

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 avoiding the attractor in n or more steps after leaving it is exponentially bounded in n . Examples of such Markov chains are those induced by probabilistic lossy channel systems and similar systems. 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, arbitrarily close approximations of the limiting average expected reward, with respect to state-based bounded reward functions, are 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 [1, 5, 7, 8, 10, 22, 25], and *probabilistic vector addition systems*, the probabilistic extension of vector addition systems (Petri nets) which models concurrency and synchronization [2, 3].

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 formulas) in the underlying infinite transition system.

We recently [2, 3] 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 given set of 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* [3] if it satisfies a slightly stronger condition: after leaving it, the probability of returning to it in n or more steps is exponentially bounded in n . Every

finite-state Markov chain trivially has a finite eager attractor (itself), but many infinite-state Markov chains also have eager finite attractors. A sufficient condition for having an eager finite attractor is that there exists a distance measure on states such that for states sufficiently far away from a given finite subset, the probability that their immediate successor is closer to this subset is greater than $\frac{1}{2}$ [3]. 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 properties of finite-state Markov chains which do not hold for general infinite-state Markov chains. These properties include the facts that

- There is at least one, but at most finitely many, bottom strongly connected components (BSCC).
- The Markov chain does not contain any persistent null-states (i.e., for every recurrent state the expected recurrence time is finite).
- The steady state distribution exists if the Markov chain is irreducible and the expected residence time (a generalization of the steady state distribution) always exists.

We use these properties to show that the expected residence time can be effectively approximated to arbitrary precision for Markov chains with eager finite attractors. 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.

Example. Consider a different variant of PLCS where at every step there is a fixed probability of losing *all* messages in *all* channels (i.e., a total reset), but there are no individual message losses. It is easy to encode a Minsky 2-counter machine into this PLCS variant s.t. the final control-state q_{acc} is reachable from the initial configuration $q_{init}\epsilon$ (channels initially empty) in the PLCS iff it is reachable in the Minsky machine. (One needs to make sure that a total reset in any other control-state than q_{init} leads to configuration $q_{init}\epsilon$ again without visiting q_{acc} .) By adding a transition from q_{acc} back to q_{init} , one obtains the eager finite attractor $\{q_{init}\epsilon\}$. However, the reachability problem whether q_{acc} can be reached from q_{init} is undecidable.

2 Preliminaries

Transition Systems. A *transition system* is a tuple $\mathcal{T} = (S, \longrightarrow)$ 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 $s_0s_1\dots$ 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]. Given a CTL^* path-formula ϕ , we use $(s \models \phi)$ to denote the set of *s-runs* that satisfy 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 $m \sim n$ s.t. $\diamond^{=m}Q$ holds.

For all $n \geq 0$ and $Q_1, Q_2 \subseteq S$, we use $Q_1 \mathcal{U}^{=n} Q_2$ to denote a formula satisfied by a run ρ iff for all $i : 0 \leq i < n$, $\rho(i) \in (Q_1 - Q_2)$, and $\rho(n) \in Q_2$. In words, runs in $(Q_1 \mathcal{U}^{=n} Q_2)$ reach the set Q_2 for the first time in the n^{th} step, only passing through states in Q_1 .

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 common to add a self-loop to deadlock states if they occur.

A *path* π is a finite sequence s_0, \dots, 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. Given a run ρ , we use ρ^n for the path $\rho(0)\rho(1)\dots\rho(n)$. Let $\Pi_s^k = \{\rho : |\rho| = k \wedge \rho(0) = s\}$ denote the set of paths starting in s of length k . For any $s, s' \in S$ and $n \in \mathbb{N}$, let

$\Pi_{s,s'}^n(Q) := \{\pi \in \Pi_s^n : (\forall i. 1 \leq i \leq n-1 \implies \pi(i) \neq s' \wedge \pi(i) \notin Q) \wedge \pi(n) = s'\}$. Intuitively, $\Pi_{s,s'}^n(Q)$ denotes the subset of Π_s^n that visits s' for the first time in the n^{th} step without going through Q .

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 direct (one step) successors.

A transition system where every state is reachable from all other states is called *strongly connected*. In the context of Markov chains (see below) this condition is called *irreducible*.

Markov Chains. A *Markov chain* is a tuple $\mathcal{M} = (S, P)$ where S is a countable set of *states* and $P : S \times S \rightarrow [0, 1]$ 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 effective, etc., if the underlying transition system is so. Notice that in the context of Markov chains, \mathcal{M} is called irreducible if the underlying transition system is strongly connected. In particular, irreducibility is an important property of Markov chains and a key ingredient in our algorithms.

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 $(\Omega, \Delta, \mathcal{P})$ 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} is first defined on finite sequences of states $u = s_0 \dots s_n \in sS^*$ by $\mathcal{P}(u) = \prod_{i=0}^{n-1} P(s_i, s_{i+1})$ and then extended to cylindric sets by $\mathcal{P}(D_u) = \mathcal{P}(u)$; it is well-known that this measure is extended in a unique way to the entire σ -algebra. We use $\mathcal{P}(s \models \phi)$ to denote the measure of the set $(s \models \phi)$ (which is measurable by [26]). For singleton sets, we sometimes omit the braces and write s for $\{s\}$ when the meaning is clear from context.

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 $\mathcal{P}(s \models \diamond A) = 1$ for each $s \in S$. In other words, for all $s \in S$, almost all s -runs will visit 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 $\mathcal{P}(s \models \bigcirc (\diamond^{\geq n} A)) \leq \beta^n$. In other words, for every state $s \in A$, the probability of avoiding A in $n+1$ (or more) steps after leaving it 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 [3] that every system whose size is (eventually) more likely to shrink than to grow (by the same amount) in every step has a finite eager attractor. In particular, every *probabilistic lossy channel system* has a finite eager attractor that can be computed and for which the parameter can also 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. Observe that the existence of BSCCs is not guaranteed in an infinite transition system.

In a Markov chain with a finite attractor A , there exists at least one BSCC. Moreover, each BSCC must contain at least one element from the attractor. Therefore, there are only finitely many BSCCs; denote them by B_1, \dots, B_r , where r can be at most the size of A . 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, $\mathcal{M}_i := (B_i, P|_{(B_i \times B_i)})$ (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 finite eager attractor $A_i := B_i \cap A$. Let $B' = B_1 \cup \dots \cup B_r$ and similarly $A' = A_1 \cup \dots \cup A_r$.

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) $\mathcal{P}(s_{init} \models \diamond A') = 1$;
- (ii) for each BSCC B_i , $\mathcal{P}(s_{init} \models \diamond A_i) = \mathcal{P}(s_{init} \models \diamond B_i)$.

Cesàro Limits. The *Cesàro limit* of a sequence a_0, a_1, \dots is defined as $\text{clim}_{n \rightarrow \infty} a_n := \lim_{n \rightarrow \infty} \frac{1}{n+1} \sum_{i=0}^n a_i$. It is well known that if $\lim_{n \rightarrow \infty} a_n$ 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 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 average probability to be in s in the long run”. Formally, it is the solution to the following equation system, if it has a unique solution.

$$\begin{cases} \pi_s = \sum_{s' \in S} P(s', s) \cdot \pi_{s'} & \text{for each } s \in S; \\ \sum_{s \in S} \pi_s = 1. \end{cases} \quad (1)$$

A sufficient condition for this system to have a unique solution is that the Markov chain is irreducible and has a finite eager attractor (see Theorem 4.1). For finite Markov chains, the solution can be computed if it exists. We will show how to approximate it for a class of *infinite* Markov chains. Formally, we define the following computation problem.

STEADY_STATE_DISTRIBUTION

Instance

- An effective irreducible Markov chain $\mathcal{M} = (S, P)$ that has a finite 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 $|\pi_s^\epsilon - \pi_s| \leq \epsilon$.

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 $Res(s_{init}, s) := \text{clim}_{n \rightarrow \infty} \mathcal{P}(s_{init} \models \bigcirc^n s)$. This is a proper generalization of the steady state distribution. We prove in Lemma 6.1 that it always exists for Markov chains with a finite 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.

¹also known as the *limiting* or *stationary* distribution.

EXPECTED_RESIDENCE_TIME

Instance

- An effective Markov chain $\mathcal{M} = (S, P)$ that has a finite 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 $Res^\epsilon(s_{init}, s)$ such that $|Res^\epsilon(s_{init}, s) - Res(s_{init}, s)| \leq \epsilon$

Here we have 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. Observe that this condition 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 $\mathcal{M} = (S, P)$, a *reward function* is a mapping $f : S \rightarrow \mathbb{R}$ from states to real numbers. Given a reward function f , we extend it to finite paths by $f(\pi) := \sum_{i=0}^{|\pi|-1} f(\pi(i))$, the “accumulated reward” for π . The *average expected reward* in the first n steps starting from s_{init} is $E_n^{s_{init}}(f) := \frac{1}{n+1} \sum_{\pi \in \Pi_{s_{init}}^n} \mathcal{P}(\pi) \cdot f(\pi)$. We study the *limiting average expected reward*, defined as $G_{s_{init}}(f) := \lim_{n \rightarrow \infty} E_n^{s_{init}}(f)$, i.e., equivalently, $G_{s_{init}}(f) = \text{clim}_{n \rightarrow \infty} \sum_{\pi \in \Pi_{s_{init}}^n} \mathcal{P}(\pi) \cdot f(\pi)$. 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 a finite eager attractor.

We define the computation problem as follows.

LIMITING_AVERAGE_EXPECTED_REWARD

Instance

- An effective Markov chain $\mathcal{M} = (S, P)$ that has a finite 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 $f : S \rightarrow \mathbb{R}$ 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 a finite eager attractor.

In order to state the theorem, we define the *expected return time* relative to a state s as $m_s := \sum_{i=1}^{\infty} i \cdot \mathcal{P}(s \mid \circ \diamond^{=i-1} \{s\})$.


The theorem relates the steady state distribution, the expected return time, the expected residence time, and the limiting average expected reward. Observe that the theorem 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 $s \in S$, $\pi_s = \text{Res}(s', s)$, where $s' \in S$ can be chosen arbitrarily;*
- (iv) *for any initial state s_{init} and any bounded reward function f , $G_{s_{init}}(f) = \sum_{s' \in S} \pi_{s'} \cdot f(s')$.*

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

Proof. Take a state $s_A \in A$. We first prove that the expected return time for s_A is finite. Once this is done, the claims will follow from classical results. Consider the Markov chain $\mathcal{M}' = (S', P')$ which is identical to \mathcal{M} except we split s_A into two states like in the following picture.



Formally, we take $S' = S \cup \{s_B\}$ where s_B is a new state, and for all $s_0, s_1 \in S'$,

$$P'(s_0, s_1) = \begin{cases} 1 & \text{if } s_0 = s_A \text{ and } s_1 = s_B; \\ 0 & \text{if } s_0 = s_A \text{ and } s_1 \neq s_B; \\ 0 & \text{if } s_0 \neq s_A \text{ and } s_1 = s_B; \\ P(s_A, s_1) & \text{if } s_0 = s_B; \\ P(s_0, s_1) & \text{otherwise.} \end{cases}$$

Clearly, $A' := A \cup \{s_B\}$ is a finite eager attractor for \mathcal{M}' , and

$$\begin{aligned} \mathcal{P}(s_A \mid \circ \diamond^{\geq n-1} \{s_A\}) &= \mathcal{P}'(s_B \mid \circ \diamond^{\geq n-1} \{s_A\}) \\ &= \mathcal{P}'(s_B \mid \diamond^{\geq n} \{s_A\}), \end{aligned}$$

where the second equality holds since $s_B \neq s_A$.

Since we have a finite eager attractor, Theorem 6.1 of [3] with initial state s_B and final states $F = \{s_A\}$ implies that there is an $\alpha < 1$ and a constant $c \in \mathbb{R}_{>0}$ such that for all $n \in \mathbb{N}_{>0}$, $\mathcal{P}'(s_B \mid \diamond^{\geq n} \{s_A\}) \leq c\alpha^n$.

It follows that $\sum_{i=1}^{\infty} i \cdot \mathcal{P}(s_A \mid \circ \diamond^{\geq i-1} \{s_A\}) \leq c \cdot \sum_{i=1}^{\infty} i \cdot \alpha^i < \infty$, i.e., m_{s_A} (relative to \mathcal{M}) is finite. Since the Markov chain is irreducible, [23, Theorem 3.6.i, p. 81] implies that m_s is finite for every $s \in S$. A Markov chain where all expected return times are finite 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. \square

The Steady State Distribution. Algorithm 1 works in two steps.

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

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

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

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

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 defined in the previous section. It approximates the steady state distribution for *all* states, in the sense that

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

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 Algorithm 3, which solves this problem.

1. Find the intersection A_1, \dots, A_r of each BSCC of the Markov chain with the attractor. This can be done due to our assumption that $s \models \exists \diamond s'$ is decidable for all $s, s' \in A$.
2. 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 $\sum_{s \in B_i - R_i^\epsilon} \pi_s < \epsilon$.
3. If $s \in R_i^\epsilon$ for some i , do the following. First use Algorithm 1 to compute an approximation π_s^ϵ of π_s in the Markov chain induced by B_i . Then use Algorithm 2 to compute an approximation b_i^ϵ of $\mathcal{P}(s_{init} \models \diamond B_i)$. Finally, return $b_i^\epsilon \cdot \pi_s^\epsilon$.
4. If $s \notin R_i^\epsilon$ for all i , return 0.

Remark. Observe that in step 3, computing an approximation b_i^ϵ of $\mathcal{P}(s_{init} \models \diamond B_i)$ can be done by a path exploration starting in s_{init} , since the probability to reach $A \cap (B_1 \cup \dots \cup B_r)$ is 1. This is similar, but not the same, to the result in [2], since in [2] the algorithm requires that reachability is decidable for *all* pairs of states while we only require decidability in the attractor.

The Limiting Average Expected Reward. First, we compute the limiting average expected reward for irreducible Markov chains and then we extend the algorithm to non-irreducible Markov chains. This is analogous to the expected residence time: we computed the steady state distribution for irreducible Markov chains, and then extended it to the expected residence time for 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 limiting average expected reward does not depend on the initial state for such Markov chains.

Given a reward function f , 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. For a BSCC B_i , we use $G_{(i)}(f)$ to denote the limiting average expected reward of the induced Markov chain \mathcal{M}_i .

First, for each BSCC B_i , we compute an approximation b_i^ϵ of the probability to reach B_i from s_{init} . Then, for each BSCC B_i , we use Algorithm 4 to compute an approximation $G_{(i)}^\epsilon(f)$ of $G_{(i)}(f)$. 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 a sufficiently large K . If a run contains a state $s \in S - R^\epsilon$, then the last K states before s cannot be 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 $A^{\leq -1} = \emptyset$), 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 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 , we have $\sum_{s \in S - A^{\leq K}} \pi_s \leq \epsilon$, for each $\epsilon > 0$ and $K \geq \frac{\log \epsilon - 2 \log(1-\beta)}{\log \beta}$.*

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

$$\sum_{s \in A^{=k}} \pi_s = \sum_{s \in A^{=k}} \lim_{n \rightarrow \infty} \mathcal{P}(s_{init} \models \bigcirc^{=n} s) = \lim_{n \rightarrow \infty} \sum_{s \in A^{=k}} \mathcal{P}(s_{init} \models \bigcirc^{=n} s) = \lim_{n \rightarrow \infty} \mathcal{P}(s_{init} \models \bigcirc^{=n} A^{=k})$$

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 $1, \dots, n-k$, but they cannot visit A in steps $n-k+1, \dots, n$ (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 :

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

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

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

In the last equality, we use the fact that the Cesàro limit equals the usual limit if that exists. 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 in the lemma statement. \square

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 $\sum_{i=1}^N i \cdot \mathcal{P}(s \models \bigcirc \diamond^{=i-1} \{s\})$ 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 first inequality holds since (1) and (2) are satisfied and the second inequality holds by (1). 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 of π_s such that $|\pi_s^\epsilon - \pi_s| \leq \epsilon$.

Constants

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

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

Variables

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

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

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

Notice that for a given m , both $A^{\leq m}$ and $\mathcal{P}(s' \models \bigcirc(\diamond^{=m} s'))$ can be computed: since the Markov chain is effective, we can just enumerate all paths of length m starting from s .

We first show termination. As the number of iterations tends to infinity, m'_s converges from below to m_s by definition. Hence, $\sum_{s \in R^\epsilon} 1/m'_s$ converges from above to $\sum_{s \in R^\epsilon} \pi_s \leq 1$. Thus, the termination condition on line 6 is satisfied after a finite number of iterations.

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 since $m'_s \leq m_s$, 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 give an algorithm to approximate the expected residence time for arbitrary Markov chains with

finite eager attractors (not necessarily irreducible). Throughout this section, we fix an effective Markov chain that has a finite eager attractor A with parameter β and use the notation 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*

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

Proof. For any $N \geq 0$, we have

$$\begin{aligned} Res(s_{init}, s) &= \lim_{n \rightarrow \infty} \mathcal{P}(s_{init} \models \bigcirc^{\leq n} s) = \\ &= \lim_{n \rightarrow \infty} \mathcal{P}(s_{init} \models \bigcirc^{\leq n} s | s_{init} \models \diamond^{\leq N} A') \cdot \mathcal{P}(s_{init} \models \diamond^{\leq N} A') \\ &+ \lim_{n \rightarrow \infty} \mathcal{P}(s_{init} \models \bigcirc^{\leq n} s | s_{init} \models \diamond^{> N} A') \cdot \mathcal{P}(s_{init} \models \diamond^{> N} A') \\ &+ \lim_{n \rightarrow \infty} \mathcal{P}(s_{init} \models \bigcirc^{\leq n} s | s_{init} \models \square \neg A') \cdot \mathcal{P}(s_{init} \models \square \neg A') \end{aligned}$$

In this expression, the first term will be important. Denote it by $Res^{\leq N}(s_{init}, s)$. The third term equals zero by Lemma 2.1. Since the series $\sum_{i=0}^{\infty} \mathcal{P}(s_{init} \models \diamond^i A')$ converges, we must have $\lim_{N \rightarrow \infty} \mathcal{P}(s_{init} \models \diamond^{> N} A') = 0$. Thus, for any $\varepsilon > 0$, there exists an N such that

$$0 \leq Res(s_{init}, s) - Res^{\leq N}(s_{init}, s) \leq \varepsilon. \quad (5)$$

We now prove the two cases of the lemma separately.

Case $s \notin B'$. Then $\lim_{n \rightarrow \infty} \mathcal{P}(s_{init} \models \bigcirc^{\leq n} s | s_{init} \models \diamond^{\leq N} A') = 0$ because s can only be reached in the first N steps by runs in $(s_{init} \models \diamond^{\leq N} A')$. Hence, $Res^{\leq N}(s_{init}, s) = 0$, and (5) reduces to

$$0 \leq Res(s_{init}, s) \leq \varepsilon.$$

Since this holds for all $\varepsilon > 0$, we must have $Res(s_{init}, s) = 0$.

Case $s \in B_i$. Since $\mathcal{P}(s_{init} \models \bigcirc^{\leq n} s | s_{init} \models \diamond^{\leq N} A_j) = 0$ if $j \neq i$, we have

$$\begin{aligned} Res^{\leq N}(s_{init}, s) &= \lim_{n \rightarrow \infty} \mathcal{P}(s_{init} \models \bigcirc^{\leq n} s | s_{init} \models \diamond^{\leq N} A') \cdot \mathcal{P}(s_{init} \models \diamond^{\leq N} A') \\ &= \lim_{n \rightarrow \infty} \sum_{k=0}^N \sum_{s' \in A'} \mathcal{P}(s_{init} \models \bigcirc^{\leq n} s | s_{init} \models (\neg A') \mathcal{U}^k s') \\ &\quad \cdot \mathcal{P}(s_{init} \models (\neg A') \mathcal{U}^k s'). \\ &= \lim_{n \rightarrow \infty} \sum_{k=0}^N \sum_{s' \in A_i} \mathcal{P}(s_{init} \models \bigcirc^{\leq n} s | s_{init} \models (\neg A') \mathcal{U}^k s') \\ &\quad \cdot \mathcal{P}(s_{init} \models (\neg A') \mathcal{U}^k s'). \end{aligned}$$

We now concentrate on the first factor inside the sums. For any $k \geq 0$ and $s' \in A_i$, we have

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

Observe that in the second equality, the first term does not depend on n . Therefore, it vanishes as n goes to infinity. The last equality follows from Theorem 4.1(iii).

We insert the result into the previous equation and obtain

$$\begin{aligned} Res^{\leq N}(s_{init}, s) &= \lim_{n \rightarrow \infty} \sum_{k=0}^N \sum_{s' \in A_i} \pi_s \cdot \mathcal{P}(s_{init} \models (\neg A') \mathcal{U}^k s') \\ &= \pi_s \cdot \mathcal{P}(s_{init} \models \diamond^{\leq N} A_i). \end{aligned}$$

We combine this with (5) and obtain that for all $\varepsilon > 0$ there is an N such that

$$0 \leq Res(s_{init}, s) - \pi_s \cdot \mathcal{P}(s_{init} \models \diamond^{\leq N} A_i) \leq \varepsilon$$

Moreover, for any $\varepsilon > 0$ we can choose an N such that

$$0 \leq \mathcal{P}(s_{init} \models \diamond A_i) - \mathcal{P}(s_{init} \models \diamond^{\leq N} A_i) \leq \varepsilon.$$

It follows that we must have

$$Res(s_{init}, s) = \pi_s \cdot \mathcal{P}(s_{init} \models \diamond A_i). \quad \square$$

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 unlike the others, Algorithm 2 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 $\{A_1, \dots, A_r\}$ of the attractor with each BSCC, an initial state $s_{init} \in S$, and an error threshold $\epsilon \in \mathbb{R}_{>0}$.

Return value

Lower approximations $b_1^\epsilon, \dots, b_r^\epsilon$ with $b_i^\epsilon \leq \mathcal{P}(s_{init} \models \diamond B_i)$, such that $\sum_{i=1}^r |b_i^\epsilon - \mathcal{P}(s_{init} \models \diamond B_i)| \leq \epsilon$.

Variables

$n : \mathbb{N}$ (initially set to 0)
 $b_1^\epsilon, \dots, b_r^\epsilon : \mathbb{R}$ (initially all are set to 0)

1. **repeat**
2. **for** $i \leftarrow 1$ to r
3. $b_i^\epsilon \leftarrow b_i^\epsilon + \mathcal{P}(s_{init} \models \diamond^{=n} A_i)$
4. $n \leftarrow n + 1$
5. **until** $\sum_{i=1}^r b_i^\epsilon \geq 1 - \epsilon$
6. **return** $(b_1^\epsilon, \dots, b_r^\epsilon)$

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), for all $i : 1 \leq i \leq r$,

$$\begin{aligned} b_i^\epsilon &= \mathcal{P}(s_{init} \models \diamond^{\leq n} A_i) \leq \mathcal{P}(s_{init} \models \diamond A_i) \\ &= \mathcal{P}(s_{init} \models \diamond B_i), \end{aligned}$$

where the last equality follows from Lemma 2.1(ii). Therefore, the termination condition guarantees that

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

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

1. Compute the BSCCs A_1, \dots, 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. $\epsilon_2 \leftarrow 3\epsilon / (4r)$
4. **for** $i \leftarrow 1 \dots r$
5. Use the method of Algorithm 1 to compute a set $R_i^{\epsilon_2}$ for the Markov chain induced by B_i such that $\sum_{s' \in B_i - R_i^{\epsilon_2}} \pi_{s'} \leq \epsilon_2 / 3$.
6. **if** $s \in R_i^{\epsilon_2}$
7. Use the method of Algorithm 1 to compute approximations $\pi_{s'}^{\epsilon_2}$ for all $\pi_{s'}$ where $s' \in R_i^{\epsilon_2}$ in the Markov chain induced by B_i , such that $\sum_{s' \in R_i^{\epsilon_2}} |\pi_{s'}^{\epsilon_2} - \pi_{s'}| \leq 2\epsilon_2 / 3$.
8. Use Algorithm 2 to compute approximations $b_j^{\epsilon_1}$ of $\mathcal{P}(s_{init} \models \diamond B_j)$ for all j , such that $\sum_{j=1}^r |b_j^{\epsilon_1} - \mathcal{P}(s_{init} \models \diamond B_j)| \leq \epsilon_1$.
9. **return** $Res^\epsilon(s_{init}, s) = b_i^{\epsilon_1} \cdot \pi_s^{\epsilon_2}$
10. **return** $Res^\epsilon(s_{init}, s) = 0$

Similarly to the previous section, we give a slightly stronger result than required. In fact, Algorithm 3 approximates the expected residence time for all states in the sense that

$$\sum_{s \in S} |Res^\epsilon(s_{init}, s) - Res(s_{init}, s)| \leq \epsilon. \quad (6)$$

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

$$\begin{aligned} & \sum_{s \in R_i^{\epsilon_2}} |Res^\epsilon(s_{init}, s) - Res(s_{init}, s)| \\ &= \sum_{s \in R_i^{\epsilon_2}} |b_i^{\epsilon_1} \cdot \pi_s^{\epsilon_2} - \mathcal{P}(s_{init} \models \diamond B_i) \cdot \pi_s| \\ &= \sum_{s \in R_i^{\epsilon_2}} |b_i^{\epsilon_1} \cdot (\pi_s^{\epsilon_2} - \pi_s) + (b_i^{\epsilon_1} - \mathcal{P}(s_{init} \models \diamond B_i)) \cdot \pi_s| \\ &\leq \sum_{s \in R_i^{\epsilon_2}} |\pi_s^{\epsilon_2} - \pi_s| + |b_i^{\epsilon_1} - \mathcal{P}(s_{init} \models \diamond B_i)| \\ &\leq 2\epsilon_2 / 3 + \epsilon_1 = 3\epsilon / (4r). \end{aligned}$$

Hence,

$$\sum_{s \in R_1^{\epsilon_2} \cup \dots \cup R_r^{\epsilon_2}} |Res^\epsilon(s_{init}, s) - Res(s_{init}, s)| \leq 3\epsilon / 4.$$

Moreover, by the condition on line 5 of the algorithm, we have

$$\sum_{s \in B' - (R_1^{\epsilon_1} \cup \dots \cup R_r^{\epsilon_1})} |Res^\epsilon(s_{init}, s) - Res(s_{init}, s)| \leq \epsilon/4.$$

For states $s \in S - B'$, the error in the approximation is 0, and hence (6) 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 Limiting Average Expected Reward

In this section, we show how to compute arbitrarily close approximations of the limiting average expected reward for a Markov chain with a finite eager 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 $\mathcal{M} = (S, P)$, a finite eager 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 $|G^\epsilon(f) - G(f)| \leq \epsilon$.

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** $\sum_{s \in R^{\epsilon_1}} \pi_s^{\epsilon_1} \cdot f(s)$

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

that the error in the approximation is

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

Non-irreducible Markov Chains. Given a Markov chain with a finite eager attractor and a reward function f , recall that for a BSCC B_i , $G_{(i)}(f)$ denotes the limiting average expected reward of the induced Markov chain \mathcal{M}_i .

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 a finite eager attractor, for any initial state s_{init} and any bounded reward function f , $G_{s_{init}}(f)$ always exists and satisfies*

$$G_{s_{init}}(f) = \sum_{i=1}^r \mathcal{P}(s_{init} \models \diamond B_i) \cdot G_{(i)}(f).$$

In order to prove Lemma 7.1, we first prove the following.

Lemma 7.2 *For any Markov Chain \mathcal{M} with an eager finite attractor A and any state s_{init} , the following holds for each BSCC B_i and each $N \geq 0$.*

$$\lim_{n \rightarrow \infty} E_n^s(f | s_{init} \models \diamond^{\leq N} A_i) = G_{(i)}(f).$$

Proof. Given a state $s_i \in A_i$, expanding the definition of $E_n^{s_{init}}(f)$, we obtain

$$\begin{aligned} & E_n^{s_{init}}(f | s_{init} \models (\neg A_i) \mathcal{U}^k s_i) \\ &= \frac{\frac{1}{n+1} \sum_{\pi \in \Pi_{s_{init}, s_i}^k(A_i)} \sum_{\pi' \in \Pi_{s_i}^{n-k}} (f(\pi) - f(s_i) + f(\pi')) P(\pi) P(\pi')}{\mathcal{P}(s_{init} \models (\neg A_i) \mathcal{U}^k s_i)} \\ &= \frac{\frac{1}{n+1} \sum_{\pi \in \Pi_{s_{init}, s_i}^k(A_i)} (f(\pi) - f(s_i)) P(\pi) \sum_{\pi' \in \Pi_{s_i}^{n-k}} P(\pi')}{\mathcal{P}(s_{init} \models (\neg A_i) \mathcal{U}^k s_i)} \\ &\quad + \frac{\frac{1}{n+1} \sum_{\pi \in \Pi_{s_{init}, s_i}^k(A_i)} P(\pi) \sum_{\pi' \in \Pi_{s_i}^{n-k}} f(\pi') P(\pi')}{\mathcal{P}(s_{init} \models (\neg A_i) \mathcal{U}^k s_i)}. \end{aligned}$$

The first term vanishes as n tends to infinity since $|f(\pi) - f(s_i)| \leq kM$. In the second term, observe that by definition we have $\mathcal{P}(s_{init} \models (\neg A_i) \mathcal{U}^k s_i) =$

$\sum_{\pi \in \Pi_{s_{init}, s_i}^k(A_i)} P(\pi)$. Thus, by simplifying, we get

$$\begin{aligned} & \lim_{n \rightarrow \infty} E_n^{s_{init}}(f | s_{init} \models (\neg A_i) \mathcal{U}^{-k} s_i) \\ &= \lim_{n \rightarrow \infty} \frac{1}{n+1} \sum_{\pi' \in \Pi_{s_i}^{n-k}} f(\pi') P(\pi') \\ &= \lim_{n \rightarrow \infty} \frac{1}{n+1+k} \sum_{\pi' \in \Pi_{s_i}^n} f(\pi') P(\pi') \\ &= G_{(i)}(f). \end{aligned}$$

Finally,

$$\begin{aligned} & \lim_{n \rightarrow \infty} E_n^{s_{init}}(f | s_{init} \models \diamond^{\leq N} A_i) \\ &= \sum_{s_i \in A_i} \sum_{k=0}^N \lim_{n \rightarrow \infty} E_n^{s_{init}}(f | s_{init} \models (\neg A_i) \mathcal{U}^{-k} s_i) \cdot \\ & \quad \mathcal{P}(s_{init} \models (\neg A_i) \mathcal{U}^{-k} s_i | s_{init} \models \diamond^{\leq N} A_i) \\ &= G_{(i)}(f) \sum_{s_i \in A_i} \sum_{k=0}^N \mathcal{P}(s_{init} \models (\neg A_i) \mathcal{U}^{-k} s_i | s_{init} \models \diamond^{\leq N} A_i) \\ &= G_{(i)}(f). \quad \square \end{aligned}$$

Proof. of Lemma 7.1.

$$\begin{aligned} & G_{s_{init}}(f) \\ &= \lim_{n \rightarrow \infty} E_n^{s_{init}}(f) \\ &= \lim_{n \rightarrow \infty} E_n^{s_{init}}(f | s_{init} \models \diamond^{\leq N} A') \cdot \mathcal{P}(s_{init} \models \diamond^{\leq N} A') \\ & \quad + \lim_{n \rightarrow \infty} E_n^{s_{init}}(f | s_{init} \models \diamond^{> N} A') \cdot \mathcal{P}(s_{init} \models \diamond^{> N} A') \\ & \quad + \lim_{n \rightarrow \infty} E_n^{s_{init}}(f | s_{init} \models \square \neg A') \cdot \mathcal{P}(s_{init} \models \square \neg A') \\ &= \lim_{n \rightarrow \infty} \sum_{i=1}^r E_n^{s_{init}}(f | s_{init} \models \diamond^{\leq N} A_i) \cdot \mathcal{P}(s_{init} \models \diamond^{\leq N} A_i) \\ & \quad + \lim_{n \rightarrow \infty} E_n^{s_{init}}(f | s_{init} \models \diamond^{> N} A') \cdot \mathcal{P}(s_{init} \models \diamond^{> N} A') \\ &= \sum_{i=1}^r \mathcal{P}(s_{init} \models \diamond^{\leq N} A_i) \cdot G_{(i)}(f) \\ & \quad + \mathcal{P}(s_{init} \models \diamond^{> N} A') \lim_{n \rightarrow \infty} E_n^{s_{init}}(f | s_{init} \models \diamond^{> N} A'). \end{aligned}$$

The first equality holds by definition and the second by basic probability theory. The third equality follows from Lemma 2.1. In the last step, we moved the limit into the sum (which is justified since the sum is finite) and then used Lemma 7.2. The claim now follows since $\mathcal{P}(s_{init} \models \diamond^{> N} A')$ can be made arbitrarily small by taking N big, while $|E_n^{s_{init}}(f | s_{init} \models \diamond^{> N} A')|$ is bounded by M . \square

Remark. Once Lemma 7.1 is proved, it can be used to obtain a shorter proof of Lemma 6.1. Given a state

$s \in S$, define the reward function $f : S \rightarrow \mathbb{R}$ by

$$f(s') = \begin{cases} 1 & \text{if } s' = s; \\ 0 & \text{otherwise.} \end{cases}$$

By unwinding the definitions, it is straightforward to verify that

- $G_{(i)}(f) = Res(s, s) = \pi_s$ if $s \in B_i$ (the second equality follows from Theorem 4.1(iii)),
- $G_{(i)}(f) = 0$ if $s \notin B_i$, and
- $Res(s_{init}, s) = G_{s_{init}}(f)$.

The claim of Lemma 6.1 now follows from Lemma 7.1. \square

The algorithm approximates $\mathcal{P}(s_{init} \models \diamond B_i)$ 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 $\mathcal{M} = (S, P)$, a finite eager 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 $G_{s_{init}}^\epsilon(f)$ of $G_{s_{init}}(f)$ such that $|G_{s_{init}}^\epsilon(f) - G_{s_{init}}(f)| \leq \epsilon$.

1. Compute the BSCCs A_1, \dots, A_r of the finite graph (A, E) where $(s, s') \in E$ iff $s \models \exists \diamond s'$
2. $\epsilon_1 \leftarrow \epsilon / (2r)$; $\epsilon_2 \leftarrow \epsilon / (2M)$
3. **for** $i \leftarrow 1$ to r
4. Use Algorithm 4 to compute an approximation $G_{(i)}^{\epsilon_1}(f)$ of $G_{(i)}(f)$, such that $|G_{(i)}^{\epsilon_1}(f) - G_{(i)}(f)| \leq \epsilon_1$
5. Use Algorithm 2 to compute lower approximations $b_1^{\epsilon_2}, \dots, b_r^{\epsilon_2}$, with $b_i^{\epsilon_2} \leq \mathcal{P}(s_{init} \models \diamond B_i)$, such that $\sum_{j=1}^r |b_j^{\epsilon_2} - \mathcal{P}(s_{init} \models \diamond B_j)| \leq \epsilon_2$
6. **return** $\sum_{i=1}^r b_i^{\epsilon_2} \cdot G_{(i)}^{\epsilon_1}(f)$

By applying Lemma 7.1 and the triangle inequality,

the error in the approximation is

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

8 Conclusions and Future Work

We have shown that, for Markov chains with an eager finite attractor, the expected residence time and the limiting average expected reward with respect to bounded reward functions exist, and that those quantities can be effectively approximated by path exploration schemes. Since these only require reachability information *inside* the finite 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. Another possible extension is to study systems with finite attractors which satisfy only weaker probability bounds on avoiding the attractor for n steps, rather than the exponential bound in our eagerness condition.

References

- [1] P. A. Abdulla, C. Baier, S. P. Iyer, and B. Jonsson. Reasoning about probabilistic lossy channel systems. In *Proc. CONCUR 2000*, volume 1877 of *LNCS*, pp. 320–333, 2000.
- [2] P. A. Abdulla, N. Ben Henda, and R. Mayr. Verifying infinite Markov chains with a finite attractor or the global coarseness property. In *Proc. LICS' 05*, pp. 127–136, 2005.
- [3] P. A. Abdulla, N. Ben Henda, R. Mayr, and S. Sandberg. Eager Markov chains. Technical Report 2006-009, Department of Information Technology, Uppsala University, 2006.
- [4] P. A. Abdulla and B. Jonsson. Verifying programs with unreliable channels. *Information and Computation*, 127(2):91–101, 1996.
- [5] P. A. Abdulla and A. Rabinovich. Verification of probabilistic systems with faulty communication. In *Proc. FOSSACS'03*, volume 2620 of *LNCS*, pp. 39–53, 2003.
- [6] A. Aziz, K. Sanwal, V. Singhal, and R. 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 *LNCS*, pp. 34–52. Springer Verlag, 1999.
- [9] C. Baier, B. Haverkort, H. Hermanns, and J. Katoen. Automated performance and dependability evaluation using model checking. In *Proc. Performance 2002*, pp. 261–289, 2002.
- [10] N. Bertrand and Ph. Schnoebelen. Model checking lossy channels systems is probably decidable. In *Proc. FOSSACS'03*, volume 2620 of *LNCS*, pp. 120–135, 2003.
- [11] T. Brázdil and A. Kučera. Computing the expected accumulated reward and gain for a subclass of infinite Markov chains. In *Proc. FSTTCS '05*, volume 3821 of *LNCS*, pp. 372–383, 2005.
- [12] G. Cécé, A. Finkel, and S. P. Iyer. Unreliable channels are easier to verify than perfect channels. *Information and Computation*, 124(1):20–31.
- [13] E. 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] L. de Alfaro, M. Kwiatkowska, G. Norman, D. Parker, and R. Segala. Symbolic model checking of probabilistic processes using mtbdds and the Kronecker representation. In *Proc. TACAS '00*, volume 1785 of *LNCS*, pp. 123–137, 2000.
- [16] J. Esparza and K. Etessami. Verifying probabilistic procedural programs. In *Proc. FSTTCS '04*, pp. 16–31, 2004.

- [17] J. Esparza, A. Kučera, and R. Mayr. Model checking probabilistic pushdown automata. In *Proc. LICS' 04*, pp. 12–21, 2004.
- [18] J. Esparza, A. Kučera, and R. Mayr. Quantitative analysis of probabilistic pushdown automata: Expectations and variances. In *Proc. LICS' 05*, pp. 117–126, 2005.
- [19] K. Etessami and M. Yannakakis. Algorithmic verification of recursive probabilistic state machines. In *Proc. TACAS '05*, volume 3440 of *LNCS*, pp. 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*, volume 2996 of *LNCS*, pp. 340–352, 2005.
- [21] K. Etessami and M. Yannakakis. Recursive Markov decision processes and recursive stochastic games. In *Proc. ICALP '05*, volume 3580 of *LNCS*, pp. 891–903, 2005.
- [22] S.P. Iyer and M. Narasimha. Probabilistic lossy channel systems. In *TAPSOFT '97: Theory and Practice of Software Development*, volume 1214 of *LNCS*, pp. 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*, volume 2719 of *LNCS*, pp. 1008–1021, 2003.
- [26] M.Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proc. FOCS '85*, pp. 327–338, 1985.