Quantitative Analysis of Probabilistic Lossy Channel Systems

Alexander Rabinovich

School of Computer Science, Tel Aviv University, Israel
rabinoa@post.tau.ac.il

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)* [3] 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 [5, 1, 4, 6] 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 $\theta > 0$.
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 $r + \theta$.

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 state transition system. The fact that Turing machines can be simulated by CFSMs [7] 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)* [3] 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. Thought an LCS defines in a natural way an infinite state transition system, it has been shown that the reachability problem for LCS is decidable [3], while progress properties are undecidable [2].

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 [10, 5, 1, 4, 6] have been made to study *probabilistic Lossy Channel Systems (PLCS)* which introduce randomization into the behaviour of LCS. The works in [10, 5, 1, 4, 6] 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 [5] 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 [4, 6] another semantics for PLCS was considered which is more realistic than that in [5, 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 [5, 1]).

For this model, the decidability of both the qualitative reachability and the qualitative model-checking problems was independently established in [4, 6].

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 $r + \theta$.

In order 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 subchain $M_n = (S_n, P_n)$ 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.

In order 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.

Outline In the next two sections we give basics of transition systems and countable Markov chains respectively. In Sect. 4 the quantitative probabilistic reachability problem over countable Markov chains is considered. In Sections 5 and 6 the semantics of lossy channel systems and probabilistic lossy channel systems are described. In Sect. 7 the algorithm for the quantitative probabilistic reachability problem over PLCS is presented and its complexity is analyzed. In Sect. 8 we generalize our results to the verification of the properties definable by ω-behavior of finite state automata.

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 $\stackrel{*}{\longrightarrow}$ to denote the reflexive transitive closure of $\longrightarrow$. We say that s_2 is *reachable* from s_1 if $s_1 \stackrel{*}{\longrightarrow} s_2$. For sets $Q_1, Q_2 \subseteq S$, we say that Q_2 is *reachable* from Q_1, denoted $Q_1 \stackrel{*}{\longrightarrow} Q_2$, if there are $s_1 \in Q_1$ and $s_2 \in Q_2$ with $s_1 \stackrel{*}{\longrightarrow} s_2$. 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 $i \geq 0$. We use $\pi(i)$ to denote s_i, and write $s \in \pi$ to denote that there is an $i \geq 0$ with $\pi(i) = s$. For states s and s', we say that π *leads* from s to s', written, $s \stackrel{\pi}{\longrightarrow} s'$, if $s = s_0$ and $s' \in \pi$.

For $Q \subseteq S$, we define the *graph* of Q, denoted $Graph(Q)$, to be the transition system $\left(Q, \longrightarrow'\right)$ where $s_1 \longrightarrow' s_2$ iff $s_1 \stackrel{*}{\longrightarrow} s_2$. A *strongly connected component (SCC)* in T is a maximal set $C \subseteq S$ such that $s_1 \stackrel{*}{\longrightarrow} s_2$ 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 $C \stackrel{*}{\longrightarrow} C_1$. 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 infinite) 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 $(S, \longrightarrow)$ 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 [12]): $\Omega = sS^\omega$ 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 $Prob(D_u) = \prod_{i=0,\ldots,n-1} P(s_i, s_{i+1})$ where $u = s_0 s_1 \ldots s_n$; 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.

Consider a set $Q \subseteq S$ of states and a path π. We say that π *reaches* Q if there is an $i \geq 0$ 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 $Prob\{\pi : \pi \text{ is a path from } s \text{ and } \pi \text{ repeatedly reaches } s\} = 1$. We say that a state s is *transient* if $Prob\{\pi : \pi \text{ is a path from } s \text{ and } \pi \text{ repeatedly reaches } s\} = 0$.

The next theorem summarizes standard properties of countable Markov chains [13].

Theorem 3.1. *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 $Prob\{\pi : \pi \text{ starts at } s_1 \text{ and repeatedly reaches every state of } C\} = 1$. 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. We introduce a central concept which we use in our solution for the probabilistic reachability problem, namely that of *attractors*.

Definition 3.2 (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 that we will reach the attractor with probability one. It is clear that an attractor has a state in every recurrent class. The Lemma below follows from Theorem 3.1 and describes properties of the BSCCs of the graph of a finite attractor A.

Lemma 3.3. *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 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 $\text{Comp}_n(s_1)$ be the set of all the computations of length n in M from s_1. Partition $\text{Comp}_n(s_1)$ into three sets:

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

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

$$\text{Endecided}_n(s_1, s_2) = \text{Comp}_n(s_1) \setminus \text{Reach}_n(s_1, s_2) \setminus \text{Escape}_n(s_1, s_2)$$

All the computations in $\text{Reach}_n(s_1, s_2)$ reach s_2, and no computation in $\text{Escape}_n(s_1, s_2)$ extends to a computation that reaches s_2. Note that $Prob(\text{Comp}_n(s_1)) = 1$. Let $p_n^+ = Prob(\text{Reach}_n(s_1, s_2))$, $p_n^- = Prob(\text{Escape}_n(s_1, s_2))$ and $p_n^? = Prob(\text{Endecided}_n(s_1, s_2))$. Observe that p_n^+ and p_n^- are increasing sequences, while $p_n^?$ is decreasing and

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

Path Enumeration (PE) scheme for approximation $Prob_M(s_1 \xrightarrow{*} s_2)$ 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 the 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

In the above problem, we do not assume that M is finite. Hence, these are not instances of an algorithmic problem. In Sect. 7 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 $\text{Escape}_n(s_1, s_2)$, $\text{Endecided}_n(s_1, s_2)$ and $p_n^?$ can be computed. Hence, in this case the scheme can be implemented. Observe that PE scheme terminates only if $\lim p_n^? < \theta$. Therefore,

Lemma 4.1. *If* $\lim p_n^? = 0$ *then 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 [15] as well.

Lemma 4.2. *If M has a finite attractor then* $\lim p_n^? = 0$.

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

Definition 4.3. *A Markov chain $M = (S, P)$ has δ-reachability property for $\delta > 0$ if* $\forall s_1, s_2 \in S(\ s_2$ *is reachable from* $s_1\ ) \Rightarrow Prob_M(s_1 \xrightarrow{*} s_2) > \delta$.

Lemma 4.4. *If M has δ-reachability property then* $\lim p_n^? = 0$.

Theorem 4.5. *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 PE scheme was suggested in [10] for the following decision problem.

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

It was claimed in [10] 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 [10]. However, assertion (a) was incorrect. Also, the Markov chains assigned to PLCSs in [10] 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 Sect. 6 (or considered in [5, 4, 6]) which have finite attractor property.

PE scheme is conceptually very simple, however, no information about the number of iterations before it terminates can be extracted from Theorem 4.5. For finite state M. C. standard algebraic methods allow to find the exact value of $Prob_M(s_1 \xrightarrow{*} s_2)$ in polynomial time; however, in this case 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 M.C. 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

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

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

Lemma 4.6. *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 4.7. *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 $D\} \leq \theta$ or (2) $\forall s \in D.\ Prob\{\pi : \pi$ starts at s and reaches $s_2\} \leq \theta$. Let $M' = M^{D,d}$. Then $Prob_{M'}(s_1 \xrightarrow{*} s_2) \leq Prob_M(s_1 \xrightarrow{*} s_2) \leq Prob_{M'}(s_1 \xrightarrow{*} s_2) + \theta$.*

In order to find D which satisfies the assumption of Lemma 4.7 we provide a reduction to a one-dimensional random walk. The following lemma is easily derived from standard properties of one dimensional random walks.

Lemma 4.8. *Let $M = (S, P)$ be a Markov chain where $S = \{0, 1, 2, 3, \ldots\}$, and*
- *$P(0, 0) = 1$.*
- *$P(i, i+1) = \nu_i$, $P(i, i-1) = \mu_i$, and $P(i, i) = 1 - \mu_i - \nu_i$, for $i \geq 1$.*
- *There is $q > 0.5$ such that $\mu_i > q$ for all $i \geq 2$.*

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

The main technical lemma for the correctness and the complexity analysis of the algorithm presented in Sect. 7 is a generalization of Lemma 4.8.

Lemma 4.9 (Main Lemma). *Consider a Markov chain $M = (S, P)$ such that*
1. *S is the union of disjoint sets S_i ($i \in \mathcal{N}$).*
2. *If $s \in S_i$, $s' \in S_j$, and $P(s, s') > 0$, then $j \leq i + 1$.*
3. *$S_0 = C \cup R$ and*
   - *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 $C \cup R$.*
4. *There is $\alpha < \frac{1}{2}$ such that $\nu_i + \gamma_i < \alpha$ for each $i \geq 2$, 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 4.8. Then, for every $s \in S_0 \cup S_1$ the probability of reaching $\bigcup_{n \geq N_0} S_n$ from s is less than θ.

Hence,

Lemma 4.10. *Let M, S_i and N_0 be as in Lemma 4.9 and assume that $s_1, s_2 \in S_0$. Let $U = \bigcup_{n \geq N_0} S_n$. Let $M' = M^{U,u}$. Then $Prob_{M'}(s_1 \xrightarrow{*} s_2) \leq Prob_M(s_1 \xrightarrow{*} u) \leq Prob_{M'}(s_1 \xrightarrow{*} s_2) + \theta$.*

5 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 i^{th} element of x where $i : 1 \leq i \leq |x|$. 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 $(S, \longrightarrow)$, 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 $\mathtt{w}_1$ is the result of removing $\mathtt{m}$ from the head of $\mathtt{w}_2(\mathtt{c})$.
- Furthermore, if $(\mathtt{s}_1, \mathtt{w}_1) \longrightarrow (\mathtt{s}_2, \mathtt{w}_2)$ according to one of the previous two rules, then $(\mathtt{s}_1, \mathtt{w}_1) \longrightarrow (\mathtt{s}'_2, \mathtt{w}'_2)$ for each $(\mathtt{s}'_2, \mathtt{w}'_2) \preceq (\mathtt{s}_2, \mathtt{w}_2)$.

In the first two cases we define $\mathtt{t}(\mathtt{s}_1, \mathtt{w}_1) = (\mathtt{s}_2, \mathtt{w}_2)$.

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 $\{s \ : \ \exists s' \in Q.\ s' \preceq s\}$.

Theorems in [3, 9] imply the following decidability results for LCS:

Lemma 5.1. *(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.*

6 Probabilistic Lossy Channel Systems

We introduce a probabilistic behaviour into LCS obtaining *Probabilistic Lossy Channel Systems*. This semantics was considered in [4, 6] and differs from that in [10, 5]

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 $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T})$. 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 [4] 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 only probabilistic properties of PLCSs which we use are summarized in the next two lemmas from [4].

Lemma 6.1. *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 $\leq (1-\lambda)^{m+1}$. The probability of the transitions from s to the set of states with m messages is $\leq m\lambda(1-\lambda)^m$.*

Lemma 6.2. *For each λ, 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 Sect. 7

Lemma 6.3. *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 i-th 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 $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T})$.*

7 Algorithm for Approximating the Probability of Reachability

Lemmas 6.2, 5.1(1) and Theorem 4.5 imply that there is an algorithm based on PE scheme for the quantitative probabilistic reachability problem. However, no information about the complexity of this algorithm can be extracted from Theorem 4.5. 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 $B_{\leq n}$ 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 $B_{\leq n}$, and then for each recurrence class D_i of M collapse the set of D_i states in $B_{\leq n}$ 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 4.9 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 $r \leq Prob_M(s_1 \xrightarrow{*} s_2) \leq r + \theta$.
Let A be the (finite) set of all states with 0 messages. A is an attractor by Lemma 6.2. By Lemma 5.1(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 $1 \leq i \leq k$ by Lemma 3.3 and by Lemma 6.3 we can compute finite sets of states V_i such that

$$\forall s \in S (s \text{ is in the recurrent class of } C_i) \text{ iff } s \text{ is not in the upward closure of } V_i \quad (2)$$

Hence, we can check whether s_1 (or any other state s) is in the i-th 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 3.1(4-5).)

Below we consider the case when s_1 is not recurrent. By Lemma 5.1(3) we can find 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 $\forall n.\ m \leq n \rightarrow (1-\lambda)^n(1-\lambda+n\lambda) < \frac{1}{3}$ i.e., the probability to move from every state with $n \geq m$ messages to the set of states with at least n messages is less than $\frac{1}{3}$ (by Lemma 6.1). Let $h = max(l, m)$+1.
Notations: Below we denote by B_i (respectively by $B_{\leq i}$) the set of states with i (respectively with at most i) messages.

For every state $s \in B_{\leq h}$ there is a 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 $B_{\leq l} \subseteq B_{\leq h}$ to a BSCC of $Graph(A)$. Let $\delta_s = Prob(\pi_s) > 0$ and let $0 < \delta = min(\delta_s : s \in B_{\leq h})$. 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 4.9 are satisfied. Let $N_0 = N(\delta, \frac{2}{3}, \theta)$, where N is the function from Lemma 4.8 and let $n = h + N_0$. Note that n depends linearly on $\ln(\frac{1}{\theta})$. By Lemma 4.9 the probability to reach from s_1 the set $U = \bigcup_{n \geq N_0} S_n$ of transient states with $\geq n$ messages is at most θ. Therefore, by Lemma 4.10 we derive that $Prob_{M'}(s_1 \xrightarrow{*} s_2) \leq Prob_M(s_1 \xrightarrow{*} s_2) \leq Prob_{M'}(s_1 \xrightarrow{*} s_2) + \theta$ 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 $|B_{\leq n}|$ 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 $O(|B_{\leq n}|^2)$ and by standard algebraic methods we can compute $Prob_{\widehat{M}}(s_1 \xrightarrow{*} s_2)$ in time $O(|B_{\leq n}|^3)$. Since n depends linearly on $\ln(\frac{1}{\theta})$, it follows that $|B_{\leq n}|$ 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 4.6 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 $(i = 1, \ldots, k)$ be the states with $\leq n$ messages, which are in the i-th recurrent class. (These sets can be computed by By Eq. (2) in time $O(|B_{\leq n}|)$. Let D be $B_{<n} \setminus \cup D_i$.

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

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

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 $i : 1 \leq i \leq k$.

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 (a 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 [14] 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.

8 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.

We consider an extension of PLCS by adding a labeling function. A *state labeled* PLCS is an PLCS together with a finite alphabet Σ and a labeling function lab from the local states to Σ. We lift the labeling to the global states: the label of every global state is the same as the label of its local state component. Similarly, with a computation $s_0, s_1, \ldots$ we associate the ω-string $lab(s_1)lab(s_2)\ldots$ over the alphabet Σ.

The next lemma reduces the Quantitative Probabilistic Model-checking Problem to the Quantitative Probabilistic Reachability Problem for PLCSs (see Sect. 1).

Lemma 8.1. *There exists an algorithm which for a state labeled PLCS $\mathcal{L}$, a global state s in $\mathcal{L}$ and a finite state ω-automaton $\mathcal{A}$ produces a PLCS $\mathcal{L}'$, a global state s' in $\mathcal{L}'$ and a set $C_1, C_2, \ldots, C_p$ of BSCC for a finite attractor of $\mathcal{L}'$ such that the following are equivalent:*

1. *The probability that a computation of $\mathcal{L}$ that starts at s is accepted by $\mathcal{A}$ is r.*
2. *The probability that a computation of $\mathcal{L}'$ that starts at s' reaches $\bigcup_{i=1}^{p} C_i$ is r.*

Theorem 8.2. *The Quantitative Probabilistic Model-checking problem can be solved in time $g(\mathcal{L}, \mathcal{A}, s_1, s_2)) \times \frac{1}{\theta^3}$. .*

Proof. First apply the algorithm from Lemma 8.1. Observe that for $i \neq j$ no path reaches both C_i and C_j. For $i = 1, \ldots, p$ choose a state $s_i \in C_i$. Lemma 3.1(4) and Lemma 3.3 imply that the probability to reach s_i is the same as the probability to reach C_i. By the algorithm of Sect. 7 compute r_i which approximates up to $\frac{\theta}{p}$ the probability to reach s_i from s' in $\mathcal{L}'$. From Lemma 8.1 it follows that $r_1 + r_2 + \cdots + r_p$ approximates up to θ the probability that a computation of $\mathcal{L}$ that starts at s is accepted by $\mathcal{A}$.

Acknowledgments We would like to thank Philippe Schnoebelen and an anonymous referee for pointing out that the path enumeration scheme terminates over Markov chains with finite attractors. We thank Parosh Abdulla, Danièle Beauquier, Philippe Schnoebelen and Anatol Slissenko for fruitful discussions and their useful comments.

References

1. P. A. Abdulla, C. Baier, P. Iyer, and B. Jonsson. Reasoning about probabilistic lossy channel systems. In *Proc. CONCUR 2000*, volume 1877 of *Lecture Notes in Computer Science*, 2000.
2. P. A. Abdulla and B. Jonsson. Undecidable verification problems for programs with unreliable channels. *Information and Computation*, 130(1):71–90, 1996.
3. P. A. Abdulla and B. Jonsson. Verifying programs with unreliable channels. *Information and Computation*, 127(2):91–101, 1996.
4. P. A. Abdulla and A. Rabinovich. Verification of probabilistic systems with faulty communication, 2003. In FOSSACS'03, volume 2620 of *LNCS*, pages 39–53. Springer Verlag, 2003.
5. C. Baier and B. Engelen. Establishing qualitative properties for probabilistic lossy channel systems. In *ARTS'99*, volume 1601 of *LNCS*, pages 34–52. Springer Verlag, 1999.
6. N. Bertrand and Ph. Schnoebelen. Model checking lossy channels systems is probably decidable. In FOSSACS'03, volume 2620 of *LNCS*, pages 120–135 Springer Verlag, 2003.
7. D. Brand and P. Zafiropulo. On communicating finite-state machines. *Journal of the ACM*, 2(5):323–342, 1983.
8. Gérard Cécé, Alain Finkel, and S. Purushothaman Iyer. Unreliable channels are easier to verify than perfect channels. *Information and Computation*, 124(1):20–31, 1996.
9. A. Finkel and Ph. Schnoebelen. Well structured transition systems everywhere!. *Theoretical Computer Science,* 256(1-2):63-92, 2001.
10. P. Iyer and M. Narasimha. Probabilistic Lossy Channel Systems. In Proc of *TAPSOFT '97* LNCS 1214, 667-681 1997.
11. S. Karlin. *A First Course in Stochastic Processes*. Academic Press, 1966.
12. J. Kemeny, J. Snell, and A. Knapp. *Denumerable Markov Chains*. D Van Nostad Co., 1966.
13. J. R. Norris. *Markov Chains*. Cambridge University Press, 1997.
14. Ph. Schnoebelen. Verifying lossy channel systems has nonprimitive recursive complexity. *Information Processing Letters,* 83(5):251-261, 2002.
15. Ph. Schnoebelen. Personal communication, Jan. 2003.