic problem. In Section 8 we consider the quantitative reachability problem when countable Markov chains are described by probabilistic lossy channel systems. For such finite descriptions we investigate the corresponding algorithmic problem.

If M has finite branching and is presented effectively, then p_n^+ is computable. Moreover, if in addition the reachability problem for the transition system underlying M is decidable, then $\text{Escape}_n(s_1,s_2)$, $\text{Undecided}_n(s_1,s_2)$ and $p_n^?$ can be computed. Hence, in this case the scheme can be implemented. Observe that the PE scheme terminates if and only if $\lim p_n^? < \theta$. Therefore,

Lemma 6. *If* $\lim p_n^? = 0$ *then the PE scheme terminates.*

It is well-known that for a finite state Markov chains $\lim p_n^? = 0$. This property holds for Markov chains with finite attractors [20] as well.

Lemma 7. *If M has a finite attractor then* $\lim p_n^? = 0$.

Proof. Let A be a finite attractor. Let A_+ (respectively, A_-) be the set of states in A from which s_2 is reachable (respectively, unreachable).

Let $B = \{\pi : \forall n.\pi(1)\cdots\pi(n) \in \text{Undecided}_n(s_1,s_2)\}$. Note that

$$Prob(B) = \lim Prob(\text{Undecided}_n(s_1,s_2)) = \lim p_n^?. \tag{2}$$

The state s_2 is reachable from each state s on $\pi \in B$. Hence no computation in B passes through a state in A_-. Therefore, by Lemma 4

$$\begin{aligned} Prob(B) &= Prob(\{\pi : \pi \in B \text{ and } \pi \text{ repeatedly reaches } A\}), \\ &= Prob(\{\pi : \pi \in B \text{ and } \pi \text{ repeatedly reaches } A_+\}). \end{aligned} \tag{3}$$

Since A is finite and $A_+ \subseteq A$ there is k such that $A_+ = \{v_1, \dots v_k\}$. Hence from (3) it follows that

$$Prob(B) \leqslant \sum_{i=1}^{k} Prob(\{\pi : \pi \in B \text{ and } \pi \text{ repeatedly reaches } v_i\}). \tag{4}$$

Note that s_2 is reachable from v_i, because $v_i \in A_+$. No computation $\pi \in B$ passes through s_2, because $\forall n.\pi(1)\cdots\pi(n) \in \text{Undecided}_n(s_1, s_2)$. Therefore, by Lemma 2

$$Prob(\{\pi : \pi \in B \text{ and } \pi \text{ repeatedly reaches } v_i\}) = 0. \tag{5}$$

From (4) and (5) we derive that $Prob(B) = 0$. Therefore, by (2) we obtain that $\lim p_n^? = 0$ and this completes the proof of the lemma. □

Another class of Markov chains for which the PE scheme terminates is the class of chains which satisfy the following property.

Definition 8. A Markov chain $M = (S, P)$ has δ-reachability property for $\delta > 0$ if

$$\forall s_1, s_2 \in S(\, s_2 \text{ is reachable from } s_1\,) \Rightarrow Prob_M(s_1 \xrightarrow{*} s_2) > \delta.$$

Lemma 9. *If M has δ-reachability property then* $\lim p_n^? = 0$.

Proof. First, we are going to show that $\forall n \exists m. p_m^? \leqslant (1 - \frac{\delta}{4}) p_n^?$.

The number of states in M is at most countable, hence $\text{Undecided}_n(s_1, s_2)$ contains at most countable number of paths and therefore there is a finite set of path $F_n \subseteq \text{Undecided}_n(s_1, s_2)$ such that

$$Prob(F_n) \geqslant \frac{Prob(\text{Undecided}_n(s_1, s_2))}{2} = \frac{p_n^?}{2}.$$

If $s = last(\pi)$ for $\pi \in F_n$ then s_2 is reachable from s. Therefore by δ-reachability property $Prob_M(s \xrightarrow{*} s_2) > \delta$. Therefore there is a finite set of paths E_s from s to s_2 such that $Prob(E_s) > \delta/2$. Let N be the set

$$N = \{\pi_1\pi_2 : \pi_1 \in F_n \wedge \pi_2 \in E_{last(\pi_1)}\},$$

where $\pi_1\pi_2$ denotes the concatenation of finite paths π_1 and π_2. Observe that

$$Prob(N) \geqslant \frac{p_n^? \delta}{4}.$$

The set N is finite. Let m be the length of the longest path in N. None of the paths in N extends a path in $\mathrm{Reach}_n(s_1,s_2)$; each path in N reaches s_2 and each path of length m that extends a path from N is in $\mathrm{Reach}_m(s_1,s_2)$. Therefore,

$$Prob(\mathrm{Reach}_m(s_1,s_2)) - Prob(\mathrm{Reach}_n(s_1,s_2)) \geqslant Prob(N) \geqslant \frac{p_n^? \delta}{4}. \tag{6}$$

Now we derive that $p_m^? \leqslant p_n^?(1-\frac{\delta}{4})$ as follows:

$$\begin{aligned}
p_m^? &= 1 - Prob(\mathrm{Reach}_m(s_1,s_2)) - Prob(\mathrm{Escape}_m(s_1,s_2))\\
&\leqslant 1 - Prob(\mathrm{Reach}_m(s_1,s_2)) - Prob(\mathrm{Escape}_n(s_1,s_2))\\
&= (1 - Prob(\mathrm{Reach}_n(s_1,s_2)) - Prob(\mathrm{Escape}_n(s_1,s_2)))\\
&\quad -(Prob(\mathrm{Reach}_m(s_1,s_2)) - Prob(\mathrm{Reach}_n(s_1,s_2)))\\
&= p_n^? - (Prob(\mathrm{Reach}_m(s_1,s_2)) - Prob(\mathrm{Reach}_n(s_1,s_2)))\\
&\leqslant p_n^?(1-\tfrac{\delta}{4}).
\end{aligned}$$

The first inequalities in the above sequence holds, because $\lambda n.Prob(\mathrm{Escape}_n(s_1,s_2))$ is increasing; the last inequality follows from Eq. (6).

We proved that

$$\forall n \exists m. p_m^? \leqslant (1-\frac{\delta}{4})p_n^? \tag{7}$$

Recall that $\{p_n^?\}_{n=0}^{\infty}$ is a decreasing sequence and $\forall n.p_n^? \geqslant 0$. These together with (7) imply that $\lim p_n^? = 0$. □

Lemma 6, Lemma 7 and Lemma 9 imply.

Theorem 10. *The PE scheme terminates over the class of Markov chains with finite attractor and over the class of Markov chains with δ-reachability property.*

A variant of the PE scheme was suggested in [14] for the following decision problem.

A decision problem for Probabilistic Reachability
Instance: A M.C. M, its states s_1, s_2, and rationals $\theta > 0$ and p.
Question: Is $p - \theta < Prob_M(s_1 \xrightarrow{*} s_2) < p + \theta$?

It was claimed in [14] that Eq. (1) implies that (a) when the scheme terminates it produces a correct answer (b) it terminates for the Markov chains defined by PLCS under the semantics of [14]. However, assertion (a) was incorrect. Also, the Markov chains assigned to PLCSs in [14] do not have finite attractor property and the termination assertion (b) is unsound. It is an open question whether the above problem is decidable for the PLCSs defined in Section 7 (or considered in [6,5,7]) which have finite attractor property.

The PE scheme is conceptually very simple, however, no information about the number of iterations before it terminates can be extracted from Theorem 10. For finite state Markov Chains standard algebraic methods allow to find the exact value of $Prob_M(s_1 \xrightarrow{*} s_2)$ in polynomial time; however, in this case the PE scheme finds an approximation in time $|M|^{\Omega(\ln(\frac{1}{\theta}))}$.

An alternative approach for approximation of $Prob_M(s_1 \xrightarrow{*} s_2)$ is to "approximate" a countable M.C. M by a finite state Markov chain M' and then to evaluate $Prob_{M'}(s_1 \xrightarrow{*} s_2)$ by standard algebraic methods. Below is a simple transformation which allows to reduce the size of Markov chains.

Let $M = (S,P)$ and let $U \subseteq S$ and let u be a new state. The chain $M' = (S',P')$ which is obtained from M by collapsing U into an absorbing state u is denoted by $M^{U,u}$ and is defined as follows: $S' = S \setminus U \cup \{u\}$ and

$$P'(s,s') = \begin{cases} \sum_{d\in U} P(s,d) & \text{if} \quad s \neq u \wedge s' = u, \\ P(s,s') & \text{if} \quad s \neq u \wedge s' \neq u, \\ 1 & \text{if} \quad s = u = s', \\ 0 & \text{otherwise.} \end{cases}$$

The following two lemmas are immediate, but useful for reductions of the size of M.C.

Lemma 11. *Let M be a M.C., let s_1,s_2 be states of M, let $u \notin S$, let C be a recurrent class such that $s_1 \notin C$ and let $M' = M^{C,u}$.*

1. If $s_2 \in C$ then $Prob_M(s_1 \xrightarrow{*} s_2) = Prob_{M'}(s_1 \xrightarrow{*} u)$.
2. If $s_2 \notin C$ then $Prob_M(s_1 \xrightarrow{*} s_2) = Prob_{M'}(s_1 \xrightarrow{*} s_2)$.

Lemma 12. *Let M be a M.C., let s_1,s_2 be states of M. Assume that $D \subseteq S \setminus \{s_1,s_2\}$ is such that either*

(1) $Prob\{\pi : \pi$ starts at s_1 and reaches $D\} \leqslant \theta$ or
(2) $\forall s \in D.Prob\{\pi{:}\pi$ starts at s and reaches $s_2\} \leqslant \theta$.

Let $M' = M^{D,d}$. Then

$$Prob_{M'}(s_1 \xrightarrow{*} s_2) \leqslant Prob_M(s_1 \xrightarrow{*} s_2) \leqslant Prob_{M'}(s_1 \xrightarrow{*} s_2) + \theta.$$

5. A generalization of one dimensional random walk

No information about the number of iterations of the Path Enumeration scheme can be extracted from Theorem 10. We consider here a generalization of one dimensional random walk. The main result of this section is Lemma 14. It is the key lemma for the correctness and the complexity analysis of the algorithm for the quantitative probabilistic reachability problem over PLCS, which is presented in Section. 8

The following lemma is easily derived from standard properties of one dimensional random walks.

Lemma 13. *Let $M = (S,P)$ be a Markov chain where $S = \{0,1,2,3,\ldots\}$, and*

1. $P(0,0) = 1$.
2. $P(i,i+1) = \nu_i$, $P(i,i-1) = \mu_i$, and $P(i,i) = 1 - \mu_i - \nu_i$, for $i \geqslant 1$.
3. There is $q > 0.5$ such that $\mu_i > q$ for all $i \geqslant 2$.

Let $N(\mu_1, q, \theta) = \ulcorner \frac{\ln(1-\mu_1)-\ln(\mu_1 \cdot \theta)}{\ln q - \ln(1-q)} \urcorner + 1$, *where* $\ulcorner x \urcorner$ *stands for the smallest integer which is greater than or equal to* x. *Then, for each* $\theta > 0$ *and* $n \geqslant N(\mu_1, q, \theta)$, *the probability of reaching a state* n *from* 1 *is less than* θ.

Proof. It is well known (see, e.g. [15] exercise 7 p. 88) that the probability of reaching state n, starting from state $k > 0$ is given by the expression $\frac{\sum_{i=0}^{k-1} \rho_i}{\sum_{i=-1}^{n-1} \rho_i}$, where

$$\rho_0 = 1 \qquad \rho_i = \frac{\mu_1 \mu_2 \cdots \mu_i}{\nu_1 \nu_2 \cdots \nu_i}.$$

This means that the probability of reaching state n from state 1 is given by

$$\frac{\rho_0}{\sum_{i=0}^{n-1} \rho_i} = \frac{1}{\sum_{i=0}^{n-1} \rho_i} \leqslant \frac{1}{\rho_{n-1}}$$

$$\leqslant \frac{\nu_1}{\mu_1 \left(\frac{q}{1-q}\right)^{n-1}} \leqslant \frac{1-\mu_1}{\mu_1 \left(\frac{q}{1-q}\right)^{n-1}}.$$

Thus, if we choose

$$n \geqslant \left(\log_{\frac{q}{1-q}} \left(\frac{1-\mu_1}{\mu_1 \cdot \theta}\right)\right) + 1$$

then the probability of reaching n from 1 is less than θ. □

The main technical lemma for the correctness and the complexity analysis of the algorithm presented in Section 8 is a generalization of Lemma 13.

Lemma 14 (Main lemma). *Consider a Markov chain* $M = (S, P)$ *such that*

(1) S *is the union of disjoint sets* S_i $(i \in Nat)$.
(2) If $s \in S_i$, $s' \in S_j$, *and* $P(s, s') > 0$, *then* $j \leqslant i + 1$.
(3) S_0 *is the union of two set of states* C *and* R *such that:*
- *For every state* $s \in R$, *only states in* R *are reachable from* s.
- *For every state* $s \in S_1$ *there is a finite path to* R *with the probability* $> \delta$ *which is inside* $C \cup R$ *(i.e., the last node of this path is in* R *and all the nodes except* s *are in* $S_0 = C \cup R$*).*

(4) There is $\alpha < \frac{1}{2}$ *such that* $\nu_i + \gamma_i < \alpha$ *for each* $i \geqslant 2$, *where*

$$\nu_i = \sup_{s \in S_i} \left(\sum_{s' \in S_{i+1}} P(s, s') \right) \text{ and } \gamma_i = \sup_{s \in S_i} \left(\sum_{s' \in S_i} P(s, s') \right).$$

Let $N_0 = N(\delta, 1 - \alpha, \theta)$, *where* N *is defined as in Lemma 13.*
Then, for every $s \in S_0 \cup S_1$ *the probability of reaching* $\bigcup_{n \geqslant N_0} S_n$ *from* s *is less than* θ.

Proof. See Appendix A. □

Lemma 12 and Lemma 14 imply.

Lemma 15. *Let M, S_i and N_0 be as in Lemma 14 and assume that $s_1, s_2 \in S_0$. Let $U = \bigcup_{n \geqslant N_0} S_n$. Let $M' = M^{U,u}$. Then $Prob_{M'}(s_1 \xrightarrow{*} s_2) \leqslant Prob_M(s_1 \xrightarrow{*} s_2) \leqslant Prob_{M'}(s_1 \xrightarrow{*} s_2) + \theta$.*

Remark. An important instance of Lemma 15 is the case when all S_i are finite sets. In this case the problem of approximating the probability of reachability for a countable Markov chain M is reduced to the problem of computing (or approximating) the probability of reachability for finite state Markov chain M'. However, the assumption of finiteness of S_i is not needed for the validity of Lemmas 14 and 15.

6. Lossy channel systems

In this section, we consider lossy channel systems: processes with a finite set of local states operating on a number of unbounded and unreliable channels.

A lossy channel system (LCS) consists of a finite state process operating on a finite set of channels each of which behaves as a FIFO buffer which is unbounded and unreliable in the sense that it can nondeterministically lose messages. Formally, a lossy channel system (LCS) $\mathcal{L}$ is a tuple $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T})$ where $\mathtt{S}$ is a finite set of *local states*, $\mathtt{C}$ is a finite set of *channels*, $\mathtt{M}$ is a finite *message alphabet*, and $\mathtt{T}$ is a set of *transitions* each of the form $(\mathtt{s}_1, \mathtt{op}, \mathtt{s}_2)$, where $\mathtt{s}_1, \mathtt{s}_2 \in \mathtt{S}$, and $\mathtt{op}$ is an *operation* of one of the forms $\mathtt{c!m}$ (sending message $\mathtt{m}$ to channel $\mathtt{c}$), or $\mathtt{c?m}$ (receiving message $\mathtt{m}$ from channel $\mathtt{c}$). A *global state s* is of the form $(\mathtt{s}, \mathtt{w})$ where $\mathtt{s} \in \mathtt{S}$ and $\mathtt{w}$ is a mapping from $\mathtt{C}$ to $\mathtt{M}^*$.

For words $x, y \in \mathtt{M}^*$, we use $x \bullet y$ to denote the concatenation of x and y. We write $x \preceq y$ to denote that x is a (not necessarily contiguous) substring of y. We use $|x|$ to denote the length of x, and use $x(i)$ to denote the ith element of x where $1 \leqslant i \leqslant |x|$. For $\mathtt{w}_1, \mathtt{w}_2 \in (\mathtt{C} \mapsto \mathtt{M}^*)$, we use $\mathtt{w}_1 \preceq \mathtt{w}_2$ to denote that $\mathtt{w}_1(\mathtt{c}) \preceq \mathtt{w}_2(\mathtt{c})$ for each $\mathtt{c} \in \mathtt{C}$, and define $|\mathtt{w}| = \sum_{\mathtt{c} \in \mathtt{C}} |\mathtt{w}(\mathtt{c})|$. We also extend $\preceq$ to a relation on $\mathtt{S} \times (\mathtt{C} \mapsto \mathtt{M}^*)$, where $(\mathtt{s}_1, \mathtt{w}_1) \preceq (\mathtt{s}_2, \mathtt{w}_2)$ iff $\mathtt{s}_1 = \mathtt{s}_2$ and $\mathtt{w}_1 \preceq \mathtt{w}_2$.

The LCS $\mathcal{L}$ induces a transition system $\left(S, \longrightarrow\right)$, where S is the set of global states, i.e., $S = (\mathtt{S} \times (\mathtt{C} \mapsto \mathtt{M}^*))$, and $(\mathtt{s}_1, \mathtt{w}_1) \longrightarrow (\mathtt{s}_2, \mathtt{w}_2)$ iff one of the following conditions are satisfied:

- There is a $\mathtt{t} \in \mathtt{T}$, where $\mathtt{t}$ is of the form $(\mathtt{s}_1, \mathtt{c!m}, \mathtt{s}_2)$ and $\mathtt{w}_2$ is the result of appending $\mathtt{m}$ to the end of $\mathtt{w}_1(\mathtt{c})$.
- There is a $\mathtt{t} \in \mathtt{T}$, where $\mathtt{t}$ is of the form $(\mathtt{s}_1, \mathtt{c?m}, \mathtt{s}_2)$ and $\mathtt{w}_1(\mathtt{c}) = \mathtt{m} \bullet x$ for some $x \in \mathtt{M}^*$ and $\mathtt{w}_2(\mathtt{c}) = x$ and $\mathtt{w}_2(\mathtt{c}') = \mathtt{w}_1(\mathtt{c}')$ for $\mathtt{c}' \neq \mathtt{c}$.
- Furthermore, if $(\mathtt{s}_1, \mathtt{w}_1) \longrightarrow (\mathtt{s}_2, \mathtt{w}_2)$ according to one of the previous two rules, then $(\mathtt{s}_1, \mathtt{w}_1) \longrightarrow \left(\mathtt{s}_2, \mathtt{w}_2'\right)$ for each $\mathtt{w}_2' \preceq \mathtt{w}_2$.

In the first two cases we define $\mathtt{t}(\mathtt{s}_1, \mathtt{w}_1) = (\mathtt{s}_2, \mathtt{w}_2)$. In the third case we say that $\left(\mathtt{s}_2, \mathtt{w}_2'\right)$ is obtained from $(\mathtt{s}_1, \mathtt{w}_1)$ by a transition which follows by losses of messages.

A transition $(\mathtt{s}_1, \mathtt{op}, \mathtt{s}_2)$ is said to be *enabled* at $(\mathtt{s}, \mathtt{w})$ if $s = s_1$ and either $\mathtt{op}$ is of the form $\mathtt{c!m}$; or $\mathtt{op}$ is of the form $\mathtt{c?m}$ and $\mathtt{w}(\mathtt{c}) = \mathtt{m} \bullet x$, for some $x \in \mathtt{M}^*$. We define *enabled* $(\mathtt{s}, \mathtt{w}) = \{\mathtt{t} : \mathtt{t} \text{ is enabled at } (\mathtt{s}, \mathtt{w})\}$.

Remark on notation. We use $\mathtt{s}$ and $\mathtt{S}$ to range over local states and sets of local states, respectively. On the other hand, s and S range over states and sets of states of the induced transition system (states of the transition system are global states of the LCS).

A set $Q \subseteq S$ is said to be *upward closed* if $s_1 \in Q$ and $s_1 \preceq s_2$ imply $s_2 \in Q$. The upward closure $Q \uparrow$ of a set Q is the set $\{s : \exists s' \in Q.\ s' \preceq s\}$.

By Higman's Lemma [13] it follows that $\preceq$ is a well quasi-ordering, i.e., for each infinite sequence $x_0, x_1, x_2, \ldots$ there are i and j with $i < j$ and $x_i \preceq x_j$. Moreover

Lemma 16 (Higman [13]). *For every set I which is upward closed wrt $\preceq$ there is a finite set V such that I is the upward closure of V.*

Theorems in [4,12] imply the following decidability results for LCS:

Lemma 17 (Decidability [4,12]). *(1) It is decidable whether a state s_2 is reachable from a state s_1. (2) It is decidable whether the upward closure of a finite set Q is reachable from a state s. (3) There is an algorithm Find-a-path$(s_1, s_2, \mathcal{L})$ which returns a path from s_1 to s_2 in the lossy channel system $\mathcal{L}$ or returns "No" if s_2 is not reachable from s_1. (4) Graph(A) is computable for every finite set of global states A of an LCS.*

7. Probabilistic lossy channel systems

We introduce a probabilistic behaviour into LCS obtaining probabilistic lossy channel systems. Probabilistic lossy channel systems were first defined in [14]. The semantics considered here was presented in [5,7] and differs from that in [14,6,1].

A probabilistic lossy channel system (PLCS) $\mathcal{L}$ is of the form $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T}, \lambda, w)$, where $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T})$ is an LCS, $\lambda \in (0, 1)$, and w is a mapping from $\mathtt{T}$ to the natural numbers. Intuitively, we derive a Markov chain from the PLCS $\mathcal{L}$ by assigning probabilities to the transitions of the (underlying) transition system of $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T})$. The probability of performing a transition $\mathtt{t}$ from a global state $(\mathtt{s}, \mathtt{w})$ is determined by the weight $w(\mathtt{t})$ of $\mathtt{t}$ compared to the weights of the other transitions which are enabled at $(\mathtt{s}, \mathtt{w})$. Furthermore, after performing each transition, each message which resides inside one of the channels may be lost, independently of the other messages, with the probability λ. This means that the probability of the transition from $(\mathtt{s}_1, \mathtt{w}_1)$ to $(\mathtt{s}_2, \mathtt{w}_2)$ is equal to (the sum over all $(s_3, \mathtt{w}_3)$ of) the probability of reaching some $(s_3, \mathtt{w}_3)$ from $(\mathtt{s}_1, \mathtt{w}_1)$ through performing a transition t of the underlying LCS, multiplied by the probability of reaching $(\mathtt{s}_2, \mathtt{w}_2)$ from $(\mathtt{s}_3, \mathtt{w}_3)$ through the loss of messages (see [5] for detailed calculations of the probabilities of the transitions).

To simplify the presentation, we assume from now on that PLCSs have no deadlock states, i.e., from every state a transition is enabled (The treatment of the PLCS with deadlock states can be reduced to the deadlock free PLCS in the same way as in [2]) . The only probabilistic properties of PLCSs which we use are summarized in the next two lemmas from [5].

Lemma 18. *Let s be a state with m messages. The probability of the transitions from s to the set of states with $> m + 1$ messages is 0. The probability of the transitions from s to the set of states with $m + 1$ messages is $\leqslant (1 - \lambda)^{m+1}$. The probability of the transitions from s to the set of states with m messages is $\leqslant \lambda m (1 - \lambda)^m$.*

Lemma 19. *For each* $\lambda, w,$ *and PLCS* $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T}, \lambda, w)$ *the set of the states with the empty set of messages is a finite attractor.*

The next lemma plays a key role in the algorithm presented in Section 8.

Lemma 20. *For each PLCS* $\mathcal{L} = (\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T}, \lambda, w)$ *there are* $V_1, \ldots, V_k$ *such that* V_i *are finite sets of global states and k is the number of the recurrent classes of* $\mathcal{L}$ *and for each state* $s : s$ *is in the ith recurrent class of* $\mathcal{L}$ *iff s is not in the upward closure of* V_i. *Moreover,* $V_1, \ldots, V_k$ *are computable from the underlying LCS* $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T})$.

Proof. Let $C_1, \ldots, C_k$ be the recurrent components of $\mathcal{L}$. By Lemma 5(4) C_i are the BSCCs of $\mathcal{L}$.

First, we show that the complement of C_i $(i = 1, \ldots, k)$ is an upward closed set. Let $s_1 \notin C_i$ and $s_1 \preceq s_2$. Assume for the purpose of the contradiction that $s_2 \in C_i$. Since s_2 is recurrent, it is repeatedly reachable from itself and therefore there exists a state $s \in C_i$ and a transition from s to s_2. Observe that s_1 can be reached from s by performing a transition to s_2 which follows by lost of messages. Since C_i is a BSCC and $s \in C_i$ and s_1 is reachable from s it follows that $s_1 \in C_i$. This contradicts the assumption that $s_1 \notin C_i$.

The complements of the sets C_i $(i = 1, \ldots, k)$ are upward closed, therefore, by Higman's Lemma 16, there are finite sets V_i such that $s \in C_i$ iff s is not in the upward closure of V_i.

Finally note that the construction of the sets V_i was based only on the underlying LCS $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T})$ and was independent from λ and w. □

8. Algorithm for approximating the probability of reachability

Lemmas 19, 17(1) and Theorem 10 imply that there is an algorithm based on the PE scheme for the quantitative probabilistic reachability problem.

The Markov chain M defined by a PLCS has finite branching and is presented effectively. Therefore the sets of paths $\mathrm{Comp}_n(s_1)$ and $\mathrm{Reach}_n(s_1, s_2)$ and the probability p_n^+ of the set $\mathrm{Reach}_n(s_1, s_2)$ can be computed. The reachability problem for the transition system underlying M is decidable by Lemma 17. Therefore, the sets $\mathrm{Escape}_n(s_1, s_2)$ of the paths of length n which escape s_2 is computable. Hence, $\mathrm{Undecided}_n(s_1, s_2)$ and $p_n^?$ can be computed. Therefore, for the Markov chain M defined by a PLCS, path enumeration scheme can be implemented. By Lemma 19 this chain has a finite attractor and therefore by Theorem 10 the algorithm for the quantitative probabilistic reachability problem based on the PE scheme terminates and outputs a correct approximation.

However, no information about the complexity of this algorithm can be extracted from Theorem 10. In this section, we provide an algorithm with a parametric complexity $f(\mathcal{L}, s_1, s_2) \times \frac{1}{\theta^3}$ for the quantitative probabilistic reachability problem.

The idea of the algorithm is to take the set $B_{\leqslant n}$ of states with at most n messages of the Markov chain M generated by PLCS $\mathcal{L}$. Construct a finite Markov chain $\widehat{M}$ by restricting the transition of M to $B_{\leqslant n}$, and then for each recurrence class D_i of M collapse the set of D_i states in $B_{\leqslant n}$ into one state of $\widehat{M}$. Finally, calculate the probability of reaching s_2 from s_1 in the finite M.C. $\widehat{M}$. The crucial fact in the correctness of our algorithm is that relying on Lemma 14 we can compute n big enough which will ensure that the probability of reaching s_2 from s_1 in the finite Markov chain $\widehat{M}$ approximates up to θ the probability of reaching s_2 from s_1 in the infinite Markov chain M.

In the rest of this section we describe the algorithm with a justification of its correctness and provide an analysis of its complexity.

Algorithm – for Quantitative Probabilistic Reachability Problem
Input: PLCS $\mathcal{L} = (\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T}, \lambda, w)$ with an underlying Markov chain $M = (S, P)$, states $s_1, s_2 \in S$, and a rational θ.
Output: a rational r such that $r \leqslant Prob_M(s_1 \xrightarrow{*} s_2) \leqslant r + \theta$.

Let A be the (finite) set of all states with 0 messages. A is an attractor by Lemma 19. By Lemma 17(4) we can construct $Graph(A)$. Then we can find the bottom strongly connected components $C_1, \ldots, C_k$ in $Graph(A)$ and for $1 \leqslant i \leqslant k$ by Lemma 5 and by Lemma 20 we can compute finite sets of states V_i such that

$$\forall s \in S (s \text{ is in the recurrent class of } C_i) \text{ iff } (s \text{ is not in the upward closure of } V_i) \tag{8}$$

Hence, we can check whether s_1 (or any other state s) is in the ith recurrent class.

In the case when s_1 is recurrent we proceed as follows: If s_2 is recurrent and in the same recurrent class as s_1 then output 1 else output 0. (The correctness of this answer follows by Lemma 1(4–5).)

Below we consider the case when s_1 is not recurrent. By Lemma 17(3) we can compute l such that for every $u, v \in A \cup \{s_1, s_2\}$ if u is reachable from v then there is a path from v to u which passes only through nodes with at most l messages. Let m be such that $\forall n. m \leqslant n \rightarrow (1-\lambda)^n(1-\lambda+n\lambda) < \frac{1}{3}$ i.e., the probability to move from every state with $n \geqslant m$ messages to the set of states with at least n messages is less than $\frac{1}{3}$ (by Lemma 18). Let $h = max(l, m)+1$.

Notations. Below we denote by B_i (respectively, by $B_{\leqslant i}$ or by $B_{<i}$) the set of states with i (respectively, with at most i or with less than i) messages.

For every state $s \in B_{\leqslant h}$ there is a finite path π_s which first chooses a lossy transition which leads to a state s' with 0 messages and then follows by a path from s' which is inside $B_{\leqslant l} \subseteq B_{\leqslant h}$ to a BSCC of $Graph\,(A)$. Let $\delta_s = Prob(\pi_s) > 0$ and let $0 < \delta = min(\delta_s : s \in B_{\leqslant h})$. Note that up to this point all our computations were independent of θ and their complexity depended only on $\mathcal{L}$, s_1 and s_2.

Observe that if we denote by R the set of recurrent states of M, by C the set of transient states with $< h$ messages; by S_0 the set $R \cup C$ and by S_i $(i > 0)$ the set of transient states with $h+i-1$ messages, then the assumptions of Lemma 14 are satisfied. Let $N_0 = N(\delta, \frac{2}{3}, \theta)$, where N is the function from Lemma 13 and let $n = h + N_0$. Note that n depends linearly on $\ln(\frac{1}{\theta})$. By Lemma 14 the probability to reach from s_1 the set $U = \bigcup_{n \geqslant N_0} S_n$ of transient states with $\geqslant n$ messages is at most θ. Therefore, by Lemma 15 we derive that $Prob_{M'}(s_1 \xrightarrow{*} s_2) \leqslant Prob_M(s_1 \xrightarrow{*} s_2) \leqslant Prob_{M'}(s_1 \xrightarrow{*} s_2) + \theta$ for $M' = M^{U,dead}$ obtained by collapsing U into a fresh state $dead$. The chain M' might be infinite. Below we are going to construct a finite state M.C. $\widehat{M}$ of size bounded by $|B_{\leqslant n}|$ such that $Prob_{M'}(s_1 \xrightarrow{*} s_2) = Prob_{\widehat{M}}(s_1 \xrightarrow{*} s_2)$. Hence, $Prob_{\widehat{M}}(s_1 \xrightarrow{*} s_2)$ will approximate up to θ the value of $Prob_M(s_1 \xrightarrow{*} s_2)$ which we are trying to compute.

The complexity of the construction of $\widehat{M}$ will be $O(|B_{\leqslant n}|^2)$ and by standard algebraic methods we can compute $Prob_{\widehat{M}}(s_1 \xrightarrow{*} s_2)$ in time $O(|B_{\leqslant n}|^3)$. Since n depends linearly on $\ln(\frac{1}{\theta})$, it follows that $|B_{\leqslant n}|$ depends linearly on $\frac{1}{\theta}$ and the complexity of the entire algorithm is $f(\mathcal{L}, s_1, s_2) \times \frac{1}{\theta^3}$.

We define $\widehat{M}$ by replacing every recurrent class of M' by an absorbing state. From Lemma 11 we will derive that this transformation preserves the probability of reaching s_2 from s_1. Formally a (finite) M. C. $\widehat{M} = (\widehat{S}, \widehat{P})$ is defined as follows.

Let D_i $(i = 1, \dots, k)$ be the states with $\leqslant n$ messages, which are in the ith recurrent class. (These sets can be computed by Eq. (8) in time $O(|B_{\leqslant n})|$. Let D be $B_{<n} \setminus \cup D_i$.

(1) $\widehat{S} = D \cup \{d_1, \dots, d_k\} \cup \{dead\}$. The states d_i correspond to the recurrent classes and the state *dead* corresponds to the set of transient states of M with $\geqslant n$ messages.
(2) The states $d_1, \dots, d_k, dead$ are absorbing, i.e., for $d \in \{d_1, \dots, d_k\} \cup \{dead\}$:

$$\widehat{P}(d, d') = \begin{cases} 1 & \text{if } d' = d, \\ 0 & \text{otherwise.} \end{cases}$$

(3) $\widehat{P}(d, d') = P(d, d')$ for $d, d' \in D$.
(4) $\widehat{P}(d, dead) = \sum_{s \in B_n \setminus \cup D_i} P(d, s)$ for $d \in D$.
(5) $\widehat{P}(d, d_i) = \sum_{d' \in D_i} P(d, d')$, for $d \in D$ and $i : 1 \leqslant i \leqslant k$.

Recall that we treated already the case when s_1 is recurrent, hence s_1 is in D. Compute the output r which approximates $Prob_M(s_1 \xrightarrow{*} s_2)$ up to θ by the following cases:

1. if $s_2 \in D$ then compute by standard algebraic methods the probability r of reaching s_2 from s_1 in (the finite Markov chain) $\widehat{M}$.
2. if $s_2 \in D_i$ then compute the probability r of reaching d_i from s_1 in $\widehat{M}$.

We completed the presentation of the algorithm, established its correctness and proved that its complexity is $f(\mathcal{L}, s_1, s_2) \times \frac{1}{\theta^3}$. It was shown in [19] that the complexity of the reachability problem for LCSs is not bounded by any primitive recursive function in the size of LCS. Therefore, f is not primitive recursive in the size of PLCS.

9. Probability of automata definable properties

In this section, we consider more general properties than reachability. Let φ be a property of computations. We will be interested in approximating

$$Prob\{\pi : \pi \text{ is a computation from } s \text{ in PLCS } \mathcal{L} \text{ and } \pi \text{ satisfies } \varphi\}.$$

We show that if the properties of computations are specified by (the ω-behavior of) finite state automata or equivalently by formulas of the monadic second-order logic of order, then the above problem is computable.

9.1. State-labeled systems and ω-automata

To specify properties of computations (ω-sequences of states) one needs to assume that the states are labeled by elements of some finite alphabet. Hence, we extend LCS's with a labeling function: a *state-labeled LCS* is an LCS together with a finite alphabet Σ and a labeling function *lab* from the local states to Σ. Throughout this section we assume that LCS's are state-labeled and will often use "LCS" for "state-labeled LCS." We lift the labeling from an LCS $\mathcal{L}$ to the *state-labeled transition system* $T = (S, \longrightarrow, \Sigma, lab\ )$ it induces: the label of every state $(\mathtt{s}, \mathtt{w})$ in T is the label $lab(\mathtt{s})$ of its local state component. When we deal with probabilistic lossy channel systems we also assume that the underlying LCS is labeled, and this labeling is lifted to the labeling of the corresponding Markov chain. In this manner we obtain *state-labeled PLCS's* inducing *state-labeled Markov chains*.

A path $s_0 s_1, \cdots$ in a state-labeled transition system give rise the ω-string $lab(s_0)lab(s_1)\cdots$ over the alphabet Σ. We consider properties of paths that are defined using automata: the ω-string of the path must be accepted by the automaton.

Recall that a *finite (Muller) automaton* $\mathcal{A}$ is a tuple $(\mathcal{Q}, \Sigma, \rightarrow, q_0, \mathcal{F})$, consisting of a finite set $\mathcal{Q}$ of *states*, a finite *alphabet* Σ, a *transition relation* $\rightarrow$ which is a subset of $\mathcal{Q} \times \Sigma \times \mathcal{Q}$, an *initial state* $q_0 \in \mathcal{Q}$, and a collection $\mathcal{F} \subseteq 2^{\mathcal{Q}}$ of *fairness conditions*. We write $q \overset{a}{\rightarrow} q'$ if $\langle q, a, q' \rangle \in\rightarrow$. We say that $\mathcal{A}$ is *deterministic* if for every state q and every letter $a \in \Sigma$ there is one and only one q' such that $q \overset{a}{\rightarrow} q'$.

A *run* of $\mathcal{A}$ is an ω-sequence $q_0 a_0 q_1 a_1 \ldots$ such that $q_i \overset{a_i}{\rightarrow} q_{i+1}$ for all i. With such a run we associate the set $\mathtt{Inf}$ of all $q \in \mathcal{Q}$ that appear infinitely many times. A run meets the fairness conditions $\mathcal{F}$ if its $\mathtt{Inf}$ set belongs to $\mathcal{F}$ (Muller acceptance). An ω-string $a_0 a_1 \ldots$ over Σ is accepted by $\mathcal{A}$ if there is a run $q_0 a_0 q_1 a_1 \ldots$ that meets the fairness conditions of $\mathcal{A}$. The ω-language *accepted* by $\mathcal{A}$ is the set of all ω-strings accepted by $\mathcal{A}$.

We recall the following classical theorem (see [22]) stating that automata have the same expressive power as the monadic logic of order:

Theorem 21. *For an ω-language L, the following conditions are equivalent*:

1. *L is accepted by a finite state Muller automaton.*
2. *L is accepted by a deterministic finite state Muller automaton.*
3. *L is definable by a MSO formula.*

9.2. Products with automata

To approximate the probability of the set of computations in the PLCS $\mathcal{L}$ that are accepted by $\mathcal{A}$ we shall built the product $\mathcal{L}'$ of $\mathcal{A}$ and $\mathcal{L}$ and then solve an appropriate quantitative reachability problem for $\mathcal{L}'$. In this section, we summarize properties of the product and in the next section it is shown how the quantitative probabilistic model checking problem is reduced to the quantitative reachability problem.

Consider an automaton $\mathcal{A} = (\mathcal{Q}, \Sigma, \rightarrow, q_0, \mathcal{F})$, and a state-labeled transition system $T = (S, \longrightarrow, \Sigma, lab\ )$. The *product* $\mathcal{A} \times T$ of $\mathcal{A}$ and T is a state-labeled transition system $T' = (S', \rightarrow', \Sigma, lab')$ defined as follows:

States: $S' = \mathcal{Q} \times S$ is the Cartesian product of the states of $\mathcal{A}$ and of T.
Labeling: A state (q, s) is labeled by $lab(s)$, i.e., it has the same label as s in T.

Transition relation: There is a transition $(q,s) \to' \left(q',s'\right)$ iff there is a transition $s \to s'$ in T and there is a transition $q \overset{lab\,(s)}{\longrightarrow} q'$ in $\mathcal{A}$.

We also define the product $\mathcal{R} = \mathcal{A} \times M$ of a *deterministic* automaton and a state-labeled Markov chain $M = \left(S, P, \Sigma, lab\right)$. Here the states and labels are as in $\mathcal{A} \times T$. The probability P' in $\mathcal{R}$ is given by

$$P'((q,s), \left(q',s'\right)) = \begin{cases} P(s,s') & \text{if } q \overset{lab\,(s)}{\to} q' in \mathcal{A}, \\ 0 & \text{otherwise.} \end{cases}$$

Observe that the requirement that $\mathcal{A}$ is deterministic ensures that the sum of probabilities of the transitions from the state (q,s) is the same as the sum of probabilities of the transitions from the state s in M, i.e. the sum is one. Hence the product is indeed a labeled Markov chain. Observe further that if T is the transition system underlying M, then $\mathcal{A} \times T$ is exactly the transition system underlying $\mathcal{A} \times M$.

As in [10,23] with a path $\pi = s_0 s_1 \ldots$ in M we associate $\pi^{\mathcal{A}}$, the only path in $\mathcal{R}$ of the form $(q_0,s_0)\,(q_1,s_1)\ldots$ with $q_i \overset{lab(s_i)}{\to} q_{i+1}$ for all i; $\pi^{\mathcal{A}}$ exists and is unique because $\mathcal{A}$ is deterministic. Moreover, for every path π' in $\mathcal{R}$ there exists a path π in M such that $\pi' = \pi^{\mathcal{A}}$. Hence the function ν that maps π to $\pi^{\mathcal{A}}$ is bijective. Moreover, it preserves the probabilistic measure of the sets of paths as explained in the following lemma.

Lemma 22. *Let s be a state and let L be a set of paths in M that start at s. The measure of L in M is equal to the measure of $L^{\mathcal{A}} = \{\pi^{\mathcal{A}} : \pi \in L\}$ in $\mathcal{A} \times M$.*

Definition 23. Let $M = (S, P, \Sigma, lab)$ be a Markov chain and $\mathcal{A} = (Q, \Sigma, \to, q_0, \mathcal{F})$ be a deterministic automaton. Let $F \in \mathcal{F}$ and let C be a subset of a recurrent class of the product $\mathcal{R} = \mathcal{A} \times M$. We say that C *satisfies* the condition F if

(1) $q \in F$ for each state (q,u) reachable from C in $\mathcal{R}$, and
(2) for each $q \in F$ there is $u \in M$ such that (q,u) is reachable from C in $\mathcal{R}$.

We say that C *satisfies* $\mathcal{F}$ if C satisfies at least one $F \in \mathcal{F}$.

Note that the property "C satisfies F" is a class property, i.e., this depends only on the recurrent class D that contains C. A recurrent class D satisfies F iff $q \in F$ for each state $(q,u) \in D$ and for each $q \in F$ there is $u \in M$ such that $(q,u) \in D$.

We say that a computation $s_1, s_2, \ldots$ is *accepted* by an automaton iff the corresponding ω-string $lab\,(s_1) lab\,(s_2) \ldots$ is accepted.

Lemma 24. *Let $\mathcal{A} = \langle Q, \Sigma, \to, q_0, \mathcal{F}\rangle$ be a deterministic automaton, let M be a labeled Markov chain, let $\mathcal{R}$ be the product of $\mathcal{A}$ and M, and let B be a finite attractor of $\mathcal{R}$. Let $C_1, C_2, \ldots, C_p$ be all the BSCC in Graph(B) which satisfy $\mathcal{F}$. Then the following are equivalent:*

(1) The probability that a computation of M that starts at s is accepted by $\mathcal{A}$ is r.
(2) The probability that a computation of $\mathcal{R}$ that starts at (q_0,s) reaches $\bigcup_{i=1}^{p} C_i$ is r.

Proof. Let π be a path in M and let $\pi^{\mathcal{A}} = \nu(\pi)$ be the corresponding path in $\mathcal{R}$. Note that the same ω-string corresponds to π and to $\pi^{\mathcal{A}}$, therefore π is accepted by $\mathcal{A}$ if and only if $\nu(\pi)$ is accepted by $\mathcal{A}$. Since ν is bijective and measure preserving, it is sufficient to prove that the probability of the set of computations of $\mathcal{R}$ that start at (q_0, s) and is accepted by $\mathcal{A}$ is r if and only if the probability of the set of computations of $\mathcal{R}$ that start at (q_0, s) and reach $\bigcup_{i=1}^{p} C_i$ is r.

Let $D_1, \ldots, D_n$ be the recurrent components of $\mathcal{R}$. We can assume that $D_1, \ldots, D_p$ are the components that satisfy $\mathcal{F}$ and that $C_i \subseteq D_i$ for $i = 1, \ldots, p$

Let X_i (respectively, Y_i) be the sets of paths that start at (q_0, s) and reach D_i (respectively, repeatedly reach every state of D_i).

By Lemma 1(4),

$$Prob(X_i) = Prob(Y_i) \tag{9}$$

and by Lemma 5(3)

$$\sum Prob(X_i) = 1. \tag{10}$$

Therefore,

$$\sum Prob(Y_i) = 1. \tag{11}$$

Note that Y_i are disjoint sets and for $\pi \in Y_i$: π is accepted by $\mathcal{A}$ iff D_i satisfies $\mathcal{F}$. Hence, the probability of the set of computations of $\mathcal{R}$ that start at (q_0, s) and is accepted by $\mathcal{A}$ is $\sum_{i=1}^{p} Prob(Y_i)$. By Lemma 1(4), $Prob(Y_i)$ is equal to the probability to reach C_i, hence the last sum is equal to the probability of the set of computations of $\mathcal{R}$ that start at (q_0, s) and reach $\bigcup_{i=1}^{p} C_i$. This completes the proof of the lemma. □

Finally, the product $\mathcal{L}' = \mathcal{A} \times \mathcal{L}$ of an automaton with an LCS is defined along the same lines: the local states are pairs $(q, \mathtt{s})$ of a state of $\mathcal{A}$ and a local state of $\mathcal{L}$. The transitions $\mathtt{T}'$ of $\mathcal{L}'$ are all $((q, \mathtt{s}), \mathtt{op}, (q', \mathtt{s}'))$ such that $(\mathtt{s}, \mathtt{op}, \mathtt{s}')$ is a transition of $\mathcal{L}$ and $q \overset{lab(\mathtt{s})}{\rightarrow} q'$ is a transition in $\mathcal{A}$. We define the product of a *deterministic* $\mathcal{A}$ and a PLCS $\mathcal{L}$ along the same lines.

A crucial property of these constructions is the following:

Lemma 25.

1. *If T is the transition system induced by an LCS $\mathcal{L}$ then $\mathcal{A} \times T$ is (isomorphic to) the transition system induced by the LCS $\mathcal{A} \times \mathcal{L}$.*
2. *If M is the Markov chain induced by a PLCS $\mathcal{L}$ then $\mathcal{A} \times M$ is (isomorphic to) the Markov chain induced by the PLCS $\mathcal{A} \times \mathcal{L}$.*

Here, the isomorphism associates $(q, (\mathtt{s}, \mathtt{w}))$, a state of $\mathcal{A} \times T$ (resp. $\mathcal{A} \times M$), with $((q, \mathtt{s}), \mathtt{w})$, a state of the transition system (respectively Markov chain) induced by $\mathcal{A} \times \mathcal{L}$.

9.3. Quantitative probabilistic model-checking

We are now ready to show that the following problem is computable.

Quantitative Probabilistic Model-checking Problem
Instance: A PLCS $\mathcal{L}$, its state s, a finite state ω-automaton $\mathcal{A}$, and a rational $\theta > 0$.
Task: Find a rational r such that the probability of the set of computations that start at s and are accepted by $\mathcal{A}$ is between r and $r + \theta$.

The Quantitative Probabilistic Model-checking Problem is reducible to the Quantitative Probabilistic Reachability Problem for PLCSs.

Theorem 26. *The Quantitative Probabilistic Model-checking Problem can be solved in time* $g(\mathcal{L}, \mathcal{A}, s)) \times \frac{1}{\theta^3}$.

Proof. Let $\mathcal{L}'$ be the product of $\mathcal{A} = (\mathcal{Q}, \Sigma, \rightarrow, q_0, \mathcal{F})$ and $\mathcal{L}$ and let s' be the global state (q_0, s) of $\mathcal{L}'$. Let $\mathcal{R}$ be the Markov chain that corresponds to $\mathcal{L}'$. We can assume that $\mathcal{A}$ is deterministic (or replace it by an equivalent deterministic automaton Theorem 21). Note that $\mathcal{R}$ is isomorphic to the product of $\mathcal{A}$ and the Markov chain M that corresponds to $\mathcal{L}$, by Lemma 25. Let A be a finite attractor of $\mathcal{R}$ (one can take for A the set of global states with the empty channels). Let $C_1, \ldots, C_p$ be the BSCC of *Graph*(A) that satisfy $\mathcal{F}$. Note that it is decidable whether a BSCC C satisfies $\mathcal{F}$, because by Lemma 17(2) for every $q \in \mathcal{Q}$ one can check whether the upward closure of the finite set $\{(q, \mathrm{s}, \epsilon) : \mathrm{s} \in \mathcal{L}\}$ is reachable from s'.

By Lemma 24, the probability of the set of computations of M that start at s and are accepted by $\mathcal{A}$ is equal to the probability of the set of computations of $\mathcal{R}$ that start at (q_0, s) and reach (a finite set) $\bigcup_{i=1}^{p} C_i$. The latter probability can be approximated by the algorithm of Section 8 in time $g(\mathcal{L}', \mathcal{A}, s)) \times \frac{1}{\theta^3}$. □

10. Conclusions and further results

This paper deals with the verification of quantitative properties of probabilistic infinite state systems. For finite state Markov chains the probability of reaching one state from another can be computed explicitly. However, even for very simple infinite state Markov chains usually there is no explicit expression which computes the probability of reaching one state from another.

Our contribution can be described as follows:

- To the best of our knowledge we were the first to formulate the quantitative reachability and model-checking problems: approximate the probability that a certain property is satisfied.
- We provided the PE scheme which reduces the quantitative reachability problem for M.C. to the reachability problem of the underlying transition system. We proved the correctness of the PE scheme for two important classes of Markov chains.
- We provided a method to reduce the quantitative reachability problem for countable Markov chains to the quantitative reachability problem for finite state Markov chains.

- We derived the decidability of the quantitative reachability and model checking problems for PLCSs.

Below we summarize our results and point out to further extensions.

10.1. Path Enumeration Scheme

We provided the Path Enumerating Scheme for the quantitative reachability problem. We proved that the PE scheme is correct for two classes of Markov chains:

(1) Markov chains with a finite attractor.
(2) Markov chains with δ-reachability property.

To approximate the probability of reaching a state s_2 from a state s_1 in a M.C. M, the PE scheme partitions $\mathrm{Comp}_n(s_1)$ - the set of all the computations of length n in M from s_1 into three sets: $\mathrm{Reach}_n(s_1,s_2)$ - all the computations in $\mathrm{Comp}_n(s_1)$ that reach s_2; $\mathrm{Escape}_n(s_1,s_2)$ - all the computations in $\mathrm{Comp}_n(s_1)$ that cannot be extended to computations that reach s_2; $\mathrm{Undecided}_n(s_1,s_2)$ - those computations in $\mathrm{Comp}_n(s_1)$ that have not reached s_2, but can be extended to computations that reach s_2. Assume that M has finite branching, all transition probabilities are rational and there exist the following algorithms:

Alg1 For every state $s \in M$ an algorithm computes the set $Next_s = \{s' :$ there is a transition $s \longrightarrow s'$ with non-zero probability $P(s,s')\}$.
Alg2 For every states $s,s' \in M$ an algorithm computes the probability $P(s,s')$.
Alg3 For every state $s \in M$ an algorithm decides whether s_2 is reachable from s.

Under the above assumptions the sets of paths $\mathrm{Reach}_n(s_1,s_2)$, $\mathrm{Escape}_n(s_1,s_2)$ and $\mathrm{Undecided}_n(s_1,s_2)$ are finite and there is an algorithm that computes the probabilities $p_n^+ = Prob(\mathrm{Reach}_n(s_1,s_2))$, $p_n^- = Prob(\mathrm{Escape}_n(s_1,s_2))$ and $p_n^? = Prob(\mathrm{Undecided}_n(s_1,s_2))$. Finally, p_n^+ approximates up to θ the probability of reaching s_1 from s_2, whenever $p_n^? < \theta$.

A variant of the PE scheme can be algorithmically implemented even for Markov chains with infinite branching. In this case we need

Alg1* an algorithm that for every state $s \in M$ and a rational $r < 1$ computes a finite set $N_{s,r} \subseteq Next_s$ such that the sum $\sum_{s' \in N_{s,r}} P(s,s')$ is greater than r.

Using such an algorithm, we can compute for every $s_1 \in M$, a natural n and a rational θ a finite set $C_n(s_1) \subseteq \mathrm{Comp}_n(s_1)$ such that $Prob(C_n(s_1)) > 1 - \frac{\theta}{2}$. Using Alg2 and Alg3, we can partition this finite set into three sets $\mathrm{Reach}_n(s_1,s_2)$, $\mathrm{Escape}_n(s_1,s_2)$ and $\mathrm{Undecided}_n(s_1,s_2)$ like above and compute the probabilities $p_n^+ = Prob(\mathrm{Reach}_n(s_1,s_2))$, $p_n^- = Prob(\mathrm{Escape}_n(s_1,s_2))$ and $p_n^? = Prob(\mathrm{Undecided}_n(s_1,s_2))$. Finally, p_n^+ approximates up to θ the probability of reaching s_1 from s_2, whenever $p_n^? < \frac{\theta}{2}$. If M has either finite attractor or δ-reachability property, then $\lim p_n^? = 0$, hence the algorithm terminates.

10.2. Reduction to finite state Markov chains

The PE scheme is conceptually very simple, however, no information about the number of iterations before it terminates can be extracted from Theorem 10. For finite state Markov chains standard algebraic methods allow to find the exact value of $Prob_M(s_1 \xrightarrow{*} s_2)$ in polynomial time; however, in this case the PE scheme finds an approximation in time $|M|^{\Omega(\ln(\frac{1}{\theta}))}$.

An alternative approach for approximation of $Prob_M(s_1 \xrightarrow{*} s_2)$ is to "approximate" a countable M.C. M by a finite state Markov chain M' and then to evaluate $Prob_{M'}(s_1 \xrightarrow{*} s_2)$ by standard algebraic methods. Lemmas 11 and 12 provide simple transformations which allow to reduce the size of Markov chains. Our main technical result, Lemma 14, identify a wide class of Markov chains for which the quantitative reachability problem can be reduced to the quantitative reachability problem for finite state Markov chains. This class of Markov chains generalize one dimensional random walks. Based on these lemmas we provided an algorithm with a parametric complexity $f(\mathcal{L}, s_1, s_2) \times \frac{1}{\theta^3}$ for the quantitative reachability and model checking problems over the PLCS.

10.3. Probabilistic lossy channel systems

We provided an algorithm for quantitative model checking for a realistic class of probabilistic lossy channel systems, where during each step of the runs of the systems, any message inside the channels may be lost with a certain predefined probability; losses occur independently of each other.

Our results extend to PLCS with the semantics defined as in [14,6] for the case when losing probability is at least 0.5; in this case the PLCSs have finite attractor. The results also extend to other variants of PLCS. For example, according to the semantics considered here, when $\mathtt{c!m}$ is performed at a state $(\mathtt{s}, \mathtt{w})$ the message $\mathtt{m}$ is appended to the contents of the channel $\mathtt{c}$ in $\mathtt{w}$ and then some messages can be lost from $\mathtt{c}$ and other channels. An alternative semantics would allow that only messages from $\mathtt{c}$ can be lost. Our algorithm can be easily modified to handle this alternative semantics.

The results can be easily generalized to PLCS with different sources of unreliability [9], such as duplication, corruption, and insertion combined with lossiness. For example, an LCS $\mathcal{L}$ with *duplication errors* is of the same form $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T})$ as an LCS. We define the behaviour of $\mathcal{L}$ as follows. For $a \in \mathtt{M}$, we use a^n to denote the concatenation of n copies of a. For $x = a_1 a_2 \ldots a_n$ with $x \in \mathtt{M}^*$, we define *duplicate* (x) to be the set $\left\{b_1 b_2 \ldots b_n : \text{ either } b_i = a_i \text{ or } b_i = a_i^2 \text{ for each } i : 1 \leqslant i \leqslant n\right\}$. In other words, we get each member of *duplicate* (x) by duplicating some of the elements of x. We extend the definition of *duplicate* to $\mathtt{S} \times (\mathtt{C} \mapsto \mathtt{M}^*)$ in a similar manner to Section 6. The transition relation of an LCS $\mathcal{L}$ with duplication errors is the enlargement of that of the corresponding standard LCS in the sense that: If $(\mathtt{s}_1, \mathtt{w}_1) \longrightarrow (\mathtt{s}_2, \mathtt{w}_2)$ according to the definition of Section 6 then $(\mathtt{s}_1, \mathtt{w}_1) \longrightarrow \left(\mathtt{s}_2', \mathtt{w}_2'\right)$ for each $\left(\mathtt{s}_2', \mathtt{w}_2'\right) \in$ *duplicate* $(\mathtt{s}_2, \mathtt{w}_2)$. In [9], it is shown that the reachability problem is decidable for LCS with duplication errors. A PLCS with *duplication errors* is of the form $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T}, \lambda, w, \lambda_D)$, where $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T}, \lambda, w)$ is a PLCS, and $\lambda_D \in [0,1]$. The value of λ_D represents the probability by which any given message is duplicated inside the channels.

The Markov chain induced by a PLCS with duplication errors, is defined in a natural way. In [5], it is shown that qualitative probabilistic model checking is decidable for PLCS with duplication errors when $\lambda_D < \lambda$. In this case the corresponding Markov chains have finite attractors. The technique of the current paper can show that there exists an efficient algorithm for the quantitative model checking for PLCS with duplication errors when $\lambda_D < \lambda$.

10.4. Open problems and related results

One of the directions for further research is approximate other important parameters of the Markov chains such as: mean return time, mean hitting time, etc.

Another important direction is to extend the verification of qualitative properties of Markov decision processes [7,21] to the verification of quantitative properties.

Finally, let us mention the following decision problem:

Problem A: *Deciding an upper bound for the probability of* ‘ *reachability*
Instance: A M.C. M, its states s_1, s_2 and a rational h.
Question: Decide whether h is greater than the probability that s_2 is reached from s_1.

It is straightforward to construct an algorithm for the quantitative reachability problem from the algorithm for the above decision problem. As the first approximation, we set $l_0 = 0$, and $u_0 = 1$. Further approximations are obtained by *binary search*: If the probability of reaching s_2 from s_1 is greater than $\frac{1}{2}(l_i + u_i)$, define $l_{i+1} = l_i$ and $u_{i+1} = \frac{1}{2}(l_i + u_i)$; otherwise define $l_{i+1} = \frac{1}{2}(l_i + u_i)$ and $u_{i+1} = u_i$. It is clear that for each $i \geq 0$, the probability of reaching s_2 from s_1 is in the interval $[l_i, u_i]$. Hence for $n = \ulcorner - log_2\theta \urcorner$, the rational number l_n approximates up to θ the probability of reaching s_2 from s_1. Therefore, the quantitative reachability problem is computable for every class C of Markov chains for which the above problem is decidable.

Recently, the reachability and model checking problems for the class of Probabilistic Pushdown Automata (PPDA) was considered in [11]. It was shown there that for a PPDA M, its states s_1, s_2 the probability that s_2 is reached from s_1 is a root of a polynom effectively computable from M, s_1 and s_2. As a consequence the authors obtained that for PPDA the above decision problem and the quantitative reachability problems are decidable. These results were extended to the quantitative model checking problem and to the model checking properties definable by Probabilistic CTL.

The decidability of Problem A for the class of PLCSs remains open.

Acknowledgments

We thank Philippe Schnoebelen for pointing out that the path enumeration scheme terminates over Markov chains with finite attractors. We thank Parosh Abdulla, Daniéle Beauquier, Philippe Schnoebelen, and Anatol Slissenko for fruitful discussions and their useful comments.

Appendix A. Proof of Lemma 14

First, we recall a general result for hitting probabilities (see Theorem 1.3.2 [17]).

Theorem 27. *Let $M = (S,P)$ be a Markov chain, Let $V \subseteq S$ and let h_s^V be the probability to reach V from s. Then the vector of probabilities $h^V = \left(h_s^V : s \in S\right)$ is the minimal non-negative solution of the system of linear equations*

$$\begin{cases} h_s^V = 1 & \text{s} \in \text{V}, \\ h_s^V = \sum_{s' \in S} P(s,s')h_{s'}^V & \text{otherwise.} \end{cases}$$

Lemma 28. *Consider Markov chains $M = (S,P)$ and $M' = \left(S',P'\right)$ such that*

(1) (a) S' is the set $Nat \quad = \{0,1,2,3,\ldots\}$ of natural numbers, and
(b) $P'(i,i+1) = \nu_i$, and $P'(i,i) = \gamma_i$ for $i \geqslant 0$, and $P'(i,i-1) = \mu_i = 1 - \nu_i - \gamma_i$ for $i \geqslant 1$.
(2) (a) S is the union of disjoint sets S_i ($i \in Nat \quad$).
(b) If $s \in S_i$, $s' \in S_j$, and $P(s,s') > 0$, then $j \leqslant i+1$.
(3) For every i
(a) $\nu_i \geqslant \sup_{s \in S_i}\left(\sum_{s' \in S_{i+1}} P(s,s')\right)$ and
(b) either $\gamma_i \geqslant \sup_{s \in S_i}\left(\sum_{s' \in S_i} P(s,s')\right)$ or $\mu_i \leqslant \inf_{s \in s_i}\left(\sum_{s' \in \bigcup_{k=0}^{i-1} S_k} P(s,s')\right)$

Then for every $i < j$ and $s \in S_i$ the probability to reach S_j from s in M is less than or equal to the probability to reach j from i in M'.

Proof. Fix j. Let $h = (h_s : s \in S)$ be the probability of reaching $\bigcup_{k \geqslant j} S_k$ in M. Let $\hat{h} = \left(\hat{h}_i : i \in Nat\right)$ be the probability of reaching $\{k : k \geqslant j\}$ in M'. Observe that if $k \leqslant i \leqslant j$ then $\hat{h}_k \leqslant \hat{h}_i$.

Define $H_s = \hat{h}_i$ for $s \in S_i$. Let H be the vector $(H_s : s \in S)$. Observe that H is a non-negative vector and $H_s = 1$ for $s \in \bigcup_{k \geqslant j} S_k$. We show that $H \geqslant P \cdot H$.

For $s \in S_i$

$$\begin{aligned} (P \cdot H)_s &= \textstyle\sum_{s' \in S} P(s,s')H_{s'} \\ &= \textstyle\sum_{s' \in S_{i+1}} P(s,s')H_{s'} + \sum_{s' \in S_i} P(s,s')H_{s'} + \sum_{s' \in \bigcup_{k<i} S_k} P(s,s')H_{s'} \\ &= \left(\textstyle\sum_{s' \in S_{i+1}} P(s,s')\right)\hat{h}_{i+1} + \left(\sum_{s' \in S_i} P(s,s')\right)\hat{h}_i + \sum_{k<i}\left(\sum_{s' \in S_k} P(s,s')\right)\hat{h}_k \\ &\leqslant \left(\textstyle\sum_{s' \in S_{i+1}} P(s,s')\right)\hat{h}_{i+1} + \left(\sum_{s' \in S_i} P(s,s')\right)\hat{h}_i + \left(\sum_{k<i}\sum_{s' \in S_k} P(s,s')\right)\hat{h}_{i-1} \\ &= \left(\textstyle\sum_{s' \in S_{i+1}} P(s,s')\right)\hat{h}_{i+1} + \left(\sum_{s' \in S_i} P(s,s')\right)\hat{h}_i + \left(\sum_{s' \in \bigcup_{k<i} S_k} P(s,s')\right)\hat{h}_{i-1} \\ &= \nu_s\hat{h}_{i+1} + \gamma_s\hat{h}_i + \mu_s\hat{h}_{i-1}, \end{aligned} \tag{A.1}$$

where $\nu_s = \left(\sum_{s' \in S_{i+1}} P(s,s')\right)$, $\gamma_s = \left(\sum_{s' \in S_i} P(s,s')\right)$ and $\mu_s = \left(\sum_{s' \in \bigcup_{k<i} S_k} P(s,s')\right)$. In (A.1) the $\leqslant$ step holds, because $\hat{h}_k$ is increasing and P is non-negative.

$$H_s = \hat{h}_i = \nu_i\hat{h}_{i+1} + \gamma_i\hat{h}_i + \mu_i\hat{h}_{i-1}, \tag{A.2}$$

$$\nu_i + \gamma_i + \mu_i = 1 \text{ and } \nu_s + \gamma_s + \mu_s = 1. \tag{A.3}$$

Therefore, from Eqs. (A.1)–(A.3) we obtain

$$H_s - (P \cdot H)_s \geqslant (\nu_i - \nu_s)(\hat{h}_{i+1} - \hat{h}_{i-1}) + (\gamma_i - \gamma_s)(\hat{h}_i - \hat{h}_{i-1}) \tag{A.4}$$

and

$$H_s - (P \cdot H)_s \geqslant (\nu_i - \nu_s)(\hat{h}_{i+1} - \hat{h}_i) + (\mu_s - \mu_i)(\hat{h}_i - \hat{h}_{i-1}). \tag{A.5}$$

The sequence $\hat{h}_k$ is increasing; therefore, from the assumptions 3(a) and 3(b) of Lemma 28 and Eqs. (A.4)–(A.5) we obtain $H_s - (P \cdot H)_s \geqslant 0$. Hence and

$$H \geqslant P \cdot H. \tag{A.6}$$

Observe that $H_s = 1 = h_s$ for $s \in V = \bigcup_{k \geqslant j} S_k$ and are non-negative. Define an operator α that maps S-tuple of reals to S-tuple of reals as follows:

$$\alpha(u)_s = \begin{cases} 1 & s \in \bigcup_{k \geqslant j} S_k, \\ \sum_{s' \in S} P(s,s') u_{s'} & \text{otherwise.} \end{cases}$$

The matrix P is non-negative. Therefore, α maps the non-negative vectors to the non-negative vectors and is monotonic (increasing) on the set of non-negative vectors. By Eq. (A.6) we obtain that $H \geqslant \alpha(H)$. Therefore, by monotonicity of α, we obtain that the sequence $\alpha^n(H)$ is decreasing and non-negative. Hence, it converges to a non-negative fixed point of α which is $\leqslant H$. From Theorem 27 and the definition of h it follows that h is the minimal non-negative fixed point of α. Hence, $h \leqslant H$. From the definition of H we conclude that $h_s \leqslant \hat{h}_i$ for $s \in S_i$. Hence, the probability to reach S_j from $s \in S_i$ in M is less than or equal to the probability to reach j from i in M'. □

The next theorem is a standard result on observing a Markov process at some subset of its state-space (see Example 1.4.4 in [17]).

Theorem 29. *Assume that $M = (S,P)$ is the transition matrix of a Markov process $(X_n)_{n \geqslant 0}$. Let $U \subset S$. Then the process obtained from $(X_n)_{n \geqslant 0}$ by observation only at states U is a Markov process and it has the transition matrix $P'(s,s') = h_s^{s'}$ for $s,s' \in U$ and where for $s' \in U$, the vector $\left(h_s^{s'} : s \in S\right)$ is the minimal non-negative solution to*

$$h_s^{s'} = P(s,s') + \sum_{s'' \notin U} P(s,s'') h_{s''}^{s'}.$$

As a consequence we derive

Lemma 30. *Assume that $M = (S,P)$ is the transition matrix of a Markov process $(X_n)_{n \geqslant 0}$.*

(1) S is the union of disjoint sets U_0, U_1 and U_2.
(2) If $s \in U_i, s' \in U_j$, and $P(s,s') > 0$, then $j \leqslant i+1$.
(3) $U_0 = C \cup R$ and

- *For every state $s \in R$, only states in R are reachable from s.*
- *For every state $s \in U_1$ there is a finite path to R with probability $> \delta$ which is inside $C \cup R$.*

Let $U = R \cup U_1 \cup U_2$. Then the process obtained from $(X_n)_{n \geqslant 0}$ by observation only at states U is a Markov process and it has the transition matrix $\widehat{P}$ which satisfies:

$$\widehat{P}(s,s') = \begin{cases} P(s,s') \ ifs' \in U_2, \\ P(s,s') \ s,s' \in R \end{cases}$$

$$\widehat{P}(s,s') \geqslant \delta \ if\, s \in U_1 \ and\, s' \in R$$

Proof. Immediately from Theorem 29. □

Now we are ready to prove Lemma 14. Let M be a Markov chain which satisfies the assumption of Lemma 14. Let $U_0 = S_0 = C \cup R$, $U_1 = S_1$ and $U_2 = \bigcup_{k \geqslant 2} S_k$. Let $\widehat{M} = \big(U, \widehat{P}\big)$ be the chain which is obtained by observing M at $U = R \cup U_1 \cup U_2 = R \cup \bigcup_{k \geqslant 1} S_k$.

$$\begin{array}{c}\text{For every } s \in U \text{ and } Z \subseteq U \\ \text{the probability to reach } Z \text{ from } s \text{ in } M \text{ is the same as} \\ \text{the probability to reach } Z \text{ from } s \text{ in } \widehat{M}.\end{array} \tag{A.7}$$

By Lemma 30, for $i \geqslant 1$ and $s \in S_i$:

$$\left(\sum_{s' \in S_{i+1}} P(s,s')\right) = \left(\sum_{s' \in S_{i+1}} \widehat{P}(s,s')\right) \tag{A.8}$$

for $i \geqslant 2$ and $s \in S_i$:

$$\left(\sum_{s' \in S_i} P(s,s')\right) = \left(\sum_{s' \in S_i} \widehat{P}(s,s')\right) \tag{A.9}$$

and for $s \in S_1$

$$\delta \leqslant \left(\sum_{s' \in R} \widehat{P}(s,s')\right) \tag{A.10}$$

From Eqs. (A.8) and (A.9) and by the assumption (4) of Lemma 14 we obtain

$$\nu_i = \sup_{s \in S_i} \left(\sum_{s' \in S_{i+1}} \widehat{P}(s,s')\right) \text{ and } \gamma_i = \sup_{s \in S_i} \left(\sum_{s' \in S_i} \widehat{P}(s,s')\right). \tag{A.11}$$

Let $M' = \big(Nat \quad , P'\big)$ be such that

- $P'(i,i+1) = \nu_i$, and $P'(i,i) = \gamma_i$ and $P'(i,i-1) = \mu_i = 1 - \nu_i - \gamma_i$ for $i \geqslant 2$.

- $P'(0,0) = 1$
- $P'(1,0) = \delta$
- $P'(1,2) = \nu_1$ and
- $P'(1,1) = 1 - P'(1,0) - P'(1,2)$.

From Lemma 28, the definition of M' and Eqs. (A.10) and (A.11) we derive that

$$\begin{array}{c}\text{For every } s \in S_1 \text{ the probability to reach } S_n \text{ in } \widehat{M}\\ \text{is less than or equal to}\\ \text{the probability to reach } n \text{ from 1 in } M'.\end{array} \tag{A.12}$$

By assumption (4) of Lemma 14 we have $\nu_i + \gamma_i < \alpha < 0.5$ for each $i \geqslant 2$. Hence, $\mu_i > 1 - \alpha > 0.5$ for $i \geqslant 2$, and therefore, by Lemma 13, for $n \geqslant N_0 = N(\delta, 1 - \alpha, \theta)$

$$\text{the probability to reach } n \text{ from 1 in } M' \text{ is } < \theta. \tag{A.13}$$

Hence, by (A.7), (A.12) and (A.13) we obtain

$$\begin{array}{c}\text{for every } s \in S_1 \text{ the probability of reaching } S_{N_0} \text{ from } s \text{ in } M\\ \text{is less than } \theta.\end{array} \tag{A.14}$$

Finally observe that for every $s \in S_0$ a computation from s which reaches S_{N_0} should pass through a state in S_1. This observation together with (A.14) implies

$$\begin{array}{c}\text{for every } s \in S_0 \text{ the probability of reaching } S_{N_0} \text{ from } s \text{ in } M\\ \text{is less than } \theta.\end{array} \tag{A.15}$$

The conclusion of Lemma 14 follows from (A.14) and (A.15).

References

[1] P.A. Abdulla, C. Baier, P. Iyer, B. Jonsson, Reasoning about probabilistic lossy channel systems, in: Proceedings of the CONCUR 2000, Lecture Notes in Computer Science, vol. 1877, 2000.

[2] P.A. Abdulla, N. Bertrand, A. Rabinovich, Ph. Schnoebelen, Verification of probabilistic systems with faulty communication, Information and Computation 202 (2) (2005) 105–228.

[3] P.A. Abdulla, B. Jonsson, Undecidable verification problems for programs with unreliable channels, Information and Computation 130 (1) (1996) 71–90.

[4] P.A. Abdulla, B. Jonsson, Verifying programs with unreliable channels, Information and Computation 127 (2) (1996) 91–101.

[5] P.A. Abdulla, A. Rabinovich, Verification of probabilistic systems with faulty communication, in: FOSSACS'03, Lecture Notes in Computer Science, vol. 2620, Springer, Berlin, 2003, pp. 39–53.

[6] C. Baier, B. Engelen, Establishing qualitative properties for probabilistic lossy channel systems, in: ARTS'99, Lecture Notes in Computer Science, vol. 1601, Springer, Berlin, 1999, pp. 34–52.

[7] N. Bertrand, Ph. Schnoebelen, Model checking lossy channels systems is probably decidable, in: FOSSACS'03, Lecture Notes in Computer Science, vol. 2620, Springer, Berlin, 2003, pp. 120–135.

[8] D. Brand, P. Zafiropulo, On communicating finite-state machines, Journal of the ACM 2 (5) (1983) 323–342.

[9] Gérard Cécé, Alain Finkel, S. Purushothaman Iyer, Unreliable channels are easier to verify than perfect channels, Information and Computation 124 (1) (1996) 20–31.

[10] C. Courcoubetis, M. Yannakakis, The complexity of probabilistic verification, Journal of the ACM 42 (1995) 857–907.

[11] J. Esparza, A. Kucera, R. Mayr, Model checking probabilistic pushdown automata, LICS04 (2004) 12–21.

[12] A. Finkel, Ph. Schnoebelen, Well structured transition systems everywhere!, Theoretical Computer Science 256 (1-2) (2001) 63–92.

[13] G. Higman, Ordering by divisibility in abstract algebras, Proceedings of the London Mathematical Society 2 (1952) 326–336.

[14] P. Iyer, M. Narasimha, Probabilistic lossy channel systems, in: Proceedings of the TAPSOFT '97, Lecture Notes in Computer Science, vol. 1214, 1997, pp. 667–681.

[15] S. Karlin, A First Course in Stochastic Processes, Academic Press, New York, 1966.

[16] J. Kemeny, J. Snell, A. Knapp, Denumerable Markov Chains, D Van Nostad Co., 1966.

[17] J.R. Norris, Markov Chains, Cambridge University Press, Cambridge, 1997.

[18] A. Rabinovich, Quantitative analysis of probabilistic lossy channel systems, in: ICALP03, Lecture Notes in Computer Sciences, vol. 2719, 2003, pp. 1008–1021.

[19] Ph. Schnoebelen, Verifying lossy channel systems has nonprimitive recursive complexity, Information Processing Letters 83 (5) (2002) 251–261.

[20] Ph. Schnoebelen, Personal communication, Jan. 2003.

[21] Ph. Schnoebelen, The verification of probabilistic lossy channel systems, in: Lecture Notes in Computer Sciences, vol. 2925, Springer, Berlin, 2004, pp. 445–465.

[22] W. Thomas, Automata on infinite objects, in: Handbook of Theoretical Computer Science, Volume B: Formal Methods and Semantics, 1990, pp. 133–192.

[23] M. Vardi, Probabilistic linear-time model checking: an overview of the automata-theoretic approach, in: Proceedings of the 5th Workshop on Formal Methods for Real-Time and Probabilsitic Systems, Lecture Notes in Computer Science, vol. 1601, 1999, pp. 265–276.