Quantitative Analysis of Infinite Markov Chains

Parosh Aziz Abdulla, Julien d'Orso, and Ahmed Rezine

Uppsala University, Dept. of Information Technology
Box 337, S-751 05 Uppsala, Sweden
{parosh,juldor,rahmed}@it.uu.se

Abstract. We present a general method for performing *quantitative analysis* of probabilistic systems with infinite state spaces: given an initial state s_{init}, a set F of final states, and a rational $\theta > 0$, compute a rational ρ such that the probability of reaching F from s_{init} is between ρ and $\rho + \theta$. The method extends the general framework of [AČJYK00,FS98] from the context of transition systems to that of Markov chains. We use an algorithm presented in [IN97], and present sufficient conditions for termination of the algorithm. We apply the algorithm to solve the quantitative analysis problem for probabilistic versions of VASS (Vector Addition Systems with States) and lossy channel systems.

1 Introduction

An interesting challenge in the area of model checking is to extend its applicability to systems with infinite state spaces. In [AČJYK00] we presented a general framework for verification of infinite-state transition systems. The framework is based on the assumption that the transition relation is monotonic with respect to a *well quasi-ordering* on the set of states (configurations). The framework has been applied to several classes of systems such as Petri nets, timed automata, lossy channel systems, relational automata, broadcast protocols, timed Petri nets, etc.

Recently, several attempts have been made to develop algorithms for model checking of infinite state systems with *probabilistic behaviours*. For instance the works in [Rab03,BS03,AR03,BE99,IN97,ABIJ00] consider *Probabilistic Lossy Channel Systems (PLCS)*: systems consisting of finite-state processes, which communicate through channels which are unbounded and unreliable in the sense that they can non-deterministically lose messages. The motivation for these works is that, since we are dealing with unreliable communication, it is relevant to take into consideration the probability by which messages are lost inside the channels.

In this paper, we extend the general framework of [AČJYK00,FS98] from the context of transition systems to that of Markov chains (with potentially infinite state spaces). We consider the *Quantitative Analysis Problem*: Given an initial state s_{init}, a set F of final states, and a rational $\theta > 0$ (corresponding to the *precision* of the analysis), compute a rational ρ such that the probability of reaching F from s_{init} is between ρ and $\rho + \theta$. To solve the Quantitative

Analysis Problem, we use a simple procedure presented in [IN97]. The idea of the procedure is to construct a reachability tree starting from s_{init} in a breadth-first manner. The edges of the tree are labeled with probabilities of the corresponding transitions. A leaf of the tree is closed if it either corresponds to a state inside F or a state from which F is not reachable. The algorithm terminates when we reach a point where the sum of probabilities of the closed nodes is at least $1-\theta$. Notice that such an algorithm will not terminate in general.

The main contribution of the paper is to provide sufficient conditions under which the reachability algorithm above is guaranteed to terminate. These conditions are

- The Markov chain is *finitely spanning*, i.e., there is a natural number K such that, if a state s_2 is reachable from a state s_1, then s_2 is reachable from s_1 within K steps.
- The Markov chain is *coarse*, i.e., there is a rational $\alpha > 0$ such that each enabled transition has a probability which is at least α.

As an application, we consider *Vector Addition Systems with States (VASS)*; a model which has the same computation power as Petri nets. We consider a probabilistic version *PVASS* of VASS, where each transition is equipped with a *weight* which is a natural number. The probability by which a transition is performed from a state is defined by the relative weight of the transition compared to the weights of other transitions which are enabled. We show that PVASSs satisfy the two conditions above, and therefore the above described reachability procedure amounts to an algorithm for computing a probabilistic version of the coverability problem (compute the probability by which a given set of states is covered from the initial state).

Next, we relax the coarseness condition and consider Markov chains which are *almost coarse*: the probability of reaching a node violating the coarseness condition from the initial state s_{init} before reaching a closed node is less than θ. We show that the reachability algorithm terminates if the Markov chain is almost coarse. We show that PLCSs, where initial states have empty channels, induce Markov chains which are almost coarse; thus implying solvability of the Quantitative Analysis Problem for PLCS.

Related Work The works in [BS03,AR03,BE99,ABIJ00] all consider PLCS. However, these papers only consider *Qualitative Analysis*: Is the probability of reaching F from s_{init} equal to one?

[IN97] considers quantitative analysis. However, no arguments are given for proving termination of the algorithm.

Recently, Rabinovich has presented an algorithm which solves the Quantitative Analysis Problem for PLCS [Rab03]. The algorithm solves both reachability and repeated reachability (in *repeated reachability* we compute the probability by which the set of final states is visited infinitely often). The algorithm of [Rab03] is different from the one we present in this paper. Furthermore, it is not easy to see how it can be extended to other classes of infinite Markov chains such as PVASS. For instance, it relies on the existence of a finite *attractor* in the Markov chain. An *attractor* is a set of states which is reachable from each state in the

Markov chain with a probability which is equal to one. In the case of PLCS, the set of states with empty channels form a finite attractor. However, a finite attractor does not necessarily exist in the case of PVASS, and therefore the algorithm of [Rab03] is not applicable. Also, our algorithm, based on a simple breadth-first procedure is much more intuitive than the algorithm presented in [Rab03]. Our algorithm does not rely on the notions of attractors or even strongly connected components as is the case with all existing decidability proofs (known to us) for infinite-state Markov chains. In fact, we regard showing that a small number of simple conditions are sufficient to guarantee termination of the breadth-first algorithm as one of the main contributions of this paper.

Outline In the next Section and Section 3, we recall some basic concepts of transition systems and Markov Chains. In Section 4, we present an algorithm for solving the Quantitative Analysis Problem, and present sufficient conditions for its termination. We apply the algorithm to Probabilistic VASS in Section 5. In Section 6 we introduce Almost Coarse Markov Chains. In Section 7 we consider PLCS, and then show in Section 8 that they are almost coarse. In Section 9, we give conclusions and directions for future work.

2 Transition Systems

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

A *transition system* $\mathcal{T}$ is a triple $(S, \longrightarrow, F)$ where S is a (potentially) infinite set of states, $\longrightarrow$ is a binary relation on S, and $F \subseteq S$ is the set of *final states*. We write $s \longrightarrow s'$ to denote that $(s, s') \in \longrightarrow$, and use $s \xrightarrow{k} s'$ to denote that there are $s_0, s_1, \ldots, s_k$ such that $s = s_0$, $s' = s_k$, and $s_i \longrightarrow s_{i+1}$ for $i : 0 \leq i < k$. We use $\xrightarrow{*}$ to denote the reflexive transitive closure of $\longrightarrow$, i.e., $\xrightarrow{*} = \bigcup_k \xrightarrow{k}$. We say that s_2 is *reachable* from s_1 if $s_1 \xrightarrow{*} s_2$. For $s \in S$, we use $Post(s)$ to denote the set $\{s' |\ s \longrightarrow s'\}$.

For sets $Q, Q' \subseteq S$, we use $Q \longrightarrow Q'$, to denote that there are $s \in Q$ and $s' \in Q'$ with $s \longrightarrow s'$. Definitions of other concepts such as $\xrightarrow{k}$, $\xrightarrow{*}$, and reachability, are extended to sets of states in a similar manner.

We define $\tilde{F}$ to be the set $\left\{ s |\ s \not\xrightarrow{*} F \right\}$. By F^* we mean the set $F \cup \tilde{F}$. For $s \in S$, we define the *distance* $dist(s)$ of s to be the minimum natural number k with $s \xrightarrow{k} F$. In case $s \in \tilde{F}$, we define $dist(s) = \infty$.

A transition system $\mathcal{T} = (S, \longrightarrow, F)$ is said to be *effective* if for each $s \in S$, we can

- compute elements of the set $Post(s)$ (notice that this implies that $\mathcal{T}$ is finitely branching); and
- check whether $s \xrightarrow{*} F$.

A transition system $\mathcal{T}$ is said to be of *span* K if for each $s \in S$, we have $dist(s) \leq K$ or $dists(s) = \infty$. We say that $\mathcal{T}$ is *finitely spanning* if $\mathcal{T}$ is of span K, for some $K \geq 0$. Now, we give a sufficient condition for being finitely spanning:

we say that $\mathcal{T}$ is *well quasi-ordered* if there is a preorder $\preceq$ (i.e. a reflexive transitive relation) on S such that for each infinite sequence $s_0, s_1, s_2, \ldots \in S$ there are $i < j$ with $s_i \preceq s_j$. We say that $\mathcal{T}$ is *monotonic* (with respect to $\preceq$) if (i) F is upward closed with respect to $\preceq$, i.e., $s_1 \in F$ and $s_1 \preceq s_2$ imply $s_2 \in F$; and (ii) $s_1 \longrightarrow s_2$ and $s_1 \preceq s_3$ imply that there is s_4 with $s_3 \longrightarrow s_4$ and $s_2 \preceq s_4$. In [AČJYK00] we show the following

Theorem 1. *If $\mathcal{T}$ is monotonic and well quasi-ordered then $\mathcal{T}$ is finitely spanning.*

3 Markov Chains

In this section, we introduce (potentially infinite state) *Markov chains.*

A *Markov chain* $\mathcal{M}$ is a tuple(S, P, F) where S is a (potentially infinite) set of *states*, P is a mapping from $S \times S$ to the set of rational numbers in the interval $[0, 1]$, such that $\sum_{s' \in S} P(s, s') = 1$, for each $s \in S$, and $F \subseteq S$ is the set of *final states.* A *computation* π *(from s_0)* of $\mathcal{M}$ is an infinite sequence $s_0, s_1, \ldots$ of states. We use $\pi(i)$ to denote s_i.

We use $(s \models \Diamond F)$ to be the set

$$\{\pi| \; (\pi(0) = s) \text{ and } (\exists i.\, \pi(i) \in F)\}$$

For $Q_1, Q_2 \subseteq S$, we use $\left(s \models Q_1 \; \underline{\mathit{Before}} \; Q_2\right)$ to be the set

$$\{\pi| \; (\pi(0) = s) \text{ and } (\exists j.\, \pi(j) \in Q_1) \text{ and } (\forall i < j.\, \pi(i) \notin Q_2)\}$$

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 $\mathcal{M}$ is $(S, \longrightarrow, F)$ where $s_1 \longrightarrow s_2$ iff $P(s_1, s_2) > 0$. In this manner, concepts defined for transition systems can be lifted to Markov chains. For instance, a Markov chain $\mathcal{M}$ is *finitely spanning* if the underlying transition system is finitely spanning.

A state s is said to be of *coarseness* α if for each $s' \in S$, $P(s, s') > 0$ implies $P(s, s') \geq \alpha$. A Markov chain $\mathcal{M} = (S, P, F)$ is said to be of *coarseness* α if each $s \in S$ is of coarseness α. We say that $\mathcal{M}$ is *coarse* if $\mathcal{M}$ is of coarseness α, for some $\alpha > 0$. Notice that if $\mathcal{M}$ is coarse then the underlying transition system is finitely branching; however, the converse is not necessarily true.

A Markov chain $\mathcal{M} = (S, P, F)$ 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 in probability context. A *σ-algebra* over Ω contains Ω and is closed under complementation and countable union. Adding to a measurable space a *probability measure* $\mathit{Prob} : \Delta \to [0, 1]$ such that $\mathit{Prob}(\Omega) = 1$ and that is countably additive, we get a *probability space* $(\Omega, \Delta, \mathit{Prob})$.

Consider a state s of a Markov chain $\mathcal{M} = (S, P, F)$. On the sets of computations that start at s, the probabilistic space is defined as follows:

Probabilistic space $(\Omega, \Delta, Prob_{\mathcal{M}})$(see [KSK66]) : $\Omega = sS^{\omega}$ is the set of all infinite 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 $Prob_{\mathcal{M}}$ is defined by $Prob_{\mathcal{M}}(D_u) = \prod_{i=0,\ldots,n-1} P(s_i, s_{i+1})$ where $u = s_0 s_1 ... s_n$; 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.

We shall Consider the following problem for Markov chains

Quantitative Analysis Problem
Instance A Markov chain $\mathcal{M} = (S, P, F)$, a state $s_{init} \in S$, and a rational θ.
Task Compute a rational ρ such that $\rho \leq Prob_{\mathcal{M}}\,(s_{init} \models \Diamond F) \leq \rho + \theta$.

4 Algorithm

In this section we present an algorithm based on [IN97] for solving the Quantitative Analysis Problem for effective coarse and finitely spanning Markov chains.

We use $\mathcal{N}$ and $\mathcal{Q}$ to denote the set of natural numbers and non-negative rational numbers, respectively.

Given an effective Markov chain $\mathcal{M} = (S, P, F)$, a state $s_{init} \in S$, and a positive $\theta \in \mathcal{Q}$, the algorithm tries to construct the "reachability tree" from s_{init}. The nodes of the tree are labeled with states in S. With each node in the tree is associated a probability. This probability is the product of the probabilities of all the transitions in the path from the root to the node. In other words, if the path from the root to a node is labeled with states $s_0, s_1, \ldots, s_n$ (where $s_0 = s_{init}$) then the probability of the node is given by $P(s_0, s_1) \cdot P(s_1, s_2) \cdots P(s_{n-1}, s_n)$. The tree is constructed in a breadth-first fashion, using a data structure in the form of a queue to store the nodes which are yet to be analyzed.

The algorithm maintains two variables *Yes* and *No* which accumulate the probabilities by which the set F is reachable (and not reachable respectively). For a leaf, labeled with s, the algorithm does one of the following:

- if s is in F, then the node is closed and the value of *Yes* is increased by the probability of the node;
- if s is in $\tilde{F}$, then the node is closed and the value of *No* is increased in a similar fashion;
- otherwise, a number of children, corresponding to the elements of $Post(s)$, are created

Notice that each step of the algorithm can be implemented due to effectiveness of $\mathcal{M}$.

The algorithm runs until we reach a point where the sum of *Yes* and *No* exceeds $1 - \theta$.

Algorithm 1 – Quantitative Analysis

Input
An effective Markov chain $\mathcal{M} = (S, P, F)$, a state $s_{init} \in S$, and a positive $\theta \in \mathcal{Q}$.
Output
$\rho \in \mathcal{Q}$ such that $\rho \leq \mathit{Prob}_{\mathcal{M}}\,(s_{init} \models \Diamond F) \leq \rho + \theta$.
Variables
Yes, *No*: $\mathcal{Q}$
store: queue with elements in $S \times \mathcal{Q}$
begin
1. *store* $:= (s_{init}, 1)$
2. **repeat**
3. remove (s, r) from *store*
4. **if** $s \in F$ **then** $\mathit{Yes} := \mathit{Yes} + r$
5. **else if** $s \in \tilde{F}$ **then** $\mathit{No} := \mathit{No} + r$
6. **else**
7. **for each** $s' \in \mathit{Post}(s)$
8. add $(s', r \cdot P(s, s'))$ to the end of *store*
9. **until** $\mathit{Yes} + \mathit{No} \geq 1 - \theta$
10. **return**(*Yes*)
end

We use $\mathit{Yes}^j(\mathcal{M}, s_{init})$ to denote the value of variable *Yes* after the algorithm has explored the "reachability tree" up to depth j (i.e. any element (s, r) in *store* is such that $s_{init} \xrightarrow{j+1} s$), when the algorithm is run on the Markov chain $\mathcal{M}$ with initial state s_{init}. We define $\mathit{No}^j(\mathcal{M}, s_{init})$ in a similar manner.

In the rest of this section we prove correctness of Algorithm 1. We do that by showing that the algorithm always terminates for coarse and finitely spanning Markov chains (Lemma 1), and from the partial correctness of the algorithm (Lemma 2).

Lemma 1. *Algorithm 1 terminates in case $\mathcal{M}$ is coarse and finitely spanning.*

Proof. Let K be the span of $\mathcal{M}$, and α its coarseness. We observe that any state s which doesn't belong to $\tilde{F}$ satisfies $\mathit{dist}(s) \leq K$. Therefore, from any such state s, it is the case that F is reachable with probability at least α^K. Hence, the set of paths avoiding $F \cup \tilde{F}$ forever has measure 0. We can thus conclude that $\mathit{Prob}_{\mathcal{M}}\left(s_{init} \models \Diamond F \cup \tilde{F}\right) = 1$. Therefore $\lim_{j \to \infty}\left(\mathit{Yes}^j + \mathit{No}^j\right) = 1$. Hence the result. $\square$

Lemma 2. *If Algorithm 1 terminates at depth j then*

$$\mathit{Yes}^j(\mathcal{M}, s_{init}) \leq \mathit{Prob}_{\mathcal{M}}\,(s_{init} \models \Diamond F) \leq \mathit{Yes}^j(\mathcal{M}, s_{init}) + \theta$$

Proof. It is straightforward to check that for each $j \geq 0$ we have

$$\mathit{Yes}^j(\mathcal{M}, s_{init}) \leq \mathit{Prob}_{\mathcal{M}}\,(s_{init} \models \Diamond F)$$

and

$$No^j(\mathcal{M}, s_{init}) \leq Prob_{\mathcal{M}} \left(s_{init} \models \tilde{F} \; \underline{Before} \; F \right)$$

We notice that

$$Prob_{\mathcal{M}} \left(s_{init} \models \Diamond F \right) \leq 1 - Prob_{\mathcal{M}} \left(s_{init} \models \tilde{F} \; \underline{Before} \; F \right)$$

It follows that

$$Yes^j(\mathcal{M}, s_{init}) \leq Prob_{\mathcal{M}} \left(s_{init} \models \Diamond F \right) \leq 1 - No^j(\mathcal{M}, s_{init})$$

The result follows from the fact that $Yes^j(\mathcal{M}, s_{init}) + No^j(\mathcal{M}, s_{init}) \geq 1 - \theta$ when the algorithm terminates. $\square$

From Lemma 1 and Lemma 2 we get the following

Theorem 2. *The Quantitative Analysis Problem is solvable for effective, coarse, and finitely spanning Markov chains*

Remark Algorithm 1 treats states in F^* as *absorbing* states, i.e., states from which there are no transitions to other states. Only successors of *open* (non-absorbing) states are generated. The set of open states is in general infinite. The reason why Algorithm 1 terminates (Lemma 1) is that each open state will reach an absorbing state within K steps with a probability which is bound from below (due to finite spanning and coarseness).

5 Probabilistic VASS

In this section, we use Algorithm 1 (Section 4) to solve the Quantitative Analysis Problem for Probabilistic VASS.

VASS A *Vector Addition System with States (VASS)* consists of a finite state process operating on a finite set of variables each of which ranges over $\mathcal{N}$. Formally, a *VASS* $\mathcal{V}$ is a tuple $(\mathtt{S}, \mathtt{X}, \mathtt{T}, \mathtt{F})$, where $\mathtt{S}$ is a finite set of *local states*, $\mathtt{F} \subseteq \mathtt{S}$ is the set of *final local states*, $\mathtt{X}$ is a finite set of *variables*, 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 a mapping from $\mathtt{X}$ to the set $\{-1, 0, 1\}$. A *(global) state* s is of the form $(\mathtt{s}, \mathtt{v})$ where $\mathtt{s} \in \mathtt{S}$ and $\mathtt{v}$ is a mapping from $\mathtt{X}$ to $\mathcal{N}$.

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, we use s and S to range over states and sets of states of the induced transition system (states of the transition system are global states of the VASS).

For $\mathtt{v}_1, \mathtt{v}_2 \in (\mathtt{X} \mapsto \mathcal{N})$, we use $\mathtt{v}_1 \preceq \mathtt{v}_2$ to denote that $\mathtt{v}_1(\mathtt{x}) \leq \mathtt{v}_2(\mathtt{x})$ for each $\mathtt{x} \in \mathtt{X}$. We extend $\preceq$ to a relation on $\mathtt{S} \times (\mathtt{X} \mapsto \mathcal{N})$, where $(\mathtt{s}_1, \mathtt{v}_1) \preceq (\mathtt{s}_2, \mathtt{v}_2)$ iff $\mathtt{s}_1 = \mathtt{s}_2$ and $\mathtt{v}_1 \preceq \mathtt{v}_2$.

A transition $\mathtt{t} = (\mathtt{s}_1, \mathtt{op}, \mathtt{s}_2)$ is said to be *enabled* at $(\mathtt{s}_1, \mathtt{v}_1)$ if $\mathtt{v}_1(\mathtt{x}) + \mathtt{op}(\mathtt{x}) \geq 0$ for each $\mathtt{x} \in \mathtt{X}$. We define $enabled(\mathtt{s}, \mathtt{v}) = \{\mathtt{t} \mid \mathtt{t} \text{ is enabled at } (\mathtt{s}, \mathtt{v})\}$. In case $\mathtt{t} = (\mathtt{s}_1, \mathtt{op}, \mathtt{s}_2)$ is enabled at $(\mathtt{s}_1, \mathtt{v}_1)$, we define $\mathtt{t}((\mathtt{s}_1, \mathtt{v}_1))$ to be $(\mathtt{s}_2, \mathtt{v}_2)$

where $\mathtt{v}_2(\mathtt{x}) = \mathtt{v}_1(\mathtt{x}) + \mathtt{op}(\mathtt{x})$ for each $\mathtt{x} \in \mathtt{X}$. The VASS $\mathcal{V}$ induces a transition system $(S, \longrightarrow, F)$, where S is the set of states, i.e., $S = (\mathtt{S} \times (\mathtt{X} \mapsto \mathcal{N}))$, $F = \{(\mathtt{s}, \mathtt{v}) \mid \mathtt{s} \in \mathtt{F}\}$, and $(\mathtt{s}_1, \mathtt{v}_1) \longrightarrow (\mathtt{s}_2, \mathtt{v}_2)$ iff there is a $\mathtt{t} \in \mathtt{T}$ with $(\mathtt{s}_2, \mathtt{v}_2) = \mathtt{t}((\mathtt{s}_1, \mathtt{v}_1))$. In the sequel, we assume, without loss of generality, that for all $(\mathtt{s}, \mathtt{v})$, the set $\mathit{enabled}(\mathtt{s}, \mathtt{v})$ is not empty, i.e., there is no deadlock. This can be guaranteed for instance by requiring that from each local state there is a self-loop not changing the values of the variables.

For $\mathtt{Q} \subseteq \mathtt{S}$, we define a $\mathtt{Q}$-*state* to be a state of the form $(\mathtt{s}, \mathtt{v})$ where $\mathtt{s} \in \mathtt{Q}$. Notice that, for any $\mathtt{Q} \subseteq \mathtt{S}$, the set of $\mathtt{Q}$-states is upward closed (with respect to $\preceq$).

Probabilistic VASS A *probabilistic VASS (PVASS)* $\mathcal{V}$ is of the form $(\mathtt{S}, \mathtt{X}, \mathtt{T}, \mathtt{F}, w)$, where $(\mathtt{S}, \mathtt{X}, \mathtt{T}, \mathtt{F})$ is a VASS and w is a mapping from $\mathtt{T}$ to the set of positive natural numbers. Intuitively, we derive a Markov chain from $\mathcal{V}$ by assigning probabilities to the transitions of the underlying transition system $(\mathtt{S}, \mathtt{X}, \mathtt{T}, \mathtt{F})$. The probability of performing a transition $\mathtt{t}$ from a state $(\mathtt{s}, \mathtt{v})$ is determined by the weight $w(\mathtt{t})$ of $\mathtt{t}$ compared to the weights of the other transitions which are enabled at $(\mathtt{s}, \mathtt{v})$. We define $w(\mathtt{s}, \mathtt{v}) = \sum_{\mathtt{t} \in \mathit{enabled}(\mathtt{s}, \mathtt{v})} w(\mathtt{t})$. The PVASS $\mathcal{V}$ induces a Markov chain (S, P, F), where S and F are defined as for a VASS, and

$$P\left((\mathtt{s}_1, \mathtt{v}_1), (\mathtt{s}_2, \mathtt{v}_2)\right) = \frac{\sum \{w(\mathtt{t}) \mid \mathtt{t}((\mathtt{s}_1, \mathtt{v}_1)) = (\mathtt{s}_2, \mathtt{v}_2)\}}{w((\mathtt{s}_1, \mathtt{v}_1))}$$

Notice that this is well-defined since $w((\mathtt{s}_1, \mathtt{v}_1)) > 0$ by the assumption that there are no deadlock states.

Properties of PVASS Coarseness of transition systems induced by VASS follows immediately from the definitions. Effectiveness is shown in [AČJYK00]. By Higman's Lemma [Hig52] it follows that $\preceq$ is a well quasi-ordering. It is straightforward to verify that, in case $\mathtt{F}$ is the set of $\mathtt{Q}$-states for some $\mathtt{Q} \subseteq \mathtt{S}$, then the transition system induced by a PVASS is monotonic. From Theorem 1 we conclude that a PVASS induces a transition system which is finitely spanning. From Theorem 2 we get the following

Theorem 3. *The Quantitative Analysis Problem is solvable for PVASS.*

6 Almost Coarse Markov Chains

In this section, we strengthen the result of Section 4. We consider Markov chains which are *almost coarse*, in the sense that the probability of reaching a state violating the coarseness condition before visiting a state in F^* (recall that $F^* = F \cup \tilde{F}$) can be made arbitrarily small.

Formally, let $\mathcal{M} = (S, P, F)$ be a Markov chain, and let $\alpha \in \mathcal{Q}$. We define $\mathit{Bad}(\mathcal{M}, \alpha)$ to be the set

$$\{s \mid \exists s'.\, 0 < P(s, s') < \alpha\}$$

In other words, the states in $\mathit{Bad}(\mathcal{M}, \alpha)$ are those which violate the coarseness condition.

Let $\alpha, \omega \in \mathcal{Q}$ with $\alpha > 0$. We define $Coarse(\mathcal{M}, \alpha, \omega)$ to be the set

$$\left\{s \middle| \; Prob_{\mathcal{M}}\left(s \models F^* \; \underline{Before} \; Bad(\mathcal{M}, \alpha)\right) \; \geq \; 1 - \omega\right\}$$

Intuitively, for a state $s \in Coarse(\mathcal{M}, \alpha, \omega)$, the probability of reaching F^* before reaching a bad state is at least $1 - \omega$. Therefore, if ω is sufficiently small then bad states which arise during the analysis will not be significant for computing the probability of reaching F.

For a Markov Chain $\mathcal{M}$ and a state s_{init}, we say that $\mathcal{M}$ is *almost coarse from* s_{init} iff for any $\omega > 0$, there exists $\alpha > 0$ such that $s_{init} \in Coarse(\mathcal{M}, \alpha, \omega)$.

We show (in Lemma 4) that Algorithm 1 terminates when given $\mathcal{M} = (S, P, F)$, $s_{init} \in S$, and $\theta \in \mathcal{Q}$, in case $\mathcal{M}$ is finitely spanning, and there are $\alpha > 0$ and $\omega < \theta$ such that $s_{init} \in Coarse(\mathcal{M}, \alpha, \omega)$. Notice that neither α nor ω is actually used in the definition of Algorithm 1. However, as we shall see below, their existence guarantees termination of Algorithm 1, and therefore also implies the main result of this section (Theorem 4).

To see that, we introduce Algorithm 2 which is a variant of Algorithm 1. We introduce a new variable B of type $\mathcal{Q}$. The variable B contains the sum of probabilities of nodes labeled with states in $Bad(\mathcal{M}, \alpha)$. Each time we generate a state $s \in Bad(\mathcal{M}, \alpha)$ we close the node and increase the value of B by the probability of the corresponding node. Furthermore, we replace the termination condition of Algorithm 1 by the condition where we check whether the value of $Yes + No + B$ exceeds $1 + \omega - \theta$.

More precisely, we derive Algorithm 2 by replacing line 5 in the code of Algorithm 1 by the following two lines:

5a. **else if** $s \in \tilde{F}$ **then** $No := No + r$
5b. **else if** $s \in Bad(\mathcal{M}, \alpha)$ **then** $B := B + r$

and replacing line 9 by:

9. **until** $Yes + No + B \geq 1 + \omega - \theta$

For Algorithm 2, we define $Yes^j_\bullet(\mathcal{M}, s)$ in a similar manner to $Yes^j(\mathcal{M}, s)$ in Algorithm 1. We define similar expressions for the variables No and B.

In a similar manner to the proof of Lemma 1, we can show that

Lemma 3. *Algorithm 2 terminates in case* $\mathcal{M}$ *is finitely spanning and* $s_{init} \in Coarse(\mathcal{M}, \alpha, \omega)$ *for some* $\alpha > 0$ *and* $\omega < \theta$.

Lemma 3 means that there is a j with

$$Yes^j_\bullet(\mathcal{M}, s_{init}) + No^j_\bullet(\mathcal{M}, s_{init}) + B^j_\bullet(\mathcal{M}, s_{init}) \geq 1 + \omega - \theta$$

Since $s_{init} \in Coarse(\mathcal{M}, \alpha, \omega)$ we know that $B^j_\bullet(\mathcal{M}, s_{init}) \; \leq \; \omega$. Furthermore, we observe that

$$Yes^j_\bullet(\mathcal{M}, s_{init}) \leq Yes^j(\mathcal{M}, s_{init}) \quad No^j_\bullet(\mathcal{M}, s_{init}) \leq No^j(\mathcal{M}, s_{init})$$

Consequently, we get the following

Lemma 4. *Algorithm 1 terminates in case* $\mathcal{M}$ *is finitely spanning and* $s_{init} \in Coarse(\mathcal{M}, \alpha, \omega)$ *for some* $\alpha > 0$ *and* $\omega < \theta$.

From Lemma 4 and Lemma 2 we get the following

Theorem 4. *The Quantitative Analysis Problem is solvable in case $\mathcal{M}$ is effective, finitely spanning, and almost coarse from s_{init}.*

Remarks
1. Algorithm 2 treats states in F^* and $Bad(\mathcal{M}, \alpha)$ as absorbing states (in the sense of Section 4). Notice that by decreasing the value of α, the set of bad states will shrink and therefore the set of open states will grow. This means that the probability of reaching F^* before reaching $Bad(\mathcal{M}, \alpha)$ will increase with decreasing alpha. However, this does not imply almost coarseness, since the above mentioned probability will not necessarily approach one (even in the case where we have a finitely spanning Markov chain). For instance, let us consider the Markov chain $(\{0, 1, 2, 3, \ldots\}, P, \{0\})$, where the probability function is given by $P(i, 0) = 10^{-i}$ and $P(i, i+1) = 1 - 10^{-i}$, for $i > 0$. In this example, the Markov chain is of span 1 and the set $\tilde{F}$ is empty and hence $F^* = F$. Suppose that we start from state 1. Then, by decreasing α, the probability of reaching F^* before the bad states will increase. However the probability will not approach one. In fact, Algorithm 1 would not terminate if applied to the above Markov chain. In Section 8, we show that the probability of reaching F^* before the bad states will indeed approach one in the case of Markov chains induced by PLCS.
2. Algorithm 2 requires that $\theta > \omega$, which means that the precision allowed by Algorithm 2 can never get better than ω. However, the definition of almost coarseness ensures that there is an α for any such ω. In particular, this is true if we take ω to be (say) $\frac{\theta}{2}$. Thus, we indeed allow to fix an arbitrary precision θ, and Algorithm 2 will still terminate. Since not discarding bad states can only make Algorithm 1 converge faster, we conclude (in Theorem 4) that we don't have to worry about bad states, and can therefore use Algorithm 1 directly. Also, we observe that neither α nor the set of bad states is used in Algorithm 1; they are only used as part of the argument that Algorithm 1 terminates for almost coarse Markov chains.

7 Probabilistic Lossy Channel Systems

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

Lossy Channel Systems A *Lossy Channel System* 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 $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T}, \mathtt{F})$ where $\mathtt{S}$ is a finite set of *local states*, $\mathtt{F} \subseteq \mathtt{S}$ is the set of *final 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 *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 define the *size* $\#x$ of x to be 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 extend $\preceq$ and the size function $\#$ to $(\mathtt{S} \times (\mathtt{C} \mapsto \mathtt{M}^*))$: $(\mathtt{s}_1, \mathtt{w}_1) \preceq (\mathtt{s}_2, \mathtt{w}_2)$ iff $\mathtt{s}_1 = \mathtt{s}_2$ and $\mathtt{w}_1 \preceq \mathtt{w}_2$; and $\#(\mathtt{s}, \mathtt{w}) = \#\mathtt{w}$. Notice that the size $\#(\mathtt{s}, \mathtt{w})$ is the total number of messages inside the channels in the state $(\mathtt{s}, \mathtt{w})$.

A state $(\mathtt{s}, \mathtt{w})$ is *initial* if $\#(\mathtt{s}, \mathtt{w}) = 0$, i.e., if all the channels are empty.

In a similar manner to VASS, an LCS $\mathcal{L}$ induces a transition system $(S, \longrightarrow, F)$, where S is the set of (global) states, i.e., $S = (\mathtt{S} \times (\mathtt{C} \mapsto \mathtt{M}^*))$, $F = \{(\mathtt{s}, \mathtt{w}) \mid \mathtt{s} \in \mathtt{F}\}$, and $(\mathtt{s}_1, \mathtt{w}_1) \longrightarrow (\mathtt{s}_2, \mathtt{w}_2)$ iff one of the following conditions is satisfied

- There is a $\mathtt{t} \in \mathtt{T}$, where $\mathtt{t}$ is of the form $(\mathtt{s}_1, \mathtt{c}!\mathtt{m}, \mathtt{s}_2)$ 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 $(\mathtt{s}_1, \mathtt{c}?\mathtt{m}, \mathtt{s}_2)$ and $\mathtt{w}_1$ is the result of appending $\mathtt{m}$ to 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 $\mathtt{s} = \mathtt{s}_1$ and either

- $\mathtt{op}$ is of the form $\mathtt{c}!\mathtt{m}$; or
- $\mathtt{op}$ is of the form $\mathtt{c}?\mathtt{m}$ and $\mathtt{w}(\mathtt{c}) = \mathtt{m} \bullet x$, for some $x \in \mathtt{M}^*$.

The notions of enabledness and $\mathtt{Q}$-state are defined in a similar manner to VASS. Also here, we assume absence of deadlock.

Probabilistic Lossy Channel Systems A *Probabilistic Lossy Channel System (PLCS)* $\mathcal{L}$ is of the form $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T}, \mathtt{F}, \lambda, w)$, where $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T}, \mathtt{F})$ is an LCS, $\lambda \in [0, 1]$, and w is a mapping from $\mathtt{T}$ to the set of positive natural numbers. In a similar manner to VASS, we derive a Markov chain from the PLCS $\mathcal{L}$ by assigning probabilities to the transitions of the underlying transition system $(\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T}, \mathtt{F})$. The probability of performing a transition $\mathtt{t}$ from a 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 $(\mathtt{s}, \mathtt{w})$. In the case of PLCS, we have also to take into consideration that, after performing each transition, each message which resides inside one of the channels may be lost with a probability λ. This means that the probability of reaching $(\mathtt{s}_2, \mathtt{w}_2)$ from $(\mathtt{s}_1, \mathtt{w}_1)$ is equal to (the sum over all $(\mathtt{s}_3, \mathtt{w}_3)$ of) the probability of reaching some $(\mathtt{s}_3, \mathtt{w}_3)$ from $(\mathtt{s}_1, \mathtt{w}_1)$ through performing a transition 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. Now, we show how to derive these probabilities from the definition of $\mathcal{L}$.

First, we compute probabilities of reaching states through the loss of messages. For $x, y \in \mathtt{M}^*$, we define $\#\,(x, y)$ to be the size of the set

$$\{(i_1, \ldots, i_n) \mid i_1 < \cdots < i_n \text{ and } x = y(i_1) \bullet \cdots \bullet y(i_n)\}$$

In other words, $\#(x,y)$ is the number of the different ways in which we can delete symbols in the word y in order to obtain x. We also define $P_L(x,y) = \#(x,y) \cdot \lambda^{\#y-\#x} \cdot (1-\lambda)^{\#x}$. For $\mathtt{w}_1, \mathtt{w}_2 \in (\mathtt{C} \mapsto \mathtt{M}^*)$, we define $P_L(\mathtt{w}_1, \mathtt{w}_2) = \prod_{\mathtt{c}\in\mathtt{C}} P_L(\mathtt{w}_1(\mathtt{c}), \mathtt{w}_2(\mathtt{c}))$. Intuitively, $P_L(\mathtt{w}_1, \mathtt{w}_2)$ defines the probability by which $\mathtt{w}_2$ can change to $\mathtt{w}_1$ through loss of messages during a single step of the execution of the system. Notice that $P_L(\mathtt{w}_1, \mathtt{w}_2) = 0$ in case $\mathtt{w}_1 \not\preceq \mathtt{w}_2$. We take $P_L((\mathtt{s}_1, \mathtt{w}_1), (\mathtt{s}_2, \mathtt{w}_2)) = P_L(\mathtt{w}_1, \mathtt{w}_2)$ if $\mathtt{s}_1 = \mathtt{s}_2$, and $P_L((\mathtt{s}_1, \mathtt{w}_1), (\mathtt{s}_2, \mathtt{w}_2)) = 0$ otherwise. We define $w(\mathtt{s}, \mathtt{w}) = \sum_{\mathtt{t}\in enabled(\mathtt{s},\mathtt{w})} w(\mathtt{t})$.

The PLCS $\mathcal{L}$ induces a Markov chain (S, P, F), where S and F are defined as for LCS, and

$$P((\mathtt{s}_1, \mathtt{w}_1), (\mathtt{s}_2, \mathtt{w}_2)) = \sum_{\mathtt{t}\in\mathtt{T}} ((w(\mathtt{t})/w(\mathtt{s}_1, \mathtt{w}_1)) \cdot P_L((\mathtt{s}_2, \mathtt{w}_2), \mathtt{t}(\mathtt{s}_1, \mathtt{w}_1)))$$

Properties of PLCS Effectiveness and finite spanning of Markov chains induced by PLCS follow in a similar manner to the case with VASS (Section 5). In Section 8 (Theorem 6) we show that Markov chains induced by PLCS are almost coarse. From Theorem 4 we get the following

Theorem 5. *The Quantitative Analysis Problem is solvable for PLCS.*

8 Almost Coarseness of PLCS

In this section we show that PLCSs induce almost coarse Markov chains. Throughout this section, we assume a PLCS $\mathcal{L} = (\mathtt{S}, \mathtt{C}, \mathtt{M}, \mathtt{T}, \mathtt{F}, \lambda, w)$ and the Markov chain $\mathcal{M} = (S, P, F)$ induced by $\mathcal{L}$. Almost coarseness of $\mathcal{M}$ is stated as follows:

Theorem 6. *For an initial state s_{init} and any $\omega > 0$, there exists $\alpha > 0$ such that $s_{init} \in Coarse(\mathcal{M}, \alpha, \omega)$.*

Below, we describe informally the main proof ingredients.

For any $j \in \mathcal{N}$, we let S_j denote the set of states $\{s \mid \#s = j\}$.

The proof relies on the following two properties:

1. Since each message inside a channel may be lost with a fixed probability λ during each step, the probability of losing messages increases with the size of the state. Starting from a state $s \in S_N$ for some $N \in \mathcal{N}$, the probability that we reach a state $s' \in S_{N'}$ with $N' \geq N$ decreases exponentially with N. This means that starting from any state, a computation tends to visit states with small sizes before reaching a state with a large size. In fact, starting from a state $s \in S_M$ for some $M \in \mathcal{N}$, the probability of reaching an initial state before reaching a state in S_N can be made arbitrarily close to one by increasing the value of N.
2. Let K be the span of $\mathcal{M}$. We know that, for each initial state s_{init}, either $s_{init} \in \tilde{F}$ or s_{init} is is at distance (at most) K from F^*. Therefore, there is a positive $\mu \in \mathcal{Q}$ such that any computation starting from s_{init} will visit F^* before visiting S_K with a probability which is at least μ.

Property 1 above implies that a computation with a high probability will return to the set of initial states "many times" before reaching a state with a large size. Property 2 means that each time a computation returns to an initial state, the probability of reaching F^* before reaching a large state (with size greater than K) is bounded from below.

Together, the two properties imply that, starting from an initial state, we can make the probability of reaching F^* before S_N arbitrarily large by increasing the value of N. This gives the result, since we can now define α to be the smallest positive rational such that there are states s, s' with $s \in \cup_{0 \leq i \leq N} S_i$ and $P(s, s') = \alpha$.

More formally, the proof reflects the two properties described earlier in this section.

The following Lemma shows the first property, namely for a state $s \in S_M$, we can make the probability of reaching an initial state before S_N arbitrarily close to one by increasing the value of N.

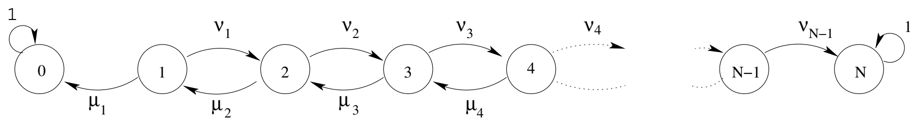

Fig. 1. The Markov Chain $\mathcal{M}_A$.

Lemma 5. *Consider $M \in \mathcal{N}$, $s \in S_M$, and $\omega \in \mathcal{Q}$ with $\omega > 0$. There is an $N \in \mathcal{N}$ such that*

$$Prob_{\mathcal{M}}\left(s \models S_0 \, \underline{Before} \, S_N\right) \quad \geq \quad 1 - \omega$$

Proof. We consider an abstraction of $\mathcal{M}$ (Figure 1) which is a Markov chain $\mathcal{M}_A = (S_A, P_A, F_A)$, where $S_A = \{0, 1, 2, \ldots\}$ and P_A is defined as follows:

- $P_A(i, i+1) = \nu_i$ for $i > 0$, where ν_i is the largest rational such that there is a state $s \in S_i$ with
$$\sum_{s' \in S_i \cup S_{i+1}} P(s, s') \quad = \quad \nu_i$$
In other words ν_i is the largest probability by which, if we start from a state in S_i, then we do not decrease the size of the state in the next step.
- $P_A(i, i-1) = \mu_i = 1 - \nu_i$ for $i > 0$.
- $P_A(i, j) = 0$, otherwise.

Intuitively, a state i in $\mathcal{M}_A$ represents an abstraction of the set S_i in $\mathcal{M}$. The claim of the lemma follows from the following two properties of $\mathcal{M}_A$:

- For natural numbers $M_1 < M_2 < M_3$ and state $s \in S_{M_2}$, we have
$$Prob_{\mathcal{M}}\left(s \models S_{M_1} \, \underline{Before} \, S_{M_3}\right) \quad \geq \quad Prob_{\mathcal{M}_A}\left(M_2 \models M_1 \, \underline{Before} \, M_3\right)$$

In other words, the Markov chain $\mathcal{M}_A$ tends to move to the left less than $\mathcal{M}$. This follows immediately from the definition of $\mathcal{M}_A$.

- There is an N such that

$$Prob_{\mathcal{M}_A}\left(M \ \models\ 0 \ \underline{Before}\ N\right) \quad \geq \quad 1-\omega$$

To show this, consider $s \in S_i$. We know that the probability of not decreasing the size of the channel is at most

$$(1-\lambda)^{i+1} \ + \ (i+1)\cdot\lambda\cdot(1-\lambda)^{i}$$

Therefore, we know that for larger values of i, the ratio $\frac{\mu_i}{\nu_i}$ gets arbitrarily large. We let L be the least index for which $i \geq L$ implies $\frac{\mu_i}{\nu_i} \geq 2$.
In [Kar66] it is shown that

$$Prob_{\mathcal{M}_A}\left(M \ \models\ 0 \ \underline{Before}\ N\right) \ = \ \frac{\sum_{i=M}^{N-1}\rho_i}{\sum_{i=0}^{N-1}\rho_i}$$

where $\rho_i = \frac{\mu_1\cdots\mu_i}{\nu_1\cdots\nu_i}$.
We notice that for $i \geq L$, we have $\rho_i \geq \gamma\cdot 2^{i-L}$ where $\gamma = \frac{\mu_1\cdots\mu_L}{\nu_1\cdots\nu_L}$
We conclude therefore that by choosing N to be sufficiently large, the value of $\sum_{i=L}^{N-1}\rho_i$ can be made arbitrarily large, i.e. we can get $Prob_{\mathcal{M}_A}\left(M \ \models\ 0 \ \underline{Before}\ N\right)$ arbitrarily close to one.

$\square$

Proof of Theorem 6

Let μ be a positive rational such that, for each initial state s, there is a sequence $s_0, s_1, \ldots, s_\ell$, where $0 \leq \ell \leq K$, $\mu \leq \prod_{i=0}^{\ell-1} P(s_i, s_{i+1})$, $s = s_0$, and $s_\ell \in F^*$. The rational μ exists since K is the span of $\mathcal{M}$. Notice that $s_i \notin S_K$ for $i : 0 \leq i < \ell$, and that μ is a lower bound on the measure of computations which originate from an initial state and reach F^* before reaching S_K.

In the proof we consider four sets of states, namely F^*, S_0, S_K and S_N, where $N \in \mathcal{N}$ and $N > K$. Consider an initial state $s \in S_0$. The state s satisfies the following properties:

- $Prob_{\mathcal{M}}\left(s \ \models\ \Diamond(F^* \cup S_K)\right) \ = \ 1$: a computation which starts from s and does not visit S_K will visit F^* by probability one by a similar reasoning to Lemma 1.
- $Prob_{\mathcal{M}}\left(s \ \models\ F^* \ \underline{Before}\ S_K\right) \ \geq \ \mu$: follows from the definition of μ.
- $Prob_{\mathcal{M}}\left(s \ \models\ S_N \ \underline{Before}\ S_K\right) \ = \ 0$: since $N > K$.

Let $\nu > 0$ be a rational such that $Prob_{\mathcal{M}}\left(s \ \models\ S_0 \ \underline{Before}\ S_N\right) \ \geq \ \nu$ for each $s \in S_K$. Such a ν exists, since S_K is finite and since for each state $s \in S_K$ the property $\left(s \ \models\ S_0 \ \underline{Before}\ S_N\right)$ holds with a positive probability. This means that each $s \in S_K$ satisfies the following properties:

- $Prob_{\mathcal{M}}\left(s \ \models\ \Diamond(S_0 \cup S_N)\right) \ = \ 1$: by absence of deadlock and the fact that we can always lose messages.

– $Prob_{\mathcal{M}}\left(s \models S_0 \underline{\mathit{Before}}\, S_N\right) \geq \nu$: by definition of ν.

From the properties stated above, we define a new abstract Markov chain (Figure 2) $\mathcal{M}_B = (S_B, P_B, F_B)$, where $S_B = \{0, 1, 2, 3\}$. Intuitively the states $0, 1, 2, 3$ represent abstractions of the sets F^*, S_0, S_K, and S_N respectively. The probability function P is defined by $P(1, 0) = \mu$, $P(1, 2) = 1 - \mu$, $P(2, 1) = \nu$, and $P(2, 3) = 1 - \nu$. From properties of the sets of states S_0 and S_K described

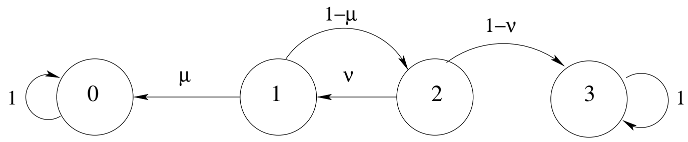

Fig. 2. Simplified Markov Chain.

above, it follows immediately that:

$$Prob_{\mathcal{M}_B}\left(1 \models 0 \underline{\mathit{Before}}\, 3\right) \leq Prob_{\mathcal{M}}\left(s \models F^* \underline{\mathit{Before}}\, S_N\right)$$

for each $s \in S_0$.

From the definition of $\mathcal{M}_B$, it is straightforward to verify that

$$Prob_{\mathcal{M}_B}\left(1 \models 0 \underline{\mathit{Before}}\, 3\right) = \frac{\mu}{1 - \nu \cdot (1 - \mu)}$$

By Lemma 5 and the definition of ν, we can choose N to make ν arbitrarily close to one. In particular, we can choose N sufficiently large such that

$$\frac{\mu}{1 - \nu \cdot (1 - \mu)} \geq 1 - \omega$$

We can now define α to be the smallest positive rational such that there are states s, s' with $s \in \cup_{0 \leq i \leq N} S_i$ and $P(s, s') = \alpha$.

Remarks

1. The proof of Theorem 6 not only shows the existence of N, but can also be used to compute the value of N. Our algorithm prunes subtrees in the reachability graph which originate from bad states (bad states are treated as absorbing states and their successors are not considered). This means that we have essentially transformed the problem into verification of a finite-state Markov chain (the set of states with sizes bounded by N is finite). It is therefore feasible to apply possibly more efficient algorithms designed for finite-state Markov chains. Notice that the finiteness property is a side effect of Theorem 6 (and not the other way round). In fact, almost coarseness does not in general enable us to *a priori* compute a bound on the size of states to be analyzed. For instance, we know (Section 5) that Markov chains induced by PVASS are

coarse and therefore also almost coarse. However, the above argument on the size of states does not apply in this case.
2. The proof of Theorem 6 can be easily modified to prove almost coarseness starting from any state (and not only from states with empty channels). Therefore, termination of Algorithm 1 can be guaranteed even in the case where the input state s_{init} has non-empty channels.

9 Conclusion, Discussion, and Future Work

We have presented a general methodology for solving the Quantitative Analysis Problem for a class of infinite-state Markov chains. The method is based on a simple breadth-first procedure presented in [IN97]. We give sufficient conditions which guarantee termination of the procedure. Providing sufficient conditions is a well-known method for deciding properties of infinite-state transition systems. For instance, the existence of a finite partitioning is used to extract an algorithm for verification of real-time automata [ACD90]. Also, the existence of a well quasi-ordered state space [AČJYK00] is used to extract an algorithm for uniform verification of lossy channel systems, Petri nets, timed Petri nets, relational automata, etc. The conditions in [ACD90,AČJYK00] and those in this paper may naturally be difficult (or impossible) to prove for some classes of transition systems (or Markov chains). The idea is to give classes which are non-trivial (and interesting) and which satisfy these conditions. In the paper we give two such classes namely PVASS and PLCS.

We are currently implementing our algorithm and intend to carry out a verification of a probabilistic model of the PGM Reliable Transport Protocol [ea00].

One direction for future research is to extend the method in order to carry out quantitative analysis of liveness properties rather than safety properties.

References

[ABIJ00] Parosh Aziz Abdulla, Christel Baier, Purushothaman Iyer, and Bengt Jonsson. Reasoning about probabilistic lossy channel systems. In C. Palamidessi, editor, *Proc. CONCUR 2000,* 11^{th} *Int. Conf. on Concurrency Theory*, volume 1877 of *Lecture Notes in Computer Science*, 2000.

[ACD90] R. Alur, C. Courcoubetis, and D. Dill. Model-checking for real-time systems. In *Proc. LICS' 90,* 5^{th} *IEEE Int. Symp. on Logic in Computer Science*, pages 414–425, Philadelphia, 1990.

[AČJYK00] Parosh Aziz Abdulla, Karlis Čerāns, Bengt Jonsson, and Tsay Yih-Kuen. Algorithmic analysis of programs with well quasi-ordered domains. *Information and Computation*, 160:109–127, 2000.

[AR03] Parosh Aziz Abdulla and Alexander Rabinovich. Verification of probabilistic systems with faulty communication, 2003. To appear in Proc. FOSSACS'03: Foundations of Software Science and Computation Structures.

[BE99] C. Baier and B. Engelen. Establishing qualitative properties for probabilistic lossy channel systems. In Katoen, editor, *ARTS'99, Formal Methods for Real-Time and Probabilistic Systems, 5th Int. AMAST Workshop*, volume 1601 of *Lecture Notes in Computer Science*, pages 34–52. Springer Verlag, 1999.

[BS03] N. Bertrand and Ph. Schnoebelen. Model checking lossy channels systems is probably decidable. In *Proc. FOSSACS03, Conf. on Foundations of Software Science and Computation Structures*, 2003.

[ea00] T. Speakman et al. PGM reliable transport protocol specification, 2000. RFC 3208 of the IEFT. Available at http://www.ietf.org/rfc.html.

[FS98] A. Finkel and Ph. Schnoebelen. Well-structured transition systems everywhere. Technical Report LSV-98-4, Ecole Normale Supérieure de Cachan, April 1998.

[Hig52] G. Higman. Ordering by divisibility in abstract algebras. *Proc. London Math. Soc.*, 2:326–336, 1952.

[IN97] Purush Iyer and Murali Narasimha. Probabilistic lossy channel systems. In Michel Bidoit and Max Dauchet, editors, *TAPSOFT '97: Theory and Practice of Software Development*, volume 1214 of *Lecture Notes in Computer Science*, pages 667–681. Springer-Verlag, 1997.

[Kar66] S. Karlin. *A First Course in Stochastic Processes.* Academic Press, 1966.

[KSK66] J.G. Kemeny, J.L. Snell, and A.W. Knapp. *Denumerable Markov Chains.* D Van Nostad Co., 1966.

[Rab03] A. Rabinovich. Quantitative analysis of probabilistic lossy channel systems, 2003. To appear in Proc. ICALP'03.