DECISIVE MARKOV CHAINS

PAROSH ABDULLA, NOOMENE BEN HENDA, AND RICHARD MAYR

Uppsala University, Department of Information Technology, Box 337, SE-751 05 Uppsala, Sweden
e-mail address: parosh@it.uu.se

Uppsala University, Department of Information Technology, Box 337, SE-751 05 Uppsala, Sweden
e-mail address: noomene.benhenda@it.uu.se

North Carolina State University, Department of Computer Science, Campus Box 8206, Raleigh, NC 27695, USA
e-mail address: mayr@csc.ncsu.edu

ABSTRACT. We consider qualitative and quantitative verification problems for infinite-state Markov chains. We call a Markov chain *decisive* w.r.t. a given set of target states F if it almost certainly eventually reaches either F or a state from which F can no longer be reached. While all finite Markov chains are trivially decisive (for every set F), this also holds for many classes of infinite Markov chains.

Infinite Markov chains which contain a finite attractor are decisive w.r.t. every set F. In particular, all Markov chains induced by probabilistic lossy channel systems (PLCS) contain a finite attractor and are thus decisive. Furthermore, all globally coarse Markov chains are decisive. The class of globally coarse Markov chains includes, e.g., those induced by probabilistic vector addition systems (PVASS) with upward-closed sets F, and all Markov chains induced by probabilistic noisy Turing machines (PNTM) (a generalization of the noisy Turing machines (NTM) of Asarin and Collins).

We consider both safety- and liveness problems for decisive Markov chains. Safety: What is the probability that a given set of states F is eventually reached. Liveness: What is the probability that a given set of states is reached *infinitely often*. There are three variants of these questions. (1) The qualitative problem, i.e., deciding if the probability is one (or zero); (2) the approximate quantitative problem, i.e., computing the probability up-to arbitrary precision; (3) the exact quantitative problem, i.e., computing probabilities exactly.

1. We express the qualitative problem in abstract terms for decisive Markov chains, and show an almost complete picture of its decidability for PLCS, PVASS and PNTM.

2. We also show that the path enumeration algorithm of Iyer and Narasimha terminates for decisive Markov chains and can thus be used to solve the approximate quantitative safety problem. A modified variant of this algorithm can be used to solve the approximate quantitative liveness problem.

3. Finally, we show that the exact probability of (repeatedly) reaching F cannot be effectively expressed (in a uniform way) in Tarski-algebra for either PLCS, PVASS or (P)NTM (unlike for probabilistic pushdown automata).

2000 ACM Subject Classification: G3, D2.4, F4.1.
Key words and phrases: Infinite Markov Chains, Model-Checking, Verification.
An extended abstract of an earlier version of this paper has appeared in the proceedings of LICS 2005 [AHM05].

LOGICAL METHODS
IN COMPUTER SCIENCE DOI:10.2168/LMCS-???

1. Introduction

Verification of infinite systems. 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.

Probabilistic systems. In a parallel development, methods have been designed for the analysis of models with stochastic behaviors (e.g. [LS82, HS84, Var85, CY88, CY95, HK97, CSS03]). 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 Φ holds with probability one (or zero).
- The approximate quantitative problem: compute the probability p of satisfying a given property Φ up-to arbitrary precision, i.e., for any pre-defined error margin $\epsilon > 0$, compute a value p' s.t. $p' \leq p \leq p' + \epsilon$.
- The exact quantitative problem: compute the probability p of satisfying a given property Φ exactly and decide exact questions, e.g., if $p \geq 0.5$.

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 [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 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 [EKM04, EKM05, EY05b, EY05a, EE04, EY05c] consider *probabilistic pushdown automata* (recursive state machines) which are natural models for probabilistic sequential programs with recursive procedures.

Our contribution. Here we consider more abstract conditions on infinite Markov chains. We show how verification problems can be solved for Markov chains with these conditions and that several infinite-state probabilistic process models satisfy them. In particular, we consider probabilistic lossy channel systems (PLCS), probabilistic vector addition systems with states (PVASS) and probabilistic noisy Turing machines (PNTM).

Let F be a given set of target states in a Markov chain, and $\widetilde{F}$ the set of states from which F cannot be reached, i.e., $\widetilde{F} := \{s \mid s \not\xrightarrow{*} F\}$. We call a Markov chain *decisive* w.r.t. a given set of target states F if it almost certainly (i.e., with probability 1) eventually reaches either F or $\widetilde{F}$. In other words, decisiveness means that if F is always reachable then it will almost certainly be reached.

While all finite Markov chains are trivially decisive (for every set F), this also holds for several classes of infinite-state Markov chains.

It is not a meaningful question if the decisiveness property is decidable for general Markov chains. For finite Markov chains the answer is always yes, and for general infinite Markov chains the problem instance is not finitely given, unless one restricts to a particular subclass. For some such subclasses decisiveness always holds, while for others (e.g., probabilistic pushdown automata (PPDA)) it is decidable (see below).

- Markov chains which contain a finite *attractor*. An *attractor* is a set of states which is eventually reached with probability one from every state in the Markov chain. Examples of Markov chains with finite attractors are all Markov chains induced by probabilistic lossy channel systems (PLCS).

  We show that infinite Markov chains which contain a finite attractor are decisive w.r.t. *every* set F.
- Markov chains which are *globally coarse*. A Markov chain is *globally coarse* w.r.t. F if there exists some $\alpha > 0$ such that, from every state, the probability of eventually reaching the set F is either zero or $\geq \alpha$. Global coarseness w.r.t. F also implies decisiveness w.r.t. F. We consider two probabilistic process models which induce globally coarse Markov chains.
  - Any probabilistic vector addition system with states (PVASS) with an upward closed set of final states F induces a globally coarse Markov chain.
  - We consider probabilistic noisy Turing machines (PNTM), a generalization of the noisy Turing machines (NTM) of Asarin and Collins [AC05]. Probabilistic noisy Turing machines induce globally coarse Markov chains w.r.t. *every* set F defined by a set of control-states.
- Another subclass of infinite Markov chains are those induced by probabilistic pushdown automata (PPDA; also called recursive state machines) [EKM04, EKM05, EY05b, EY05a, EE04, EY05c]. These infinite Markov chains are not decisive in general. However, it follows directly from the results in [EKM04] that decisiveness is decidable for PPDA, provided that the set of final states F is effectively regular.

The focus of this paper are the classes PLCS, PVASS and PNTM, not PPDA. We strive to be as general as possible and use only the weak condition of decisiveness. We do not advocate the use of our algorithms for PPDA, even for those instances which are decisive. Since PPDA is a special class with a particular structure, specialized algorithms like those described in [EKM04, EKM05, EY05b, EY05a, EE04, EY05c] are more suitable for it. However, we show in Section 9 that the techniques used for analyzing PPDA cannot possibly be applied to PLCS, PVASS or PNTM.

We consider both qualitative and quantitative analysis for decisive Markov chains. The main contributions of the paper are the following.

- The qualitative reachability problem, i.e., the question if F is reached with probability 1 (or 0). For decisive Markov chains, this problem is equivalent to a question about the underlying (non-probabilistic) transition system.

  For PVASS, the decidability of this question depends on the set of target states F. It is decidable if F is defined by a set of control-states, but undecidable if F is a more general upward-closed set of configurations. (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).

For both PLCS and PNTM, the qualitative reachability problem is generally decidable. In particular for PLCS, although this was already shown in [AR03, BS03], our construction is more abstract and simpler. In particular, our algorithm does not require explicit construction of the attractor as in [AR03, BS03].

- The qualitative repeated reachability problem.

  If a Markov chain is decisive w.r.t. F then the question whether F will be visited infinitely often with probability 1 is equivalent to a simple question about the underlying transition graph, which is decidable for PVASS, PLCS and PNTM. (For PVASS, the decidability of probabilistic repeated reachability is surprising, given the undecidability of probabilistic simple reachability above.)

  If a Markov chain is decisive w.r.t. both F and $\widetilde{F}$ then the question whether F will be visited infinitely often with probability 0 is equivalent to another question about the underlying transition graph. The precondition holds for all Markov chains with a finite attractor (such a PLCS) since they are decisive w.r.t. every set, and the question is decidable for PLCS. For PNTM, we show that if F is defined by a set of control-states then so is $\widetilde{F}$. Since PNTM induce globally coarse Markov chains w.r.t. any set defined by control-states, the question is also decidable.

  However, for PVASS, decisiveness w.r.t. F does not generally imply decisiveness w.r.t. $\widetilde{F}$ and thus our algorithm is not always applicable. For PVASS, decidability of the question whether F is visited infinitely often with probability 0 is an open problem.
- To approximate the probability of eventually reaching F, we recall an algorithm from [IN97] which was also used in [Rab03] for PLCS. We show that the algorithm can be used to solve the problem for all decisive Markov chains (in particular also for both PVASS and PNTM).

  Furthermore, we show that a minor modification of the algorithm yields an algorithm for approximating the probability of visiting F infinitely often for all Markov chains which are decisive w.r.t. F and $\widetilde{F}$. In particular this works for all Markov chains with a finite attractor, such as PLCS. This is a more abstract, general and simpler solution than the result for PLCS in [Rab03]. However, it does not yield precise complexity bounds as [Rab03].
- The question if the exact probability of (either eventually, or infinitely often) reaching F in PLCS is expressible by standard mathematical functions was stated as an open problem in [Rab03]. We provide a partial answer by showing that for PVASS, PLCS and (P)NTM, this probability cannot be effectively expressed (in a uniform way) in Tarski-algebra, the first-order theory of the reals $(\mathbb{R}, +, *, \leq)$. (By 'in a uniform way' we mean that quantitative parameters in the system should be reflected directly by constants in the Tarski-algebra-formula.) This is in contrast to the situation for probabilistic pushdown automata for which these probabilities can be effectively expressed, in a uniform way, in $(\mathbb{R}, +, *, \leq)$ [EKM04, EY05b, EE04].

2. TRANSITION SYSTEMS AND MARKOV CHAINS

We introduce some basic concepts for transition systems and Markov chains. Let $\mathbb{N}$ and $\mathbb{Q}_{\geq 0}$ denote the set of natural numbers (including 0) and non-negative rational numbers, respectively.

2.1. **Transition Systems.** A *transition system* $\mathcal{T}$ is a tuple $(S, \longrightarrow)$ where S is a (potentially) infinite set of states, and $\longrightarrow$ is a binary relation on S. We write $s \longrightarrow s'$ for $(s, s') \in\longrightarrow$ and let $Post(s) := \{s' |\ s \longrightarrow s'\}$. A *run* ρ *(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 $i \geq 0$. We use $\rho(i)$ to denote s_i and say that ρ is an s-run if $\rho(0) = s$. A *path* is a finite prefix of a run. We assume familiarity with the syntax and semantics of the temporal logic CTL^* [CGP99]. We use $(s \models \phi)$ to denote the set of s-runs that satisfy the CTL^* path-formula ϕ. 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. Given a set $F \subseteq S$, we define $\widetilde{F} := \{s|\ s \not\models \exists\Diamond F\}$, i.e., the set of states from which F is not reachable. For $s \in S$ and $F \subseteq S$, we define the *distance* $dist_F(s)$ of s to F to be the minimal natural number n with $s \xrightarrow{n} F$. In other words, $dist_F(s)$ is the length of the shortest path leading from s to F. In case $s \in \widetilde{F}$, we define $dist_F(s) = \infty$. A transition system $\mathcal{T}$ is said to be of *span* N with respect to a given set F if for each $s \in S$ we either have $dist_F(s) \leq N$ or $dist_F(s) = \infty$. We say that $\mathcal{T}$ is *finitely spanning* with respect to a given set F if $\mathcal{T}$ is of span N w.r.t. F for some $N \geq 0$. A transition system $\mathcal{T} = (S, \longrightarrow)$ is said to be *effective* w.r.t. a given set F 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 $s \models \exists\Diamond F$.

2.2. **Markov Chains.** A *Markov chain* $\mathcal{M}$ is a tuple (S, P) where S is a (potentially infinite) set of *states*, and $P : S \times S \rightarrow [0, 1]$, such that $\sum_{s' \in S} P(s, s') = 1$, for each $s \in S$. A Markov chain induces a transition system, where the transition relation consists of pairs of states related by positive probabilities. 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* w.r.t. given set F if the underlying transition system is finitely spanning w.r.t. F, etc.

Consider a state s_0 of a Markov chain $\mathcal{M} = (S, P)$. On the sets of s_0-runs, the probability space $(\Omega, \Delta, Prob_{\mathcal{M}})$ is defined as follows (see also [KSK66]): $\Omega = s_0 S^\omega$ 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 $u \in s_0 S^*$, 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.

Given a CTL^* path-formula ϕ, we use $(s \models \phi)$ to denote the set of s-runs that satisfy ϕ. We use $Prob_{\mathcal{M}}\,(s \models \phi)$ to denote the measure of the set of s-runs $(s \models \phi)$ (which is measurable by [Var85]). For instance, given a set $F \subseteq S$, $Prob_{\mathcal{M}}\,(s \models \Diamond F)$ is the measure of s-runs 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 $Prob_{\mathcal{M}}\,(s \models \phi) = 1$.

3. Classes of Markov Chains

In this section we define several abstract properties of infinite-state Markov chains: decisiveness, the existence of a finite attractor, and global coarseness. We show that both the existence of a finite attractor and global coarseness imply decisiveness. In particular, all three properties hold trivially for finite Markov chains.

In the rest of this section, we assume a Markov chain $\mathcal{M} = (S, P)$.

3.1. **Decisive Markov Chains.**

Definition 3.1. Given a Markov chain $\mathcal{M} = (S, P)$ and a set $F \subseteq S$, $\mathcal{M}$ is said to be *decisive* w.r.t. F iff $Prob_{\mathcal{M}}(s \models \Diamond F \vee \Diamond \widetilde{F}) = 1$, for each $s \in S$.

In other words, the set of runs, along which F is always reachable but which never reach F, is almost empty.

Similarly, $\mathcal{M}$ is said to be *strongly decisive* w.r.t. F if $Prob_{\mathcal{M}}(s \models \Diamond \widetilde{F} \vee \Box\Diamond F) = 1$. Intuitively, this means that the set of runs along which F is always reachable and which visit F only finitely many times is almost empty.

Lemma 3.2. *Given a Markov chain $\mathcal{M} = (S, P)$ and a set $F \subseteq S$, $\mathcal{M}$ is decisive w.r.t. F iff it is strongly decisive w.r.t. F.*

Proof. Assume a Markov chain $\mathcal{M} = (S, P)$ and a set $F \subseteq S$. We want to prove that $\forall s \in S, Prob_{\mathcal{M}}(s \models \Diamond F \vee \Diamond \widetilde{F}) = 1 \iff \forall s \in S, Prob_{\mathcal{M}}(s \models \Diamond \widetilde{F} \vee \Box\Diamond F) = 1$. This is equivalent to proving that

$$\forall s \in S, Prob_{\mathcal{M}}(s \models \Box\neg F \wedge \Box\neg\widetilde{F}) = 0 \iff \forall s \in S, Prob_{\mathcal{M}}(s \models \Diamond\Box\neg F \wedge \Box\neg\widetilde{F}) = 0$$

Let U be a set of sequences of states. U is called *proper* if no sequence in U is a prefix of another sequence in U. Assume that all sequences in U are finite and start at the same state, we will denote by $P(U) = Prob_{\mathcal{M}}(D_U)$ where $D_U = \{uS^{\omega} |\, u \in U\}$. Given a proper set of finite sequences (namely paths) U ending all in the same state s_c and a proper set V of possibly infinite sequences (runs) starting all from s_c, we define $U \bullet V$ to be the set of all sequences us_cv where $us_c \in U$ and $s_cv \in V$.

We now prove both implications of the required equivalence above.

($\Longleftarrow$) is straightforward since $\forall s \in S\, (s \models \Box\neg F) \subseteq (s \models \Diamond\Box\neg F)$.

($\Longrightarrow$) Given a state $s \in S$,we define the following sets of paths:

$$\Pi^s_{if} := \{\pi \,|\, \pi = s(\overline{F}^* F)^{i-1}\overline{F}^* f\} \text{ where } i > 0 \text{ and } f \in F.$$

Now, consider the following sets of runs:

- $\forall i > 0, \Gamma^s_{iF} := \bigcup_{f\in F} \Gamma^s_{if}$ where $\forall f \in F$, $\Gamma^s_{if} := \Pi^s_{if} \bullet (f \models \Box\neg F \wedge \Box\neg\widetilde{F})$
- $\forall i > 0, \Delta^s_{iF} := \bigcup_{f\in F} \Delta^s_{if}$ where $\forall f \in F$, $\Delta^s_{if} := \Pi^s_{if} \bullet (fS^{\omega})$.

Intuitively, Δ^s_{iF} is the set of s-runs which visit F at least i times while Γ^s_{iF} is the set of all s-runs which visit F exactly i times and then never visit neither F nor $\widetilde{F}$. In particular for $i = 0$, we define $\Gamma^s_{0F} = (s \models \Box\neg F \wedge \Box\neg\widetilde{F})$. It is straightforward to check that:

- $(\forall i \in \mathbb{N})(\forall f_1, f_2 \in F \wedge f_1 \neq f_2)(\Gamma^s_{if_1} \cap \Gamma^s_{if_2} = \emptyset)$
- $(\forall i, j \in \mathbb{N} \wedge i \neq j)(\Gamma^s_{iF} \cap \Gamma^s_{jF} = \emptyset)$
- $(\forall i \in \mathbb{N})(\forall f \in F)(Prob_{\mathcal{M}}(\Gamma^s_{if}) = P(\Pi^s_{if})Prob_{\mathcal{M}}(f \models \Box\neg F \wedge \Box\neg\widetilde{F}))$

Since $\mathcal{M}$ is decisive w.r.t. F, for all $i \in \mathbb{N}$, $Prob_{\mathcal{M}}(\Gamma^s_{iF}) = \sum_{f\in F} Prob_{\mathcal{M}}(\Gamma^s_{if}) = 0$.

We notice that $\bigcap_{i=1}^{\infty} \Delta^s_{iF} \subseteq (s \models \Box\Diamond F)$. Therefore, $(s \models \Diamond\Box\neg F) \subseteq \bigcup_{i=1}^{\infty} \neg\Delta^s_{iF}$ where for all $i > 0$, $\neg\Delta^s_{iF}$ is the set of s-runs visiting F at most $i-1$ times. For all $i > 0$, we have $(\neg\Delta^s_{iF} \cap (s \models \Box\neg\widetilde{F})) \subseteq \bigcup_{k=0}^{i-1} \Gamma^s_{kF}$. Thus, we obtain $Prob_{\mathcal{M}}(s \models \Diamond\Box\neg F \wedge \Box\neg\widetilde{F}) \leq Prob_{\mathcal{M}}(\bigcup_{i=0}^{\infty} \Gamma^s_{iF}) = \sum_{i=0}^{\infty} Prob_{\mathcal{M}}(\Gamma^s_{iF}) = 0$. $\square$

3.2. **Markov Chains with a Finite Attractor.**

Definition 3.3. Given a Markov chain $\mathcal{M} = (S, P)$, 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.

Lemma 3.4. *A Markov chain $\mathcal{M}$ which has a finite attractor is decisive w.r.t. every set $F \subseteq S$.*

Proof. Assume a Markov chain $\mathcal{M} = (S, P)$ that has a finite attractor A, a state s and a set $F \subseteq S$. All states s' visited by runs in $R_s := (s \models \Box\neg F \wedge \Box\neg\widetilde{F})$ satisfy $s' \models \exists\Diamond F$, since $s' \notin \widetilde{F}$. In particular this holds for the finitely many different $s'' \in A$ visited by those runs. Let $A' := \{s'' \in A \,|\, \exists\rho \in R_s\, \exists i.\, \rho(i) = s''\}$. For every $s'' \in A'$ we define $\alpha_{s''} := Prob_{\mathcal{M}}(s'' \models \Diamond F)$ and obtain $\alpha_{s''} > 0$. Since A is finite we get $\alpha := \min_{s'' \in A'} \alpha_{s''} > 0$. Almost every run in R_s must visit A (and thus A') infinitely often. Thus, $Prob_{\mathcal{M}}(s \models \Box\neg F \wedge \Box\neg\widetilde{F}) \leq (1-\alpha)^\infty = 0$. Finally, we obtain for each $s \in S$, $Prob_{\mathcal{M}}(s \models \Diamond F \vee \Diamond\widetilde{F}) = 1 - Prob_{\mathcal{M}}(s \models \Box\neg F \wedge \Box\neg\widetilde{F}) = 1$. □

3.3. **Globally Coarse Markov Chains.**

Definition 3.5. 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 $\mathcal{M} = (S, P)$ 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. Given a Markov chain $\mathcal{M} = (S, P)$ and a set $F \subseteq S$. We say that a Markov chain $\mathcal{M} = (S, P)$ is *globally coarse* w.r.t. F if there exists some $\alpha > 0$ s.t. $\forall s \in S.\, (s \xrightarrow{*} F) \Rightarrow Prob_{\mathcal{M}}(s \models \Diamond F) \geq \alpha$.

Lemma 3.6. *If a Markov chain is coarse and finitely spanning w.r.t. a set F then it is globally coarse w.r.t. F.*

Proof. If a Markov chain is coarse (of coarseness $\beta > 0$) and finitely spanning w.r.t. a given set F (of span N) then it is globally coarse w.r.t. the same set F (define $\alpha := \beta^N$). □

Lemma 3.7. *Given a Markov chain $\mathcal{M}$ and a set F such that $\mathcal{M}$ is globally coarse w.r.t. F, then $\mathcal{M}$ is decisive w.r.t. F.*

Proof. Assume a Markov chain $\mathcal{M} = (S, P)$, a state s and a set $F \subseteq S$ such that $\mathcal{M}$ is globally coarse w.r.t. F. All states s' visited by runs in $(s \models \Box\neg F \wedge \Box\neg\widetilde{F})$ satisfy $s' \models \exists\Diamond F$, because $s' \notin \widetilde{F}$. Since $\mathcal{M}$ is globally coarse w.r.t. F 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 those runs. Therefore, $Prob_{\mathcal{M}}(s \models \Box\neg F \wedge \Box\neg\widetilde{F}) \leq (1-\alpha)^\infty = 0$. Finally, we obtain for each $s \in S$, $Prob_{\mathcal{M}}(s \models \Diamond F \vee \Diamond\widetilde{F}) = 1 - Prob_{\mathcal{M}}(s \models \Box\neg F \wedge \Box\neg\widetilde{F}) = 1$. □

4. System Models and their Properties

We define three classes of infinite-state probabilistic system models and describe the induced Markov chains.

4.1. **Vector Addition Systems.** A *Vector Addition System with States (VASS)* consists of a finite-state process operating on a finite set of unbounded variables each of which ranges over $\mathbb{N}$. Formally, a *VASS* $\mathcal{V}$ is a tuple $(\mathtt{S}, \mathtt{X}, \mathtt{T})$, where $\mathtt{S}$ is a finite set of *control-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 $\mathbb{N}$.

Remark on notation. We use $\mathtt{s}$ and $\mathtt{S}$ to range over control-states and sets of control-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 $\mathtt{v}_1, \mathtt{v}_2 \in (\mathtt{X} \mapsto \mathbb{N})$, 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 $\mathtt{S} \times (\mathtt{X} \mapsto \mathbb{N})$, 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 $\mathit{enabled}(\mathtt{s}, \mathtt{v}) = \{\mathtt{t} |\ \mathtt{t} \text{ is enabled at } (\mathtt{s}, \mathtt{v})\}$. 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 $\mathcal{V}$ induces a transition system $(S, \longrightarrow)$, where S is the set of states, i.e., $S = (\mathtt{S} \times (\mathtt{X} \mapsto \mathbb{N}))$, 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 $\mathit{enabled}(\mathtt{s}, \mathtt{v})$ is not empty, i.e., there is no deadlock. This can be guaranteed by requiring that from each control-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 $w(\mathtt{s}, \mathtt{v}) = \sum_{\mathtt{t} \in \mathit{enabled}(\mathtt{s}, \mathtt{v})} w(\mathtt{t})$. The PVASS $\mathcal{V}$ induces a Markov chain (S, P), where S is defined as for a VASS, and

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

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 [AČJYK00] (Section 4 and 7.2) that each Markov chain induced by a PVASS is effective and finitely spanning w.r.t. any upward closed set of final markings F. VASS induce well-structured systems in the sense of [AČJYK00] 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 w.r.t. F of our Markov chain derived from a PVASS. By applying Lemma 3.6 we obtain

Theorem 4.1. *Each Markov chain $\mathcal{M}$, induced by a PVASS satisfies the following:*

- *$\mathcal{M}$ is coarse.*
- *$\mathcal{M}$ is effective, finitely spanning and globally coarse w.r.t. any upward closed set of final states.*

This combined with Lemma 3.2 and Lemma 3.7 yields the following corollary.

Corollary 4.2. *Each Markov chain, induced by a PVASS is decisive (and thus, by Lemma 3.2, strongly decisive) w.r.t. any upward closed set of final states.*

4.2. **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 [AJ96, CFI96] 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. 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.

By default, we assume the relatively realistic PLCS model, similar to the one from [AR03, BS03, Rab03], 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, the definition of PLCS in [AR03, BS03, Rab03] assumes that messages can be lost only after discrete steps, but not before them. Thus, since no messages can be lost before the first discrete step, the set $\{s \in S : s \models \exists \Diamond F\}$ of predecessors of a given set F of target states is generally not upward-closed.

Here we assume a more realistic PLCS model where messages can be lost before and after discrete steps. This PLCS model is also closer to the classic non-probabilistic LCS model where also messages can be lost before and after discrete steps [AJ96, CFI96]. So we obtain that the set $\{s \in S : s \models \exists \Diamond F\}$ is always upward-closed. 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 [Cou91].) In [AČJYK00] it is shown that a Markov chain, induced by a PLCS is effective w.r.t. any effectively representable set F.

However, many of our 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. In [AR03], it is shown that each Markov chain induced by a PLCS contains a finite attractor.

Theorem 4.3. *Each Markov chain induced by a PLCS contains a finite attractor and is effective w.r.t. any effectively representable set of global states F.*

From this result, Lemma 3.2 and Lemma 3.4, we obtain the following corollary.

Corollary 4.4. *Each Markov chain induced by a PLCS is decisive w.r.t. every set F and thus strongly decisive w.r.t. every set F.*

The PLCS models used here (and in [AR03, BS03, Rab03]) differ from the more unrealistic models considered previously in [ABIJ00, BE99]. In [BE99] at most one message could be lost during any step and in [ABIJ00] 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 [BS04]. 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, but a Markov decision process. Another difference is that the model of [BS04] allows (and in some cases requires) idle transitions which are not present in our PLCS model. However, for any scheduler, the PLCS model of [BS04] also has a finite attractor (w.r.t. the system-state, though not necessarily w.r.t. the state of the scheduler).

4.3. **Noisy Turing Machines.** *Noisy Turing Machines (NTM)* were introduced in [AC05]. They are *Turing Machines* augmented by an additional parameter $\epsilon > 0$ giving the *noise level*. Each transition of an NTM consists of two steps. First, in the *noisy step* the tape cells are subjected to noise. In this manner, each symbol in each tape may change independently and uniformly with probability ϵ to any other symbol in the tape alphabet (possibly the same as before). Then, in the *normal step*, the NTM proceeds like a normal Turing machine.

Probabilistic Turing Machines (PTM) [dLMSS56], which are Turing machines where transitions are random choices among finitely many alternatives, are more general than the model of [AC05]. In fact, any NTM can be simulated by a PTM by adding extra steps where the machine makes a pass over the tapes changing the symbols randomly. However, as described below, general PTM do not satisfy our conditions.

Probabilistic NTM. In this paper, we adopt the model of *Probabilistic Noisy Turing Machines (PNTM)* which are a generalization of NTM. In a PNTM, the transitions are similar to those of an NTM except that normal steps are subject to probabilistic choices. Formally, a PNTM $\mathcal{N}$ is a tuple $(\mathtt{S}, \Sigma, \Gamma, \mathtt{M}, \mathtt{T}, \epsilon, w)$ where $\mathtt{S}$ is a finite set of control-states, Σ is the input alphabet, $\Gamma \supseteq \Sigma \cup \sharp$ (where $\sharp$ is the blank symbol) is the tape alphabet, $\mathtt{M}$ is the number of tapes, $\mathtt{T} \subseteq \mathtt{S} \times \Gamma^{\mathtt{M}} \times \mathtt{S} \times \Gamma^{\mathtt{M}} \times \{-1, 0, 1\}^{\mathtt{M}}$ is the transition relation, ϵ is the noise level and $\mathtt{w} : \mathtt{T} \to \mathbb{N}_{>0}$ is the weight function. The probability of a transition $\mathtt{t} \in \mathtt{T}$ is given by comparing the weight $\mathtt{w}(\mathtt{t})$ to the weights of all possible alternatives.

Assume a PNTM $\mathcal{N} = (\mathtt{S}, \Sigma, \Gamma, \mathtt{M}, \mathtt{T}, \epsilon, w)$. A *global state* of $\mathcal{N}$ can be represented by a triple: (i) the control-state, (ii) the current time, and (iii) an $\mathtt{M}$-tuple of *tape configurations*. A tape configuration is a triple: (i) the head position; (ii) a finite word $\omega \in \Gamma^*$ representing the content of all cells visited by the head so far; and (iii) a $|\omega|$-tuple of natural numbers, each representing the last point in time the head visited the corresponding cell. For a set $\mathtt{Q} \subseteq \mathtt{S}$, we let $\mathtt{Q}$-states denotes the set of all global states whose control-states are in $\mathtt{Q}$.

For a PNTM $\mathcal{N} = (\mathtt{S}, \Sigma, \Gamma, \mathtt{M}, \mathtt{T}, \epsilon, w)$, we use $\mathcal{G}(\mathcal{N})$ to denote the graph obtained from $\mathcal{N}$ by abstracting away the memory tapes. Formally, $\mathcal{G}(\mathcal{N})$ is the tuple $(\mathtt{S}, \mathtt{T}')$ where $\mathtt{S}$ is the set of control-states of the underlying PNTM $\mathcal{N}$, and $\mathtt{T}' \subseteq \mathtt{S} \times \mathtt{S}$ is obtained form the transition relation of $\mathcal{N}$ by projection. Observe that any path in $\mathcal{G}(\mathcal{N})$ corresponds to a possible sequence of transitions in $\mathcal{N}$ since in each step, symbols under the reading heads can always change enabling the desired transition. Such statements are not possible for general PTM, since the reachability of any control-state still depends on the tape configurations and thus cannot be reduced to a reachability question in the induced graph. Nevertheless, for PNTM the following holds.

Lemma 4.5. *Given a PNTM $\mathcal{N} = (\mathtt{S}, \Sigma, \Gamma, \mathtt{M}, \mathtt{T}, \epsilon, w)$, for any CTL^* formula ϕ over sets $F_1 \cdots F_n$ of $\mathtt{Q}_1$-states $\cdots$ $\mathtt{Q}_n$-states for some $\mathtt{Q}_1 \cdots \mathtt{Q}_n \subseteq \mathtt{S}$, the set of global states s satisfying $s \models \phi$ is computable.*

Proof. Observe that checking $s \models \phi$ is equivalent to checking, in $\mathcal{G}(\mathcal{N})$, $\mathtt{s} \models \phi'$ where $\mathtt{s}$ is the control-state in s and ϕ' is the formula obtained from ϕ by replacing all occurrences of $F_1 \cdots F_n$ by $\mathtt{Q}_1 \cdots \mathtt{Q}_n$ respectively. It follows that the set of global states satisfying ϕ is exactly the set of $\mathtt{Q}$-states such that for any $\mathtt{s} \in \mathtt{Q}$, $\mathtt{s} \models \phi'$. Since $\mathcal{G}(\mathcal{N})$ is finite, the result follows by decidability of CTL^*model-checking in finite-state systems ([CGP99]). ☐

A PNTM $\mathcal{N}$ induces a Markov chain $\mathcal{M} = (S, P)$ on the set of global states. Each transition in $\mathcal{M}$ is also a combination of a noisy step followed by a normal step. However, in the noisy steps, we assume that cells not under the reading heads are not subjected to noise. Observe that this is different than the way noise is added in the model of [AC05] where, for instance, all cells are subject to noise. Intuitively, the noise doesn't affect the computations of the underlying Turing machine unless it changes a cell which is going to be visited by the reading head. Now, whether the content of that cell changes when the reading head reaches it or has changed in the previous steps; the resulting computation is the same.

In order to compensate for the missing noise, we assume a higher noise probability for the cell under the head. If the cell was last visited k time units ago, then we increase the noise probability to $1-(1-\epsilon)^k$. The probability of a transition in the induced Markov chain is obtained by multiplying the noise probability by the probability of the normal step described earlier.

Theorem 4.6. *Each Markov chain induced by a PNTM $\mathcal{N} = (\mathtt{S}, \Sigma, \Gamma, \mathtt{M}, \mathtt{T}, \epsilon, w)$ is coarse, effective and finitely spanning with respect to any set of $\mathtt{Q}$-states for some $\mathtt{Q} \subseteq \mathtt{S}$.*

Proof. Assume a PNTM $\mathcal{N} = (\mathtt{S}, \Sigma, \Gamma, \mathtt{M}, \mathtt{T}, \epsilon, w)$, a set $\mathtt{Q} \subseteq \mathtt{S}$ and the induced Markov chain $\mathcal{M}$. Let F be the set of $\mathtt{Q}$-states. Effectiveness of $\mathcal{M}$ w.r.t. F follows from the definition and Lemma 4.5. For any state $s \in S$, if $s \models \exists\Diamond F$ then there is a path in $\mathcal{G}(\mathcal{N})$ from the control-state of s to a control-state in $\mathtt{Q}$. Such a path has length at most $N = |\mathcal{G}(\mathcal{N})|$. Thus $\mathcal{M}$ has span N with respect to F. Along this path, it is possible that, in each step, each symbol under a reading head is subject to noise. Since in each step, $\mathtt{M}$ cells are subject to noise and each happens with probability $\geq \epsilon$, it follows that the probability of each successor is $\geq (\epsilon/|\Gamma|)^{\mathtt{M}}$. This gives the coarseness of $\mathcal{M}$. $\square$

This, combined with Lemma 3.6 and Lemma 3.7, yields the following corollary.

Corollary 4.7. *Each Markov chain induced by a PNTM $\mathcal{N} = (\mathtt{S}, \Sigma, \Gamma, \mathtt{M}, \mathtt{T}, \epsilon, w)$ is decisive and thus (by Lemma 3.2) strongly decisive with respect to any set of $\mathtt{Q}$-states for some $\mathtt{Q} \subseteq \mathtt{S}$.*

5. 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 $\mathcal{M} = (S, P)$
- A state $s_{init} \in S$
- A set of final states $F \subseteq S$

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 and any set of states F.

Lemma 5.1. $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) = 1$ *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 $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) < 1$. $\square$

The reverse implication of Lemma 5.1 holds only for Markov chains which satisfy certain conditions.

Lemma 5.2. *Given a Markov chain $\mathcal{M}$ and a set F such that $\mathcal{M}$ is decisive w.r.t. F, then we have that $s_{init} \not\models \widetilde{F}\ \underline{Before}\ F$ implies $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) = 1$.*

Proof. From $s_{init} \not\models \widetilde{F}\ \underline{Before}\ F$ it follows that $(s_{init} \models \Diamond\widetilde{F}) \subseteq (s_{init} \models \Diamond F)$. Therefore $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) = Prob_{\mathcal{M}}(s_{init} \models \Diamond F \vee \Diamond\widetilde{F})$. Since $\mathcal{M}$ is decisive w.r.t. F, it follows that $1 = Prob_{\mathcal{M}}(s_{init} \models \Diamond F \vee \Diamond\widetilde{F}) = Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$. □

Lemma 5.2 does not hold for general Markov chains; see Remark 6.3 in Section 6.

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

Theorem 5.3. *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 F of $\mathtt{Q}$-states is upward closed, we obtain from Corollary 4.2 that the Markov chain derived from our PVASS is decisive w.r.t. such F. Thus, by Lemma 5.1 and Lemma 5.2, 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 [May84, Jan90]. □

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

We will need the following definition for the proof.

Definition 5.5. We define a PVASS which weakly simulates a Minsky [Min67] 2-counter machine. Since this construction will be used in several proofs (Theorem 5.4 and Theorem 9.1), it contains a parameter $x > 0$ which will be instantiated as needed.

Consider a deterministic Minsky 2-counter machine M with a set of control-states K, initial control-state k_0, final accepting state k_{acc}, two counters c_1 and c_2 which are initially zero and the usual instructions of increment and test-for-zero-decrement. For technical reasons we require the following conditions on the behavior of M.

- Either M terminates in control-state k_{acc}, or
- M does not terminate. In this case we require that in its infinite run it infinitely often tests a counter for zero in a configuration where the tested counter contains a non-zero value.

We call a counter machine that satisfies these conditions an IT-2-counter machine (IT for 'infinitely testing'). Any 2-counter machine M' can be effectively transformed into an equivalent IT-2-counter machine M by the following operations. After every instruction of M' we add two new instructions: First increment c_1 by 1 (thus it is now certainly nonzero). Then test c_1 for zero (this test always yields answer 'no'), decrement it by 1 (so it has its original value again), and then continue with the next instruction of M'. So M is infinitely testing and accepts if and only if M' accepts. Since acceptance is undecidable for 2-counter machines [Min67], it follows that acceptance is also undecidable for IT-2-counter machines.

We construct a PVASS $\mathcal{V} = (\mathtt{S}, \mathtt{X}, \mathtt{T}, w)$ that weakly simulates M as follows. $\mathtt{S} = K \cup \{k^i \mid k \in K, i \in \{1, 2\}\} \cup \{err\}$ and $\mathtt{X} = \{c_1, c_2\}$. For every instruction $k_1 : c_i := c_i + 1;$ goto k_2 we add a transition $(k_1, \mathtt{op}, k_2)$ to T, where $\mathtt{op}(c_i) = 1$ and $\mathtt{op}(c_j) = 0$ for $j \neq i$ and $w((k_1, \mathtt{op}, k_2)) := 1$. For every instruction k_1 : If $c_i = 0$ then goto k_2 else $c_i := c_i - 1;$ goto k_3 we add the following transitions to T.

- α**:** $(k_1, \mathtt{op}_1, k_3)$ with $\mathtt{op}_1(c_i) = -1$ and $\mathtt{op}_1(c_j) = 0$ for $j \neq i$ and $w((k_1, \mathtt{op}_1, k_3)) = 1$.
- β**:** $(k_1, \mathtt{op}_2, k_2^i)$ with $\mathtt{op}_2(c_j) = 0$ for $j = 1, 2$ and $w((k_1, \mathtt{op}_2, k_2^i)) = x$ ($x > 0$ is a parameter of w).
- γ**:** $(k_2^i, \mathtt{op}_a, k_2)$ with $\mathtt{op}_a(c_j) = 0$ for $j = 1, 2$ and $w((k_2^i, \mathtt{op}_a, k_2)) = 1$.
- δ**:** $(k_2^i, \mathtt{op}_b, err)$ with $\mathtt{op}_b(c_i) = -1$ and $\mathtt{op}_b(c_j) = 0$ for $j \neq i$ and $w((k_2^i, \mathtt{op}_b, err)) = 1$.

Finally, to avoid deadlocks in $\mathcal{V}$, we add two self-loops (k_{acc}, op_l, k_{acc}) and (err, op_l, err) with $\mathtt{op}_l(c_j) = 0$ for $j = 1, 2$ and weight 1.

Proof. (of Theorem 5.4) Since F is upward closed, we obtain from Corollary 4.2 that the Markov chain derived from our PVASS is decisive w.r.t. F. Thus, by Lemma 5.1 and Lemma 5.2, we have $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) < 1 \iff s_{init} \models \widetilde{F}\ \underline{Before}\ F$. Now we show that the condition $s_{init} \models \widetilde{F}\ \underline{Before}\ F$ is undecidable if F is a general upward closed set. We use the IT-2-counter machine M and the PVASS $\mathcal{V}$ from Def. 5.5 and instantiate parameter $x := 1$. Let F be the set of configurations where transitions of type δ are enabled. This set is upward-closed, because of the monotonicity of VASS, and effectively constructible (i.e., its finitely many minimal elements). It follows directly from the construction in Def. 5.5 that a transition of type δ is enabled if and only if the PVASS has been unfaithful in the simulation of the 2-counter machine, i.e., if a counter was nonzero and a 'zero' transition (of type β) has wrongly been taken instead of the correct 'decrement' transition (of type α).

If the 2-counter machine M accepts then there is a run in the PVASS $\mathcal{V}$ which faithfully simulates the run of M and thus never enables transitions of type δ and thus avoids the set F. Since the k_{acc}-states have no outgoing transitions (except for the self-loop), they are trivially contained in $\widetilde{F}$. Thus $s_{init} \models \widetilde{F}\ \underline{Before}\ F$.

If the 2-counter machine M does not accept then its run is infinite. By our convention in Def. 5.5, M is an IT-2-counter machine and every infinite run must contain infinitely many nontrivial tests for zero. Thus in our PVASS $\mathcal{V}$, the set F is reachable from every reachable state s' which was reached in a faithful simulation of M, i.e., without visiting F before. Therefore in $\mathcal{V}$ the set $\widetilde{F}$ cannot be reached unless F is visited first, and so we get $s_{init} \not\models \widetilde{F}\ \underline{Before}\ F$.

We obtain that M accepts iff $s_{init} \models \widetilde{F}\ \underline{Before}\ F$ iff $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) < 1$. This proves the undecidability of the problem for the case of $\rho = 1$.

To show the undecidability for general $\rho \in (0, 1]$ we modify the construction as follows. Consider a new PVASS $\mathcal{V}'$ which with probability ρ does the same as $\mathcal{V}$ defined above and with probability $1-\rho$ immediately goes to the accepting state k_{acc}. Then the IT-2-CM M accepts iff $Prob_{\mathcal{M}'}\left(s_{init} \models \Diamond F\right) \neq \rho$. □

Notice the difference between Theorem 5.3 and Theorem 5.4 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 5.6. *Consider a PLCS $\mathcal{L}$ and an effectively representable set of final states F. Then the question $Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) = 1$ is decidable.*

Proof. By Corollary 4.4, the Markov chain induced by $\mathcal{L}$ is decisive w.r.t. such F. Thus we obtain from Lemma 5.1 and Lemma 5.2 that $Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) = 1$ iff $s_{init} \not\models \widetilde{F}\ \underline{Before}\ F$. This condition can be checked with a standard construction for LCS (from [AR03]) as follows. First one can effectively compute the set $\widetilde{F} = \overline{Pre^*(F)}$ using the techniques from, e.g., [AJ96]. 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 [Hig52]. We get that $X = X_n$ is effectively constructible. Finally we have that $Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) = 1$ iff $s_{init} \notin X$, which can be effectively checked. □

Notice that, unlike in earlier work [AR03, BS03], it is not necessary to compute the finite attractor of the PLCS-induced Markov chain for Theorem 5.6. It suffices to know that it exists.

Theorem 5.7. *For a PNTM $\mathcal{N} = (\mathtt{S}, \Sigma, \Gamma, \mathtt{M}, \mathtt{T}, \epsilon, w)$, the question $Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) = 1$ is decidable for any set F of $\mathtt{Q}$-states for some $\mathtt{Q} \subseteq \mathtt{S}$.*

Proof. By Corollary 4.7, we obtain that the Markov chain $\mathcal{M}$ derived from $\mathcal{N}$ is decisive w.r.t. F. This combined with Lemma 5.1 and Lemma 5.2 yields $Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) < 1 \iff s_{init} \models \widetilde{F}\ \underline{Before}\ F$. Observe that since F is a set of $\mathtt{Q}$-states, we obtain by Lemma 4.5 that we can compute a set $\mathtt{Q}' \subseteq \mathtt{S}$ such that $\widetilde{F} = \mathtt{Q}'$-states. Since F and $\widetilde{F}$ are sets of $\mathtt{Q}$-states and $\mathtt{Q}'$-states respectively, it follows by Lemma 4.5 that the question $s_{init} \models \widetilde{F}\ \underline{Before}\ F$ is decidable. This gives the result. □

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

Lemma 5.8. *$Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) = 0$ iff $s_{init} \not\models \exists \Diamond F$.*

Since the (submarking) reachability problem for Petri nets is decidable [May84, Jan90], we get the following consequence of Lemma 5.8.

Theorem 5.9. *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* [Jan90] *(in particular any upward closed set, any finite set, and their complements), then the question $Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) = 0$ is decidable.*

From Lemma 5.8 and the result that for LCS the set of all predecessors of any effectively representable set can be effectively constructed (e.g., [AJ96]), we get the following.

Theorem 5.10. *Given a PLCS $\mathcal{L}$ and a set of final states F which is effectively representable, then the question $Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) = 0$ is decidable.*

By Lemma 4.5, we obtain the following.

Theorem 5.11. *Given a PNTM $\mathcal{N} = (\mathtt{S}, \Sigma, \Gamma, \mathtt{M}, \mathtt{T}, \epsilon, w)$ and a set F of $\mathtt{Q}$-states for some $\mathtt{Q} \subseteq \mathtt{S}$, then the question $Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) = 0$ is decidable.*

6. 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 $\mathcal{M} = (S, P)$
• A state $s_{init} \in S$
• A set of final states $F \subseteq S$
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 and any set of states F.

Lemma 6.1. *Let $\mathcal{M} = (S, P)$ be a Markov chain and $F \subseteq S$. Then $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1$ 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 starting from s_{init} leading to a state s s.t. $s \models \forall\Box\neg F$. The set of all s_{init}-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 $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) < 1$. $\square$

The reverse implication of Lemma 6.1 does not hold generally, but it is true for strongly decisive Markov chains.

Lemma 6.2. *Given a Markov chain $\mathcal{M}$ and a set F such that $\mathcal{M}$ is strongly decisive w.r.t. F, then we have that $s_{init} \models \forall\Box\exists\Diamond F$ implies $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1$.*

Proof. We show that $Prob_{\mathcal{M}}(s_{init} \models \Diamond\Box\neg F) = 0$ which implies the result.

If $s_{init} \models \forall\Box\exists\Diamond F$ then every state s reached by runs in $(s_{init} \models \Diamond\Box\neg F)$ satisfies $s \models \exists\Diamond F$. Therefore $(s_{init} \models \Diamond\Box\neg F) \subseteq (s_{init} \models \Box\neg\widetilde{F})$.

Since the Markov chain is strongly decisive, $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F} \vee \Box\Diamond F) = 1$ and thus $Prob_{\mathcal{M}}(s_{init} \models \Box\neg\widetilde{F} \wedge \Diamond\Box\neg F) = 0$. It follows from the inclusion above that $0 = Prob_{\mathcal{M}}(s_{init} \models \Box\neg\widetilde{F} \wedge \Diamond\Box\neg F) = Prob_{\mathcal{M}}(s_{init} \models \Box\neg\widetilde{F}) \geq Prob_{\mathcal{M}}(s_{init} \models \Diamond\Box\neg F)$. $\square$

Remark 6.3. Neither Lemma 5.2 nor Lemma 6.2 hold for general Markov chains. A counterexample is the Markov chain $\mathcal{M} = (S, P)$ of the "gambler's ruin" problem [Fel66] where $S = \mathbb{N}$, $P(i, i+1) := x$, $P(i, i-1) := 1 - x$ 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 6.4. *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 Corollary 4.2 that the Markov chain derived from our PVASS is strongly decisive w.r.t. F. Thus it follows from Lemma 6.1 and Lemma 6.2 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 [May84, Jan90]. $\Box$

Notice the surprising contrast of the decidability of repeated reachability of Theorem 6.4 to the undecidability of simple reachability in Theorem 5.4.

Theorem 6.5. *Consider a PLCS $\mathcal{L}$ and an effectively representable set of final states F. Then the question $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1$ is decidable.*

Proof. By Corollary 4.4, $\mathcal{L}$ induces a strongly decisive Markov chain. Thus, we obtain from Lemma 6.1 and Lemma 6.2 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. $\Box$

Similarly as in Theorem 5.6, it is not necessary to compute the finite attractor of the PLCS-induced Markov chain for Theorem 6.5.

Theorem 6.6. *Let $\mathcal{N} = (\mathtt{S}, \Sigma, \Gamma, \mathtt{M}, \mathtt{T}, \epsilon, w)$ be a PNTM and F a set of $\mathtt{Q}$-states for some $\mathtt{Q} \subseteq \mathtt{S}$. Then the question $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1$ is decidable.*

Proof. Since F is a set of $\mathtt{Q}$-states, we obtain from Corollary 4.7 that the Markov chain derived from $\mathcal{N}$ is strongly decisive w.r.t. F. Thus it follows from Lemma 6.1 and Lemma 6.2 that $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1 \iff s_{init} \models \forall\Box\exists\Diamond F$. We can check this condition by Lemma 4.5. $\Box$

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

Lemma 6.7. *For any Markov chain and any set of states F, we have $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) \neq 0$ implies $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) \neq 1$*

The following lemma implies that the reverse implication holds for strongly decisive Markov chains.

Lemma 6.8. *Given a Markov chain $\mathcal{M}$ which is strongly decisive w.r.t. a given set F, then we have that $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 1 - Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F})$.*

Proof. $1 - Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) = Prob_{\mathcal{M}}(s_{init} \models \Box\neg\widetilde{F}) = Prob_{\mathcal{M}}(s_{init} \models \Box\neg\widetilde{F} \wedge \Box\Diamond F) + Prob_{\mathcal{M}}(s_{init} \models \Box\neg\widetilde{F} \wedge \Diamond\Box\neg F)$. Since the Markov chain is strongly decisive, it follows that $Prob_{\mathcal{M}}(s_{init} \models \Box\neg\widetilde{F} \wedge \Diamond\Box\neg F) = 0$. Moreover, we have $(s_{init} \models \Box\Diamond F) \subseteq (s_{init} \models \Box\neg\widetilde{F})$. Therefore, $1 - Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) = Prob_{\mathcal{M}}(s_{init} \models \Box\neg\widetilde{F}) = Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F)$. $\Box$

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

Lemma 6.9. *For any Markov chain $\mathcal{M}$ and any set of states F, $s_{init} \models \exists\Diamond\widetilde{\widetilde{F}}$ implies $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) \neq 1$.*

Proof. If $s_{init} \models \exists\Diamond\widetilde{\widetilde{F}}$ then there exists a finite path π starting from s_{init} and leading to some state $s \in \widetilde{\widetilde{F}}$ and thus $s \models \forall\Box\exists\Diamond F$. Therefore $s \not\models \exists\Diamond\widetilde{F}$. Consider the set R of s_{init}-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 $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) \leq 1 - Prob_{\mathcal{M}}(R) < 1$. ☐

If $\mathcal{M}$ is decisive w.r.t. $\widetilde{F}$ then the reverse implication holds.

Lemma 6.10. *Given a Markov chain $\mathcal{M}$ and a set of states F s.t. $\mathcal{M}$ is decisive w.r.t. $\widetilde{F}$, then the condition $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) \neq 1$ implies $s_{init} \models \exists\Diamond\widetilde{\widetilde{F}}$.*

Proof. As $\mathcal{M}$ is decisive w.r.t. $\widetilde{F}$, it follows that $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F} \vee \Diamond\widetilde{\widetilde{F}}) = 1$. Since by assumption $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{F}) \neq 1$, it follows that $Prob_{\mathcal{M}}(s_{init} \models \Diamond\widetilde{\widetilde{F}}) > 0$. The result follows immediately. ☐

Theorem 6.11. *For any Markov chain $\mathcal{M}$ and any set F s.t. $\mathcal{M}$ is decisive w.r.t. F*

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

For any Markov chain $\mathcal{M}$ and any set F s.t. $\mathcal{M}$ is decisive w.r.t. F and w.r.t. $\widetilde{F}$

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

Proof. Directly from Lemma 6.7, 6.8, 6.9 and 6.10. ☐

Remark 6.12. Observe that decisiveness w.r.t. a given set F does not imply decisiveness w.r.t. $\widetilde{F}$. Therefore the reverse implication of Lemma 6.9 does not hold in general. In particular, it holds for Markov chains with a finite attractor (since they are decisive w.r.t. every set), but not generally when we have global coarseness. This is because global coarseness depends on the set of final states. Global coarseness of a Markov chain w.r.t. a certain set F does not imply global coarseness w.r.t. the set $\widetilde{F}$.

Theorem 6.13. *Consider a PLCS $\mathcal{L}$ and an effectively representable set of final states F. Then the question $Prob_{\mathcal{M}}\,(s_{init} \models \Box\Diamond F) = 0$ is decidable.*

Proof. By Corollary 4.4, $\mathcal{L}$ induces a strongly decisive Markov chain w.r.t. every set of states, in particular it is decisive w.r.t. F and $\widetilde{F}$. Therefore, by Theorem 6.11, it suffices to check if $s_{init} \not\models \exists\Diamond\widetilde{\widetilde{F}}$. Since the upward closure of F is effectively constructible, one can effectively construct a symbolic representation of the set of all states which satisfy $\exists\Diamond\widetilde{\widetilde{F}}$ (using the techniques from, e.g., [AJ96]) and check if s_{init} is not in this set. ☐

Theorem 6.14. *Given a PNTM $\mathcal{N} = (\mathtt{S}, \Sigma, \Gamma, \mathtt{M}, \mathtt{T}, \epsilon, w)$ and a set F of $\mathtt{Q}$-states for some $\mathtt{Q} \subseteq \mathtt{S}$, the question $Prob_{\mathcal{M}}\,(s_{init} \models \Box\Diamond F) = 0$ is decidable.*

Proof. Observe that since F is a set of $\mathtt{Q}$-states, by Lemma 4.5, we can construct the set $\mathtt{Q}' \subseteq \mathtt{S}$ such that $\widetilde{F}$ is exactly the set of $\mathtt{Q}'$-states. By Corollary 4.7, the Markov chain induced by $\mathcal{N}$ is decisive w.r.t. any set of $\mathtt{Q}''$-states for some $\mathtt{Q}'' \subseteq \mathtt{S}$. In particular, it is decisive w.r.t. $\mathtt{Q}'$-states $= \widetilde{F}$. By Theorem 6.11, it suffices to check if $s_{init} \not\models \exists\Diamond\widetilde{\widetilde{F}}$ which is again decidable by Lemma 4.5. ☐

For PVASS, decidability of $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) = 0$ 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 5.3 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 6.3), but let now $F := \{1, 2, \dots\}$ (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 $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F)$ cannot be effectively expressed in the first-order theory of the reals $(\mathbb{R}, +, *, \leq)$, as shown in Section 9, Remark 9.5.

7. Approximate Quantitative Reachability

In this section we consider the approximate quantitative reachability problem.

APPROX_QUANT_REACH
Instance
- A Markov chain $\mathcal{M} = (S, P)$
- A state $s_{init} \in S$
- A set of states $F \subseteq S$
- A rational $\varepsilon > 0$

Task Compute a rational θ such that $\theta \leq Prob_{\mathcal{M}}(s_{init} \models \Diamond F) \leq \theta + \varepsilon$.

First, we present a path enumeration algorithm, based on [IN97], for solving the problem, and then we show that the algorithm is guaranteed to terminate for all instances where $\mathcal{M}$ is decisive w.r.t. F [1].

Given an effective Markov chain $\mathcal{M} = (S, P)$, a state $s_{init} \in S$, a set F and a positive $\varepsilon \in \mathbb{Q}_{>0}$, 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 certainly reached (and certainly not reached, 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 $1 - \varepsilon$.

[1] Termination of the algorithm is stated (without giving a proof) in [Rab03] for the case of Markov chains which contain a finite attractor.

Algorithm 1 – APPROX_QUANT_REACH
Input
An effective Markov chain $\mathcal{M} = (S, P)$, a state $s_{init} \in S$, a set $F \subseteq S$ and a positive $\varepsilon \in \mathbb{Q}_{>0}$.
Return value
A rational θ such that $\theta \leq Prob_{\mathcal{M}}\left(s_{init} \models \Box\Diamond F\right) \leq \theta + \varepsilon$
Variables
Yes, No: $\mathbb{Q}$ (initially all are set to 0)
$store$: queue with elements in $S \times \mathbb{Q}$

begin
$store := (s_{init}, 1)$
repeat
  remove (s, r) from $store$
  if $s \in F$ **then** $Yes := Yes + r$
  else if $s \in \widetilde{F}$ **then** $No := No + r$
  else **for each** $s' \in Post(s)$
    add $(s', r \cdot P(s, s'))$ to the end of $store$
until $Yes + No \geq 1 - \varepsilon$
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 7.1. *If Algorithm 1 terminates at depth j then*
$$Yes^j(\mathcal{M}, s_{init}) \leq Prob_{\mathcal{M}}(s_{init} \models \Diamond F) \leq Yes^j(\mathcal{M}, s_{init}) + \varepsilon$$

Proof. It is straightforward to check that for each $j \geq 0$ we have
$$Yes^j(\mathcal{M}, s_{init}) \leq Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$$
and
$$No^j(\mathcal{M}, s_{init}) \leq Prob_{\mathcal{M}}(s_{init} \models \widetilde{F} \ \underline{Before}\ F)$$
We notice that
$$Prob_{\mathcal{M}}(s_{init} \models \Diamond F) \leq 1 - Prob_{\mathcal{M}}(s_{init} \models \widetilde{F} \ \underline{Before}\ F)$$
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 $Yes^j(\mathcal{M}, s_{init}) + No^j(\mathcal{M}, s_{init}) \geq 1 - \varepsilon$ when the algorithm terminates. $\square$

Lemma 7.2. *Algorithm 1 terminates in case the Markov chain $\mathcal{M}$ is decisive w.r.t. F.*

Proof. Since $\mathcal{M}$ is decisive we have $Prob_{\mathcal{M}}(s_{init} \models \Diamond F \vee \Diamond \widetilde{F}) = 1$. Therefore $\lim_{j\to\infty}(Yes^j + No^j) = 1$, which implies termination of the algorithm. $\square$

From Lemma 7.1 and Lemma 7.2 it follows that APPROX_QUANT_REACH is solvable for Markov chains which are globally coarse w.r.t. the target set and for Markov chains which contain a finite attractor. This, together with Theorem 4.1 and Theorem 4.3, yield the following theorems.

Theorem 7.3. APPROX_QUANT_REACH *is solvable for PVASS in case F is upward closed.*

Theorem 7.4. APPROX_QUANT_REACH *is solvable for PLCS in case F is effectively representable.*

Theorem 7.5. APPROX_QUANT_REACH *is solvable for PNTM in case F is a set of* $\mathtt{Q}$*-states.*

8. APPROXIMATE QUANTITATIVE REPEATED REACHABILITY

In this section we approximate the probability of reaching a given set of states infinitely often, i.e., we compute arbitrarily close approximations of $Prob_{\mathcal{M}}\left(s_{init} \models \Box\Diamond F\right)$.

APPROX_QUANT_REP_REACH
Instance
- A Markov chain $\mathcal{M} = (S, P)$
- A state $s_{init} \in S$
- A set $F \subseteq S$
- A rational $\varepsilon > 0$

Task Compute a rational θ such that $\theta \leq Prob_{\mathcal{M}}\left(s_{init} \models \Box\Diamond F\right) \leq \theta + \varepsilon$.

We present an algorithm which is a modification of Algorithm 1 (in Section 7) and show that it is guaranteed to terminate for all Markov chains which are decisive w.r.t. both F and $\widetilde{F}$.

Algorithm 2 – APPROX_QUANT_REP_REACH
Input
An effective Markov chain $\mathcal{M} = (S, P)$, a state $s_{init} \in S$, a set $F \subseteq S$ and a positive $\varepsilon \in \mathbb{Q}_{>0}$.
Return value
A rational θ such that $\theta \leq Prob_{\mathcal{M}}\left(s_{init} \models \Box\Diamond F\right) \leq \theta + \varepsilon$
Variables
Yes, No: $\mathbb{Q}$ (initially all are set to 0)
$store$: queue with elements in $S \times \mathbb{Q}$

begin
$store := (s_{init}, 1)$
repeat
  remove (s, r) from $store$
  if $s \in \widetilde{\widetilde{F}}$ **then** $Yes := Yes + r$
  else if $s \in \widetilde{F}$ **then** $No := No + r$
  else for each $s' \in Post(s)$
    add $(s', r \cdot P(s, s'))$ to the end of $store$
until $Yes + No \geq 1 - \varepsilon$
return Yes
end

We define $Yes^j(\mathcal{M}, s_{init})$ and $No^j(\mathcal{M}, s_{init})$ as the values of the variables Yes and No after the algorithm has explored the "reachability tree" with root s_{init} up-to depth j, similarly as for Algorithm 1. The following Lemma states partial correctness of Algorithm 2.

Lemma 8.1. *For a Markov chain $\mathcal{M}$ and a set F such that $\mathcal{M}$ is strongly decisive w.r.t. F, if Algorithm 2 terminates at depth j then*

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

Proof. If Algorithm 2 reaches some state $s \in \widetilde{\widetilde{F}}$ (at line 4.) then we have $s \models \forall\Box\exists\Diamond F$. Since $\mathcal{M}$ is strongly decisive, it follows from Lemma 6.2 that $Prob_{\mathcal{M}}(s \models \Box\Diamond F) = 1$. Thus, for each $j \ge 0$, we have $Yes^j(\mathcal{M}, s_{init}) \le Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F)$. Similarly, if the algorithm reaches some state $s \in \widetilde{F}$ (at line 5.) then we have $(s \models \Box\Diamond F) = \emptyset$. Thus, for each $j \ge 0$, we have $No^j(\mathcal{M}, s_{init}) \le Prob_{\mathcal{M}}(s_{init} \not\models \Box\Diamond F) = 1 - Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F)$. It follows that $Yes^j(\mathcal{M}, s_{init}) \le Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F) \le 1 - No^j(\mathcal{M}, s_{init})$. The result follows from the fact that $Yes^j(\mathcal{M}, s_{init}) + No^j(\mathcal{M}, s_{init}) \ge 1 - \varepsilon$ when the algorithm terminates. □

Lemma 8.2. *Algorithm 2 terminates if $\mathcal{M}$ is decisive w.r.t. $\widetilde{F}$.*

Proof. Since $\mathcal{M}$ is decisive w.r.t. $\widetilde{F}$, we have $Prob_{\mathcal{M}}\left(s_{init} \models \Diamond\widetilde{F} \vee \Diamond\widetilde{\widetilde{F}}\right) = 1$. It follows that $\lim_{j\to\infty}(Yes^j + No^j) = 1$, which implies termination. □

Note that Algorithm 2 only works for Markov chains which are decisive w.r.t. both F and $\widetilde{F}$. Decisiveness w.r.t. $\widetilde{F}$ is required for termination (Lemma 8.2), while decisiveness w.r.t. F is required for correctness (Lemma 8.1).

Theorem 8.3. APPROX_QUANT_REP_REACH *is solvable for PLCS in case F is effectively representable.*

Proof. By Theorem 4.3, a Markov chain induced by a PLCS is decisive w.r.t. every set, in particular w.r.t. F and $\widetilde{F}$. Thus it follows from Lemma 8.1 and Lemma 8.2, that Algorithm 2 solves APPROX_QUANT_REP_REACH for PLCS. □

Theorem 8.4. APPROX_QUANT_REP_REACH *is solvable for PNTM in case F is a set of* $\mathtt{Q}$*-states.*

Proof. Assume a PNTM $\mathcal{N} = (\mathtt{S}, \Sigma, \Gamma, \mathtt{M}, \mathtt{T}, \epsilon, w)$ and a set F of $\mathtt{Q}$-states for some $\mathtt{Q} \subseteq \mathtt{S}$. Since F is the set of $\mathtt{Q}$-states, it follows by Lemma 4.5 that $\widetilde{F}$ is also a set of $\mathtt{Q}'$-states where $\mathtt{Q}' \subseteq \mathtt{S}$. By Theorem 4.6, the Markov chain induced by a PNTM is decisive w.r.t. every set of $\mathtt{Q}''$-states, in particular w.r.t. F and $\widetilde{F}$. Thus it follows from Lemma 8.1 and Lemma 8.2, that Algorithm 2 solves APPROX_QUANT_REP_REACH for PNTM. □

9. Exact Quantitative Analysis

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

EXACT_QUANT_REACH
Instance • A Markov chain $\mathcal{M} = (S, P)$ • A state $s_{init} \in S$ • A set of states $F \subseteq S$ • A rational ν
Task Check whether $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) \ge \nu$.

By Theorem 5.4, EXACT_QUANT_REACH is undecidable for PVASS and upward closed sets F. If F is a set of $\mathtt{Q}$-states then decidability of EXACT_QUANT_REACH is open for PVASS, PLCS and PNTM.

However, for PVASS, PLCS and PNTM, 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 a uniform way in the first-order theory of the reals $(\mathbb{R}, +, *, \leq)$, or any other decidable logical theory with first-order quantification. By 'expressed in a uniform way' we mean that parameters of the system (e.g., transition weights for PVASS, the message loss probability for PLCS or the noise parameter for PNTM) should be reflected directly by constants in the constructed formula (see the remarks at the end of this section for details). This negative result for PVASS, PLCS and PNTM is in contrast to the situation for probabilistic pushdown automata (probabilistic recursive systems) for which this probability can be effectively expressed in $(\mathbb{R}, +, *, \leq)$ [EKM04, EY05b, EE04].

Theorem 9.1. *Let $\mathcal{V} = (\mathtt{S}, \mathtt{X}, \mathtt{T}, w)$ be a PVASS and $w_1, \dots, w_n \in \mathbb{Q}$ 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 $\Phi(p, w_1, \dots, w_n)$ with parameters w_i and p which expresses the probability $Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$, i.e., $\Phi(p, w_1, \dots, w_n) = true$ iff $p = Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$.

Proof. We assume the contrary and derive a contradiction. Consider the IT-2-CM M and the PVASS $\mathcal{V}$ with parameter x from Def. 5.5, i.e., let $w_1 = x$ and $w_i = 1$ for $i > 1$. Let F be the set of err-states. Suppose that one could effectively construct the $(\mathbb{R}, +, *, \leq)$-formula $\Phi(p, w_1, \dots, w_n) = \Phi(p, x)$ with the required properties.

If a counter is tested for zero then, in our weak simulation by the PVASS, there are two cases:

- If the counter contains zero, then only one transition (the one of type β) is enabled and the simulation is faithful. After firing transition β, only transition γ (but not δ) is enabled.
- If the counter does not contain zero then two transitions, α and β, are enabled. The probability of choosing α, the faithful simulation, is $1/(1 + x)$ and the probability of choosing β (the wrong transition in this case) is $x/(1 + x)$. If β is fired then both γ and δ are enabled.

If the IT-2-counter machine M accepts after a finite number L of steps, then it can make at most L tests for zero. Thus the probability of ever choosing the wrong transition is bounded from above by $1 - 1/(1 + x)^L$. A transition of type δ, leading to the err-state can only be taken if a wrong transition has been taken first. Thus the probability of reaching the err-state is also bounded from above by $1 - 1/(1 + x)^L$. For $x \to 0$ this probability converges to 0. Thus we have $\exists x > 0\, \exists p\, (\Phi(p, x) \wedge p < 1/10)$.

If the IT-2-counter machine M does not accept then its infinite run will contain infinitely many tests for zero on counters with are non-zero by Def. 5.5. In each of those tests, the chance of firing the wrong transition (type β) is $x/(1 + x) > 0$. Thus it will happen eventually with probability 1. If the wrong transition has been fired then the probability of going to the err-state by the next transition is $1/2$ (competing enabled transitions γ and δ of weight 1 each). Thus the probability of reaching the err-state is at least $1/2$, i.e., $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) \geq 1/2$, regardless of the value of x (provided $x > 0$).

It follows that M accepts if and only if

$$\exists x > 0\, \exists p\, (\Phi(p, x) \wedge p < 1/10)$$

If the formula $\Phi(p, x)$ was effectively constructible, then the formula $\exists x > 0\, \exists p\, (\Phi(p, x) \wedge p < 1/10)$ would also be effectively constructible and the question would be decidable, because of the decidability of $(\mathbb{R}, +, *, \leq)$ [Tar51]. This is a contradiction, since acceptance of M is undecidable. □

Remark 9.2. In Theorem 9.1 the set F is the set of err-states, i.e., an upward- and downward-closed set. However, the result holds just as well if F is a single configuration. To show this, it suffices to modify the construction as follows. Add two new transitions $(err, \mathtt{op}_i, err)$ with weight 1 and $\mathtt{op}_i(c_i) = -1$ and $\mathtt{op}_i(c_j) = 0$ for $i \neq j$ and $i, j \in \{1, 2\}$. The only other possible transition in err-states is the self-loop which changes nothing. So almost every run starting in an err-state will eventually reach $(err, 0, 0)$. Thus the probability of eventually reaching configuration $(err, 0, 0)$ is the same as that of eventually reaching some err-state.

For PLCS, the result corresponding to Theorem 9.1 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 [BZ83] (unlike for VASS [May84]). The following theorem shows a stronger non-expressibility result even for the restricted case of $\lambda > 0$.

Theorem 9.3. *Let $\mathcal{L}$ be a PLCS with message loss probability $\lambda > 0$ and F a set of $\mathtt{Q}$-states for some $\mathtt{Q} \subseteq \mathtt{S}$. Then 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., such that for any $\lambda > 0$ one has $\Phi(p, \lambda) = true$ iff $p = Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$.*

Proof. We assume the contrary and derive a contradiction. It is known that the termination problem for deterministic non-lossy non-probabilistic FIFO-channel systems $\mathcal{L}'$ is undecidable [BZ83]. Given such a FIFO-channel system $\mathcal{L}'$ (and the contrary of our theorem) we will construct a PLCS $\mathcal{L}$ and a $(\mathbb{R}, +, *, \leq)$ formula Ψ s.t. Ψ is true if and only if $\mathcal{L}'$ terminates.

Consider a deterministic non-lossy non-probabilistic FIFO-channel system $\mathcal{L}'$ which starts in control-state s_{init} and the empty channel. Let k_{acc} be the final accepting control-state of $\mathcal{L}'$. $\mathcal{L}'$ either eventually reaches the final accepting control-state k_{acc} and terminates, or continues forever. We construct the PLCS $\mathcal{L}$ by modifying $\mathcal{L}'$ as follows. First we add a new control-state err to the system. Then, for every transition t_1 of $\mathcal{L}'$ we add an additional transition t_2 to $\mathcal{L}$. Transition t_2 is almost identical to t_1, except that its target control-state is err. We assign the same transition weight 1 to every transition in $\mathcal{L}$. Thus, in every single step, the system $\mathcal{L}$ has probability $1/2$ of going to an err-state. In order to avoid deadlocks, we add self-loops to the control-states k_{acc} and err. These are the only possible transitions from k_{acc} and err. In particular, it is impossible to get to an err-state from any k_{acc}-state or vice-versa. Finally, we add the message loss probability $\lambda > 0$ to $\mathcal{L}$, i.e., in every step any message in transit is lost with probability λ. Let F be the set of k_{acc}-states.

If $\mathcal{L}'$ terminates, then it terminates in a finite number L of steps. Since one starts with the empty channel, the maximal number of messages in transit at any time in this run is bounded from above by L. The PLCS $\mathcal{L}$ can imitate this run of $\mathcal{L}'$ (however, $\mathcal{L}$ also has many other possible runs). The probability that none of the (at most L) messages is lost in any single step is bounded from below by $(1-\lambda)^L$. Thus $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) \geq \left((1-\lambda)^L \cdot \frac{1}{2}\right)^L$. This is the probability that $\mathcal{L}$ faithfully simulates the behavior of $\mathcal{L}'$. No messages are lost and all transitions leading to err-states are avoided. Now let $\epsilon := (0.9)^{L^2}(0.5)^L$. It follows that

$$\exists \epsilon > 0. \forall \lambda : 0 < \lambda \leq 0.1 \, Prob_{\mathcal{M}}(s_{init} \models \Diamond F) \geq \epsilon$$

Now we consider the case that $\mathcal{L}'$ does not terminate. Let G be the set of all err-states. In particular, $G \subseteq \widetilde{F}$. Since G is not reachable from F either, we have $Prob_{\mathcal{M}}(s_{init} \models \neg\Diamond F) \geq Prob_{\mathcal{M}}(s_{init} \models \Diamond G)$. Consider all those runs of $\mathcal{L}$ which do not lose any messages in the first

N steps and reach G after at most N steps. These runs faithfully simulate the first N steps of $\mathcal{L}'$, unless they go to an err-state. In particular, they do not reach k_{acc}, because $\mathcal{L}'$ does not reach k_{acc}, since $\mathcal{L}'$ does not terminate. We have

$$Prob_{\mathcal{M}}\left(s_{init} \models \Diamond G\right) \geq (1-\lambda)^{N^2}\left(1-(0.5)^N\right)$$

Now let $N := \left(\frac{1}{\lambda}\right)^{1/4}$. We obtain

$$\lim_{\lambda \to 0} Prob_{\mathcal{M}}\left(s_{init} \models \Diamond G\right) = 1$$

and thus

$$\lim_{\lambda \to 0} Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) = 0$$

and finally

$$\neg\left(\exists \epsilon > 0.\, \forall \lambda : 0 < \lambda \leq 0.1\, Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) \geq \epsilon\right)$$

Therefore, $\mathcal{L}'$ terminates if and only if

$$\exists \epsilon > 0.\, \forall \lambda : 0 < \lambda \leq 0.1\, Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) \geq \epsilon$$

Assume that for $\mathcal{L}$ the $(\mathbb{R}, +, *, \leq)$-formula $\Phi(p, \lambda)$ was effectively constructible. Then $\mathcal{L}'$ terminates if and only if

$$\Psi := \exists \epsilon > 0.\, \forall \lambda : 0 < \lambda \leq 0.1\, \exists p\, (\Phi(p, \lambda) \wedge p \geq \epsilon)$$

Since $(\mathbb{R}, +, *, \leq)$ is decidable [Tar51] one can decide if Ψ is true and thus if $\mathcal{L}'$ terminates. Contradiction. □

The proof of the corresponding result for PNTM is similar to the one for PLCS. Note that the following Theorem 9.4 also holds for NTM, since the proof does not use the extensions of PNTM over NTM.

Theorem 9.4. *Let $\mathcal{N}$ be a (P)NTM with noise probability $\epsilon > 0$ and F a set of $\mathtt{Q}$-states for some $\mathtt{Q} \subseteq \mathtt{S}$. Then it is impossible to effectively construct a $(\mathbb{R}, +, *, \leq)$ formula $\Phi(p, \epsilon)$ with parameters $\epsilon > 0$ and p which expresses the probability $Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right)$, i.e., such that for any $\epsilon > 0$ one has $\Phi(p, \epsilon) = true$ iff $p = Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right)$.*

Proof. We assume the contrary and derive a contradiction. In the proof, we start form a Turing machine $\mathcal{N}'$. Then we construct an PNTM $\mathcal{N}$ and a $(\mathbb{R}, +, *, \leq)$-formula Ψ s.t. Ψ is true if and only if $\mathcal{N}'$ halts. The contradiction follows from the undecidability of the halting problem for Turing-machines [HU79].

Assume a single-tape Deterministic Turing machine $\mathcal{N}'$ that starts in an initial control-state $\mathtt{s}_0$ with an input word ω_0. Once started, the machine either reaches a control-state $\mathtt{s}_h$ and halts, or runs forever. We derive a PNTM $\mathcal{N}$ from $\mathcal{N}'$ as follows. First, we add a control-state $\mathtt{s}_e$ to the system and a special symbol σ to the tape alphabet. Then, from every state, we add a transition that reads σ and moves to $\mathtt{s}_e$ where it halts. We add loops to $\mathtt{s}_e$ and $\mathtt{s}_h$ to avoid deadlocks. Finally, we assume that $\mathcal{N}$ has a noise probability $\epsilon > 0$ and all the transitions have weight 1.

Assume that the runs of $\mathcal{N}$ start all from an initial (global) state s_{init} where the control-state is $\mathtt{s}_0$ and the initial word is ω_0. Let F be the set of $\mathtt{s}_h$-states. If $\mathcal{N}'$ halts then it halts after a finite number of steps L. The runs which faithfully imitate those of $\mathcal{N}'$ avoid noise in each step. Assume that the computation reached a step where the cell under the head was visited k time units ago. Then, the probability to avoid the noise and move to the next configuration is exactly

$1 - \sum_{\sigma' \in \Gamma} \frac{1-(1-\epsilon)^k}{|\Gamma|} = (1-\epsilon)^k$. Observe that this probability is bounded from below by $(1-\epsilon)^L$. Therefore, it follows that $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) \geq (1-\epsilon)^{L^2}$. Thus, by taking $\theta = 0.9^{L^2}$ we obtain

$$\exists \theta > 0. \forall \epsilon : 0 < \epsilon \leq 0.1\, Prob_{\mathcal{M}}(s_{init} \models \Diamond F) \geq \theta.$$

If $\mathcal{N}'$ does not halt and continues running forever. Let G be the set of $\mathtt{s}_e$-states. Observe that F is not reachable from G. Similarly, by reasoning about runs of $\mathcal{N}$ which faithfully imitate those of $\mathcal{N}'$ the first N steps, we obtain that $Prob_{\mathcal{M}}(s_{init} \models \neg\Diamond F) \geq Prob_{\mathcal{M}}(s_{init} \models \Diamond G) \geq (1-\epsilon)^{N^2}$. For $N = (\frac{1}{\epsilon})^{1/4}$, it follows that $\lim_{\epsilon \to 0} Prob_{\mathcal{M}}(s_{init} \models \Diamond G) = 1$. As a consequence, we obtain $\lim_{\epsilon \to 0} Prob_{\mathcal{M}}(s_{init} \models \Diamond F) = 0$. Therefore, we have

$$\neg \left(\exists \theta > 0. \forall \epsilon : 0 < \epsilon \leq 0.1\, Prob_{\mathcal{M}}(s_{init} \models \Diamond F) \geq \theta\right).$$

Finally we obtain that $\mathcal{N}'$ halts if and only if

$$\exists \theta > 0. \forall \epsilon : 0 < \epsilon \leq 0.1\, Prob_{\mathcal{M}}(s_{init} \models \Diamond F) \geq \theta.$$

Suppose that for $\mathcal{N}$, there exists a $(\mathbb{R}, +, *, \leq)$-formula $\Phi(p, \epsilon)$. Then $\mathcal{N}'$ terminates if and only if

$$\Psi := \exists \theta > 0. \forall \epsilon : 0 < \epsilon \leq 0.1\, \exists p \left(\Phi(p, \epsilon) \wedge p \geq \theta\right)$$

which is a contradiction. □

Remark 9.5. In the constructions for Theorem 9.1, Theorem 9.3 and Theorem 9.4 only states in F can be reached from F. Therefore, $Prob_{\mathcal{M}}(s_{init} \models \Diamond F) = Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F)$ in these constructions. Thus, for PVASS, PLCS and PNTM, the probability $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F)$ is not effectively expressible in $(\mathbb{R}, +, *, \leq)$ either.

To summarize, the results of this section only show that for PVASS, PLCS and (P)NTM the probabilities $Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$ and $Prob_{\mathcal{M}}(s_{init} \models \Box\Diamond F)$ cannot be effectively expressed in $(\mathbb{R}, +, *, \leq)$ in the same uniform way as for probabilistic pushdown automata (see [EKM04, EY05b, EE04]).

The following three scenarios are not ruled out for PVASS, PLCS, or PNTM. It is an open question which of them is true for each class of models.

(1) The probability $Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$ is effectively expressible in $(\mathbb{R}, +, *, \leq)$, but not in a uniform way. For example, the parameters of the (PVASS, PLCS or PNTM) system would influence the formula in a complex (but still computable) way such that it is not possible to quantify over them in the logic.
(2) The probability $Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$ is expressible in $(\mathbb{R}, +, *, \leq)$, but the formula is not effectively constructible.
(3) The probability $Prob_{\mathcal{M}}(s_{init} \models \Diamond F)$ is not generally expressible in $(\mathbb{R}, +, *, \leq)$, i.e., there is no corresponding formula at all.

10. Conclusions and Future Work

We have defined the 'decisiveness' property for infinite-state Markov chains. Some other abstract conditions like the existence of a finite attractor or global coarseness imply the decisiveness property. In particular, several classes of infinite Markov chains which are derived from program-like probabilistic system models like PVASS, PLCS and PNTM satisfy these properties.

We have studied the qualitative and quantitative (repeated) reachability problem for decisive Markov chains. Our main results are that qualitative questions (i.e., if some probability is 0 or 1) can often be reduced to questions about the underlying infinite transition graph of the systems, while a simple path enumeration algorithm can be used to approximate the probabilities arbitrarily closely.

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 5). Furthermore, for probabilistic systems, reachability is not always easier to decide than repeated reachability (Theorems 5.4 and 6.4).

Open questions for future work are the decidability of qualitative reachability problems for Markov chains with downward closed sets of final states, and an algorithm to approximate quantitative repeated reachability in PVASS. Furthermore, the decidability of exact quantitative questions like $Prob_{\mathcal{M}}\left(s_{init} \models \Diamond F\right) \geq 0.5$ is still open for PVASS, PLCS and (P)NTM.

References

[ABIJ00] P. Abdulla, C. Baier, P. Iyer, and B. Jonsson. Reasoning about probabilistic lossy channel systems. In *Proc. of CONCUR 2000*, volume 1877 of *LNCS*, 2000.

[AC05] E. Asarin and P. Collins. Noisy Turing machines. In *Proc. ICALP '05, 32. International Colloquium on Automata, Languages, and Programming*, pages 1031–1042, 2005.

[AČJYK00] P. 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.

[AHM05] P. Abdulla, N.B. Henda, and R. Mayr. Verifying infinite markov chains with a finite attractor or the global coarseness property. In *Twentieth Annual IEEE Symposium on Logic in Computer Science (LICS 2005)*. IEEE, 2005.

[AJ96] P. Abdulla and B. Jonsson. Verifying programs with unreliable channels. *Information and Computation*, 127(2):91–101, 1996.

[AR03] P. 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.

[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 *LNCS*, 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*, volume 2620, 2003.

[BS04] N. Bertrand and Ph. 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.

[BZ83] D. Brand and P. Zafiropulo. On communicating finite-state machines. *Journal of the ACM*, 2(5):323–342, April 1983.

[CFI96] G. Cécé, A. Finkel, and P. Iyer. Unreliable channels are easier to verify than perfect channels. *Information and Computation*, 124(1):20–31, 10 January 1996.

[CGP99] E.M. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, Dec. 1999.

[Cou91] B. Courcelle. On constructing obstruction sets of words. *Bulletin of the EATCS*, 44:178–185, June 1991.

[CSS03] J.-M. Couvreur, N. Saheb, and G. Sutre. An optimal automata approach to LTL model checking of probabilistic systems. In *Proc. LPAR'2003*, volume 2850 of *LNCS*, pages 361–375, 2003.

[CY88] C. Courcoubetis and M. Yannakakis. Verifying temporal properties of finite-state probabilistci programs. In *Proc. 29. Annual Symp. Foundations of Computer Science*, pages 338–345, 1988.

[CY95] C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, 1995.

[dLMSS56] K. de Leeuw, E.F. Moore, C.E. Shannon, and N. Shapiro. Computability by probabilistic machines. In C.E. Shannon and J. McCarthy, editors, *Annals of Mathematics Studies,* Automata Studies, volume 34, pages 183–212. Princeton University Press, 1956.

[EE04] J. Esparza and K. Etessami. Verifying probabilistic procedural programs. In *Proc. of FST&TCS 2004*, volume 3328 of *LNCS*. Springer Verlag, 2004.

[EKM04] J. Esparza, A. Kucera, and R. Mayr. Model checking probabilistic pushdown automata. In *Proc. LICS' 04, 20. IEEE Int. Symp. on Logic in Computer Science*, pages 12–21, 2004.

[EKM05] J. Esparza, A. Kučera, and R. Mayr. Quantitative analysis of probabilistic pushdown automata: Expectations and variances. In *Proc. LICS' 05, 21. IEEE Int. Symp. on Logic in Computer Science*, pages 117–126, 2005.

[EY05a] K. Etessami and M. Yannakakis. Algorithmic verification of recursive probabilistic state machines. In *Proc. TACAS '05, 11. Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, volume 3440 of *LNCS*, pages 253–270, 2005.

[EY05b] 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 3404 of *LNCS*. Springer Verlag, 2005.

[EY05c] K. Etessami and M. Yannakakis. Recursive Markov decision processes and recursive stochastic games. In *Proc. ICALP '05, 32. International Colloquium on Automata, Languages, and Programming*, volume 3580 of *LNCS*, pages 891–903, 2005.

[Fel66] W. Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. Wiley & Sons, second edition, 1966.

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

[HK97] M. Huth and M. Kwiatkowska. Quantitative analysis and model checking. In *Proc. LICS'97, 12. IEEE Int. Symp. on Logic in Computer Science*, pages 111–122, 1997.

[HS84] S. Hart and M. Sharir. Probabilistic temporal logics for finite and bounded models. In *Proc. 16. ACM Symp. on Theory of Computing*, pages 1–13, 1984.

[HU79] J.E. Hopcroft and J.D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 1979.

[IN97] P. Iyer and M. Narasimha. Probabilistic lossy channel systems. In *TAPSOFT '97: Theory and Practice of Software Development*, volume 1214 of *LNCS*, pages 667–681, 1997.

[Jan90] P. Jančar. Decidability of a temporal logic problem for Petri nets. *Theoretical Computer Science*, 74:71–93, 1990.

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

[LS82] D. Lehmann and S. Shelah. Reasoning with time and chance. *Information and Computation*, 53:165–198, 1982.

[May84] E. Mayr. An algorithm for the general Petri net reachability problem. *SIAM Journal of Computing*, 13:441–460, 1984.

[Min67] M.L. Minsky. *Computation: Finite and Infinite Machines*. Prentice-Hall, 1967.

[Rab03] A. Rabinovich. Quantitative analysis of probabilistic lossy channel systems. In *Proc. ICALP 2003, 30. International Colloquium on Automata, Languages, and Programming*, volume 2719 of *LNCS*, pages 1008–1021, 2003.

[Tar51] A. Tarski. *A Decision Method for Elementary Algebra and Geometry*. Univ. of California Press, Berkeley, 1951.

[Var85] M.Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proc. 26. Annual Symp. Foundations of Computer Science*, pages 327–338, 1985.