Limiting Behavior of Markov Chains with Eager Attractors

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

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

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

Sven Sandberg
Uppsala University, Sweden.
svens@it.uu.se

Abstract

We consider discrete infinite-state Markov chains which contain an eager finite attractor. A finite attractor is a finite subset of states that is eventually reached with probability 1 from every other state, and the eagerness condition requires that the probability of returning to the attractor in n or more steps is exponentially bounded in n. Examples of such Markov chains are those induced by probabilistic lossy channel systems and other probabilistic systems which are (eventually) more likely to shrink than to grow in every single step. This class of infinite-state Markov chains preserves certain important characteristics of finite-state Markov chains. We show that the expected residence time (a generalization of the steady state distribution) exists for Markov chains with eager attractors and that it can be effectively approximated to arbitrary precision. Furthermore, the limiting average expected reward, with respect to state-based bounded reward functions, is also computable.

1 Introduction

Overview. Probabilistic models can be used to capture the behaviors of systems with uncertainty, such as programs with unreliable channels, randomized algorithms, and fault-tolerant systems. The goal is to develop algorithms to analyze quantitative aspects of their behavior such as performance and dependability. In those cases where the underlying semantics of a system is defined as a *finite-state* Markov chain, techniques based on extensions of finite-state model checking can be used to carry out verification [14, 26, 15, 6, 9, 24]. However, many systems that arise in computer applications can only be faithfully modeled as Markov chains which have *infinite* state spaces. Examples include *probabilistic pushdown automata* (recursive state machines) which are natural models for probabilistic sequential programs with recursive procedures [17, 18, 20, 19, 16, 21], *probabilistic lossy channel systems* (PLCS) which consist of finite-state processes communicating through unreliable and unbounded channels in which messages are lost with a certain probability [2, 5, 7, 8, 10, 22, 25], and *probabilistic vector addition systems*, the probabilistic extension of vector addition systems (Petri nets) which model concurrency and synchronization [3, 1].

Related Work. A method for analyzing the limiting behavior of certain classes of infinite Markov chains (including PLCS) has recently been presented by Brázdil and Kučera in [11]. The main idea in [11] is to approximate an infinite-state Markov chain by a sequence of effectively constructible finite-state Markov chains such that the obtained solutions for the finite-state Markov chains converge toward the solution for the original infinite-state Markov chain. The infinite Markov chain needs to satisfy certain preconditions to ensure this convergence. In particular, the method requires decidability of the reachability problem (and even of model checking with certain path formulae) in the underlying infinite transition system.

We recently [3, 1] defined weak abstract conditions on infinite-state Markov chains which are sufficient to make many verification problems computable. Among those are decision problems ("Is a set a final states reached eventually (or infinitely often) with probability 1?"), and approximation problems ("Compute the expected cost/reward of all runs until they reach some final state."). One such sufficient condition is the existence of an eager finite attractor. An attractor is a subset of states that is eventually reached with probability 1 from every other state. We call an attractor *eager* [1] if it satisfies a slightly stronger condition: once outside it, the probability of returning to it in n or more steps is exponentially bounded in n. Every finite-state Markov chain trivially has an eager finite attractor (itself), but many infinite-state Markov chains also have eager finite attractors. A sufficient condition for having an eager finite attractor is that the system size is (eventually) more likely to shrink than to grow (by the same amount) in every step [1]. For example, probabilistic lossy channel systems (PLCS) always satisfy this condition. The condition that an eager finite attractor exists is generally incomparable to the conditions in [11], but classic PLCS satisfy both.

Our contribution. We show that infinite-state Markov chains that contain an eager finite attractor retain many important properties of finite-state Markov chains, particularly those concerning the behavior in the limit. The expected residence time (a generalization of the steady state distribution) exists for Markov chains with eager finite attractors and can be effectively approximated to arbitrary precision. In a similar way, one can compute arbitrarily close approximations to the limiting average expected reward with respect to state-based bounded reward functions.

In contrast to [11], our method is a pure path exploration scheme which computes approximate solutions for the original infinite-state Markov chain directly. We do not require decidability of the general reachability problem, but only information about the mutual reachability of states *inside* some eager finite attractor (but not necessarily inside every finite attractor). This weaker condition can be satisfied even if general reachability is undecidable, e.g., if the eager finite attractor is known to be strongly connected or just a single point. Thus, our method is applicable not only to classic PLCS (where every message in transit can be lost at any moment, and reachability is decidable [4, 12]) but also to more general and realistic models of unreliable communication where the pattern of message loss can depend on complex conditions (burst disturbances; interdependencies of conditions which cause interference) and where general reachability is undecidable.

Outline of the Paper. Section 2 gives fundamental definitions that we will need throughout the paper. Section 3 defines the problems that we are going to solve, and Section 4 gives an overview of our algorithms. Then, Sections 5, 6, and 7 give algorithms to solve the problems. Finally, Section 8 summarizes the paper and points out some directions of future research.

The proofs of lemmas are moved to the appendix.

2 Preliminaries

Transition Systems. A *transition system* is a triple $\mathcal{T} = (S, \longrightarrow)$ where S is a countable set of *states* and $\longrightarrow \subseteq S \times S$ is the *transition relation*. We write $s \longrightarrow s'$ to denote that $(s, s') \in \longrightarrow$.

A *run* ρ is an infinite sequence $s_0 s_1 \ldots$ of states satisfying $s_i \longrightarrow s_{i+1}$ for all $i \geq 0$. We use $\rho(i)$ to denote s_i and say that ρ is an s-run if $\rho(0) = s$. We assume familiarity with the syntax and semantics of the temporal logic *CTL** [13]. We use $(s \models \phi)$ to denote the set of s-runs that satisfy the *CTL** formula ϕ. For instance, if $Q \subseteq S$, $(s \models \bigcirc Q)$ and $(s \models \Diamond Q)$ are the sets of s-runs that visit Q in the next state resp. eventually reach Q. For a natural number n, $\bigcirc^{=n} Q$ denotes a formula which is satisfied by a run ρ iff $\rho(n) \in Q$. We use $\Diamond^{=n} Q$ to denote a formula which is satisfied by ρ iff ρ reaches Q first in its n^{th} step, i.e., $\rho(n) \in Q$ and $\rho(i) \notin Q$ when $0 \leq i < n$. Similarly, for $\sim \in \{<, \leq, \geq, >\}$, $\Diamond^{\sim n} Q$ holds for a run ρ if there is an $m \in \mathbb{N}$ with $m \sim n$ such that $\Diamond^{=m} Q$ holds.

The properties we consider are defined on (infinite) runs. Thus, we assume transition systems that are deadlock-free, i.e., each state has at least one successor. It is usually natural to add a self-loop to deadlock states if they occur.

A *path* π is a finite sequence $s_0, \ldots, s_n$ of states such that $s_i \longrightarrow s_{i+1}$ for all $i : 0 \leq i < n$. We let $|\pi| := n$ denote the length (number of transitions) in a path. Note that a path is a prefix of a run. We use ρ^n for the path $\rho(0)\rho(1) \cdots \rho(n)$. Let $\Pi_s^k = \{\rho : |\rho| = k \wedge \rho(0) = s\}$ denote the set of paths starting in s of length k.

A transition system is said to be *effective* if (1) it is finitely branching, and (2) for each state, we can explicitly compute all its successors.

A transition system where every state is reachable from all other states is called *irreducible*.

Markov Chains. A *Markov chain* is a tuple $\mathcal{M} = (S, P)$ where S is a countable set of *states* and $P : S \times S \to [0, 1]$ is the *probability distribution*, satisfying $\forall s \in S.\ \sum_{s' \in S} P(s, s') = 1$.

A Markov chain induces a transition system, where the transition relation consists of pairs of states related by a positive probability. Formally, the *underlying transition system* of $\mathcal{M}$ is $(S, \longrightarrow)$ where $s_1 \longrightarrow s_2$ iff $P(s_1, s_2) > 0$. In this manner, concepts defined for transition systems can be lifted to Markov chains. For instance, a run or path in a Markov chain $\mathcal{M}$ is a run or path in the underlying transition system, and $\mathcal{M}$ is irreducible, effective, etc, if the underlying transition system is so. In particular, irreducibility is an important property of Markov chains and a key ingredient in our algorithms, as Theorem 4.1 will show.

A Markov chain $\mathcal{M} = (S, P)$ and a state s induce a probability space on the set of runs that start at s. The probability space $(\Omega, \Delta, \mathcal{P}_\mathcal{M})$ is defined as follows: $\Omega = sS^\omega$ is the set of all infinite sequences of states starting from s and Δ is the σ-algebra generated by the basic cylindric sets $\{D_u = us^\omega : u \in sS^*\}$. The probability measure $\mathcal{P}_\mathcal{M}$ is first defined on finite sequences of states $u = s_0 \dots s_n \in sS^*$ by $\mathcal{P}_\mathcal{M}(u) = \prod_{i=0}^{n-1} P(s_i, s_{i+1})$ and then extended to cylindric sets by $\mathcal{P}_\mathcal{M}(D_u) = \mathcal{P}_\mathcal{M}(u)$; it is well-known that this measure is extended in a unique way to the entire σ-algebra. We use $\mathcal{P}_\mathcal{M}(s \models \phi)$ to denote the measure of the set $(s \models \phi)$.

We say that a property of runs holds *almost certainly* (or *for almost all runs*) if it holds with probability 1.

Eager Attractors. A set $A \subseteq S$ is said to be an *attractor* if $\mathcal{P}_\mathcal{M}(s \models \Diamond A) = 1$ for each $s \in S$. In other words, for all $s \in S$, almost all s-runs will return back to A. We will only work with attractors that are *finite*; therefore we assume finiteness (even when not explicitly mentioned) for all the attractors in the sequel. We say that an attractor $A \subseteq S$ is *eager* if there is a $\beta < 1$ such that for each $s \in A$ and $n \geq 0$ it is the case that $\mathcal{P}_\mathcal{M}\left(s \models \bigcirc\left(\Diamond^{\geq n} A\right)\right) \leq \beta^n$. In other words, for every state $s \in A$, the probability of first returning to A in $n+1$ (or more) steps is exponentially bounded in n. We call β the *parameter* of A. Notice that it is not a restriction to have β independent of s, since A is finite. We showed in [1] that every *probabilistic lossy channel system* has an eager attractor that can be computed and for which the parameter can be computed.

Bottom Strongly Connected Components. Consider the directed acyclic graph (DAG) of maximal strongly connected components (SCCs) of the transition system. An SCC is called a *bottom SCC (BSCC)* if no other SCC is reachable from it.

In a Markov chain with a finite attractor, each BSCC must contain at least one element from the attractor. In particular, there are finitely many BSCCs; denote them by $B_1, \dots, B_r$. If $s \models \exists\Diamond s'$ is decidable for all $s, s' \in A$, we can compute the sets $A_1 = B_1 \cap A, \dots, A_r = B_r \cap A$ (they are the BSCCs of the finite directed graph (A, E) where $(s, s') \in E \iff s \models \exists\Diamond s'$).

Note that a run that enters a BSCC never leaves it. Thus, $\mathcal{M}_i := \left(B_i, P|_{(B_i \times B_i)}\right)$ (where the second component is the restriction of P to $B_i \times B_i$) is a Markov chain on its own; call it the Markov chain *induced by* B_i. The Markov chain induced by a BSCC B_i is irreducible and has the eager attractor $A_i := B_i \cap A$. Let $B' = B_1 \cup \cdots \cup B_r$ anb similarly $A' = A_1 \cup \cdots \cup A_r$.

The following Lemma from [5, 10] implies that almost all runs reach a BSCC.

Lemma 2.1 *For any Markov chain with a finite attractor A and for any initial state s_{init},*

(i) $\mathcal{P}_\mathcal{M}(s_{init} \models \Diamond A') = 1$;

(ii) for each BSCC B_i, $\mathcal{P}_\mathcal{M}(s_{init} \models \Diamond A_i) = \mathcal{P}_\mathcal{M}(s_{init} \models \Diamond B_i)$.

Cesàro Limits. Define the *Cesàro limit* for a sequence $a_0, a_1, \dots$ as $\operatorname{clim}_{n\to\infty} a_n := \lim_{n\to\infty} \frac{1}{n+1} \sum_{i=0}^{\infty} a_i$. It is well known that if $\lim_{n\to\infty} a_n$ exists, then the Cesàro limit exists and equals the limit. Cesàro limits are therefore a natural generalization of the usual limit that can be used when the limit does not exist. For instance, although the sequence $\{1, 0, 1, 0, \dots\}$ does not have a limit in the usual sense, it has the Cesàro limit $\frac{1}{2}$.

3 Problem Statements

In this section, we give the mathematical definitions of the problems we want to solve, as well as the associated computational approximation problems.

The Steady State Distribution. The *steady state distribution*[1] of a Markov chain is a probability distribution over states. For a state $s \in S$, the steady state distribution of s, denoted by π_s, expresses "the probability to be in s in the long run". Consider the following equation system:

$$\begin{cases} \pi_s = \displaystyle\sum_{s' \in S} P(s', s) \cdot \pi_{s'}, s \in S \\ \displaystyle\sum_{s \in S} \pi_s = 1. \end{cases} \tag{1}$$

A sufficient condition for this system to have a unique solution is that the Markov chain is irreducible and has an eager attractor (see Theorem 4.1). In this case, the solution is called the *steady state distribution* for $\mathcal{M}$.

If the solution exists, it can be computed for finite Markov chains. We show how to approximate it for a class of *infinite* Markov chains. Formally, we define the following computation problem.

[1] also known as the *limiting* or *stationary* distribution

Steady_State_Distribution
Instance
- An effective irreducible Markov chain $\mathcal{M} = (S, P)$ that has an eager attractor A with parameter β.
- A state s.
- An error tolerance $\epsilon \in \mathbb{R}_{>0}$.

Task Compute a number $\pi_s^\epsilon \in \mathbb{R}$ such that
$$|\pi_s^\epsilon - \pi_s| < \epsilon.$$

The Expected Residence Time. Given a Markov chain, an initial state s_{init} and a state s, define *the expected residence time* in s when starting from s_{init} as $Res(s_{init}, s) := \mathrm{clim}_{n\to\infty} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=n} s)$. This is a proper generalization of the steady state distribution. We prove in Lemma 6.1 that it always exists for Markov chains with an eager attractor, as opposed to the steady state distribution. When the steady state distribution exists, the two quantities are equal (see Theorem 4.1).

The associated computation problem is as follows.

Expected_Residence_Time
Instance
- An effective Markov chain $\mathcal{M} = (S, P)$ that has an eager attractor A with parameter β and where it is decidable for all states $s, s' \in A$ whether $s \models \exists\Diamond s'$.
- An initial state s_{init} and a state s.
- An error tolerance $\epsilon \in \mathbb{R}_{>0}$.

Task Compute a number $Res(s_{init}, s)^\epsilon$ such that
$$|Res(s_{init}, s)^\epsilon - Res(s_{init}, s)| \leq \epsilon$$

Here, we introduced the requirement that reachability is computable for states in the attractor. In our algorithms, this will be used to compute the BSCCs of the Markov chain. Note that it is much weaker than requiring decidable reachability for all pairs of states; in particular, it only requires a correct yes/no answer to finitely many questions.

The Limiting Average Expected Reward. Given a Markov chain and a function $f : S \to \mathbb{R}$ from states to real numbers which assigns a "reward" to each step of a run, extend it to finite paths by $f(\pi) := \sum_{i=0}^{|\pi|} f(\pi(i))$, the "accumulated reward" for π. The *expected average reward* in the first n steps starting from s_{init} is $E_n^{s_{init}}(f) := \frac{1}{n+1} \sum_{\pi\in\Pi_{s_{init}}^n} \mathcal{P}_{\mathcal{M}}(\pi) \cdot f(\pi)$. We study the *limiting average expected reward*, defined as $G_{s_{init}}(f) := \lim_{n\to\infty} E_n^{s_{init}}(f)$, i.e., equivalently, $G_{s_{init}}(f) = \mathrm{clim}_{n\to\infty} \sum_{\pi\in\Pi_{s_{init}}^n} \mathcal{P}_{\mathcal{M}}(\pi) \cdot f(\pi)$. Intuitively, this quantity expresses the average reward per step in the long run.

Throughout this paper, we assume f is computable and *bounded*, meaning that $\exists M.\forall s \in S.|f(s)| \leq M$. Under this assumption, we show in Lemma 7.1 that the limiting average expected reward exists for all Markov chains with an eager attractor.

We define the computation problem as follows.

Limiting_Average_Expected_Reward
Instance
- An effective Markov chain $\mathcal{M} = (S, P)$ that has an eager attractor A with parameter β and where it is decidable for all states $s, s' \in A$ whether $s \models \exists\Diamond s'$.
- An initial state s_{init}.
- A computable reward function $f : S \to \mathbb{R}$ bounded by M.
- An error tolerance $\epsilon \in \mathbb{R}_{>0}$

Task Compute a number $G_{s_{init}}^\epsilon(f) \in \mathbb{R}$ such that
$$|G_{s_{init}}^\epsilon(f) - G_{s_{init}}(f)| \leq \epsilon.$$

4 Overview of the Algorithms

In this section, we give intuitive descriptions of the algorithms which are formally stated in the following sections. We start with a key theorem that lists important properties of irreducible Markov chains with an eager attractor.

In order to state this, we define the *expected return time* relative to a state s as $m_s := \sum_{i=1}^{\infty} i \cdot \mathcal{P}_{\mathcal{M}}(s \models \bigcirc\Diamond^{=i-1}\{s\})$. Note that the expected return time for a state in an eager attractor is finite.

The theorem connects the steady state distribution, the expected return time, the expected residence time, and the limiting average expected reward. Note that it only characterizes these quantities without indicating how to compute them. The topic for the remainder of this paper is to show that they can be approximated to arbitrary precisions.

Theorem 4.1 *The following holds for an irreducible Markov chain with a finite eager attractor.*

(i) The linear equation system (1) has a unique solution;
(ii) the solution is given by $\pi_s = 1/m_s$, for all $s \in S$;
(iii) for all $s \in S$, $\pi_s = Res(s', s)$, where $s' \in S$ can be chosen arbitrarily;
(iv) for any bounded reward function f, $G_s(f) = \sum_{s'\in S} \pi_{s'} \cdot f(s')$.

In particular, the limiting average expected reward does not depend on the initial state. We thus simply write $G(f)$ instead of $G_s(f)$ when the Markov chain is irreducible and has an eager attractor.

Proof. Take some $s \in A$ and apply Theorem 6.1 of [1] with $F = \{s\}$ to obtain that there is an $\alpha < 1$ such that for all $n \in \mathbb{N}_{>0}$, $\mathcal{P}_{\mathcal{M}}(s \models \bigcirc\Diamond^{\geq n-1}\{s\}) \leq \alpha^n$. It follows that $\sum_{i=1}^{\infty} i \cdot \mathcal{P}_{\mathcal{M}}(s \models \bigcirc\Diamond^{\geq n-1}\{s\}) \leq \sum_{i=1}^{\infty} i \cdot \alpha^n < \infty$, i.e., m_s is finite. Since the Markov chain is irreducible, [23, Theorem 3.6.i, p. 81] implies that m_s is finite for every $s \in S$. (Such a Markov chain is called *positive recurrent*.) Now, (i), (ii), (iii), and (iv) follow from Theorem 3.18 (p. 111), the second equality of equation (3.144) (p. 108), Theorem 3.17 (p. 109), and Theorem 3.23 (p. 140) of [23], respectively. □

The Steady State Distribution. The algorithm works in two steps.

1. It computes a *finite* set R^ϵ of states such that

$$\sum_{s \in S - R^\epsilon} \pi_s < \frac{\epsilon}{3}. \tag{2}$$

   We take R^ϵ as the set of states reachable from some state in the attractor in K steps, for sufficiently large K. Lemma 5.1 shows how to use the parameter β of the eager attractor to find K. The steady state probability for states s outside R^ϵ can thus be approximated by $\pi_s^\epsilon = 0$.

2. For each state $s \in R^\epsilon$, it computes an approximation π_s^ϵ such that

$$\sum_{s \in R^\epsilon} |\pi_s^\epsilon - \pi_s| < \frac{2\epsilon}{3}. \tag{3}$$

   We approximate m_s, and apply Theorem 4.1(ii) to obtain the approximation π_s^ϵ of π_s.

By combining (2) and (3), we see that the algorithm solves a more general problem than the one given in the previous section. It approximates the steady state distribution for *all* states, in the sense that

$$\sum_{s \in S} |\pi_s^\epsilon - \pi_s| \leq \epsilon. \tag{4}$$

The Expected Residence Time. We show that the expected residence time for s when starting in s_{init} is 0 if s is not in a BSCC, while if $s \in B_i$, it is the steady state probability of s with respect to the Markov chain induced by B_i, weighted by the probability to reach B_i from s_{init}. Here is an outline of our algorithm to solve this problem.

1. Find the intersection $A_1, \ldots, A_r$ of each BSCC of the Markov chain with the attractor. This can be done due to our assumption that $s \models \Diamond s'$ is decidable for all $s, s' \in A$.

2. For each BSCC B_i, compute an approximation b_i^ϵ of $\mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_i)$. This can be done by a path exploration starting in s_{init}, since the probability to reach $A \cap (B_1 \cup \cdots \cup B_r)$ is 1. This is inspired by but not the same as the result in [11] since they require that reachability is decidable for *all* pairs of states while we only require it in the attractor.

3. For each BSCC B_i, apply the method of Algorithm 1 on the Markov chain induced by B_i, to find a set $R_i^\epsilon \subseteq B_i$ such that $\sum_{s \in B_i - R_i^\epsilon} \pi_s < \epsilon$.

4. If $s \in R_i^\epsilon$ for some i, use Algorithm 1 to compute an approximation π_s^ϵ of π_s in the Markov chain induced by B_i, and return $b_i^\epsilon \cdot \pi_s^\epsilon$.

5. If $s \notin R_i^\epsilon$ for all i, return 0.

The Limiting Average Expected Reward. In a similar manner to the expected residence time, we solve this problem in two steps. First, we compute the limiting average expected reward for irreducible Markov chains and then we extend the algorithm to non-irreducible Markov chains.

1. Algorithm 4 solves the problem under the assumption that $\mathcal{M}$ is irreducible. Recall from Theorem 4.1 that the expected reward does not depend on the initial state for such Markov chains.

   Recall that f is bounded by M and let $\epsilon_1 = \epsilon/M$. First, the algorithm finds the set R^{ϵ_1} and the approximation $\pi_s^{\epsilon_1}$ for all $s \in R^{\epsilon_1}$ as in Algorithm 1. Then, it returns $\sum_{s \in R^{\epsilon_1}} \pi_s^{\epsilon_1} \cdot f(s)$.

2. Next, in Algorithm 5 we remove the assumption that $\mathcal{M}$ is irreducible. The method is similar to the method we used to compute the expected residence time in terms of the steady state distribution.

   First, for each BSCC B_i, we compute an approximation b_i^ϵ of the probability to reach B_i from s_{init}. Then, we use Algorithm 4 to compute for each BSCC B_i an approximation of the limiting average expected reward of the induced Markov chain $\mathcal{M}_i$, denoted by $G_{(i)}^\epsilon(f)$. Finally, we return $\sum_{i=1}^{r} b_i^\epsilon \cdot G_{(i)}^\epsilon(f)$.

5 The Steady State Distribution

In this section, we give an algorithm to solve STEADY_STATE_DISTRIBUTION. We first show how to find the set R^ϵ such that (2) is satisfied and then how to compute the approximation π_s^ϵ so that (3) holds.

Computing R^ϵ. Take R^ϵ as the set of states reachable in at most K steps from some state in the attractor, for sufficiently large K. If a run contains a state $s \in S - R^\epsilon$, then the last K states before s were not in A. Intuitively, such "long" sequences of states outside the attractor occur "seldom" because the attractor is eager, and thus the steady state probability for states outside R^ϵ is "small".

For all $k \in \mathbb{N}$, let $A^{\leq k} := \{s \in S : \exists s' \in A. s' \models \exists \Diamond^{\leq k} s\}$. We define $A^{=k} := A^{\leq k} - A^{\leq k-1}$ (where $A^{-1} = \emptyset$), i.e., $A^{=k}$ consists of all states that can be reached in k steps from some state in A but not in less than k steps from any state in A. In particular, $A^{\leq 0} = A^{=0} = A$. Note that $A^{=k}$ is finite for all k since the Markov chain is finitely branching, it is computable since the Markov chain is effective, and $\bigcup_{k=0}^{\infty} A^{=k} = S$.

Lemma 5.1 *Given an irreducible Markov chain that has a finite eager attractor A with parameter β, for each $\epsilon > 0$ and $K > \frac{\log \epsilon - 2\log(1-\beta)}{\log \beta}$, we have $\sum_{s \in S - A^{\leq K}} \pi_s \leq \epsilon$.*

Approximating π_s for a state $s \in R^\epsilon$. For the case when $s \in R^\epsilon$, we use Theorem 4.1(ii), and obtain π_s^ϵ by approximating m_s. By definition, the finite sum $\sum_{i=1}^{N} i \cdot \mathcal{P}_{\mathcal{M}}(s \models \bigcirc \Diamond^{=i-1}\{s\})$ converges to m_s as N tends to infinity. Our algorithm computes this sum for a sufficiently large N.

The convergence rate is not known in advance, i.e., we do not know beforehand how large N must be for a given ϵ. However, we observe that $1 - \epsilon/3 \leq \sum_{s \in R^\epsilon} \pi_s \leq 1$, where the second inequality holds by (1) and the first inequality holds since (2) is satisfied. Since our approximation of m_s increases with N, the approximation of $\pi_s = 1/m_s$ decreases with N. We can thus approximate π_s for all $s \in R^\epsilon$ simultaneously, and terminate when the sum over $s \in R^\epsilon$ of our approximations becomes less than $1 + \epsilon/3$. It is not guaranteed to reach 1 in finite time.

Algorithm 1 – STEADY_STATE_DISTRIBUTION

Input

An effective irreducible Markov chain $\mathcal{M} = (S, P)$, a finite eager attractor A with parameter β, a state $s \in S$, and an error tolerance $\epsilon \in \mathbb{R}_{>0}$.

Return value

An approximation π_s^ϵ such that $\sum_{s \in S} |\pi_s^\epsilon - \pi_s| \leq \epsilon$.

Constants

$K := \left\lceil \frac{\log(\epsilon/3) - 2\log(1-\beta)}{\log \beta} \right\rceil$

$R^\epsilon := A^{\leq K}$

Variables

$n : \mathbb{N}$ (initially set to 0)

$\{m'_s : \mathbb{R}\}_{s \in R^\epsilon}$ (initially all are set to 0)

begin

1. **if** $s \in S - R^\epsilon$ **return** 0
2. **repeat**
3. **for each** $s' \in R^\epsilon$
4. $m'_{s'} \leftarrow m'_{s'} + \mathcal{P}_{\mathcal{M}}\left(s' \models \bigcirc(\Diamond^{=n-1} s')\right) \cdot n$
5. $n \leftarrow n + 1$
6. **until** $\sum_{s' \in R^\epsilon} \frac{1}{m'_{s'}} \leq 1 + \epsilon/3$
7. **return** $1/m'_s$

end

Notice that since the Markov chain is effective, for each $n \geq 1$ and $s \in S$ we can compute $\mathcal{P}_{\mathcal{M}}(s \models \bigcirc(\Diamond^{=n-1} s))$ by enumerating all paths of length n starting from s. Termination is guaranteed since as the number of iterations tends to infinity, m'_s converges from below to m_s by definition, so that $\sum_{s \in R^\epsilon} 1/m'_s$ converges from above to $\sum_{s \in R^\epsilon} \pi_s \leq 1$. Thus, after a finite number of iterations, $\sum_{s \in R^\epsilon} 1/m'_s \leq 1 + \epsilon/3$.

It remains to show that the return value is a correct approximation of π_s.

If $s \in S - R^\epsilon$, then (2) is satisfied by the choice of K and Lemma 5.1.

Otherwise, by Lemma 5.1 together with the choice of R^ϵ, $1 - \epsilon/3 \leq \sum_{s \in R^\epsilon} \pi_s$. By the termination condition on line 6, $\sum_{s \in R^\epsilon} \frac{1}{m'_s} \leq 1 + \epsilon/3$. Combining these inequalities gives $\sum_{s \in R^\epsilon} \frac{1}{m'_s} - \sum_{s \in R^\epsilon} \pi_s \leq 2\epsilon/3$. By Theorem 4.1(ii) and because $m'_s \leq m_s$, we thus have

$$\sum_{s \in R^\epsilon} \left| \frac{1}{m'_s} - \pi_s \right| = \sum_{s \in R^\epsilon} \left(\frac{1}{m'_s} - \pi_s \right) \leq \frac{2\epsilon}{3}.$$

Thus, (3) and hence also (4) are satisfied. In other words, the algorithm returns a value for π_s^ϵ such that the sum of errors over all states does not exceed ϵ.

6 The Expected Residence Time

We now give an algorithm to approximate the expected residence for arbitrary Markov chains with eager attractors (not necessarily irreducible). Throughout this section, we fix an effective Markov chain that has an eager attractor A with parameter β and use the same definitions from section 2 (paragraph Bottom Strongly Connected Components). For all $s \in B'$, let π_s denote the steady state probability of s relative to the Markov chain induced by the BSCC to which s belongs. We are now ready to state a key lemma used in this section.

Lemma 6.1 *In a Markov chain with a finite eager attractor, for any initial state s_{init}, the expected residence time $Res(s_{init}, s)$ always exists and satisfies*

$$Res(s_{init}, s) = \begin{cases} \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_i) \cdot \pi_s, & \textit{if } s \in B_i; \\ 0 & \textit{if } s \notin B'. \end{cases}$$

This result indicates how our algorithm works. Roughly speaking, we approximate the probability to reach each BSCC, we approximate π_s if $s \in B'$, and we return the product of these quantities.

The Probability to Reach a BSCC. We first give a path exploration algorithm that approximates the probability to reach each BSCC. Since we do not require that reachability is decidable, it is not possible to check whether $s \in B_i$. However, it suffices to check whether $s \in A_i$, which is possible since A_i is finite and can be computed explicitly. Note that Algorithm 2, unlike the others, does not require that the attractor is eager.

Algorithm 2 – PROBABILITY_TO_REACH_BSCC
Input
An effective Markov chain $\mathcal{M} = (S, P)$ with a finite attractor; the intersections $\{A_1, \ldots, A_r\}$ of the attractor with each BSCC, an initial state $s_{init} \in S$, an index $i \in \{1, \ldots, r\}$ of a BSCC, and an error threshold $\epsilon \in \mathbb{R}_{>0}$.
Return value
A lower approximation $b_i^\epsilon \leq \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_i)$ such that $\sum_{i=1}^{r} |b_i^\epsilon - \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_i)| \leq \epsilon$
Variables
$n : \mathbb{N}$ (initially set to 0)
$b_1^\epsilon, \ldots, b_r^\epsilon : \mathbb{R}$ (initially all are set to 0)
begin
1. **repeat**
2. **for** $i \leftarrow 1$ to r
3. $b_i^\epsilon \leftarrow b_i^\epsilon + \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond^{=n} A_i)$
4. $n \leftarrow n + 1$
5. **until** $\sum_{i=1}^{r} b_i^\epsilon \geq 1 - \epsilon$
6. **return** b_i^ϵ
end

It is easy to see that the algorithm returns a correct value if it terminates: each time the algorithm reaches line 4 (but has not yet executed it), we clearly have: for all $i : 1 \leq i \leq r$, $b_i^\epsilon = \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond^{\leq n} A_i) \leq \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond A_i) = \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_i)$ where the last equality follows from Lemma 2.1(ii). Therefore, the termination condition guarantees that

$$\sum_{i=1}^{r} |b_i^\epsilon - \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_i)| \leq \epsilon.$$

It remains to show that the algorithm actually terminates. By Lemma 2.1(i), almost all runs reach A', so $\sum_{n=0}^{\infty} \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond^{=n} A') = \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond A_i) = 1$. By the definition of a convergent sum, there is an N such that $\sum_{n=0}^{N} \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond^{=n} A') \geq 1 - \epsilon$, and hence the algorithm terminates.

The Expected Residence Time. We are now ready to state the algorithm.

Algorithm 3 – EXPECTED_RESIDENCE_TIME
Input
An effective Markov chain $\mathcal{M} = (S, P)$, a finite eager attractor A with parameter β, an initial state $s_{init} \in S$, a state $s \in S$, and an error tolerance $\epsilon \in \mathbb{R}_{>0}$.
Return value
An approximation $Res(s_{init}, s)^\epsilon$ of $Res(s_{init}, s)$ such that $|Res(s_{init}, s)^\epsilon - Res(s_{init}, s)| \leq \epsilon$.
begin
1. Compute the BSCCs $A_1, \ldots, A_r$ of the finite graph (A, E) where $(s', s'') \in E$ iff $s' \models \exists \Diamond s''$
2. $\epsilon_1 \leftarrow \epsilon/(4r)$
3. **for** $i \leftarrow 1 \ldots r$
4. Use the method of Algorithm 1 to compute a set $R_i^{3\epsilon_1}$ for the Markov chain induced by B_i such that $\sum_{s' \in B_i - R_i^{3\epsilon_1}} \pi_{s'} \leq \epsilon_1$.
5. **if** $s \in R_i^{3\epsilon_1}$
6. Use the method of Algorithm 1 to compute approximations $\pi_{s'}^{3\epsilon_1}$ for all $\pi_{s'}$ where $s' \in R_i^{3\epsilon_1}$ in the Markov chain induced by B_i, such that $\sum_{s' \in R_i^{3\epsilon_1}} |\pi_{s'}^{3\epsilon_1} - \pi_{s'}| \leq 2\epsilon_1$.
7. Use Algorithm 2 to compute approximations $b_j^{\epsilon_1}$ of $\mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_j)$ for all j, such that $\sum_{j=1}^{r} |b_j^{\epsilon_1} - \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_j)| \leq \epsilon_1$.
8. **return** $Res(s_{init}, s)^\epsilon = b_i^{\epsilon_1} \cdot \pi_s^{3\epsilon_1}$
9. **return** $Res(s_{init}, s)^\epsilon = 0$
end

Similarly to the previous section, we give a slightly stronger result than required. The error in the approximation is not just small for one single state; in fact we will show that

$$\sum_{s \in S} |Res(s_{init}, s)^\epsilon - Res(s_{init}, s)| \leq \epsilon. \tag{5}$$

For any $i : 1 \leq i \leq r$, Lemma 6.1 implies

$$\sum_{s \in R_i^\epsilon} |Res(s_{init}, s)^\epsilon - Res(s_{init}, s)|$$
$$= \sum_{s \in R_i^\epsilon} |b_i^{\epsilon_1} \cdot \pi_s^{3\epsilon_1} - \mathcal{P}_\mathcal{M}(s_{init} \models \Diamond B_i) \cdot \pi_s|$$
$$= \sum_{s \in R_i^\epsilon} |b_i^{\epsilon_1} \cdot (\pi_s^{3\epsilon_1} - \pi_s) + (b_i^{\epsilon_1} - \mathcal{P}_\mathcal{M}(s_{init} \models \Diamond B_i)) \cdot \pi_s|$$
$$\leq \sum_{s \in R_i^\epsilon} |\pi_s^{3\epsilon_1} - \pi_s| + |b_i^{\epsilon_1} - \mathcal{P}_\mathcal{M}(s_{init} \models \Diamond B_i)|$$
$$\leq 2\epsilon_1 + \epsilon_1 = 3\epsilon/(4r).$$

Hence,

$$\sum_{s \in R_1^\epsilon \cup \cdots \cup R_r^\epsilon} |Res(s_{init}, s)^\epsilon - Res(s_{init}, s)| \leq 3\epsilon/4.$$

Moreover, by the condition on line 4 of the algorithm, we have

$$\sum_{s \in B' - (R_1^\epsilon \cup \cdots \cup R_r^\epsilon)} |Res(s_{init}, s)^\epsilon - Res(s_{init}, s)| \leq \epsilon/4.$$

For states $s \in S - B'$, the error in the approximation is 0, and hence (5) follows.

Remark. In Algorithm 2, we can replace A_i by any subset of B_i, since each state of B_i is reached with probability 1 if B_i is reached. (This holds because the attractor is reached infinitely often, and each state is reachable from the attractor with some positive probability.) The larger this set is, the faster Algorithm 2 will converge. In our case, we have already computed the set R_i^ϵ for some i. Since it satisfies $A_i \subseteq R_i^\epsilon \subseteq B_i$, we can re-use it here instead of A_i.

7 The Limiting Average Expected Reward

We now show how to compute arbitrarily close approximations of the limiting average expected reward for a Markov chain with an eager finite attractor.

First, Algorithm 4 relies on Theorem 4.1(iv) to compute the limiting average expected reward for an irreducible Markov chain. Recall that the limiting average expected reward in an irreducible Markov chain is independent of the initial state.

Then, Algorithm 5 combines outputs from Algorithm 2 and Algorithm 4 in order to approximate the limiting average expected reward in a non-irreducible Markov chain.

Algorithm 4 – LIMITING_AVERAGE_EXPECTED_REWARD–IRREDUCIBLE

Input

An effective irreducible Markov chain $\mathcal{M} = (S, P)$, an eager finite attractor A with parameter β, a computable reward function f bounded by M, and an error tolerance $\epsilon \in \mathbb{R}_{>0}$.

Return value

An approximation $G^\epsilon(f)$ of $G(f)$ such that $|G^\epsilon(f) - G(f)| \leq \epsilon$.

begin

1. $\epsilon_1 \leftarrow \epsilon/M$
2. Use methods from Algorithm 1 to compute the set R^{ϵ_1} and the approximations $\{\pi_s^{\epsilon_1}\}_{s \in R^{\epsilon_1}}$ such that $\sum_{s \in S - R^{\epsilon_1}} \pi_s^{\epsilon_1} < \epsilon_1/3$ and $\sum_{s \in R^{\epsilon_1}} |\pi_s^{\epsilon_1} - \pi_s| < (2\epsilon_1)/3$.
3. **return** $\sum_{s \in R^{\epsilon_1}} \pi_s^{\epsilon_1} \cdot f(s)$

end

We now show correctness. By applying Theorem 4.1(iv), the triangle inequality, and (4), we see that the error in the approximation is

$$\left| \sum_{s \in R^{\epsilon_1}} \pi_s^{\epsilon_1} \cdot f(s) - G(f) \right|$$
$$= \left| \sum_{s \in R^{\epsilon_1}} (\pi_s^{\epsilon_1} - \pi_s) \cdot f(s) - \sum_{s \in S - R^{\epsilon_1}} \pi_s \cdot f(s) \right|$$
$$\leq \sum_{s \in R^{\epsilon_1}} |\pi_s^{\epsilon_1} - \pi_s| \cdot M + \sum_{s \in S - R^{\epsilon_1}} \pi_s \cdot M$$
$$\leq \frac{2\epsilon_1}{3} \cdot M + \frac{\epsilon_1}{3} \cdot M = \epsilon.$$

Non-irreducible Markov Chains. The following lemma is used analogously to the way Lemma 6.1 was used in Section 6.

Lemma 7.1 *For any Markov chain with an eager finite attractor and for any initial state s_{init}, $G_{s_{init}}(f)$ always exists and satisfies*

$$G_{s_{init}}(f) = \sum_{i=1}^{r} \mathcal{P}_\mathcal{M}(s_{init} \models \Diamond B_i) \cdot G_{(i)}(f),$$

The algorithm approximates $\mathcal{P}_\mathcal{M}(s_{init} \models \Diamond B_i)$ and $G_{(i)}(f)$ for all BSCCs. Then it returns the sum over all BSCCs of the products of these approximations.

Algorithm 5 – LIMITING_AVERAGE_EXPECTED_REWARD

Input

An effective Markov chain $\mathcal{M} = (S, P)$, an eager finite attractor A with parameter β, a computable reward function f bounded by M, an initial state s_{init}, and an error tolerance $\epsilon \in \mathbb{R}_{>0}$.

Return value

An approximation $G^{\epsilon}_{s_{init}}(f)$ of $G_{s_{init}}(f)$ such that $|G^{\epsilon}_{s_{init}}(f) - G_{s_{init}}(f)| \leq \epsilon$.

begin

1. Compute the BSCCs $A_1, \ldots, A_r$ of the finite graph (A, E) where $(s, s') \in E$ iff $s \models \exists\Diamond s'$
2. $\epsilon_1 \leftarrow \epsilon/(2r)$
3. $\epsilon_2 \leftarrow \epsilon/(2M)$
4. **for** $i \leftarrow 1$ to r
5. Use Algorithm 4 to compute approximations $G^{\epsilon_1}_{(i)}(f)$ of $G_{(i)}(f)$, such that $\forall i \in \{1, \ldots, r\}. |G^{\epsilon_1}_{(i)}(f) - G_{(i)}(f)| \leq \epsilon_1$
6. Use Algorithm 2 to compute a lower approximation $b^{\epsilon_2}_i \leq \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_i)$, such that $\sum_{j=1}^{r} |b^{\epsilon_2}_j - \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_j)| \leq \epsilon_2$
7. **return** $\sum_{i=1}^{r} b^{\epsilon_2}_i \cdot G^{\epsilon_1}_{(i)}(f)$

end

The correctness proof is similar to that for Algorithm 3. By applying Lemma 7.1 and the triangle inequality, we see that the error in the approximation is

$$\begin{aligned}
&\left|\sum_{i=1}^{r} b^{\epsilon_2}_i \cdot G^{\epsilon_1}_{(i)}(f) - G_{s_{init}}(f)\right| \\
&= \left|\sum_{i=1}^{r} b^{\epsilon_2}_i \cdot G^{\epsilon_1}_{(i)}(f) - \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_i) \cdot G_{(i)}(f)\right| \\
&= \left|\sum_{i=1}^{r} b^{\epsilon_2}_i \cdot (G^{\epsilon_1}_{(i)}(f) - G_{(i)}(f))\right. \\
&\qquad \left. + (b^{\epsilon_2}_i - \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_i)) \cdot G_{(i)}(f)\right| \\
&\leq (\max_{1 \leq i \leq r} b^{\epsilon_2}_i) \sum_{i=1}^{r} |G^{\epsilon_1}_{(i)}(f) - G_{(i)}(f)| \\
&\qquad + (\max_{1 \leq i \leq r} G_{(i)}(f)) \sum_{i=1}^{r} |b^{\epsilon_2}_i - \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_i)| \\
&\leq 1 \cdot r \cdot \epsilon_1 + M \cdot \epsilon_2 = \epsilon/2 + \epsilon/2 = \epsilon
\end{aligned}$$

8 Conclusions and Future Work

We have shown that infinite-state Markov chains with an eager finite attractor have a similar limiting behavior as finite-state Markov chains. The expected residence time and limiting average expected reward w.r.t. bounded reward functions can be effectively approximated by path exploration schemes. Since these only require reachability information *inside* the attractor, they are applicable even to some systems where general reachability is undecidable.

One direction for future work is to further weaken the required preconditions, in order to handle larger classes of systems. For example, the finiteness condition of the attractor can possibly be replaced by a weaker condition that symbolic representations of sufficiently likely parts of some infinite attractor can be effectively constructed.

References

[1] P. A. Abdulla, N. Ben Henda, R. Mayr, and S. Sandberg. Eager markov chains. Technical Report 2006-009, Department of Information Technology, Uppsala University, March 2006.

[2] Parosh Aziz Abdulla, Christel Baier, Purushothaman Iyer, and Bengt Jonsson. Reasoning about probabilistic lossy channel systems. In *Proc. CONCUR 2000, 11^{th} Int. Conf. on Concurrency Theory*, volume 1877 of *Lecture Notes in Computer Science*, pages 320–333, 2000.

[3] Parosh Aziz Abdulla, Noomene Ben Henda, and Richard Mayr. Verifying infinite Markov chains with a finite attractor or the global coarseness property. In *Proc. LICS' 05 21^{st} IEEE Int. Symp. on Logic in Computer Science*, pages 127–136, 2005.

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

[5] Parosh Aziz Abdulla and Alexander Rabinovich. Verification of probabilistic systems with faulty communication. In *Proc. FOSSACS03, Conf. on Foundations of Software Science and Computation Structures*, volume 2620 of *Lecture Notes in Computer Science*, pages 39–53, 2003.

[6] Adnan Aziz, Kumud Sanwal, Vigyan Singhal, and Robert Brayton. Model-checking continuous-time Markov chains. *ACM Trans. on Computational Logic*, 1(1):162–170, 2000.

[7] C. Baier, N. Bertrand, and Ph. Schnoebelen. A note on the attractor-property of infinite-state Markov chains. *Information Processing Letters*, 97(2):58–63, 2006.

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

[9] Christel Baier, Boudewijn R. Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Automated performance and dependability evaluation using model checking. In *Proc. Performance 2002*, pages 261–289, 2002.

[10] 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 of *Lecture Notes in Computer Science*, pages 120–135, 2003.

[11] Tomás Brázdil and Antonn Kučera. Computing the expected accumulated reward and gain for a subclass of infinite Markov chains. In *Proc. FSTTCS '05, 25^{th} Conf. on Foundations of Software Technology and Theoretical Computer Science*, volume 3821 of *Lecture Notes in Computer Science*, pages 372–383, 2005.

[12] Gérard Cécé, Alain Finkel, and S. Purushothaman Iyer. Unreliable channels are easier to verify than perfect channels. *Information and Computation*, 124(1):20–31, 10 January 1996.

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

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

[15] Luca de Alfaro, Marta Z. Kwiatkowska, Gethin Norman, David Parker, and Roberto Segala. Symbolic model checking of probabilistic processes using mtbdds and the Kronecker representation. In *Proc. TACAS '00, 6^{th} Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, volume 1785 of *Lecture Notes in Computer Science*, pages 123–137, 2000.

[16] J. Esparza and K. Etessami. Verifying probabilistic procedural programs. In *Proc. FSTTCS '04, 24^{th} Conf. on Foundations of Software Technology and Theoretical Computer Science*, pages 16–31, 2004.

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

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

[19] K. Etessami and M. Yannakakis. Algorithmic verification of recursive probabilistic state machines. In *Proc. TACAS '05, 11^{th} Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, volume 3440 of *Lecture Notes in Computer Science*, pages 253–270, 2005.

[20] K. Etessami and M. Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of non-linear equations. In *Proc. STACS'05, 22nd Int. Symp. on Theoretical Aspects of Computer Science*, volume 2996 of *Lecture Notes in Computer Science*, pages 340–352, 2005.

[21] K. Etessami and M. Yannakakis. Recursive Markov decision processes and recursive stochastic games. In *Proc. ICALP '05, 32^{nd} International Colloquium on Automata, Languages, and Programming*, volume 3580 of *Lecture Notes in Computer Science*, pages 891–903, 2005.

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

[23] V. G. Kulkarni. *Modeling and Analysis of Stochastic Systems*. Texts in Statistical Science Series. Chapman and Hall Ltd., London, 1995.

[24] M. Kwiatkowska, G. Norman, and D. Parker. Probabilistic model checking in practice: Case studies with PRISM. *ACM Performance Evaluation Review*, 32(2):16–21, 2005.

[25] A. Rabinovich. Quantitative analysis of probabilistic lossy channel systems. In *Proc. ICALP '03, 30^{th} International Colloquium on Automata, Languages, and Programming*, volume 2719 of *Lecture Notes in Computer Science*, pages 1008–1021, 2003.

[26] M.Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proc. FOCS '85, 26^{th} Annual Symp. Foundations of Computer Science*, pages 327–338, 1985.

A Proof of Lemmas

Lemma 5.1 For any $s_{init} \in A$ and $k \geq 1$, we have by Theorem 4.1(iii)

$$\begin{aligned}\sum_{s \in A^{=k}} \pi_s &= \sum_{s \in A^{=k}} \operatorname*{clim}_{n\to\infty} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=n} s) \\ &= \operatorname*{clim}_{n\to\infty} \sum_{s \in A^{=k}} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=n} s) \\ &= \operatorname*{clim}_{n\to\infty} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=n} A^{=k}),\end{aligned}$$

where the sum and limit commute because the sum is finite. The runs in $(s_{init} \models \bigcirc^{=n} A^{=k})$ visit A for sure in step 0 (since $s_{init} \in A$), they may visit A in steps $1, \ldots, n-k$, but they cannot visit A in steps $n-k+1, \ldots, n$ (by the definition of $A^{=k}$). Let i be the step in which A is last visited before the n^{th} step and let $s' \in A$ be the state visited at that point. Graphically, any run in $(s_{init} \models \bigcirc^{=n} A^{=k})$ looks as follows:

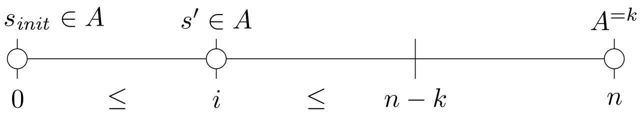

We split into disjoint cases and sum over all possible values for s' and i:

$$\begin{aligned}&\mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=n} A^{=k}) \\ &\leq \sum_{i=0}^{n-k} \sum_{s' \in A} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=i} s') \cdot \mathcal{P}_{\mathcal{M}}(s' \models \bigcirc(\Diamond^{\geq n-i} A)) \\ &\leq \sum_{i=0}^{n-k} \sum_{s' \in A} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=i} s') \cdot \beta^{n-i} \\ &\leq \sum_{i=0}^{n-k} \beta^{n-i} = \frac{\beta^k - \beta^{n+1}}{1-\beta},\end{aligned}$$

where β is the parameter of eagerness and the last inequality holds because $\sum_{s' \in A} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=i} s') \leq 1$. Combining the two equations above, we obtain

$$\begin{aligned}\sum_{s \in A^{=k}} \pi_s &= \lim_{n\to\infty} \frac{1}{n} \sum_{m=1}^{n} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=m} A^{=k}) \\ &\leq \lim_{n\to\infty} \frac{1}{n} \sum_{m=1}^{n} \frac{\beta^k - \beta^{m+1}}{1-\beta} = \frac{\beta^k}{1-\beta}.\end{aligned}$$

In the last equality, $\frac{1}{n}\sum_{m=1}^{n} \beta^{m+1}$ vanishes because $\sum_{m=1}^{n} \beta^{m+1}$ is bounded. We now sum the above inequality over all $k > K$:

$$\begin{aligned}\sum_{s \in A^{>K}} \pi_s &= \sum_{k=K+1}^{\infty} \sum_{s \in A^{=k}} \pi_s \leq \sum_{k=K+1}^{\infty} \beta^k/(1-\beta) \\ &= \frac{\beta^{K+1}}{(1-\beta)^2} \leq \epsilon,\end{aligned}$$

where the last inequality follows from the choice of K. $\square$

Definition A.1 *For all $n \geq 0$ and $Q_1, Q_2 \subseteq S$, we use $Q_1\, \mathcal{U}^{=n}\, Q_2$ to denote the formula satisfied by runs which reach the set Q_2 the first time in the n^{th} step; only through states which belong to Q_1. For example given $s, s' \in S$, $(s \models (\neg A)\, \mathcal{U}^{=n}\, s')$ denotes the set of runs ρ such that $\rho(n) = s'$ and for all $i : 0 \leq i < n$, $\rho(i) \notin A$.*

Lemma 6.1 First consider the case $s \notin B'$. For any $N \geq 0$, we have

$$\begin{aligned}&Res(s_{init}, s) = \operatorname*{clim}_{n\to\infty} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=n} s) \\ &= \operatorname*{clim}_{n\to\infty} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=n} s | s_{init} \models \Diamond^{\leq N} A') \cdot \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond^{\leq N} A') \\ &+ \operatorname*{clim}_{n\to\infty} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=n} s | s_{init} \models \Diamond^{>N} A') \cdot \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond^{>N} A').\end{aligned}$$

The first term equals 0 because s can only be reached in the first N steps by runs in $(s_{init} \models \Diamond^{\leq N} A')$. Since for any $s_{init} \in S$, $\mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond A') = 1$, we can make $\mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond^{>N} A')$ arbitrarily small by taking N sufficiently large. Thus, for any $\varepsilon > 0$, there exists an N such that $Res(s_{init}, s) < \varepsilon$. Therefore, we must have $Res(s_{init}, s) = 0$.

We now prove the case $s \in B_i$.

For any $k \geq 0$ and $s' \in A_i$, using Definition A.1, we obtain

$$\begin{aligned}&\operatorname*{clim}_{n\to\infty} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=n} s | s_{init} \models (\neg A')\, \mathcal{U}^{=k}\, s') \\ &= \lim_{n\to\infty} \frac{1}{n} \sum_{m=1}^{n} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=m} s | s_{init} \models (\neg A')\, \mathcal{U}^{=k}\, s') \\ &= \lim_{n\to\infty} \frac{1}{n} \sum_{m=1}^{k} \frac{\mathcal{P}_{\mathcal{M}}(s_{init} \models (\neg A')\, \mathcal{U}^{=m}\, s) \cdot \mathcal{P}_{\mathcal{M}}(s \models (\neg A')\, \mathcal{U}^{=k-m}\, s')}{\mathcal{P}_{\mathcal{M}}(s_{init} \models (\neg A')\, \mathcal{U}^{=k}\, s')} \\ &+ \lim_{n\to\infty} \frac{1}{n} \sum_{m=k+1}^{n} \frac{\mathcal{P}_{\mathcal{M}}(s_{init} \models (\neg A')\, \mathcal{U}^{=k}\, s') \cdot \mathcal{P}_{\mathcal{M}}(s' \models \bigcirc^{=m-k} s)}{\mathcal{P}_{\mathcal{M}}(s_{init} \models (\neg A')\, \mathcal{U}^{=k}\, s')} \\ &= \operatorname*{clim}_{n\to\infty} \mathcal{P}_{\mathcal{M}}(s' \models \bigcirc^{=n-k} s) = Res(s', s) = \pi_s.\end{aligned}$$

Observ that in the second equality, the first term doesn't depend on n. Therefore, it vanishes as n goes to infinity. Moreover, the last equality follows from Theorem 4.1(iii).

Now, since $s \in B_i$, Lemma 2.1 implies that for any $N \geq 0$:

$$\begin{aligned}
&Res(s_{init}, s)\\
&= \operatorname*{clim}_{n\to\infty} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=n} s | s_{init} \models \Diamond^{\leq N} A_i) \cdot \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond^{\leq N} A_i)\\
&+ \operatorname*{clim}_{n\to\infty} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=n} s | s_{init} \models \Diamond^{> N} A_i) \cdot \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond^{> N} A_i).
\end{aligned}$$

Let us denote by $Res(s_{init}, s)'(N)$ and $Res(s_{init}, s)''(N)$ resp. the first and the second term of the sum. By applying the same reasoning as in the first case, we can make $Res(s_{init}, s)''(N)$ arbitrarily small by taking N large enough. Moreover, if we consider the first term, we obtain:

$$\begin{aligned}
&Res(s_{init}, s)'(N)\\
&= \operatorname*{clim}_{n\to\infty} \sum_{s' \in A_i} \sum_{k=0}^{N} \mathcal{P}_{\mathcal{M}}(s_{init} \models \bigcirc^{=n} s | s_{init} \models (\neg A')\, \mathcal{U}^{=k}\, s')\\
&\qquad \cdot \mathcal{P}_{\mathcal{M}}(s_{init} \models (\neg A')\, \mathcal{U}^{=k}\, s')\\
&= \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond^{\leq N} A_i) \cdot \pi_s\\
&= (\mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond A_i) - \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond^{> N} A_i)) \cdot \pi_s.
\end{aligned}$$

where moving the limit into the sum is justified (since the sum is finite). In the last step, again we can make $\mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond^{>N} A_i)$ arbitrarily small, and thus we obtain

$$\begin{aligned}
Res(s_{init}, s) &= \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond A_i) \cdot \pi_s\\
&= \mathcal{P}_{\mathcal{M}}(s_{init} \models \Diamond B_i) \cdot \pi_s
\end{aligned}$$

by Lemma 2.1 *(ii)*. □

Definition A.2 *For any $s, s' \in S$ and $k \in \mathbb{N}$, let $\Pi^k_{s,s'}(Q) := \{\pi \in \Pi^k_s : (\forall i. 1 \leq i \leq k-1 \implies \pi(k) \neq s' \wedge \pi(k) \notin Q \wedge \pi(k) = s'\}$ denote the subset of Π^k_s that visits s' for the first time in the k^{th} step without "going" through Q.*

Lemma A.3 *For any Markov Chain $\mathcal{M}$ with an eager finite attractor A and any state s, the following holds:*

$$\begin{aligned}
&\forall N \geq 0, \text{ for each BSCC } A_i \text{ in } A,\\
&\lim_{n\to\infty} E^s_n(f | s \models \Diamond^{\leq N} A_i) = G_{(i)}(f).
\end{aligned}$$

Proof. Given a state $s_i \in A_i$, using Definition A.2 and Definition A.1:

$$\begin{aligned}
&E^s_n(f | s \models (\neg A_i)\, \mathcal{U}^{=k}\, s_i) =\\
&\frac{\frac{1}{n+1} \sum_{\pi \in \Pi^k_{s,s_i}(A_i)} \sum_{\pi' \in \Pi^{n-k}_{s_i}} (f(\pi) - f(s_i) + f(\pi')) P(\pi) P(\pi')}{\mathcal{P}_{\mathcal{M}}(s \models (\neg A_i)\, \mathcal{U}^{=k}\, s_i)}\\
&= \frac{\frac{1}{n+1} \sum_{\pi \in \Pi^k_{s,s_i}(A_i)} (f(\pi) - f(s_i)) P(\pi) \sum_{\pi' \in \Pi^{n-k}_{s_i}} P(\pi')}{\mathcal{P}_{\mathcal{M}}(s \models (\neg A_i)\, \mathcal{U}^{=k}\, s_i)}\\
&\quad + \frac{\frac{1}{n+1} \sum_{\pi \in \Pi^k_{s,s_i}(A_i)} P(\pi) \sum_{\pi' \in \Pi^{n-k}_{s_i}} f(\pi') P(\pi')}{\mathcal{P}_{\mathcal{M}}(s \models (\neg A_i)\, \mathcal{U}^{=k}\, s_i)}.
\end{aligned}$$

The first term vanishes as n tends to infinity since $|f(\pi) - f(s_i)| \leq kM$. In the second term, observ that by definition we have$\mathcal{P}_{\mathcal{M}}(s \models (\neg A_i)\ \mathcal{U}^{=k}\ s_i) = \sum_{\pi \in \Pi^k_{s,s_i}(A_i)} P(\pi)$, and thus by simplifying

$$\begin{aligned}
&\lim_{n\to\infty} E^s_n(f | s \models (\neg A_i)\, \mathcal{U}^{=k}\, s_i)\\
&= \lim_{n\to\infty} \frac{1}{n+1} \sum_{\pi' \in \Pi^{n-k}_{s_i}} f(\pi') P(\pi')\\
&= \lim_{n\to\infty} \frac{1}{n+1+k} \sum_{\pi' \in \Pi^{n}_{s_i}} f(\pi') P(\pi')\\
&= G_{(i)}(f).
\end{aligned}$$

Finally,

$$\begin{aligned}
&\lim_{n\to\infty} E^s_n(f | s \models \Diamond^{\leq N} A_i)\\
&= \sum_{s_i \in A_i} \sum_{k=0}^{N} \lim_{n\to\infty} E^s_n(f | s \models (\neg A_i)\, \mathcal{U}^{=k}\, s_i) \cdot\\
&\qquad \mathcal{P}_{\mathcal{M}}(s \models (\neg A_i)\, \mathcal{U}^{=k}\, s_i | s \models \Diamond^{\leq N} A_i)\\
&= G_{(i)}(f) \sum_{s_i \in A_i} \sum_{k=0}^{N} \mathcal{P}_{\mathcal{M}}(s \models (\neg A_i)\, \mathcal{U}^{=k}\, s_i | s \models \Diamond^{\leq N} A_i)\\
&= G_{(i)}(f).
\end{aligned}$$

□

Lemma 7.1 By definition, we have

$$\begin{aligned}
G_s(f) &= \lim_{n\to\infty} E^s_n(f)\\
&= \lim_{n\to\infty} E^s_n(f | s \models \Diamond^{\leq N} A') \cdot \mathcal{P}_{\mathcal{M}}(s \models \Diamond^{\leq N} A')\\
&+ \lim_{n\to\infty} E^s_n(f | s \models \Diamond^{> N} A') \cdot \mathcal{P}_{\mathcal{M}}(s \models \Diamond^{> N} A')\\
&= \lim_{n\to\infty} \sum_{i=1}^{r} E^s_n(f | s \models \Diamond^{\leq N} A_i) \cdot \mathcal{P}_{\mathcal{M}}(s \models \Diamond^{\leq N} A_i)\\
&+ \lim_{n\to\infty} E^s_n(f | s \models \Diamond^{> N} A') \cdot \mathcal{P}_{\mathcal{M}}(s \models \Diamond^{> N} A')\\
&= \sum_{i=1}^{r} \mathcal{P}_{\mathcal{M}}(s \models \Diamond^{\leq N} A_i) \cdot G_{(i)}(f)\\
&\quad + \mathcal{P}_{\mathcal{M}}(s \models \Diamond^{> N} A') \lim_{n\to\infty} E^s_n(f | s \models \Diamond^{> N} A').
\end{aligned}$$

In the last step, we moved the limit into the sum (which is justified since the sum is finite) and then used Lemma A.3. The claim now follows since $\mathcal{P}_{\mathcal{M}}(s \models \Diamond^{>N} A')$ can be made arbitrarily smaller than ϵ/M for any given $\epsilon > 0$ and since $|E^s_n(f | s \models \Diamond^{>N} A')| \leq M$. □