Verifying Infinite Markov Chains with a Finite Attractor or the Global Coarseness Property

Parosh Aziz Abdulla
Uppsala University, Sweden.
parosh@it.uu.se

Noomene Ben Henda
Uppsala University, Sweden.
Noomene.BenHenda@it.uu.se

Richard Mayr
NC State University, USA.
mayr@csc.ncsu.edu

Abstract

We consider infinite Markov chains which either have a finite attractor or satisfy the global coarseness property. Markov chains derived from probabilistic lossy channel systems (PLCS) or probabilistic vector addition systems with states (PVASS) are classic examples for these types, respectively. We consider three different variants of the reachability problem and the repeated reachability problem: The qualitative problem, i.e., deciding if the probability is one (or zero); the approximate quantitative problem, i.e., computing the probability up-to arbitrary precision; the exact quantitative problem, i.e., computing probabilities exactly.

We express the qualitative problem in abstract terms for Markov chains with a finite attractor and for globally coarse Markov chains, and show an almost complete picture of its decidability of PLCS and PVASS.

We also show that the path enumeration algorithm of [20] terminates for our types of Markov chain and can thus be used to solve the approximate quantitative reachability problem. Furthermore, a modified variant of this algorithm can solve the approximate quantitative repeated reachability problem for Markov chains with a finite attractor.

Finally, we show that the exact probability of (repeated) reachability cannot be effectively expressed in the first-order theory of the reals $(\mathbb{R}, +, *, \leq)$ for either PLCS or PVASS (unlike for other probabilistic models, e.g., probabilistic pushdown automata [15, 16, 14]).

1 Introduction

The aim of model checking is to decide algorithmically whether a transition system satisfies a specification. Specifications which are formulated as *reachability* or *repeated reachability* of a given set of target states are of particular interest since they allow to analyze *safety* and *progress* properties respectively. In particular, model checking problems w.r.t. ω-regular specifications are reducible to the repeated reachability problem.

A main challenge has been to extend the applicability of model checking to systems with infinite state spaces. Algorithms have been developed for numerous models such as timed automata, Petri nets, pushdown systems, lossy channel systems, parameterized systems, etc.

In a parallel development, methods have been designed for the analysis of models with stochastic behaviors (e.g. [23, 17, 27, 11, 12, 19, 13]). The motivation is to capture the behaviors of systems with uncertainty such as programs with unreliable channels, randomized algorithms, and fault-tolerant systems. The underlying semantics for such models is often that of a *Markov chain*. In a Markov chain, each transition is assigned a *probability* by which the transition is performed from a state of the system. In probabilistic model checking, three classes of problems are relevant:

- The qualitative problem: check whether a certain property holds with probability one (or zero).
- The approximate quantitative problem: compute the probability up-to arbitrary precision.
- The exact quantitative problem: compute the probabilities exactly.

Recently, several attempts have been made to consider systems which combine the above two features, i.e., systems which are infinite-state and which exhibit probabilistic behavior. For instance the works in [25, 7, 5, 6, 20, 1] 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 spontaneously 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. The papers [15, 16, 14] consider *probabilistic pushdown automata* which are natural models for probabilistic sequential programs with recursive procedures.

In this paper, we identify two general classes of infinite-state Markov chains, namely:

- Markov chains which contain a finite *attractor*. An *attractor* is a set of states which is reached with probability one from each state in the Markov chain. Examples of Markov chains with finite attractors are PLCS.

- Markov chains which are *globally coarse*. A Markov chain is *globally coarse* if there exists some $\alpha > 0$ such that the following property is satisfied: from every state, the probability of eventually reaching the set of final (target) states is either zero or $\geq \alpha$. For example, any PVASS with an upward closed set of final states induces a globally coarse Markov chain.

We consider both qualitative and quantitative analysis for the above two classes. For globally coarse Markov chains, we show several new decidability and undecidability results. Some of these are quite surprising as they deviate from usual patterns known in model checking of both probabilistic and non-probabilistic systems. For Markov chains with finite attractors, we give simpler constructions for several decidability results previously presented in the literature. The main contributions of the paper are the following:

- We show decidability of qualitative reachability problem for PVASS when the set of target states is given by a set of local states of the automaton. Decidability is shown through a reduction to model checking of a CTL-formula interpreted over the underlying (non-probabilistic) transition system. On the other hand, we show that the problem is undecidable in case the set of target states is an arbitrary upward closed set of states. This is in contrast to all known decidability results for other models such as non-probabilistic VASS, and PLCS, where the two problems can effectively be reduced to each other.

- We prove decidability of qualitative reachability problem for PLCS. This was the main result of [5, 7]. In a similar manner to those papers, the proof in this paper relies on the existence of a finite attractor. However, our construction is much simpler than those of [5, 7]. In particular, our algorithm does not require explicit construction of the attractor as is the case in [5, 7].

- We show decidability of qualitative repeated reachability problem for PVASS when the set of target states is an arbitrary upward closed set of states. This is surprising since, as mentioned above, the corresponding reachability problem is undecidable. In all models known to us from the literature (both probabilistic and non-probabilistic) repeated reachability is at least as difficult as reachability. We also show decidability of the problem for PLCS. Again, in contrast to [5, 7], our algorithm does not need to construct a finite attractor (it only relies on its existence).

- We recall an algorithm from [20] which was used in [25] to perform approximate quantitative reachability analysis for PLCS. We show that the algorithm can be used to solve the same problem for PVASS. Furthermore, we show that a minor modification of the algorithm yields an algorithm for solving approximate quantitative repeated reachability analysis for PLCS. This is a much simpler solution than the rather complicated construction which is the main result of [25].

- The question if the exact probability of (repeated) reachability for PLCS is expressible by standard mathematical functions was stated as an open problem in [25]. We provide a partial answer by showing that for both PVASS and PLCS the probability cannot be effectively expressed in the first-order theory of the reals $(\mathbb{R}, +, *, \leq)$. This is in contrast to the situation for probabilistic pushdown automata for which it can be effectively expressed in $(\mathbb{R}, +, *, \leq)$ [15, 16, 14].

Due to space constraints, many proofs had to be shifted into an appendix. They can be found in the full version of the paper [3].

2 Transition Systems and Markov Chains

We introduce some basic concepts for transition systems and Markov chains. Let $\mathcal{N}$ and $\mathcal{Q}$ denote the set of natural numbers and non-negative rational numbers, respectively.
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 $s \longrightarrow s'$ for $(s, s') \in \longrightarrow$ and let $Post(s) := \{s' |\ s \longrightarrow s'\}$. A *run* π *(from s_0)* of $\mathcal{T}$ is an infinite sequence $s_0, s_1, \ldots$ of states such that $s_i \longrightarrow s_{i+1}$ for $i \geq 0$. Let $\pi(i) := s_i$. A *path* is a finite prefix of a run. We assume familiarity with the syntax and semantics of the temporal logic CTL^* [9]. For $s \in S$ and $Q \subseteq S$, we say that Q is *reachable* from s if $s \models \exists \Diamond Q$. For $Q_1, Q_2 \subseteq S$, we use Q_1 $\underline{Before}$ Q_2 to denote the CTL formula $\exists (\neg Q_2 \mathcal{U} Q_1)$, i.e., there exists a run which reaches a state in Q_1 without having previously passed through any state in Q_2. We define $\widetilde{F} := \{s |\ s \not\models \exists \Diamond F\}$, i.e., the set of states from which F is not reachable. For $s \in S$, we define the *distance* $dist(s)$ of s to be the minimal natural number n with $s \xrightarrow{n} F$. In other words, $dist(s)$ is the length of the shortest path leading from s to F. In case $s \in \widetilde{F}$, we define $dist(s) = \infty$. A transition system $\mathcal{T}$ is said to be of *span* N if for each $s \in S$ we have $dist(s) \leq N$ or $dists(s) = \infty$. We say that $\mathcal{T}$ is *finitely spanning* if $\mathcal{T}$ is of span N for some $N \geq 0$. A transition system $\mathcal{T} = (S, \longrightarrow, F)$ is said to be *effective* if for each $s \in S$, we can (1) compute elements of the set $Post(s)$ (notice that this implies that $\mathcal{T}$ is finitely branching); and (2) check whether $s \models \exists \Diamond F$.

A *Markov chain* $\mathcal{M}$ is a tuple (S, P, F) where S is a (potentially infinite) set of *states*, $P : S \times S \to [0, 1]$, 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 Markov chain induces a transition system, where the transition relation consists of pairs of states related by positive probabilities. In this manner, concepts defined for transition systems can be lifted to Markov chains. For instance, for a Markov chain $\mathcal{M}$, a run of $\mathcal{M}$ is a run in the underlying transition system, and $\mathcal{M}$ is *finitely spanning* if the underlying transition system is finitely spanning, etc.

Consider a state s of a Markov chain $\mathcal{M} = (S, P, F)$. On the sets of runs that start at state s, the probability space $(\Omega, \Delta, Prob_{\mathcal{M}})$ is defined as follows (see also [22]): $\Omega = sS^{\omega}$ 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 $u = s_0 s_1 \ldots s_n$; this measure is extended in a unique way to the elements of the σ-algebra generated by the basic cylindric sets.

We use $Prob_{\mathcal{M}} (s \models \phi)$ to denote the measure of the set of runs in $\mathcal{M}$ which start from s and satisfy the formula ϕ. For instance, $Prob_{\mathcal{M}} (s \models \Diamond F)$ is the measure of runs from s which eventually reach F. In other words, it is the probability by which s satisfies $\Diamond F$. We say that *almost all* runs of a Markov chain satisfy a given property ϕ if $Prob_{\mathcal{M}} (s \models \phi) = 1$.

A set $A \subseteq S$ is said to be an *attractor*, if for each $s \in S$, we have $Prob_{\mathcal{M}} (s \models \Diamond A) = 1$, i.e., the set A is reached from s with probability one. If A is finite then this condition also implies $Prob_{\mathcal{M}} (s \models \Box\Diamond A) = 1$.

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

We say that a Markov chain $\mathcal{M} = (S, P, F)$ is *globally coarse* if there exists some $\alpha > 0$ s.t. $\forall s \in S. (s \xrightarrow{*} F) \Rightarrow Prob_{\mathcal{M}}(s \models \Diamond F) \geq \alpha$. If a Markov chain is coarse (of coarseness β) and finitely spanning (of span N) then it is globally coarse (define $\alpha := \beta^N$).

Lemma 1 *If a Markov chain is coarse and finitely spanning then it is globally coarse.*

3 System Models and their Properties

We give two examples of models and describe the induced Markov chains.

3.1 Vector Addition Systems

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 $(\mathtt{S}, \mathtt{X}, \mathtt{T})$, where $\mathtt{S}$ is a finite set of *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 $enabled(\mathtt{s}, \mathtt{v}) = \{\mathtt{t} | \ \mathtt{t} \text{ is enabled at } (\mathtt{s}, \mathtt{v})\}$. 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 $\mathcal{V}$, together with a set F of global states, induces a transition system $(S, \longrightarrow, F)$, where S is the set of states, i.e., $S = (\mathtt{S} \times (\mathtt{X} \mapsto \mathcal{N}))$, 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 $enabled(\mathtt{s}, \mathtt{v})$ is not empty, i.e., there is no deadlock. This can be guaranteed 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 and downward closed with respect to $\preceq$. When we speak of an upward closed set of VASS configurations, we assume that it is represented by its finitely many minimal elements.

Probabilistic VASS A *probabilistic VASS (PVASS)* $\mathcal{V}$ is of the form $(\mathtt{S}, \mathtt{X}, \mathtt{T}, w)$, where $(\mathtt{S}, \mathtt{X}, \mathtt{T})$ 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. 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 $w(\mathtt{s}, \mathtt{v}) = \sum_{\mathtt{t} \in enabled(\mathtt{s}, \mathtt{v})} w(\mathtt{t})$. The PVASS $\mathcal{V}$, together with a set F of global states, induces a Markov chain (S, P, F), where S is defined as for a VASS, and

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

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

Coarseness of Markov chains induced by PVASS follows immediately from the definitions. It follows from results in [2] (Section 4 and 7.2) that each Markov chain induced by a PVASS and an upward closed set of final markings F is effective and finitely spanning. VASS induce well-structured systems in the sense of [2] and the computation of the set of predecessors of an ideal (here this means an upward closed set) converges after some finite number k of steps. This yields the finite span k of our Markov chain derived from a PVASS. By applying Lemma 1 we obtain

Theorem 2 *A Markov chain induced by a PVASS and an upward closed set of final states F is effective, coarse, and finitely spanning and thus globally coarse.*

3.2 Probabilistic Lossy Channel Systems

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 spontaneously lose messages (see [4] for a precise definition).

Probabilistic lossy channel systems (PLCS) are a generalization of LCS to a probabilistic model for message loss and choice of transitions. There exist several variants of PLCS which differ in how many messages can be lost, with which probabilities, and in which situations, and whether normal transitions are subject to non-deterministic or probabilistic choice.

By default, we assume the relatively realistic PLCS model from [5, 7, 25] where each message in transit independently has the probability $\lambda > 0$ of being lost in every step, and the transitions are subject to probabilistic choice in a similar way as for PVASS. However, our decidability results do not strongly depend on a particular PLCS model. The only crucial aspects are the existence of a finite attractor in the induced Markov chain (most PLCS models have it) and the standard decidability results of the underlying non-probabilistic LCS (which is almost the same for all PLCS models). We will discuss other PLCS models below.

We say that a set of target states F is *effectively representable* if a finite set F' can be computed s.t. $F' \uparrow = F \uparrow$, i.e., their upward closures are equivalent. (For instance, any context-free language is effectively representable [10].) In [2] it is shown that a Markov chain, induced by a PLCS and an effectively representable set F is effective. In [5], it is shown that each Markov chain induced by a PLCS contains a finite attractor. We consider a partial order on channel contents, defined by $w_1 \leq w_2$ iff w_1 is a (not necessarily continuous) substring of w_2.

Theorem 3 *Each Markov chains, induced by a PLCS and an effectively representable set of global states, is effective and contains a finite attractor.*

The PLCS model used here and in [5, 7, 25] differs from the more unrealistic models considered previously in [1, 6]. In [6] at most one message could be lost during any step and in [1] messages could be lost only during send operations. If one assumes a sufficiently high probability (> 0.5) of message loss for these models then they also contain a finite attractor. Another different PLCS model was studied in [8]. It has the same kind of probabilistic message loss as our PLCS, but differs in having nondeterministic choice (subject to external schedulers) instead of probabilistic choice for the transitions, and thus does not yield a Markov chain. Another difference is that the model of [8] allows (and in some cases requires) idle transitions which are not present in our PLCS model. However, the PLCS model of [8] also has a finite attractor.

4 Qualitative Reachability

We consider the qualitative reachability problem for Markov chains, i.e., the problem if a given set of final states is eventually reached with probability 1, or probability 0, respectively.

QUAL_REACH
Instance A Markov chain $\mathcal{M} = (S, P, F)$ and a state $s_{init} \in S$.
Task Decide if $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) = 1$ (or $= 0$).

First we consider the problem if $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) = 1$. The following Lemma holds for any Markov chain.

Lemma 4 *$Prob_{\mathcal{M}}(s_{init} \models \Diamond F) = 1$ implies $s_{init} \not\models \widetilde{F}$ $\underline{Before}$ F.*

Proof If $s_{init} \models \widetilde{F}$ $\underline{Before}$ F then there is a path π of finite length from s_{init} to some state in $\widetilde{F}$ s.t. F is not visited in π. The set of all continuation runs of the form $\pi\pi'$ thus has a non-zero probability and never visits F. Thus $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) < 1$. □

The reverse implication of Lemma 4 holds only for Markov chains which satisfy certain conditions.

Lemma 5 *Given a Markov chain which either*

- *is globally coarse, or*
- *contains a finite attractor.*

Then we have that $s_{init} \not\models \widetilde{F}$ $\underline{Before}$ F implies $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) = 1$.

Proof The set $runs(s_{init})$ of all runs starting at s can be partitioned into $R_F := \{\pi \mid \pi(0) = s \wedge \exists i. \pi(i) \in F\}$ and $R_{\neg F} := runs(s) - R_F$. We show $Prob_{\mathcal{M}}(R_{\neg F}) = 0$ which implies the result.

Consider the case that the Markov chain is globally coarse. From $s_{init} \not\models \widetilde{F}\ \underline{Before}\ F$ it follows that all states s visited by runs in $R_{\neg F}$ satisfy $s \notin \widetilde{F}$ and thus $s \models \exists\Diamond F$. Since the Markov chain is globally coarse there exists some universal constant $\alpha > 0$ s.t. $Prob_{\mathcal{M}}\,(s \models \Diamond F) \geq \alpha$ for any s which is visited by runs in $R_{\neg F}$. Therefore $Prob_{\mathcal{M}}(R_{\neg F}) \leq (1-\alpha)^{\infty} = 0$.

Now consider the case that the Markov chain has a finite attractor A. From $s_{init} \not\models \widetilde{F}\ \underline{Before}\ F$ it follows that all states s visited by runs in $R_{\neg F}$ satisfy $s \models \exists\Diamond F$. In particular this holds for the finitely many $s \in A$ which are visited by runs in $R_{\neg F}$. Let $A' := \{s \in A \,|\, \exists \pi \in R_{\neg F}\, \exists i.\, \pi(i) = s\}$. For every $s \in A'$ we define $\alpha_s := Prob_{\mathcal{M}}\,(s \models \Diamond F)$ and obtain $\alpha_s > 0$. Since A (and A') is finite, we get $\alpha := min_{s \in A'} \alpha_s > 0$. As A is a finite attractor, almost every run in $R_{\neg F}$ must visit A (and thus A') infinitely often. Thus $Prob_{\mathcal{M}}(R_{\neg F}) \leq (1-\alpha)^{\infty} = 0$. □

Lemma 5 does not hold for general Markov chains; see Remark 1 in Section 5.

Now we apply these results to Markov chains derived from PVASS. Interestingly, decidability depends on whether F is a set of $\mathtt{Q}$-states for some $\mathtt{Q} \subseteq \mathtt{S}$ or a general upward closed set.

Theorem 6 *Given a PVASS* $(\mathtt{S}, \mathtt{X}, \mathtt{T}, w)$ *and a set of final states F which is the set of* $\mathtt{Q}$*-states for some* $\mathtt{Q} \subseteq \mathtt{S}$*. Then the question* $Prob_{\mathcal{M}}\,(s_{init} \models \Diamond F) = 1$ *is decidable.*

Proof Since any set of $\mathtt{Q}$-states is upward closed, we obtain from Theorem 2 that the Markov chain derived from our PVASS is globally coarse. Thus, by Lemma 4 and Lemma 5, we obtain $Prob_{\mathcal{M}}\,(s_{init} \models \Diamond F) < 1 \iff s_{init} \models \widetilde{F}\ \underline{Before}\ F$. To decide the question $s_{init} \models \widetilde{F}\ \underline{Before}\ F$, we construct a modified PVASS $(\mathtt{S}, \mathtt{X}, \mathtt{T}', w')$ by removing all outgoing transitions from states $q \in \mathtt{Q}$. Formally, $\mathtt{T}'$ contains all transitions of the form $(\mathtt{s}_1, \mathtt{op}, \mathtt{s}_2) \in \mathtt{T}$ with $\mathtt{s}_1 \notin \mathtt{Q}$ and $w'(\mathtt{t}) = w(\mathtt{t})$ for $\mathtt{t} \in \mathtt{T} \cap \mathtt{T}'$. Furthermore, to avoid deadlocks, we add to each state in $\mathtt{Q}$ a self-loop which does not change the values of the variables and whose weight is equal to one. It follows that $s_{init} \models \widetilde{F}\ \underline{Before}\ F$ in $(\mathtt{S}, \mathtt{X}, \mathtt{T}, w)$ iff $s_{init} \models \exists\Diamond\widetilde{F}$ in $(\mathtt{S}, \mathtt{X}, \mathtt{T}', w')$.

So we obtain that $Prob_{\mathcal{M}}\,(s_{init} \models \Diamond F) = 1$ in $(\mathtt{S}, \mathtt{X}, \mathtt{T}, w)$ iff $\widetilde{F}$ is not reachable in the VASS $(\mathtt{S}, \mathtt{X}, \mathtt{T}')$. Since $\widetilde{F}$ is downward closed and effectively constructible (due to the monotonicity of VASS), decidability follows from the decidability of the (submarking) reachability problem for Petri nets [24, 21]. □

The situation changes if one considers not a set of $\mathtt{Q}$-states as final states F, but rather some general upward-closed set F (described by its finitely many minimal elements). In this case one cannot effectively check the condition $s_{init} \models \widetilde{F}\ \underline{Before}\ F$.

Theorem 7 *Given a PVASS* $\mathcal{V} = (\mathtt{S}, \mathtt{X}, \mathtt{T}, w)$ *and an upward closed set of final states F, then the question* $Prob_{\mathcal{M}}\,(s_{init} \models \Diamond F) = \rho$ *is undecidable for any* $\rho \in (0, 1]$.

Notice the difference between Theorem 6 and Theorem 7 in the case of $\rho = 1$. Unlike for non-probabilistic VASS, reachability of control-states and reachability of upward-closed sets cannot be effectively expressed in terms of each other for PVASS.

Theorem 8 *Given a PLCS $\mathcal{L}$ where F is effectively representable. Then the question* $Prob_{\mathcal{M}}\,(s_{init} \models \Diamond F) = 1$ *is decidable.*

Proof By Theorem 3, $\mathcal{L}$ has a finite attractor. Thus we obtain from Lemma 4 and Lemma 5 that $Prob_{\mathcal{M}}\,(s_{init} \models \Diamond F) = 1$ iff $s_{init} \not\models \widetilde{F}\ \underline{Before}\ F$. This condition can be checked with a standard construction for LCS (from [5]) as follows. First one can effectively compute the set $\widetilde{F} = \overline{Pre^*(F)}$ using the techniques from, e.g., [4]. Next one computes the set X of all configurations from which it is possible to reach $\widetilde{F}$ without passing through F. This is done as follows. Let $X_0 := \widetilde{F}$ and $X_{i+1} := X_i \uparrow\ \cup (Pre(X_i) \cap \overline{F}) \uparrow$. Since all X_i are upward closed, this construction converges at some finite index n, by Higman's Lemma [18]. We get that $X = X_n$ is effectively constructible. Finally we have that $Prob_{\mathcal{M}}\,(s_{init} \models \Diamond F) = 1$ iff $s_{init} \notin X$, which can be effectively checked. □

Notice that, unlike in earlier work [5, 7], it is not necessary to compute the finite attractor for Theorem 8. It suffices to know that it exists.

Now we consider the question $Prob_{\mathcal{M}}\,(s_{init} \models \Diamond F) = 0$. The following property trivially holds for all Markov chains.

Lemma 9 $Prob_{\mathcal{M}}\,(s_{init} \models \Diamond F) = 0$ *iff* $s_{init} \not\models \exists\Diamond F$.

Since the (submarking) reachability problem for Petri nets is decidable [24, 21], we get the following consequence of Lemma 9.

Theorem 10 *Given a PVASS* $\mathcal{V} = (\mathtt{S}, \mathtt{X}, \mathtt{T}, w)$ *and a set of final states F which is expressible in the constraint logic of [21] (in particular any upward closed set, any finite set, and their complements), then the question* $Prob_{\mathcal{M}}\,(s_{init} \models \Diamond F) = 0$ *is decidable.*

From Lemma 9 and the result that for LCS the set of all predecessors of an upward closed set can be effectively constructed (e.g., [4]), we get the following.

Theorem 11 *Given a PLCS $\mathcal{L}$ where F is effectively representable, then the question* $Prob_{\mathcal{M}}\,(s_{init} \models \Diamond F) = 0$ *is decidable.*

5 Qualitative Repeated Reachability

Here we consider the qualitative repeated reachability problem for Markov chains, i.e., the problem if a given set of final states F is visited infinitely often with probability 1, or probability 0, respectively.

QUAL_REP_REACH
Instance A Markov chain $\mathcal{M} = (S, P, F)$ and a state $s_{init} \in S$.
Task Decide if $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1$ (or $= 0$).

First we consider the problem if $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1$. The following Lemma holds for any Markov chain.

Lemma 12 *$Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1$ implies $s_{init} \models \forall\Box\exists\Diamond F$.*

Proof Suppose that $s_{init} \not\models \forall\Box\exists\Diamond F$. Then $s_{init} \models \exists\Diamond\forall\Box\neg F$. Thus there exists a finite path π leading to a state s s.t. $s \models \forall\Box\neg F$. The set of all runs of the form $\pi\pi'$ (for any π') has non-zero probability and they all satisfy $\neg\Box\Diamond F$. So we get that $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) < 1$. $\Box$

The reverse implication of Lemma 12 holds only for Markov chains which satisfy certain conditions.

Lemma 13 *Given a Markov chain which either*

- *is globally coarse, or*
- *contains a finite attractor.*

Then we have that $s_{init} \models \forall\Box\exists\Diamond F$ implies $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1$.

Proof Let $R(s_{init})$ be the set of all states which are reachable from s_{init}. If $s_{init} \models \forall\Box\exists\Diamond F$ then every state $s \in R(s_{init})$ satisfies $s \models \exists\Diamond F$.

Consider the case that the Markov chain is globally coarse. Then there exists some universal constant $\alpha > 0$ (independent of $s \in R(s_{init})$) s.t. $\forall s \in R(s_{init})$. $Prob_{\mathcal{M}}(s \models \Diamond F) \geq \alpha$. Thus the set of runs which visit F only finitely often has probability $\leq (1-\alpha)^{\infty} = 0$ and therefore $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1$.

Now consider the case that the Markov chain has a finite attractor A. We have that every state $s \in R(s_{init})$ satisfies $s \models \exists\Diamond F$. In particular, this holds for the finitely many $s \in A \cap R(s_{init})$. For every $s \in A \cap R(s_{init})$ we define $\alpha_s := Prob_{\mathcal{M}}(s \models \Diamond F)$ and obtain $\alpha_s > 0$. Since A is finite, $A \cap R(s_{init})$ is finite and so we get $\alpha := min_{s \in A \cap R(s_{init})} \alpha_s > 0$. As A is a finite attractor, almost every infinite run starting at s_{init} must visit $A \cap R(s_{init})$ infinitely often. Thus the set of runs starting at s_{init} which visit F only finitely often has probability $\leq (1-\alpha)^{\infty} = 0$ and so we obtain $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1$. $\Box$

Remark 1 *Neither Lemma 5 nor Lemma 13 hold for general Markov chains. A counterexample is the Markov chain $\mathcal{M} = (S, P, F)$ of the "gambler's ruin" problem where $S = \mathcal{N}$, $P(i, i+1) := x$, $P(i, i-1) := 1-x$ for $i \geq 1$ and $P(0,0) = 1$ and $F := \{0\}$, for some parameter $x \in [0,1]$. It follows that $\widetilde{F} = \emptyset$ for $x < 1$. If $x \in [0, 1/2]$ then $Prob_{\mathcal{M}}(1 \models \Diamond F) = 1$ and for $x > 1/2$ one has $Prob_{\mathcal{M}}(1 \models \Diamond F) = (1-x)/x$.*

For $x \in (1/2, 1)$ one has that $1 \not\models \widetilde{F}$ $\underline{Before}$ F, but $Prob_{\mathcal{M}}(1 \models \Diamond F) = (1-x)/x < 1$. Similarly, although $1 \models \forall\Box\exists\Diamond F$, one still has $Prob_{\mathcal{M}}(1 \models \Box\Diamond F) \leq Prob_{\mathcal{M}}(1 \models \Diamond F) = (1-x)/x < 1$.

Theorem 14 *Let $\mathcal{V} = (\mathtt{S}, \mathtt{X}, \mathtt{T}, w)$ be a PVASS and F an upward closed set of final states. Then the question $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1$ is decidable.*

Proof Since F is upward closed, we obtain from Theorem 2 that the Markov chain derived from our PVASS is globally coarse. Thus it follows from Lemma 12 and Lemma 13 that $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1 \iff s_{init} \models \forall\Box\exists\Diamond F$. This condition can be checked as follows. Since F is upward closed and represented by its finitely many minimal elements, the set $Pre^*(F)$ is upward-closed and effectively constructible. Then $\widetilde{F} = \overline{Pre^*(F)}$ is downward-closed and effectively constructible. We get that $s_{init} \models \forall\Box\exists\Diamond F$ iff $s_{init} \not\models \exists\Diamond\widetilde{F}$ which is decidable by reduction to the (submarking) reachability problem for Petri nets [24, 21]. $\Box$

Notice the surprising contrast of the decidability of repeated reachability of Theorem 14 to the undecidability of simple reachability in Theorem 7.

Theorem 15 *Given a PLCS $\mathcal{L}$ where F is effectively representable, then the question $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1$ is decidable.*

Proof By Theorem 3, $\mathcal{L}$ has a finite attractor. Thus, we obtain from Lemma 12 and Lemma 13 that $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1 \iff s_{init} \models \forall\Box\exists\Diamond F$. This condition can be checked as follows. First one can effectively compute the set $\widetilde{F} = \overline{Pre^*(F)}$. Next one computes the set X of all configurations from which it is possible to reach $\widetilde{F}$, i.e., $X := Pre^*(\widetilde{F})$. Finally we have that $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) = 1$ iff $s_{init} \models \forall\Box\exists\Diamond F$ iff $s_{init} \notin X$, which can be effectively checked. $\Box$

Notice that, unlike in earlier work [5, 7], it is not necessary to compute the finite attractor for Theorem 15. It suffices to know that it exists.

Now we consider the question $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 0$. From the definitions we get the following.

Lemma 16 *For any Markov chain* $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) \neq 0$ *implies* $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) \neq 1$

The reverse implication holds only under certain conditions.

Lemma 17 *Given a Markov chain which either*

- *is globally coarse, or*
- *contains a finite attractor.*

Then we have that $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1 - Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F})$*.*

Proof Let R be the set of all runs starting from s_{init} and $R_{\widetilde{F}} \subseteq R$ be the subset of runs which satisfy $\Diamond\widetilde{F}$. Let $R_{\neg\widetilde{F}} := R - R_{\widetilde{F}}$. It follows that every run $\pi \in R_{\neg\widetilde{F}}$ satisfies $\Box(\exists\Diamond F)$. Let $R_F \subseteq R_{\neg\widetilde{F}}$ be the subset of runs in $R_{\neg\widetilde{F}}$ which satisfy $\Box\Diamond F$ and $R_{\neg F} := R_{\neg\widetilde{F}} - R_F$, i.e., all runs in $R_{\neg F}$ satisfy $\neg\Box\Diamond F$. Now we show that $Prob_{\mathcal{M}}(R_{\neg F}) = 0$.

If the Markov chain is globally coarse then there exists some universal constant $\alpha > 0$ s.t. for every state s visited by any run in $R_{\neg\widetilde{F}}$ we have $Prob_{\mathcal{M}}(s \models \Diamond F) \geq \alpha$. Therefore, $Prob_{\mathcal{M}}(R_{\neg F}) \leq (1-\alpha)^\infty = 0$.

If the Markov chain has a finite attractor A then almost all (the rest has probability measure 0) runs in $R_{\neg\widetilde{F}}$ must visit A infinitely often. Formally, $R_{\neg\widetilde{F}}$ can be partitioned into R_A and $R_{\neg A}$ where all runs in R_A visit A infinitely often and $Prob_{\mathcal{M}}(R_{\neg A}) = 0$. Let $A' \subseteq A$ be the states in A which are visited by runs in R_A. It follows that $\alpha_s := Prob_{\mathcal{M}}(s \models \Diamond F) > 0$ for any $s \in A'$. Since A (and A') are finite, we obtain that $\alpha := min_{s\in A'}\alpha_s > 0$. So all runs in R_A pass infinitely often through states $s \in A'$ which have a probability $\geq \alpha$ to visit F. Therefore, $Prob_{\mathcal{M}}(R_{\neg F}) = Prob_{\mathcal{M}}(R_{\neg F} \cap R_{\neg\widetilde{F}}) = Prob_{\mathcal{M}}(R_{\neg F} \cap (R_A \cup R_{\neg A})) = Prob_{\mathcal{M}}(R_{\neg F} \cap R_A) + Prob_{\mathcal{M}}(R_{\neg F} \cap R_{\neg A}) \leq (1-\alpha)^\infty + Prob_{\mathcal{M}}(R_{\neg A}) = 0 + 0 = 0$.

So in both cases $Prob_{\mathcal{M}}(R_{\neg F}) = 0$. Thus, $1 = Prob_{\mathcal{M}}(R) = Prob_{\mathcal{M}}(R_{\widetilde{F}}) + Prob_{\mathcal{M}}(R_{\neg\widetilde{F}}) = Prob_{\mathcal{M}}(R_{\widetilde{F}}) + Prob_{\mathcal{M}}(R_F) + Prob_{\mathcal{M}}(R_{\neg F}) = Prob_{\mathcal{M}}(R_{\widetilde{F}}) + Prob_{\mathcal{M}}(R_F) = Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) + Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F)$. □

There is also a correspondence of $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) \neq 0$ to a property of the underlying transition graph.

Lemma 18 *For any Markov chain* $s_{init} \models \exists\Diamond\forall\Box\exists\Diamond F$ *implies* $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) \neq 1$*.*

Proof If $s_{init} \models \exists\Diamond\forall\Box\exists\Diamond F$ then there exists a finite path π leading to some state s s.t. $s \models \forall\Box\exists\Diamond F$. Therefore $s \not\models \exists\Diamond\widetilde{F}$. Consider the set R of runs of the form $\pi\pi'$ for any π'. This set has a non-zero measure and all runs in it satisfy $\neg\Diamond\widetilde{F}$. Thus $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) \leq 1 - Prob_{\mathcal{M}}(R) < 1$. □

The reverse implication of Lemma 18 holds only for Markov chains with a finite attractor, but not generally for globally coarse Markov chains. This is because global coarseness depends on the set of final states. Global coarseness of a Markov chain (S, P, F) does not imply global coarseness of $\left(S, P, \widetilde{F}\right)$.

Lemma 19 *Given a Markov chain which contains a finite attractor, then the condition* $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) \neq 1$ *implies* $s_{init} \models \exists\Diamond\forall\Box\exists\Diamond F$*.*

Proof If $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) \neq 1$ then there exists a set R of runs starting at s_{init} s.t. $Prob_{\mathcal{M}}(R) > 0$ and every run $\pi \in R$ satisfies $\neg\Diamond\widetilde{F}$ and thus $\Box(\exists\Diamond F)$. Since the Markov chain has some finite attractor A, almost every run in R visits A infinitely often (the rest has probability measure 0). Therefore, there exists a set of runs $R' \subseteq R$ s.t. every run in R' visits A infinitely often and $Prob_{\mathcal{M}}(R') = Prob_{\mathcal{M}}(R) > 0$. Since A is finite, there must exist at least one state $a \in A$ which is visited infinitely often by a subset of runs in R' with non-zero probability measure, i.e., there exists some $R'' \subseteq R'$ s.t. $Prob_{\mathcal{M}}(R'') > 0$ and every run in R'' visits state a infinitely often.

We now show that $a \models \forall\Box\exists\Diamond F$. We assume the contrary and derive a contradiction. If $a \not\models \forall\Box\exists\Diamond F$ then $a \models \exists\Diamond\widetilde{F}$. It follows that $\alpha := Prob_{\mathcal{M}}(a \models \Diamond\widetilde{F}) > 0$. Therefore the set of runs which infinitely often visit state a, but which do not satisfy $\Diamond\widetilde{F}$ has probability measure $\leq (1-\alpha)^\infty = 0$, and thus $Prob_{\mathcal{M}}(R'') = 0$. Contradiction.

So we get $a \models \forall\Box\exists\Diamond F$ and therefore $s_{init} \models \exists\Diamond\forall\Box\exists\Diamond F$, because a is reachable from s_{init}. □

Notice that the finite attractor of the Markov chain need not be known or constructible. It suffices to know that a finite attractor exists.

Theorem 20 *For any globally coarse Markov chain*

$$s_{init} \models \exists\Diamond\forall\Box\exists\Diamond F \Rightarrow Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) \neq 1 \iff Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) \neq 0$$

For any Markov chain which contains a finite attractor

$$s_{init} \models \exists\Diamond\forall\Box\exists\Diamond F \iff Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) \neq 1 \iff Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) \neq 0$$

Proof Directly from Lemma 16, 17, 18 and 19. □

Theorem 21 *Given a PLCS* $\mathcal{L}$ *where* F *is effectively representable, then the question* $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 0$ *is decidable.*

Proof By Theorem 3, PLCS have a finite attractor. Therefore, by Theorem 20, it suffices to check if $s_{init} \not\models \exists\Diamond\forall\Box\exists\Diamond F$. Since the upward closure of F is effectively constructible, one can effectively construct a symbolic representation of the set of all states which satisfy $\exists\Diamond\forall\Box\exists\Diamond F$ (using the techniques from, e.g., [4]) and check if s_{init} is not in this set. $\Box$

For PVASS, decidability of $Prob_{\mathcal{M}}\left(s_{init} \models \Box\Diamond F\right) = 0$ and the equivalent question $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) = 1$ is open. For an upward closed set F the set $\widetilde{F}$ is downward-closed, but in general not a set of Q-states, and thus Theorem 6 does not always apply. The question $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) = 1$ can certainly not be reduced to purely structural questions about the underlying transition system (unlike for PLCS), because it depends on the exact values of the probabilities, i.e., on the transition weights. Consider a PVASS model of the gambler's ruin problem (see Remark 1), but let now $F := \{1, 2, \ldots\}$ (upward closed) and thus $\widetilde{F} = \{0\}$. We have $Prob_{\mathcal{M}}(1 \models \Diamond\widetilde{F}) = (1-x)/x < 1$ for $x > 1/2$ and $Prob_{\mathcal{M}}(1 \models \Diamond\widetilde{F}) = 1$ otherwise. Furthermore, for PVASS, the probability $Prob_{\mathcal{M}}\left(s_{init} \models \Box\Diamond F\right)$ cannot be effectively expressed in the first-order theory of the reals $(\mathbb{R}, +, *, \leq)$, as shown in Section 8, Remark 3.

6 Approximate Quantitative Reachability

In this section we consider the approximate quantitative reachability problem.

> APPROX_QUANT_REACH
> **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 $\rho \leq Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) \leq \rho + \theta$.

First, we present a path enumeration algorithm based on [20] for solving the problem, and then show that the algorithm is guaranteed to terminate for Markov chains which are globally coarse, and for Markov chains which contain a finite attractor [1].

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 constructs (a prefix of) the "reachability tree" from s_{init}, in a breadth-first fashion. The nodes of the tree are labeled with pairs (s, r) where $s \in S$ and r is the probability of traversing the path from the root to the current node. Every node in the tree is labeled with a probability. This probability is the product of the probabilities of all the transitions in the path from the root to the node. The algorithm maintains two variables Yes and No which accumulate the probabilities by which the set F is reachable (and not reachable respectively). Each step of the algorithm can be implemented due to the effectiveness of $\mathcal{M}$. The algorithm runs until we reach a point where the sum of Yes and No exceeds $1 - \theta$.

[1] Termination of the algorithm is stated (without giving a proof) in [25] for the case of Markov chains which contain a finite attractor.

Algorithm 1 – APPROX_QUANT_REACH
Input
An effective Markov chain $\mathcal{M} = (S, P, F)$, a state $s_{init} \in S$, and a positive $\theta \in \mathcal{Q}$.
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** $Yes := Yes + r$
   5. **else if** $s \in \widetilde{F}$ **then** $No := No + r$
   6. **else** **for each** $s' \in Post(s)$
      7. add $(s', r \cdot P(s, s'))$ to the end of $store$
8. **until** $Yes + No \geq 1 - \theta$
9. **return**(Yes)
end

Let $Yes^j(\mathcal{M}, s_{init})$ denote the value of variable Yes after the algorithm has explored the "reachability tree" with root s_{init} up to depth j (i.e. any element (s, r) in *store* is such that $s_{init} \xrightarrow{\leq j+1} s$). We define $No^j(\mathcal{M}, s_{init})$ in a similar manner. First we show partial correctness of Algorithm 1.

Lemma 22 *If Algorithm 1 terminates at depth j then*

$$Yes^j(\mathcal{M}, s_{init}) \leq Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) \leq Yes^j(\mathcal{M}, s_{init}) + \theta$$

Notice that if $Prob_{\mathcal{M}}\left(s_{init} \models \Diamond\left(F \cup \widetilde{F}\right)\right) = 1$ then $\lim_{j\to\infty}\left(Yes^j + No^j\right) = 1$, which implies termination of the algorithm.

Lemma 23 *Algorithm 1 terminates in case $\mathcal{M}$ either*

- *is globally coarse, or*
- *contains a finite attractor [25].*

From Lemma 22 and Lemma 23 it follows that APPROX_QUANT_REACH is solvable for Markov chains which are globally coarse and for Markov chains which contain a finite attractor. This, together with Theorem 2 and Theorem 3, yield the following theorems.

Theorem 24 APPROX_QUANT_REACH *is solvable for PVASS in case F is upward closed.*

Theorem 25 APPROX_QUANT_REACH *is solvable for PLCS in case F is effectively representable.*

7 Approximate Quantitative Repeated Reachability

In this section we consider the approximate quantitative repeated reachability problem.

APPROX_QUANT_REP_REACH
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 $\rho \leq Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) \leq \rho + \theta$.

We present an algorithm which is a modification of Algorithm 1 (in Section 6) and show that it is guaranteed to terminate for Markov chains which contain a finite attractor.

Let $F^\bullet$ be the set $\{s | \; s \models \forall\Box \; \exists\Diamond \; F\}$, i.e., $F^\bullet = \widetilde{\left(\widetilde{F}\right)}$. We modify Algorithm 1, and replace line 4 by the following

Algorithm 2 – APPROX_QUANT_REP_REACH
...
4. **if** $s \in F^\bullet$ **then** $Yes := Yes + r$
...

We define $Yes^j(\mathcal{M}, s_{init}) \leq Prob_{\mathcal{M}}$ and $No^j(\mathcal{M}, s_{init}) \leq Prob_{\mathcal{M}}$ in a similar manner to Algorithm 1. The following Lemma states partial correctness of Algorithm 2 for Markov chains with finite attractors.

Lemma 26 *For a Markov chain $\mathcal{M}$ with a finite attractor, if Algorithm 2 terminates at depth j then*

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

Proof From Lemma 13 it follows that, for each $j \geq 0$, we have $Yes^j(\mathcal{M}, s_{init}) \leq Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F)$. It is straightforward to check that, for each $j \geq 0$, we have $No^j(\mathcal{M}, s_{init}) \leq Prob_{\mathcal{M}}(s_{init} \not\models \Box\Diamond F) = 1 - (s_{init} \models \Box\Diamond F)$. It follows that $Yes^j(\mathcal{M}, s_{init}) \leq Prob_{\mathcal{M}}(s_{init} \models \Diamond F) \leq 1 - No^j(\mathcal{M}, s_{init})$. 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. $\Box$

Lemma 27 *Algorithm 2 terminates in case $\mathcal{M}$ contains a finite attractor.*

Proof Let A be the finite attractor. The set $runs(s)$ of all runs starting at s can be partitioned into $R_{\widetilde{F}} := \{\pi \,|\, \pi(0) = s \wedge \exists i.\, \pi(i) \in \widetilde{F}\}$ and $R_{\neg\widetilde{F}} := runs(s) - R_{\widetilde{F}}$. The set $R_{\neg\widetilde{F}}$ can be partitioned into $R_{F^\bullet} := \{\pi \,|\, \pi \in R_{\neg\widetilde{F}} \wedge \exists i.\, \pi(i) \in F^\bullet\}$ and $R_{\neg F^\bullet} := R_{\neg\widetilde{F}} - R_{F^\bullet}$. Below, we show that $Prob_{\mathcal{M}}(R_{\neg F^\bullet}) = 0$ which implies the result.

Since $R_{\neg F^\bullet} \cap R_{F^\bullet} = \emptyset$, it follows that all states s' visited by runs in $R_{\neg F^\bullet}$ satisfy $s' \models \exists\Diamond\widetilde{F}$. In particular this holds for the finitely many $s' \in A$ which are visited by runs in $R_{\neg F^\bullet}$. Let $A' := \{s' \in A \,|\, \exists\pi \in R_{\neg F^\bullet}\, \exists i.\, \pi(i) = s'\}$. For every $s' \in A'$ we define $\alpha_{s'} := Prob_{\mathcal{M}}\left(s' \models \Diamond\widetilde{F}\right)$ and obtain $\alpha_{s'} > 0$. By finiteness of A (and A') we get $\alpha := min_{s' \in A'} \alpha_{s'} > 0$. As A is a finite attractor, almost every run in $R_{\neg F^\bullet}$ must visit A (and thus A') infinitely often. Since $R_{\neg F^\bullet} \cap R_{\widetilde{F}} = \emptyset$ it follows that the set of runs in $R_{\neg F^\bullet}$ has measure $\leq (1-\alpha)^\infty = 0$. $\Box$

From Lemma 26 and Lemma 27, it follows that APPROX_QUANT_REP_REACH is solvable for Markov chains with finite attractors. This, together with Theorem 3, gives the following theorem.

Theorem 28 APPROX_QUANT_REP_REACH *is solvable for PLCS in case F is effectively representable.*

Remark 2 *Algorithm 2 does not generally terminate for globally coarse Markov chains (S, P, F). However, it is easy to show that it does terminate in the special case where $\left(S, P, \widetilde{F}\right)$ is also globally coarse.*

8 Exact Quantitative Analysis

In this section we consider the *Exact Quantitative Reachability Analysis Problem*, defined as follows.

EXACT_QUANT_REACH
Instance A Markov chain $\mathcal{M} = (S, P, F)$, a state $s_{init} \in S$, and a rational ρ.
Task Check whether $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) \geq \rho$.

By Theorem 7, EXACT_QUANT_REACH is undecidable for PVASS and upward closed F. If F is a set of $\mathtt{Q}$-states then decidability of EXACT_QUANT_REACH is open for PVASS and PLCS.

However, for both PVASS and PLCS, we show that the probability $Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$ (and thus the question of EXACT_QUANT_REACH) cannot be effectively expressed in the first-order theory of the reals $(\mathbb{R}, +, *, \leq)$. This is in contrast to the situation for probabilistic pushdown automata for which this probability can be effectively expressed in $(\mathbb{R}, +, *, \leq)$ [15, 16, 14].

Theorem 29 *Let $\mathcal{V} = (\mathtt{S}, \mathtt{X}, \mathtt{T}, w)$ be a PVASS and $w_1, \ldots, w_n \in \mathcal{Q}$ the constants used in the transition weight function w. Let F be the set of $\mathtt{Q}$-states for some $\mathtt{Q} \subseteq \mathtt{S}$.*

It is impossible to effectively construct a $(\mathbb{R}, +, *, \leq)$ formula $\Phi(p, w_1, \ldots, w_n)$ with parameters w_i and p which expresses the probability $Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$, i.e., $\Phi(p, w_1, \ldots, w_n) = true$ iff $p = Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$.

For PLCS, the result corresponding to Theorem 29 would be trivial if one allowed the case that the message loss probability λ is zero, because the reachability problem for deterministic non-lossy non-probabilistic FIFO-channel systems is undecidable (unlike for VASS). The following theorem shows a stronger non-expressibility result even for the restricted case of $\lambda > 0$.

Theorem 30 *Let $\mathcal{L}$ be a PLCS with message loss probability $\lambda > 0$ and F the set of $\mathtt{Q}$-states for some $\mathtt{Q} \subseteq \mathtt{S}$. It is impossible to effectively construct a $(\mathbb{R}, +, *, \leq)$ formula $\Phi(p, \lambda)$ with parameters $\lambda > 0$ and p which expresses the probability $Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$, i.e., for any $\lambda > 0$ one has $\Phi(p, \lambda) = true$ iff $p = Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$.*

Remark 3 *In the constructions for Theorem 29 and Theorem 30 only states in F can be reached from F. Therefore, in these constructions, one gets that $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) = Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F)$. Thus, for PVASS and PLCS, the probability $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F)$ is not effectively expressible in $(\mathbb{R}, +, *, \leq)$ either.*

9 Conclusions and Future Work

We have studied the qualitative and quantitative (repeated) reachability problem for Markov chains with the global coarseness property or a finite attractor. A surprising result was that reachability of control-states and reachability of upward closed sets cannot be effectively expressed in terms of each other for PVASS, unlike for normal VASS (Section 4). Furthermore, for probabilistic systems, reachability is not always easier to decide than repeated reachability (Theorems 7 and 14). Finally, a simple path enumeration algorithm can solve the approximate quantitative repeated reachability problem for all Markov chains with a finite attractor. In particular it can solve the same problem for PLCS as the more complex construction of [25], although, unlike [25], it does not yield any precise complexity bound.

References

[1] P. A. Abdulla, C. Baier, P. Iyer, and B. Jonsson. Reasoning about probabilistic lossy channel systems. In *Proc. CONCUR 2000, 11^{th} Int. Conf. on Concurrency Theory*, volume 1877 of *Lecture Notes in Computer Science*, 2000.

[2] P. A. Abdulla, K. Čerāns, B. Jonsson, and T. Yih-Kuen. Algorithmic analysis of programs with well quasi-ordered domains. *Information and Computation*, 160:109–127, 2000.

[3] P. A. Abdulla, N. B. Henda, and R. Mayr. Verifying infinite markov chains with a finite attractor or the global coarseness property. Technical report, Dept. of Information Technology, Uppsala University, Sweden, 2005.

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

[5] P. A. Abdulla and A. Rabinovich. Verification of probabilistic systems with faulty communication. In *Proc. FOSSACS03, Conf. on Foundations of Software Science and Computation Structures*, volume 2620, 2003.

[6] 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.

[7] N. Bertrand and P. Schnoebelen. Model checking lossy channels systems is probably decidable. In *Proc. FOSSACS03, Conf. on Foundations of Software Science and Computation Structures*, volume 2620, 2003.

[8] N. Bertrand and P. Schnoebelen. Verifying nondeterministic channel systems with probabilistic message losses. In R. Bharadwaj, editor, *Proc. of the 3rd Int. Workshop on Automated Verification of Infinite-State Systems (AVIS'04)*, ENTCS. Elsevier, 2004.

[9] E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, Dec. 1999.

[10] B. Courcelle. On constructing obstruction sets of words. *Bulletin of the EATCS*, 44:178–185, June 1991.

[11] C. Courcoubetis and M. Yannakakis. Verifying temporal properties of finite-state probabilistci programs. In *Proc. 29^{th} Annual Symp. Foundations of Computer Science*, pages 338–345, 1988.

[12] C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, 1995.

[13] J. Couvreur, N. Saheb, and G. Sutre. An optimal automata approach to LTL model checking of probabilistic systems. In *Proc. LPAR'2003*, volume 2850 of *Lecture Notes in Computer Science*, pages 361–375, 2003.

[14] J. Esparza and K. Etessami. Verifying probabilistic procedural programs. In *Proc. 24^{th} Conf. on Foundations of Software Technology and Theoretical Computer Science*, 2004.

[15] J. Esparza, A. Kucera, and R. Mayr. Model checking probabilistic pushdown automata. In *Proc. LICS' 04 20^{th} IEEE Int. Symp. on Logic in Computer Science*, pages 12–21, 2004.

[16] K. Etessami and M. Yannakakis. Recursive markov chains, stochastic grammars, and monotone systems of non-linear equations. In *22nd Int. Symp. on Theoretical Aspects of Computer Science (STACS'05)*, volume 2996 of *Lecture Notes in Computer Science*. Springer Verlag, 2005.

[17] S. Hart and M. Sharir. Probabilistic temporal logics for finite and bounded models. In *Proc. 16^{th} ACM Symp. on Theory of Computing*, pages 1–13, 1984.

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

[19] M. Huth and M. Kwiatkowska. Quantitative analysis and model checking. In *Proc. LICS' 97 12^{th} IEEE Int. Symp. on Logic in Computer Science*, pages 111–122, 1997.

[20] P. Iyer and M. Narasimha. Probabilistic lossy channel systems. In *TAPSOFT '97: Theory and Practice of Software Development*, volume 1214 of *Lecture Notes in Computer Science*, pages 667–681, 1997.

[21] P. Jančar. Decidability of a temporal logic problem for Petri nets. *Theoretical Computer Science*, 74:71–93, 1990.

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

[23] D. Lehmann and S. Shelah. Reasoning with time and chance. *Information and Control*, 53:165–198, 1982.

[24] E. Mayr. An algorithm for the general Petri net reachability problem. *SIAM Journal of Computing*, 13:441–460, 1984.

[25] A. Rabinovich. Quantitative analysis of probabilistic lossy channel systems. In *Proc. ICALP '2003, 30^{th} International Colloquium on Automata, Languages, and Programming*, volume 2719 of *Lecture Notes in Computer Science*, pages 1008–1021, 2003.

[26] A. Tarski. *A Decision Method for Elementary Algebra and Geometry*. Univ. of California Press, Berkeley, 1951.

[27] M. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proc. 26^{th} Annual Symp. Foundations of Computer Science*, pages 327–338, 1985.