

Accurately Computing Expected Visiting Times and Stationary Distributions in Markov Chains

Anonymous authors

Anonymous institute

Abstract. We study the accurate and efficient computation of the expected number of times each state is visited in discrete- and continuous-time Markov chains. To obtain sound accuracy guarantees efficiently, we lift interval iteration and topological approaches known from the computation of reachability probabilities and expected rewards. We further study applications of expected visiting times, including the sound computation of the stationary distribution and expected rewards conditioned on reaching multiple goal states. The implementation of our methods in the probabilistic model checker **Storm** scales to large systems with millions of states. Our experiments on the quantitative verification benchmark set show that the computation of stationary distributions via expected visiting times consistently outperforms existing approaches — sometimes by several orders of magnitude.

1 Introduction

Expected visiting times. Common questions for the quantitative analysis of Markov chains include reachability probabilities, stationary distributions, and expected rewards [36]. Many authors [38,24,56,25,46,49,20] have recognized the importance of another quantity namely *expected visiting times (EVTs)*, which describe the expected time a system spends in each state. EVT_s are characterized as the unique solution of a linear equation system [38]. They are not only relevant in their own right, but also useful to obtain various other quantities, including the ones mentioned above. This applies particularly to *forward analyses* which aim at computing, e.g., the distribution over terminal states given an initial distribution.

Sound approximation of EVT_s. In the context of (probabilistic) model checking, the two main requirements for any numeric procedure are *scalability* and *soundness*, i.e., the error in the reported result has to be bounded by a predefined threshold. Scalability is typically achieved via numerically robust iterative methods [53,57,58] such as the *Jacobi* or *Gauss-Seidel method* [57]. In general, these methods do not converge to the exact solution after a finite number of iterations. Thus, the procedure is usually stopped as soon as a termination criterion is satisfied [53]. However, standard stopping criteria such as small difference of consecutive iterations are not sound in the above sense: They do not actually indicate how close the approximation is to the true solution. Since the correctness of results in model checking, especially for safety-critical systems, is crucial,

several authors have proposed sound iterative algorithms [27,7,51,31]. While these works focus on computing quantities such as reachability probabilities and expected rewards, the sound computation of EVTs has not yet been studied.

Motivating example: Verifying sampling algorithms. To illustrate the use of EVTs in probabilistic verification tasks, consider the Markov chain in Figure 1. It is a finite-state model of a program — the *Fast Dice Roller* [44] — which takes as input an integer $N \geq 1$ and produces a uniformly distributed output in $\{1, \dots, N\}$ using unbiased coin flips only. The Fast Dice Roller thus solves a generalized *Bernoulli Factory* problem [37]. Our model in Figure 1 is for the case where $N = 6$ is fixed. How can we establish that each of the terminal states $\blacksquare, \dots, \blacksquare$ is indeed reached with probability sufficiently close to or exactly $\frac{1}{6}$?

The standard approach for answering this question is to solve N linear equation systems, one for each terminal state [6, Ch. 10]. An alternative (and seemingly less well-known) method is to compute the EVTs of each state in the Markov chain, which requires solving just a single linear equation system. All N desired probabilities can then easily be derived from the EVTs [38]: For instance, the states s_3, s_4, s_6 can all be shown¹ to have EVT $\frac{1}{3}$, and thus the reachability probabilities of the terminal states are all $\frac{1}{6} = \frac{1}{2} \cdot \frac{1}{3}$. Similarly, EVTs are useful for computing conditional expected rewards.




Fig. 1: Fast Dice Roller

For Bernoulli Factories, this allows us to examine if some outcomes take longer to compute on expectation than others, which is important to analyze possible side channel attacks in a security context. Furthermore, we show in this paper how computing stationary distributions reduces to EVTs. Such distributions provide insight into a system’s long-term behaviour; applications include the *mean-payoff* of a given policy in a Markov decision process (MDP) [50, Pr. 8.1.1], or the distribution computed by a Chemical Reaction Network [13].

Contributions. In summary, the contributions of this paper are as follows:

- We describe, analyze, and implement the *first sound numerical approximation algorithm for EVTs* in finite discrete- and continuous-time Markov chains. Our algorithm is an adaption of the known *Interval Iteration (II)* [45,28,7].
- We show that computing (sound bounds on) a Markov chain’s *stationary distribution* reduces to EVT computations. The resulting algorithm significantly outperforms preexisting techniques [43,47] for stationary distributions.
- Similarly, we show how the *conditional expected rewards* until reaching each of the, say, M absorbing states of a Markov chain can be obtained by computing the EVTs and solving a second linear equation system — this is in contrast to the standard approach which requires solving M equation systems [6, Ch. 10].
- We implement our algorithm in the probabilistic model checker **Storm** [33] and demonstrate its scalability on various benchmarks.

¹ See Section 4 for how to compute these numbers. An EVT of $\frac{1}{3}$ means that *on average*, the state is visited $\frac{1}{3}$ times. In general, EVTs of reachable recurrent states are ∞ , and the EVT of reachable transient states can take any value in $(0, \infty)$.

Outline. We define general notation and EVTs in Sections 2 and 3, respectively. In Section 4, we present our sound iterative algorithms for computing EVTs approximately. Sections 5 and 6 present the reductions of stationary distributions and conditional expected rewards to EVTs. We report on the experimental evaluation of our algorithms in Section 7 and summarize related work in Section 8.

2 Background

Let \mathbb{N} denote the set of non-negative integers and $\overline{\mathbb{R}} = \mathbb{R} \cup \{\infty, -\infty\}$ the set of extended real numbers. We equip finite sets $S \neq \emptyset$ with an arbitrary indexing $S = \{s_1, \dots, s_n\}$ and identify functions of type $\mathbf{v}: S \rightarrow \overline{\mathbb{R}}$ and $\mathbf{A}: S \times S' \rightarrow \overline{\mathbb{R}}$ with (column) vectors $\mathbf{v} \in \overline{\mathbb{R}}^{|S|}$ and matrices $\mathbf{A} \in \overline{\mathbb{R}}^{|S| \times |S'|}$, respectively. \mathbf{I} denotes the identity matrix. Vectors are compared component-wise, i.e., $\mathbf{v} \leq \mathbf{v}'$ iff for all $s \in S$, $\mathbf{v}(s) \leq \mathbf{v}'(s)$. *Iverson brackets* $\llbracket B \rrbracket$ cast the truth value of a Boolean expression B to a numerical value 1 or 0, such that $\llbracket B \rrbracket = 1$ iff B is true.

Definition 1. A discrete-time Markov chain (DTMC) is defined as a triple $\mathcal{D} = (S^{\mathcal{D}}, \mathbf{P}^{\mathcal{D}}, \iota_{\text{init}}^{\mathcal{D}})$, where $S^{\mathcal{D}}$ is a finite set of states, $\mathbf{P}^{\mathcal{D}}: S^{\mathcal{D}} \times S^{\mathcal{D}} \rightarrow [0, 1]$ is the transition probability function satisfying $\sum_{t \in S^{\mathcal{D}}} \mathbf{P}^{\mathcal{D}}(s, t) = 1$ for all $s \in S$, and $\iota_{\text{init}}^{\mathcal{D}}: S \rightarrow [0, 1]$ is the initial distribution with $\sum_{s \in S} \iota_{\text{init}}^{\mathcal{D}}(s) = 1$.

We often omit the superscript from objects associated with a DTMC \mathcal{D} whenever this is clear from context, e.g., we write \mathbf{P} rather than $\mathbf{P}^{\mathcal{D}}$. An *infinite path* $\pi = s_0 s_1 \dots \in S^{\omega}$ in a DTMC $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}})$ is a sequence of states such that $\mathbf{P}(s_i, s_{i+1}) > 0$ for all $i \in \mathbb{N}$. We use $\pi[i] = s_i$ to refer to the i -th state. $\text{Paths}^{\mathcal{D}}$ denotes the set of all infinite paths in \mathcal{D} . The probability measure $\text{Pr}^{\mathcal{D}}$ over measurable subsets of $\text{Paths}^{\mathcal{D}}$ is obtained by a standard construction: For finite path $\hat{\pi}$ we set $\text{Pr}^{\mathcal{D}}(\text{Cyl}(\hat{\pi})) = \iota_{\text{init}}(\hat{\pi}[0]) \cdot \prod_{k=0}^{|\hat{\pi}|-1} \mathbf{P}(\hat{\pi}[k], \hat{\pi}[k+1])$, where the cylinder set $\text{Cyl}(\hat{\pi}) = \{\pi \in \text{Paths}^{\mathcal{D}} \mid \forall i \in \{0, \dots, |\hat{\pi}|\}: \pi[i] = \hat{\pi}[i]\}$ contains all possible infinite continuations of $\hat{\pi}$. We write $\text{Pr}_s^{\mathcal{D}}$ for the probability measure induced by \mathcal{D} with the initial distribution assigning probability 1 to $s \in S$. We use LTL-style notation for measurable sets of infinite paths. For $R, T \subseteq S$ and $k \in \mathbb{N}$, let $R\mathbf{U}^k T = \{\pi \in \text{Paths}^{\mathcal{D}} \mid \pi[k] \in T \wedge \forall i < k: \pi[i] \in R\}$ be the set of infinite paths that visit a state $s \in T$ in the k -th step while only visiting states in R before. We also define $R\mathbf{U} T = \bigcup_{k \geq 0} R\mathbf{U}^k T$, $\Diamond T = S\mathbf{U} T$ and $\Diamond^{\leq k} T = S\mathbf{U}^k T$.

Expected Rewards. A (non-negative) *random variable* over the probability space induced by \mathcal{D} is a measurable function $\mathbf{v}: \text{Paths}^{\mathcal{D}} \rightarrow \overline{\mathbb{R}}_{\geq 0}$. Its *expected value* is given by the Lebesgue integral $\mathbb{E}^{\mathcal{D}}[\mathbf{v}] = \int_{\text{Paths}^{\mathcal{D}}} \mathbf{v} d\text{Pr}^{\mathcal{D}}$. We write $\mathbb{E}_s^{\mathcal{D}}$ for the expectation obtained under $\text{Pr}_s^{\mathcal{D}}$. The *total reward* w.r.t. a reward structure $\text{rew}: S \rightarrow \mathbb{R}_{\geq 0}$ is defined by the random variable $\text{tr}_{\text{rew}}: \text{Paths}^{\mathcal{D}} \rightarrow \overline{\mathbb{R}}_{\geq 0}$ with $\text{tr}_{\text{rew}}(\pi) = \sum_{k=0}^{\infty} \text{rew}(\pi[k])$; the *expected total reward* is $\mathbb{E}^{\mathcal{D}}[\text{tr}_{\text{rew}}]$.

Connectivity in DTMCs. A *strongly connected component (SCC)* of a DTMC \mathcal{D} is a set of states $C \subseteq S$ such that any $s, t \in C$ are mutually reachable, i.e., $\text{Pr}_s(\Diamond\{t\}) > 0$ and $\text{Pr}_t(\Diamond\{s\}) > 0$, and there is no proper subset of C that satisfies this property. An SCC C is called *bottom SCC (BSCC)* if no state outside C




Fig. 2: Running example DTMC. The individual EVTs are below the states.

is reachable from C . In the following, $\text{SCC}^{\mathcal{D}}$ denotes the set of SCCs of \mathcal{D} . We call a DTMC *absorbing* if all its BSCCs are singleton sets, and *irreducible* if $\text{SCC}^{\mathcal{D}} = \{S\}$. The SCCs are ordered by a strict partial order \hookrightarrow based on the topology of the DTMC, where $C' \hookrightarrow C$ if and only if $\Pr_{s'}(\Diamond C) > 0$ for some $s' \in C'$ and $C \neq C'$. An *SCC chain* of \mathcal{D} is a sequence $\kappa = C_0 \hookrightarrow C_1 \hookrightarrow \dots \hookrightarrow C_k$ of SCCs $C_0, C_1, \dots, C_k \in \text{SCC}^{\mathcal{D}}$, where $k \geq 0$. The set of all SCC chains in \mathcal{D} is denoted by $\text{Chains}^{\mathcal{D}}$, and $\text{Chains}_{\text{tr}}^{\mathcal{D}}$ denotes the set of SCC chains that do not contain a BSCC. The *length* of SCC chain $\kappa = C_0 \hookrightarrow C_1 \hookrightarrow \dots \hookrightarrow C_k$ is $|\kappa| = k$.

A state s is called *transient* if the probability that the DTMC, starting from s , will ever return to s is strictly less than one, otherwise, s is a *recurrent* state. Thus, in a finite MC, recurrent states are precisely the states contained in the BSCCs whereas the transient states coincide with non-BSCC states. The sets of recurrent and transient states in a DTMC are denoted by S_{re} and S_{tr} , respectively.

Example 1. In the DTMC \mathcal{D} depicted in Figure 2, states s_1, s_2, s_3, s_4 are transient, and s_5, s_6, s_7 are recurrent. Also, $\text{SCC}^{\mathcal{D}} = \{\{s_1, s_2\}, \{s_3\}, \{s_4\}, \{s_5, s_6\}, \{s_7\}\}$, where only $\{s_5, s_6\}$ and $\{s_7\}$ are BSCCs. An example SCC chain is $\{s_3\} \hookrightarrow \{s_1, s_2\} \hookrightarrow \{s_4\} \hookrightarrow \{s_7\}$; its length is 3.

Stationary distributions. The stationary distribution (also referred to as steady-state or long-run distribution) is a probability distribution that specifies the fraction of time spent in each state in the long run (see, e.g., [6, Def. 10.79]).

Definition 2. The stationary distribution of DTMC \mathcal{D} is given by $\theta^{\mathcal{D}} \in [0, 1]^{|S|}$ with $\theta^{\mathcal{D}}(s) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \Pr(\Diamond^{\leq k} \{s\})$.

If \mathcal{D} is irreducible, the stationary distribution is given by the unique eigenvector $\theta^{\mathcal{D}}$ satisfying $\theta^{\mathcal{D}} = \theta^{\mathcal{D}} \cdot \mathbf{P}$ and $\sum_{s \in S} \theta^{\mathcal{D}}(s) = 1$, see, e.g., [41, Thm. 4.18]. If \mathcal{D} is reducible, $\theta^{\mathcal{D}}$ can be obtained by combining a reachability analysis and the eigenvector computation for each BSCC individually, see Section 5.

Definition 3. A continuous-time Markov chain (CTMC) is a quadruple $\mathcal{C} = (S, \mathbf{P}, \iota_{\text{init}}, \mathbf{r})$, where $(S, \mathbf{P}, \iota_{\text{init}})$ is a DTMC and $\mathbf{r}: S \rightarrow \mathbb{R}_{>0}$ defines exit rates.

CTMCs extend DTMCs by assigning the rate of an exponentially distributed residence time to each state $s \in S$. We denote the *embedded DTMC* of a CTMC \mathcal{C} by $\text{emb}(\mathcal{C}) = (S, \mathbf{P}, \iota_{\text{init}})$. The semantics are defined in the usual way (see, e.g., [4, 42]). An *infinite timed path* in a CTMC $\mathcal{C} = (S, \mathbf{P}, \iota_{\text{init}}, \mathbf{r})$ is a sequence

$\pi = s_0 \xrightarrow{\tau_0} s_1 \xrightarrow{\tau_1} \dots$ consisting of states $s_0, s_1, \dots \in S$ and time instances $\tau_0, \tau_1 \dots \in \mathbb{R}_{\geq 0}$, such that $\mathbf{P}(s_i, s_{i+1}) > 0$ for all $i \in \mathbb{N}$. We denote the time τ_i spent in state s_i by $time_i(\pi)$. Notations $\pi[i]$ and Paths^C are as for DTMCs. The probability measure of \mathcal{C} over infinite timed paths [4] is denoted by Pr^C . In a CTMC, the *total reward* w.r.t. a reward structure $\text{rew}: S \rightarrow \mathbb{R}_{\geq 0}$ is the random variable $\text{tr}_{\text{rew}}: \text{Paths}^C \rightarrow \mathbb{R}_{\geq 0}$ with $\text{tr}_{\text{rew}}(\pi) = \sum_{i=0}^{\infty} \text{rew}(\pi[i]) \cdot time_i(\pi)$.

3 Expected Visiting Times

We provide characterizations of expected visiting times for a fixed DTMC $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}})$. Omitted proofs are in Appendix A.1.

Definition 4. *The expected visiting time (EVT) of a state $s \in S$ is the expected value $\mathbb{E}^{\mathcal{D}}[\text{vt}_s]$ of the random variable vt_s with $\text{vt}_s(\pi) = \sum_{k=0}^{\infty} \llbracket \pi[k] = s \rrbracket$.*

Example 2. The EVTs of the DTMC from Figure 2 are depicted below its states.

Intuitively, random variable vt_s *counts* the number of times state s occurs on an infinite path. Consequently, the EVTs of unreachable states and reachable recurrent states in a DTMC are always 0 and ∞ , respectively. For this reason we focus on the EVTs of the transient states S_{tr} . The following lemma provides an alternative characterization of EVTs in terms of expected total rewards.

Lemma 1. *For a fixed $s \in S$ and $x \in \mathbb{R}_{>0}$ the reward structure $\text{rew}: S \rightarrow \mathbb{R}_{\geq 0}$ given by $\text{rew}(t) = x \cdot \llbracket t = s \rrbracket$ satisfies $\mathbb{E}[\text{vt}_s] = \frac{1}{x} \cdot \mathbb{E}[\text{tr}_{\text{rew}}]$.*

By Lemma 1, EVTs can be obtained using existing algorithms for expected total rewards. This approach is, however, inefficient for computing the EVTs of multiple states since it requires solving an equation system for each single state.

Next, we elaborate on EVTs for multiple states as a solution of a single linear equation system. In [38, Def. 3.2.2], EVTs are defined using the so-called *fundamental matrix* for absorbing DTMCs. The fundamental matrix contains as its coefficients for each possible start and target state s and t the EVT $\mathbb{E}_s[\text{vt}_t]$. Computing the fundamental matrix explicitly becomes infeasible for large models as it requires determining the inverse of a $|S_{\text{tr}}| \times |S_{\text{tr}}|$ matrix. To obtain the vector $(\mathbb{E}^{\mathcal{D}}[\text{vt}_s])_{s \in S_{\text{tr}}}$ of EVTs that take the initial distribution of \mathcal{D} into account, it suffices to solve an equation system which is linear in the size of the DTMC. The same equation system arises by applying the dual linear program for expected rewards in MDPs [50, Ch. 7.2.7] to the special case of DTMCs.

Theorem 1 ([38, Cor. 3.3.6]). *$(\mathbb{E}[\text{vt}_s])_{s \in S_{\text{tr}}}$ is the unique solution $(\mathbf{x}(s))_{s \in S_{\text{tr}}}$ of the following equation system: $\forall s \in S_{\text{tr}}: \mathbf{x}(s) = \iota_{\text{init}}(s) + \sum_{t \in S_{\text{tr}}} \mathbf{P}(t, s) \cdot \mathbf{x}(t)$.*

Intuitively, this equation system shows that a state s can be visited initially and that it can receive visits from its predecessor states, i.e., the EVT is computed by considering the *incoming* transitions to a state. As a consequence, we obtain that the EVTs of the transient states are always finite. In particular, if $s \in S_{\text{tr}}$ is reachable, then $\mathbb{E}[\text{vt}_s] \in \mathbb{R}_{>0}$, and otherwise $\mathbb{E}[\text{vt}_s] = 0$.

Example 3. Reconsider the DTMC from Figure 2 with transient states $S_{\text{tr}} = \{s_1, s_2, s_3, s_4\}$. The EVTs of S_{tr} are the unique solution $(\mathbf{x}(s))_{s \in S_{\text{tr}}}$ of

$$\begin{aligned} \mathbf{x}(s_1) &= 0.4 + 1.0 \cdot \mathbf{x}(s_2) + 0.7 \cdot \mathbf{x}(s_3) & \mathbf{x}(s_2) &= 0.5 \cdot \mathbf{x}(s_1) \\ \mathbf{x}(s_4) &= 0.5 \cdot \mathbf{x}(s_1) + 0.3 \cdot \mathbf{x}(s_3) + 0.8 \cdot \mathbf{x}(s_4) & \mathbf{x}(s_3) &= 0.6 \end{aligned}$$

Expected visiting times in CTMCs. Following [39], we define the EVT of a state $s \in S^C$ of a CTMC \mathcal{C} as the expected value $\mathbb{E}^{\mathcal{C}}[\mathbf{vt}_s]$ of the random variable $\mathbf{vt}_s: \text{Paths}^{\mathcal{C}} \rightarrow \mathbb{R}_{\geq 0}$ with $\mathbf{vt}_s(\pi) = \sum_{k=0}^{\infty} \mathbb{I}[\pi[k] = s] \cdot \text{time}_k(\pi)$. Intuitively, \mathbf{vt}_s considers the total time the system spends in state s . Computing EVTs in CTMCs reduces to the discrete-time case: The EVT of state s coincides with the EVT in the embedded DTMC weighted by the expected residence time $\frac{1}{r(s)}$ in s :

Theorem 2. *For all states $s \in S$, it holds that $\mathbb{E}^{\mathcal{C}}[\mathbf{vt}_s] = \frac{1}{r(s)} \cdot \mathbb{E}^{\text{emb}(\mathcal{C})}[\mathbf{vt}_s]$.*

Theorem 2 implies that all results and algorithms to compute EVTs in DTMCs are readily applicable to CTMCs, too. We thus focus on DTMCs in the remainder.

4 Accurately Computing EVTs

In this section, we discuss algorithms to compute EVTs approximately: An unsound value iteration algorithm (Section 4.1), its sound interval iteration extension (Section 4.2), and finally a topological, i.e., SCC-wise algorithm (Section 4.3). Since the EVTs for recurrent states are always either 0 or ∞ , we focus on the EVTs of the transient states. Omitted proofs are in Appendix A.2.

4.1 Value Iteration

Value Iteration (VI) was originally introduced to approximate expected rewards in MDPs [8]. In a broader sense, VI simply refers to iterating a function $f: \mathbb{R}^{|S|} \rightarrow \mathbb{R}^{|S|}$ (called *Bellman operator* in the MDP setting) from some given initial vector $\mathbf{x}^{(0)}$, i.e., to compute the sequence $\mathbf{x}^{(1)} = f(\mathbf{x}^{(0)})$, $\mathbf{x}^{(2)} = f(\mathbf{x}^{(1)})$, etc. Instances of VI are usually set up such that the sequence converges to a (generally non-unique) fixed point $\mathbf{x} = f(\mathbf{x})$. In this paper, we only consider VI for the case where f is a linear function $f(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b}$, where \mathbf{A} and \mathbf{b} are a matrix and a vector, respectively. A fixed point \mathbf{x} of f is then a solution of the linear equation system $(\mathbf{I} - \mathbf{A})\mathbf{x} = \mathbf{b}$. Other iterative methods for solving linear equation systems such as the Jacobi or Gauss-Seidel method can be considered optimized variants of VI, and are applicable in our setting as well, see Appendix C.

Value iteration for EVTs. For EVTs, the function iterated during VI is as follows:

Definition 5. *The EVTs-operator $\Phi: \mathbb{R}^{|S_{\text{tr}}|} \rightarrow \mathbb{R}^{|S_{\text{tr}}|}$ for DTMC \mathcal{D} is defined as*

$$\Phi(\mathbf{x}) = \left(\iota_{\text{init}}(s) + \sum_{t \in S_{\text{tr}}} \mathbf{x}(t) \cdot \mathbf{P}(t, s) \right)_{s \in S_{\text{tr}}}.$$

The above definition is motivated by Theorem 1. The following result, which is analogous to [50, Thm. 6.3.1], means that VI for EVTs (stated explicitly as Algorithm 1 for the sake of concreteness) works for arbitrary initial vectors.

Algorithm 1: Value iteration for EVTs without precision guarantee.**Input:** DTMC \mathcal{D} , $\mathbf{x}^{(0)} \in \mathbb{R}^{|S_{\text{tr}}|}$, $\text{crit} \in \{\text{abs}, \text{rel}\}$, $\epsilon > 0$ **Output:** $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$

```

1 for  $k = 1, 2, 3, \dots$  do
2    $\mathbf{x}^{(k)} \leftarrow \Phi(\mathbf{x}^{(k-1)})$  //  $\mathbf{x}^{(k)}(s) = \iota_{\text{init}}(s) + \sum_{t \in S_{\text{tr}}} \mathbf{x}^{(k-1)}(t) \cdot \mathbf{P}(t, s)$ ,  $s \in S_{\text{tr}}$ 
3   if  $\text{diff}^{\text{crit}}(\mathbf{x}^{(k-1)}, \mathbf{x}^{(k)}) \leq \epsilon$  then return  $\mathbf{x}^{(k)}$ 

```

Theorem 3. *The EVTs-operator from Definition 5 has the following properties:*

- (i) $(\mathbb{E}[\mathbf{v}_t]_{s \in S_{\text{tr}}})$ is the unique fixed point of Φ .
- (ii) For all $\mathbf{x}^{(0)} \in \mathbb{R}^{|S_{\text{tr}}|}$ we have $\lim_{k \rightarrow \infty} \Phi^{(k)}(\mathbf{x}^{(0)}) = (\mathbb{E}[\mathbf{v}_t]_{s \in S_{\text{tr}}})$.

When to stop VI? A general issue with value iteration is that even if the generated sequence converges to the desired fixed point in the limit, it is not easy to determine how many iterations are necessary to obtain an ϵ -precise result. An ad hoc solution, which is implemented in probabilistic model checkers such as Storm [33], prism [43], and mcsta [30], is to stop the iteration once the difference between two consecutive approximations is small, i.e., the number of iterations is the smallest $k > 0$ such that $\text{diff}(\mathbf{x}^{(k)}, \mathbf{x}^{(k-1)}) < \epsilon$ for some predefined fixed $\epsilon > 0$. Common choices for the distance diff between vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^{|S|}$ are the *absolute difference* $\text{diff}^{\text{abs}}(\mathbf{x}, \mathbf{y}) = \max_{s \in S} |\mathbf{x}(s) - \mathbf{y}(s)|$, and the *relative difference*

$$\text{diff}^{\text{rel}}(\mathbf{x}, \mathbf{y}) = \max_{s \in S} \left| \frac{\mathbf{x}(s) - \mathbf{y}(s)}{\mathbf{y}(s)} \right|,$$

where by convention $0/0 = 0$ and $a/0 = \infty$ for $a \neq 0$. As pointed out by various authors [58, 28, 7, 45, 51], there exist instances where the iteration terminates with a result which vastly differs from the true fixed point, even if ϵ is small (e.g. $\epsilon = 10^{-6}$). Appendix B considers an example of this for the EVT variant of VI.

4.2 Interval Iteration

Interval iteration (II) [45, 28, 7] is an extension of VI that formally guarantees ϵ -close results for all possible inputs. The general idea of II is to construct *two* sequences of vectors $(\mathbf{l}^{(k)})_{k \in \mathbb{N}}$ and $(\mathbf{u}^{(k)})_{k \in \mathbb{N}}$ such that for all $k \in \mathbb{N}$ we have $\mathbf{l}^{(k)} \leq \mathbf{x} \leq \mathbf{u}^{(k)}$, where \mathbf{x} is the desired fixed point solution. II can be stopped with precision guarantee ϵ once it detects that $\text{diff}(\mathbf{l}^{(k)}, \mathbf{u}^{(k)}) \leq \epsilon$.

Initial bounds for II. In general, II requires initial vectors $\mathbf{l}^{(0)}$ and $\mathbf{u}^{(0)}$ which are already sound (but perhaps very crude) lower and upper bounds on the solution. In the case of EVTs, we can use $\mathbf{l}^0 = \mathbf{0}$. Finding an upper bound $\mathbf{u}^{(0)} \geq (\mathbb{E}[\mathbf{v}_t]_{s \in S_{\text{tr}}})$ is more involved since EVTs may be unboundedly large in general. We solve this issue using a technique from [7].

II for EVTs. In Lemma 2 below we show that once we have found initial bounds $\mathbf{l}^{(0)}$ and $\mathbf{u}^{(0)}$, we can readily perform a sound II for EVTs by simply iterating the operator Φ from Definition 5 on $\mathbf{l}^{(0)}$ and $\mathbf{u}^{(0)}$ in parallel. Inspired by [7], we propose the following optimization to speed up convergence: Whenever Φ

Algorithm 2: Interval iteration for EVTs with precision guarantee.

Input: DTMC \mathcal{D} , $\mathbf{l}^{(0)} \leq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}$, $\mathbf{u}^{(0)} \geq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}$, $\text{crit} \in \{\text{abs}, \text{rel}\}$, $\epsilon > 0$

Output: $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$ with $\text{diff}^{\text{crit}}(\mathbf{x}, (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}) \leq \epsilon$

```

1 for  $k = 1, 2, \dots$  do
2    $\mathbf{l}^{(k)} \leftarrow \Phi_{\max}(\mathbf{l}^{(k-1)})$  ;  $\mathbf{u}^{(k)} \leftarrow \Phi_{\min}(\mathbf{u}^{(k-1)})$ 
3   if  $\text{diff}^{\text{crit}}(\mathbf{u}^{(k)}, \mathbf{l}^{(k)}) \leq 2 \cdot \epsilon$  then return  $\frac{1}{2}(\mathbf{l}^{(k)} + \mathbf{u}^{(k)})$ 
```

decreases the current lower bound in some entries, we retain the old values for these entries (and similar for upper bounds). The next definition formalizes this.

Definition 6. The Max and Min EVTs-operators $\Phi_{\max}, \Phi_{\min} : \mathbb{R}^{|S_{\text{tr}}|} \rightarrow \mathbb{R}^{|S_{\text{tr}}|}$ are defined by $\Phi_{\max}(\mathbf{x}) = \max\{\mathbf{x}, \Phi(\mathbf{x})\} = (\max\{\mathbf{x}(s), (\Phi(\mathbf{x}))(s)\})_{s \in S_{\text{tr}}}$ and $\Phi_{\min}(\mathbf{x}) = \min\{\mathbf{x}, \Phi(\mathbf{x})\} = (\min\{\mathbf{x}(s), (\Phi(\mathbf{x}))(s)\})_{s \in S_{\text{tr}}}$.

The following result is analogous to [7, Lem. 3.3]:

Lemma 2. Let $\mathbf{u}, \mathbf{l} \in \mathbb{R}^{|S_{\text{tr}}|}$ with $\mathbf{l} \leq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}} \leq \mathbf{u}$. Then,

- (i) $\Phi_{\max}^{(k)}(\mathbf{l}) \leq \Phi_{\max}^{(k+1)}(\mathbf{l})$ and $\Phi_{\min}^{(k)}(\mathbf{u}) \geq \Phi_{\min}^{(k+1)}(\mathbf{u})$ for all $k \in \mathbb{N}$.
- (ii) $\Phi^{(k)}(\mathbf{l}) \leq \Phi_{\max}^{(k)}(\mathbf{l}) \leq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}} \leq \Phi_{\min}^{(k)}(\mathbf{u}) \leq \Phi^{(k)}(\mathbf{u})$ for all $k \in \mathbb{N}$.
- (iii) $\lim_{k \rightarrow \infty} \Phi_{\max}^{(k)}(\mathbf{l}) = \lim_{k \rightarrow \infty} \Phi_{\min}^{(k)}(\mathbf{u}) = (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}$.

The resulting II algorithm for EVTs is presented as Algorithm 2. Note the following additional optimization: The algorithm stops as soon as $\text{diff}^{\text{crit}}(\mathbf{u}^{(k)}, \mathbf{l}^{(k)}) \leq 2\epsilon$ and returns the mean of $\mathbf{u}^{(k)}$ and $\mathbf{l}^{(k)}$, ensuring that the absolute or relative difference between $(\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}$ and the returned result is at most ϵ .

Example 4. We illustrate a run of Algorithm 2 on the DTMC from Figure 2, with $\text{crit} = \text{abs}$ and $\epsilon = 0.05$ (the following numbers are rounded to 4 decimal digits):

k	$\mathbf{l}^{(k)}$	$\mathbf{u}^{(k)}$	$\text{diff}^{\text{abs}}(\mathbf{l}^{(k)}, \mathbf{u}^{(k)})$
0	(0.000, 0.000, 0.000, 0.000)	(2.000, 2.000, 1.000, 5.000)	5.000
1	(0.400, 0.000, 0.600, 0.000)	(2.000, 1.000, 0.600, 5.000)	5.000
2	(0.820, 0.200, 0.600, 0.380)	(1.820, 1.000, 0.600, 5.000)	4.602

22	(1.639, 0.819, 0.600, 4.899)	(1.640, 0.820, 0.600, 5.000)	0.101
23	(1.639, 0.819, 0.600, 4.919)	(1.640, 0.820, 0.600, 5.000)	0.081

After $k = 23$ iterations, the algorithm stops as $\text{diff}^{\text{abs}}(\mathbf{l}^{(k)}, \mathbf{u}^{(k)}) = 0.081 \leq 2 \cdot \epsilon$ and outputs the mean $\frac{1}{2}(\mathbf{l}^{(23)} + \mathbf{u}^{(23)})$.

Theorem 4 (Correctness of Algorithm 2). Given an input DTMC \mathcal{D} , initial vectors $\mathbf{l}^{(0)}, \mathbf{u}^{(0)}$ with $\mathbf{l}^{(0)} \leq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}} \leq \mathbf{u}^{(0)}$, $\text{crit} \in \{\text{abs}, \text{rel}\}$, and a threshold $\epsilon > 0$, Algorithm 2 terminates and returns a vector $\mathbf{x}^{\text{res}} \in \mathbb{R}^{|S_{\text{tr}}|}$ satisfying $\text{diff}^{\text{crit}}(\mathbf{x}^{\text{res}}, (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}) \leq \epsilon$.

Remark 1. The monotonicity of the sequences $(\Phi_{\max}^{(k)}(\mathbf{l}))_{k \in \mathbb{N}}$ and $(\Phi_{\min}^{(k)}(\mathbf{u}))_{k \in \mathbb{N}}$ (see Lemma 2 (i)) is not used in the proof of Theorem 4. By Lemma 2 (ii), we can replace Φ_{\max} and Φ_{\min} with Φ in Algorithm 2 and still obtain ϵ -sound results. However, using Φ_{\max} and Φ_{\min} instead of Φ can lead to faster convergence.

4.3 Topological Algorithm

To increase the efficiency of VI for the analysis of rewards and probabilities in MDPs, several authors have proposed topological VI [16,17], which is also known as blockwise VI [15]. The idea is to avoid the analysis of the complete model at once and instead consider the strongly connected components (SCCs) sequentially based on the order relation \hookrightarrow . We lift this approach to EVTs and in particular consider error propagations when approximative methods are used.

SCC Restrictions. To formalize the topological approach, we introduce the *SCC restriction* $\mathcal{D}|_C[\mathbf{x}]$ of DTMC \mathcal{D} to $C \in \text{SCC}^\mathcal{D}$ with parameters $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$. Intuitively, $\mathcal{D}|_C[\mathbf{x}]$ is a DTMC-like model obtained by restricting \mathcal{D} to the states C and assigning each $s \in C$ the “initial value” $\iota_{\text{init}}^\mathcal{D}(s) + \sum_{s' \in S \setminus C} \mathbf{P}^\mathcal{D}(s', s) \cdot \mathbf{x}(s')$. The idea is that \mathbf{x} is an approximation of the EVTs of the predecessor SCCs $C' \hookrightarrow C$. We also define $\mathcal{D}|_C = \mathcal{D}|_C[\mathbf{x}]$ with $\mathbf{x}(s) = \mathbb{E}^\mathcal{D}[\mathbf{v}_t_s]$ (i.e., the *exact* EVT) for each state $s \in S$ that can reach C with positive probability in one step. See Definition 7 in Appendix A.2 for more formal details.

Example 5. The SCC restriction of the DTMC \mathcal{D} from Figure 2 to the SCC $C = \{s_1, s_2\}$ with parameters $\mathbf{x} = (\mathbb{E}^\mathcal{D}[\mathbf{v}_t_s])_{s \in S}$ is depicted in Figure 3. Note that the initial values depend only on the initial distribution $\iota_{\text{init}}^\mathcal{D}$ and the \mathbf{x} -values of the states that can reach the SCC C in one step.




Fig. 3: SCC restriction.

Lemma 3. For non-bottom SCC C and $s \in C$ we have $\mathbb{E}^\mathcal{D}[\mathbf{v}_t_s] = \mathbb{E}^{\mathcal{D}|_C}[\mathbf{v}_t_s]$.

Remark 2. Since a parametric SCC restriction $\mathcal{D}|_C[\mathbf{x}]$ is defined for arbitrary vectors $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$, the initial values do not necessarily form a probability distribution. Strictly speaking, this means that $\mathcal{D}|_C[\mathbf{x}]$ is not a DTMC in general. Thus, by abuse of notation, we define the “EVTs” in the parametric SCC restriction $\mathcal{D}|_C[\mathbf{x}]$ as the unique solution of the following linear equation system: For all $s \in C$: $\mathbf{x}(s) = \iota_{\text{init}}^{\mathcal{D}|_C[\mathbf{x}]}(s) + \sum_{t \in S_{\text{tr}}} \mathbf{P}^{\mathcal{D}|_C}(t, s) \cdot \mathbf{x}(t)$. To solve this system, we can still apply the methods described in Sections 4.1 and 4.2 without further ado.

We now describe an algorithm for computing the EVTs SCC-wise in topological order. The desired precision $\epsilon \geq 0$ is an input parameter ($\epsilon = 0$ is possible). Due to space limitations we only discuss relative precision, see Appendix D for an algorithm with absolute precision. The idea is to solve the linear equation systems tailored to the parametric SCC restrictions, each of which is constructed based on the analysis of the preceding SCCs. Algorithm 3 outlines the procedure.

The algorithm first decomposes the input DTMC \mathcal{D} into its non-bottom SCCs: The function $\text{SCC}_{\text{tr}}^\mathcal{D}$ called in Line 1 returns the set $\{C_1, \dots, C_n\}$ of SCCs of \mathcal{D} which consist of transient states, i.e., the non-bottom SCCs. Furthermore,

Algorithm 3: Topological EVT algorithm with relative precision.

Input: DTMC \mathcal{D} , $\epsilon \geq 0$
Output: $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$ with $\text{diff}^{\text{rel}}(\mathbf{x}, (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}) \leq \epsilon$

```

1  $\{C_1, \dots, C_n\} \leftarrow \text{SCC}_{\text{tr}}^{\mathcal{D}}$  // obtain the non-bottom SCCs,  $C_i \hookrightarrow C_j$  implies  $i < j$ 
2  $\mathbf{x} \leftarrow \mathbf{0}$ 
3  $\delta \leftarrow \sqrt[L+1]{1 + \epsilon} - 1$  //  $L$  is the length of a longest non-bottom SCC chain
4 for  $j = 1$  to  $n$  do
5   compute  $\hat{\mathbf{x}}$  such that  $\text{diff}^{\text{rel}}((\mathbb{E}^{\mathcal{D}|_{C_j}}[\mathbf{x}][\mathbf{vt}_s])_{s \in C_j}, \hat{\mathbf{x}}) \leq \delta$  // use e.g. Alg. 2
6   for  $s \in C_j$  do  $\mathbf{x}(s) \leftarrow \hat{\mathbf{x}}(s)$ 
7 return  $\mathbf{x}$ 
```

we assume that the SCCs are indexed such that $C_i \hookrightarrow C_j$ implies $i < j$. The algorithm considers each SCC in topological order. We assume that a “black box” can obtain the (approximate) EVTs in Line 5. Then, the vector $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$ is updated such that it contains the approximations of the EVTs in \mathcal{D} upon termination. For the analysis of an SCC C_j for some $j \in \{1, \dots, n\}$, the algorithm considers the parametric SCC restriction $\mathcal{D}|_{C_j}[\mathbf{x}]$, which is based on the result $\mathbf{x}(t)$ for states t in SCCs C_i that are topologically before C_j . For each parametric SCC restriction $\mathcal{D}|_{C_j}[\mathbf{x}]$, the algorithm computes the vector $(\mathbb{E}^{\mathcal{D}|_{C_j}}[\mathbf{x}][\mathbf{vt}_s])_{s \in C_j}$ in Line 5. Then, the algorithm updates the corresponding entries in \mathbf{x} in Line 6. After each non-bottom SCC has been considered, the algorithm terminates and returns the vector \mathbf{x} that contains the updated value for each transient state. The following lemma provides an upper bound on the error that accumulates during the topological computation.

Lemma 4. *Let $\epsilon \in [0, 1)$ and let $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$ such that for every non-bottom SCC C , $(\mathbf{x}(s))_{s \in C}$ satisfies $\text{diff}^{\text{rel}}((\mathbf{x}(s))_{s \in C}, (\mathbb{E}^{\mathcal{D}|_C}[\mathbf{x}][\mathbf{vt}_s])_{s \in C}) \leq \epsilon$. Then*

$$\text{diff}^{\text{rel}}(\mathbf{x}, (\mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}) \leq (1 + \epsilon)^{L+1} - 1,$$

where $L = \max_{\kappa \in \text{Chains}_{\text{tr}}^{\mathcal{D}}} |\kappa|$ is the largest length of a chain of non-bottom SCCs.

Theorem 5 (Correctness of Algorithm 3). *Algorithm 3 returns a vector $\mathbf{x}^{\text{res}} \in \mathbb{R}^{|S_{\text{tr}}|}$ such that $\text{diff}^{\text{rel}}(\mathbf{x}^{\text{res}}, (\mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}) \leq \epsilon$.*

5 Stationary Distributions via EVTs

We show that EVTs can be used to determine *sound* approximations of the stationary distribution $\theta^{\mathcal{D}}$ of a (reducible) DTMC \mathcal{D} . It is known that $\theta^{\mathcal{D}}$ can be computed as follows (see, e.g., [41, Thm. 4.23]). For each BSCC B :

- Compute the reachability probability $\Pr^{\mathcal{D}}(\Diamond B)$ to B .
- Determine the stationary distribution $\theta^{\mathcal{D}|_B}$ of the DTMC restricted to B .

Then, we obtain the stationary distribution for the recurrent states s in the BSCC B : $\theta^{\mathcal{D}}(s) = \Pr^{\mathcal{D}}(\Diamond B) \cdot \theta^{\mathcal{D}|_B}(s)$. For the remaining transient states, we have $\theta^{\mathcal{D}}(s) = 0$. We show how both, $\Pr^{\mathcal{D}}(\Diamond B)$ and $\theta^{\mathcal{D}|_B}$, can be computed efficiently

for every BSCC B using EVTs. We also elaborate on how relative errors propagate through the computation, allowing us to derive sound lower- and upper bounds for the stationary distribution $\theta^{\mathcal{D}}$. Omitted proofs are in Appendix A.3.

Computing BSCC reachability probabilities. The *absorption probabilities*, i.e., the probability of reaching a singleton BSCC can be computed using EVTs [38]. By collapsing the BSCCs into a single state, a slightly generalised result is obtained:

Theorem 6 ([38, Thm. 3.3.7]). *For any BSCC B of a DTMC \mathcal{D} it holds that $\Pr(\Diamond B) = \sum_{s \in B} (\iota_{\text{init}}(s) + \sum_{s' \in S_{\text{tr}}} \mathbf{P}(s', s) \cdot \mathbb{E}[\mathbf{vt}_{s'}])$.*

Applying Theorem 6, we can compute the EVTs *once* to derive the reachability probabilities for *every* BSCC. Further, when using interval iteration from Section 4.2 to obtain $\mathbf{x} \in \mathbb{R}^{|S|}$ with $\text{diff}^{\text{rel}}(\mathbf{x}, (\mathbb{E}[\mathbf{vt}_s])_{s \in S}) \leq \epsilon$ for some $\epsilon \in (0, 1)$, the relative error does not increase when deriving the reachability probabilities, i.e., $\text{diff}^{\text{rel}}(\sum_{s \in B} (\iota_{\text{init}}(s) + \sum_{s' \in S_{\text{tr}}} \mathbf{P}(s', s) \cdot \mathbf{x}(s)), \Pr(\Diamond B)) \leq \epsilon$.

Computing the stationary distribution within a BSCC. Next, we leverage EVTs to compute the stationary distribution $\theta^{\mathcal{B}}$ of an irreducible DTMC $\mathcal{B} = (S, \mathbf{P}, \iota_{\text{init}})$. This method can be applied to derive the stationary distribution $\theta^{\mathcal{D}|_B}$ of $\mathcal{D}|_B$ since the latter is an irreducible DTMC for each BSCC B . Let $v \in S$ be an arbitrary state. We construct the DTMC $\mathcal{B}^{\hat{r}^{\hat{v}}}$ in which all incoming transitions of state v are redirected to a fresh absorbing state \hat{v} . Thus, its only BSCC is $\{\hat{v}\}$, all other states are transient. Formally, $\mathcal{B}^{\hat{r}^{\hat{v}}} = (S \uplus \{\hat{v}\}, \hat{\mathbf{P}}, \hat{\iota}_{\text{init}})$, where $\hat{\mathbf{P}}(\hat{v}, \hat{v}) = 1$, $\hat{\mathbf{P}}(s, \hat{v}) = \mathbf{P}(s, v)$ and $\hat{\mathbf{P}}(s, v) = 0$ for all $s \in S$, $\hat{\mathbf{P}}(s, t) = \mathbf{P}(s, t)$ for all $s \in S, t \in S \setminus \{v\}$, and $\hat{\iota}_{\text{init}}(v) = 1$.

Theorem 7. *The stationary distribution $\theta^{\mathcal{B}}$ of an irreducible DTMC \mathcal{B} is given by $\theta^{\mathcal{B}}(s) = \frac{\mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{vt}_s]}{\sum_{t \in S} \mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{vt}_t]}$. Further, if $\text{diff}^{\text{rel}}(\mathbf{x}, (\mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{vt}_s])_{s \in S}) \leq \epsilon$ for $\mathbf{x} \in \mathbb{R}^{|S|}$ and $\epsilon \in (0, 1)$, then $\text{diff}^{\text{rel}}(\frac{\mathbf{x}}{\sum_{s \in S} \mathbf{x}(s)}, \theta^{\mathcal{B}}) \leq \frac{2\epsilon}{1-\epsilon}$.*

Combining both steps. Theorems 6 and 7 and the interval iteration method yield approximations p_B and d_s for $\Pr^{\mathcal{D}}(\Diamond B)$ and $\theta^{\mathcal{D}|_B}(s)$, respectively, where B is a BSCC $s \in B$, and $\epsilon_1, \epsilon_2 \in (0, 1)$ such that $\text{diff}^{\text{rel}}(p_B, \Pr^{\mathcal{D}}(\Diamond B)) \leq \epsilon_1$ and $\text{diff}^{\text{rel}}(d_s, \theta^{\mathcal{D}|_B}(s)) \leq \epsilon_2$. The product $p_B \cdot d_s$ approximates $\theta^{\mathcal{D}}(s) = \Pr^{\mathcal{D}}(\Diamond B) \cdot \theta^{\mathcal{D}|_B}(s)$ such that $\text{diff}^{\text{rel}}(p_B \cdot d_s, \theta^{\mathcal{D}}(s)) \leq \epsilon_1 + \epsilon_2 + \epsilon_1 \epsilon_2$.

Example 6. We compute the stationary distribution of the running example DTMC \mathcal{D} from Figure 2. Its only non-trivial BSCC is $\mathcal{B} = \mathcal{D}|_{\{s_5, s_6\}}$. The DTMC $\mathcal{B}^{\hat{r}^{\hat{s}_6}}$ along with its (exact) EVTs is depicted in Figure 4. We conclude that $\theta^{\mathcal{B}}$ is proportional to $(\frac{3}{5}, 1)$, i.e., $\theta^{\mathcal{B}} = (\frac{5}{8}, \frac{3}{8})$. Since the two BSCCs $\{s_5, s_6\}$ and $\{s_7\}$ are both reached with probability $\frac{1}{2}$, it follows that the stationary probabilities of the three recurrent states s_5, s_6, s_7 are $\frac{5}{16}, \frac{3}{16}$, and $\frac{1}{2}$, respectively.




Fig. 4: The DTMC $\mathcal{B}^{\hat{r}^{\hat{s}_6}}$.

6 Conditional Expected Rewards

We have seen in Theorem 6 that the EVTs of the transient states of a DTMC \mathcal{D} can be used to compute the *probability* to reach each individual BSCC of \mathcal{D} . We now generalize this result and show that the EVTs can also be used to compute the *total expected rewards conditioned on reaching each BSCC*.

The total expected reward conditioned on reaching a set $T \subseteq S$ of states with $\Pr^{\mathcal{D}}(\Diamond T) > 0$ is defined as:

$$\mathbb{E}^{\mathcal{D}}[\text{tr}_{\text{rew}} | \Diamond T] = \frac{1}{\Pr^{\mathcal{D}}(\Diamond T)} \int_{\Diamond T} \text{tr}_{\text{rew}} d\Pr^{\mathcal{D}}$$

Our next result asserts that given the EVTs of \mathcal{D} , all the values $\{\mathbb{E}^{\mathcal{D}}[\text{tr}_{\text{rew}} | \Diamond B] \mid B \text{ a BSCC of } \mathcal{D}\}$ can be computed by solving a *single* linear equation system (the standard approach is to solve one linear equation system *per BSCC* [6, Ch. 10]). For simplicity, we state the result only for BSCCs $\{r\}$ with a single absorbing state r , and for reward functions that assign zero reward to all recurrent (BSCC) states; this is w.l.o.g. as larger BSCCs can be collapsed, and positive reward in a (reachable) BSCC causes the conditional expected reward w.r.t. this BSCC to be ∞ , rendering numeric computations unnecessary.

Theorem 8. *Let $\text{rew} : S \rightarrow \mathbb{Q}_{\geq 0}$ with $\text{rew}(S_{\text{re}}) = \{0\}$. Then the equation system*

$$\forall s \in S_{\text{tr}} : \quad \mathbf{y}(s) = \text{rew}(s) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}_t] + \sum_{t \in S_{\text{tr}}} \mathbf{P}(t, s) \cdot \mathbf{y}(t)$$

has a unique solution $(\mathbf{y}(s))_{s \in S_{\text{tr}}}$ and for all absorbing $r \in S$, $\Pr^{\mathcal{D}}(\Diamond \{r\}) > 0$,

$$\mathbb{E}^{\mathcal{D}}[\text{tr}_{\text{rew}} | \Diamond \{r\}] = \frac{\sum_{t \in S_{\text{tr}}} \mathbf{P}(t, r) \cdot \mathbf{y}(t)}{\iota_{\text{init}}(r) + \sum_{t \in S_{\text{tr}}} \mathbf{P}(t, r) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}_t]}.$$

Theorem 8 assumes rational rewards as required in our proof in Appendix A.4.

7 Experimental Evaluation

Implementation details. We integrated the presented algorithms for EVTs and stationary distributions in the model checker **Storm** [33]. The implementation uses explicit data structures such as sparse matrices and vectors. When computing EVTs, we can use value iteration (VI) or interval iteration (II) as presented in Algorithms 1 and 2. Alternatively, the corresponding linear equation systems can be solved using LU factorization (a direct method implemented in the **Eigen** library [26]) or **gmres** (a numerical method implemented in **gmm++** [52]). Each EVT approach can be used in combination with the topological algorithm (**topo**) from Section 4.3. We use double precision floating point numbers. For II, the propagation of (relative) errors is respected in a way that the error of the end result does not exceed a user-defined threshold (here: $\epsilon = 10^{-3}$). Implementing II with safe rounding modes as in [29] is left for future work. The methods **gmres** and VI are configured with a fixed relative precision parameter (here: $\epsilon = 10^{-6}$). For LU, floating point errors are the only source of inaccuracies. We also consider an exact configuration LU^X that uses rational arithmetic instead of floats.

Stationary distributions can be computed in **Storm** using the approaches **Classic**, **EVTreach**, and **EVTfull**. The **Classic** approach computes each BSCC reachability probability separately and the stationary distributions within the BSCCs are computed using the standard equation system [41, Thm. 4.18]. **EVTreach** and **EVTfull** implement our approaches from Section 5, where **EVTreach** only considers EVTs for BSCC reachability and **EVTfull** also derives the BSCC distributions from EVTs. As for EVTs, we use $\text{LU}^{(X)}$ or **gmres** to solve linear equation systems. For the BSCC reachability probabilities, a topological algorithm can be enabled as well. Using **EVTfull** with **II** yields sound approximations.

Experimental setup. The experiments ran on an Intel® Xeon® Platinum 8160 Processor limited to 4 cores and 12 GB of memory. The time timeout was set to 30 minutes. Our implementation does not use multi-threading.

7.1 Verifying the Fast Dice Roller

Recall Lumbroso’s Fast Dice Roller [44] from Section 1. For a given parameter $N \geq 1$, we verify that the obtained distribution is indeed uniform by computing the stationary distribution of the corresponding DTMC which, for this model family, coincides with the individual BSCC reachability probabilities as each BSCC consists of a single state. The plot in Figure 5 shows for different values of N (x-axis) the runtime (y-axis) of the approaches **Classic** (using **gmres** or LU^x) and **EVTfull** (using **II** or LU^x), all using topological algorithms. Our novel **EVTfull** approach is significantly faster than the **Classic** method, enabling us to verify large instances with up to 4 800 255 states ($N = 32\,000$) within the time limit. In particular, we can compute the values using exact arithmetic as LU^x has runtimes similar to those of **II** in the **EVTfull** approach.




Fig. 5: Fast Dice Roller Results.

7.2 Performance Comparison

To evaluate the various approaches, we computed EVTs and stationary distributions for all applicable finite models of type DTMC or CTMC of the Quantitative Verification Benchmark Set (QVBS) [32]: We excluded 2 model families for which none of the tested parameter valuations allowed any algorithm to complete using the available resources and for the EVT computation we excluded 8 models that do not contain any transient states. In addition to QVBS, we included Lumbroso’s Fast Dice Roller [44] and the handcrafted models (**branch** and **loop**) introduced in [47]. We considered multiple parameter valuations yielding a total of 62 instances (including 12 CTMCs) for computing EVTs and 79 instances (including 28 CTMCs) for computing the stationary distribution.

Our experiments for stationary distributions also include the implementations of the naive and guided sampling approaches (**ap-naive** and **ap-sample**) from [47]

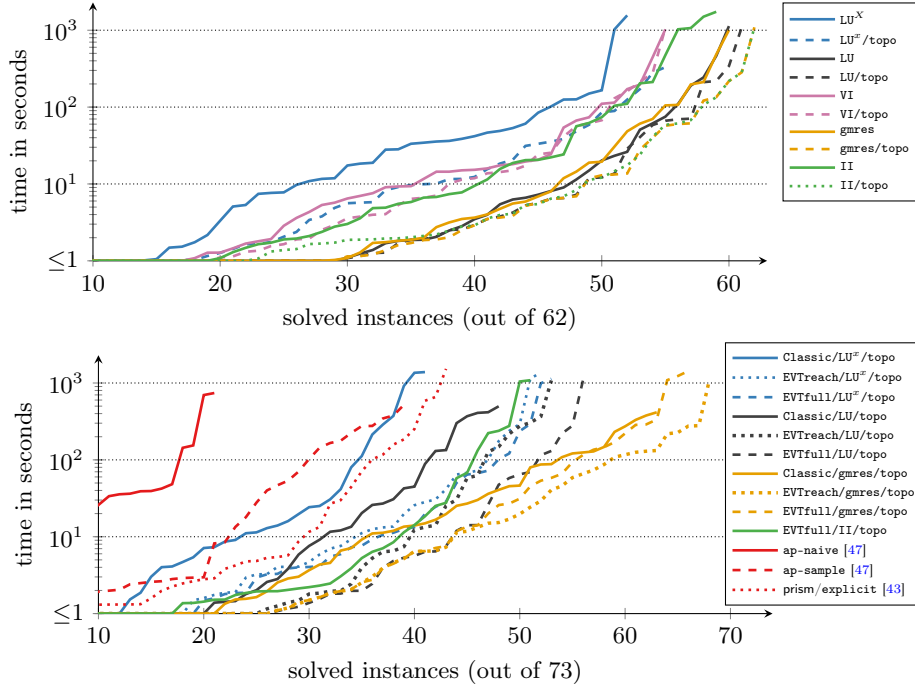


Fig. 6: Runtime comparison for EVTs (top) and stationary distributions (bottom).

as well as the implementation of the `classic` approach in `prism` [43]² as external baselines. These tools do not support the Jani models and can not compute EVTs.

We measured the runtime of the respective computation including the time for model construction. In cases where the exact results are known (using exact computations via LU^X), we consider results as *incorrect* if the relative difference to the exact value is greater than 10^{-3} . Results provided by the tool from [47] and by `prism` are not checked for correctness. We set the relative termination threshold of `gmres` and `VI` to $\epsilon = 10^{-6}$ to compensate for inaccuracies of the unsound methods. When using `II`, the implementation of [47], or `prism`, a relative precision of $\epsilon = 10^{-3}$ was requested. More experiments can be found in Appendix E.

Computing EVTs. The quantile plot at the top of Figure 6 indicates the time that is required for computing the EVTs in 62 models for different approaches. A point at position (x, y) indicates that the corresponding method solved the x^{th} fastest instance in y seconds, where only correctly solved instances are considered. The unsound methods `VI` and `gmres` produced 7 and 2 incorrect results, respectively. Furthermore, as errors accumulated, the topological variations of `VI` and `gmres` more frequently exceeded the threshold of 10^{-3} . The variants of `II` always produced correct results.

² We consider the `explicit` engine of `prism` v4.8 with the `Jacobi` method (default).




Fig. 7: Scatter plots showing the EVT computation runtime for standard and topological II (left) as well as the runtime of the EVT (right) approaches for different state space sizes $|S|$.

The plot indicates that (topological) LU and **gmres** outperform VI and II for easier instances. However, II catches up for the more intricate instances as it is more scalable than LU and always yields correct results. The exact method LU^X is significantly slower compared to the other methods. We also observe that the topological algorithms are superior to the non-topological variants. This is confirmed by the leftmost scatter plot in Figure 7 which compares the topological and non-topological variant of II. Here, each point (x, y) indicates an instance for which the methods on the x -axis and the y -axis respectively required x and y seconds to compute the EVTs. Instances that contain only singleton SCCs are depicted as triangles (\blacktriangle), whereas cycles (\bullet) represent the remaining instances. No incorrect results (INC) were obtained. The scatter plot in the middle in Figure 7 indicates that the iterative methods are more scalable than exact LU^X . We also see that models with millions of states can be solved in reasonable time.

Computing stationary distributions. The quantile plot at the bottom of Figure 6 summarizes the runtimes for the different stationary distribution approaches³. We only consider the topological variants of the approaches of **Storm** as they were consistently faster. No incorrect results were observed in this experiment.

The plot indicates that the guided sampling method (**ap-sample**) from [47] and **prism** perform significantly better than **ap-naive**. However, all algorithms provided by **Storm** — except for **Classic**/ LU^x — outperform the other implementations. For LU and **GMRES**, we observe that the **EVTreach** and **EVTfull** variants are significantly faster than the **Classic** approach. The **EVTreach** approach using **gmres** provides the fastest configuration, but is also the least reliable one in terms of accuracy. For the sound methods, we observe that the **EVTfull** approach with II is outperformed by LU combined with either **EVTreach** or **EVTfull**, where the latter shows the better performance.

³ Six Jani [11] models that **ap-naive**, **ap-sample**, and **prism** do not support are omitted.

8 Related Work

Computing stationary distributions. Other methods for the sound computation of the stationary distribution have been proposed in, e.g., [22,12,10]. In contrast to our work, they consider only subclasses of Markov chains: [22,12] introduce an (iterative) algorithm applicable to Markov chains with positive rows while [10] presents a technique limited to time-reversible Markov chains. The recent approach from [47] can also handle general Markov chains. Our technique ensures soundness with respect to both absolute and relative differences, whereas the approaches of [47] only consider absolute precision.

Other applications of EVTs. The authors of [38] have suggested to use EVTs for the *expected time to absorption*, i.e., the expected number of steps until the chain reaches an absorbing state. Indeed, this quantity is given by the sum of the vector $(\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}}$ [38, Thm. 3.3.5]. For acyclic DTMCs the EVT of a state coincides with the probability of reaching this state. This is relevant in the context of Bayesian networks [54,55] since inference queries in the network can be reduced to reachability queries by translating Bayesian networks into tree-like DTMCs. Existing procedures for multi-objective model checking of MDPs employ linear programming methods relying on the EVTs of state-action pairs [21,19,14,18]. EVTs are also employed in an algorithm proposed in [9] for LTL model checking of *interval Markov chains*. Moreover, EVTs have been leveraged for minimizing and learning DTMCs [1,2]. Further recent applications of EVTs to MDPs include verifying *cause-effect dependencies* [3], as well as an abstraction-refinement procedure that measures the importance of states based on the EVTs under a fixed policy [35]. [23] employs EVTs in the context of policy iteration in reward-robust MDPs.

9 Conclusion

We elaborated on the computation of EVTs in DTMCs and CTMCs: The EVTs in DTMCs can be determined by solving a linear equation system, while computing EVTs in CTMCs reduces to the discrete-time setting. We developed an iterative algorithm based on the value iteration [50] algorithm lacking assurance of precision. Building on interval iteration [28,7] — an algorithm for the sound computation of reachability probabilities and expected rewards — we developed an algorithm for approximating EVTs with accuracy guarantees. To enhance efficiency, we adapted a topological algorithm [16,17,15] to compute EVTs SCC-wise in topological order. We showed that EVTs enable the sound approximation of the stationary distribution and the efficient computation of conditional expected rewards. For future work, we want to extend our implementation provided in the model checker Storm [33] by symbolic computations. We also want to leverage EVTs as a heuristic to identify significant sub-regions within the state space of DTMCs, drawing inspiration from frameworks for approximate MDP verification [40].

References

1. Bacci, G., Bacci, G., Larsen, K.G., Mardare, R.: On the metric-based approximate minimization of Markov chains. In: ICALP. LIPIcs, vol. 80, pp. 104:1–104:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2017)
2. Bacci, G., Ingólfssdóttir, A., Larsen, K.G., Reynouard, R.: Active learning of Markov decision processes using Baum-Welch algorithm. In: ICMLA. pp. 1203–1208. IEEE (2021)
3. Baier, C., Funke, F., Piribauer, J., Ziemek, R.: On probability-raising causality in Markov decision processes. In: FoSSaCS. Lecture Notes in Computer Science, vol. 13242, pp. 40–60. Springer (2022)
4. Baier, C., Haverkort, B.R., Hermanns, H., Katoen, J.: Model-checking algorithms for continuous-time Markov chains. *IEEE Trans. Software Eng.* **29**(6), 524–541 (2003)
5. Baier, C., Hensel, C., Hutschenreiter, L., Junges, S., Katoen, J., Klein, J.: Parametric Markov chains: PCTL complexity and fraction-free gaussian elimination. *Inf. Comput.* **272**, 104504 (2020)
6. Baier, C., Katoen, J.: Principles of Model Checking. MIT Press (2008)
7. Baier, C., Klein, J., Leuschner, L., Parker, D., Wunderlich, S.: Ensuring the reliability of your model checker: Interval iteration for Markov decision processes. In: CAV (1). Lecture Notes in Computer Science, vol. 10426, pp. 160–180. Springer (2017)
8. Bellman, R.: A Markovian Decision Process. *Journal of Mathematics and Mechanics* **6**(5), 679–684 (1957)
9. Benedikt, M., Lenhardt, R., Worrell, J.: LTL model checking of interval Markov chains. In: TACAS. Lecture Notes in Computer Science, vol. 7795, pp. 32–46. Springer (2013)
10. Bressan, M., Peserico, E., Pretto, L.: On approximating the stationary distribution of time-reversible markov chains. *Theory Comput. Syst.* **64**(3), 444–466 (2020)
11. Budde, C.E., Dehnert, C., Hahn, E.M., Hartmanns, A., Junges, S., Turrini, A.: JANI: quantitative model and tool interaction. In: TACAS (2). Lecture Notes in Computer Science, vol. 10206, pp. 151–168 (2017)
12. Busic, A., Fourneau, J.: Iterative component-wise bounds for the steady-state distribution of a Markov chain. *Numer. Linear Algebra Appl.* **18**(6), 1031–1049 (2011)
13. Cardelli, L., Kwiatkowska, M., Laurenti, L.: Programming discrete distributions with chemical reaction networks. *Nat. Comput.* **17**(1), 131–145 (2018)
14. Chatterjee, K., Majumdar, R., Henzinger, T.A.: Markov decision processes with multiple objectives. In: STACS. Lecture Notes in Computer Science, vol. 3884, pp. 325–336. Springer (2006)
15. Ciesinski, F., Baier, C., Grötker, M., Klein, J.: Reduction techniques for model checking Markov decision processes. In: QEST. pp. 45–54. IEEE Computer Society (2008)
16. Dai, P., Goldsmith, J.: Topological value iteration algorithm for Markov decision processes. In: IJCAI. pp. 1860–1865 (2007)
17. Dai, P., Mausam, Weld, D.S., Goldsmith, J.: Topological value iteration algorithms. *J. Artif. Intell. Res.* **42**, 181–209 (2011)
18. Delgrange, F., Katoen, J., Quatmann, T., Randour, M.: Simple strategies in multi-objective MDPs. In: TACAS (1). Lecture Notes in Computer Science, vol. 12078, pp. 346–364. Springer (2020)

19. Etessami, K., Kwiatkowska, M.Z., Vardi, M.Y., Yannakakis, M.: Multi-objective model checking of Markov decision processes. *Log. Methods Comput. Sci.* **4**(4) (2008)
20. Fiondella, L., Rajasekaran, S., Gokhale, S.S.: Efficient software reliability analysis with correlated component failures. *IEEE Trans. Reliab.* **62**(1), 244–255 (2013)
21. Forejt, V., Kwiatkowska, M.Z., Norman, G., Parker, D., Qu, H.: Quantitative multi-objective verification for probabilistic systems. In: TACAS. *Lecture Notes in Computer Science*, vol. 6605, pp. 112–127. Springer (2011)
22. Fourneau, J., Quessette, F.: Some improvements for the computation of the steady-state distribution of a Markov chain by monotone sequences of vectors. In: ASMTA. *Lecture Notes in Computer Science*, vol. 7314, pp. 178–192. Springer (2012)
23. Gadot, U., Derman, E., Kumar, N., Elfatih, M.M., Levy, K., Mannor, S.: Solving non-rectangular reward-robust MDPs via frequency regularization. *CoRR abs/2309.01107* (2023)
24. Gokhale, S.S., Trivedi, K.S.: Reliability prediction and sensitivity analysis based on software architecture. In: ISSRE. pp. 64–78. IEEE Computer Society (2002)
25. Gokhale, S.S., Wong, W.E., Horgan, J.R., Trivedi, K.S.: An analytical approach to architecture-based software performance and reliability prediction. *Perform. Evaluation* **58**(4), 391–412 (2004)
26. Guennebaud, G., Jacob, B.: Eigen v3. <http://eigen.tuxfamily.org> (2010), [Accessed 10-Nov-2022]
27. Haddad, S., Monmege, B.: Interval iteration algorithm for MDPs and IMDPs. *Theor. Comput. Sci.* **735**, 111–131 (2018)
28. Haddad, S., Monmege, B.: Interval iteration algorithm for MDPs and IMDPs. *Theor. Comput. Sci.* **735**, 111–131 (2018)
29. Hartmanns, A.: Correct probabilistic model checking with floating-point arithmetic. In: TACAS (2). *Lecture Notes in Computer Science*, vol. 13244, pp. 41–59. Springer (2022)
30. Hartmanns, A., Hermanns, H.: The modest toolset: An integrated environment for quantitative modelling and verification. In: TACAS. *Lecture Notes in Computer Science*, vol. 8413, pp. 593–598. Springer (2014)
31. Hartmanns, A., Kaminski, B.L.: Optimistic value iteration. In: CAV (2). *Lecture Notes in Computer Science*, vol. 12225, pp. 488–511. Springer (2020)
32. Hartmanns, A., Klauck, M., Parker, D., Quatmann, T., Ruijters, E.: The quantitative verification benchmark set. In: TACAS (1). *Lecture Notes in Computer Science*, vol. 11427, pp. 344–350. Springer (2019)
33. Hensel, C., Junges, S., Katoen, J., Quatmann, T., Volk, M.: The probabilistic model checker Storm. *Int. J. Softw. Tools Technol. Transf.* **24**(4), 589–610 (2022)
34. Horn, R.A., Johnson, C.R.: *Matrix Analysis*, 2nd Ed. Cambridge University Press (2012)
35. Junges, S., Spaan, M.T.J.: Abstraction-refinement for hierarchical probabilistic models. In: CAV (1). *Lecture Notes in Computer Science*, vol. 13371, pp. 102–123. Springer (2022)
36. Katoen, J.: The probabilistic model checking landscape. In: LICS. pp. 31–45. ACM (2016)
37. Keane, M.S., O’Brien, G.L.: A bernoulli factory. *ACM Trans. Model. Comput. Simul.* **4**(2), 213–219 (1994)
38. Kemeny, J., Snell, J.: *Finite Markov Chains*. Undergraduate texts in mathematics, Springer (1976)
39. Kemeny, J.G., Snell, J.L.: *Finite continuous time Markov chains. Theory of Probability & Its Applications* **6**(1), 101–105 (1961)

40. Křetínský, J., Meggendorfer, T.: Of cores: A partial-exploration framework for Markov decision processes. *Log. Methods Comput. Sci.* **16**(4) (2020)
41. Kulkarni, V.: *Modeling and Analysis of Stochastic Systems*. Chapman & Hall/CRC Texts in Statistical Science, CRC Press (2020), <https://books.google.de/books?id=D2gCzgEACAAJ>
42. Kwiatkowska, M.Z., Norman, G., Parker, D.: Stochastic model checking. In: *SFM. Lecture Notes in Computer Science*, vol. 4486, pp. 220–270. Springer (2007)
43. Kwiatkowska, M.Z., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: *CAV. Lecture Notes in Computer Science*, vol. 6806, pp. 585–591. Springer (2011)
44. Lumbroso, J.O.: Optimal discrete uniform generation from coin flips, and applications. *CoRR* **abs/1304.1916** (2013)
45. McMahan, H.B., Likhachev, M., Gordon, G.J.: Bounded real-time dynamic programming: RTDP with monotone upper bounds and performance guarantees. In: *ICML. ACM International Conference Proceeding Series*, vol. 119, pp. 569–576. ACM (2005)
46. Meedeniya, I., Moser, I., Aleti, A., Grunske, L.: Architecture-based reliability evaluation under uncertainty. In: *QoSA/ISARCS*. pp. 85–94. ACM (2011)
47. Meggendorfer, T.: Correct approximation of stationary distributions. In: *TACAS (1). Lecture Notes in Computer Science*, vol. 13993, pp. 489–507. Springer (2023)
48. Neumann, C.: *Untersuchungen über das logarithmische und Newton’sche Potential*. BG Teubner (1877)
49. Pietrantuono, R., Russo, S., Trivedi, K.S.: Online monitoring of software system reliability. In: *EDCC*. pp. 209–218. IEEE Computer Society (2010)
50. Puterman, M.L.: *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics, Wiley (1994)
51. Quatmann, T., Katoen, J.: Sound value iteration. In: *CAV (1). Lecture Notes in Computer Science*, vol. 10981, pp. 643–661. Springer (2018)
52. Renard, Y.: Gmm++. <https://getfem.org/gmm.html> (2004), [Accessed 10-Nov-2022]
53. Saad, Y.: *Iterative methods for sparse linear systems*. SIAM (2003)
54. Salmani, B., Katoen, J.: Bayesian inference by symbolic model checking. In: *QEST. Lecture Notes in Computer Science*, vol. 12289, pp. 115–133. Springer (2020)
55. Salmani, B., Katoen, J.: Fine-tuning the odds in Bayesian networks. In: *ECSQARU. Lecture Notes in Computer Science*, vol. 12897, pp. 268–283. Springer (2021)
56. Sharma, V.S., Trivedi, K.S.: Reliability and performance of component based software systems with restarts, retries, reboots and repairs. In: *ISSRE*. pp. 299–310. IEEE Computer Society (2006)
57. Varga, R.S.: *Matrix Iterative Analysis*. Springer Series in Computational Mathematics, Springer Berlin Heidelberg (1999)
58. Wimmer, R., Kortus, A., Herbst, M., Becker, B.: Probabilistic model checking and reliability of results. In: *DDECS*. pp. 207–212. IEEE Computer Society (2008)

A Omitted Proofs

A.1 Proofs of Section 3

Lemma 1. *For a fixed $s \in S$ and $x \in \mathbb{R}_{>0}$ the reward structure $\text{rew}: S \rightarrow \mathbb{R}_{\geq 0}$ given by $\text{rew}(t) = x \cdot \llbracket t = s \rrbracket$ satisfies $\mathbb{E}[\text{vt}_s] = \frac{1}{x} \cdot \mathbb{E}[\text{tr}_{\text{rew}}]$.*

Proof. Let $s \in S$ be a state of \mathcal{D} and let rew be a reward structure $\text{rew}: S \rightarrow \mathbb{R}_{\geq 0}$ with $\text{rew}(t) = x \cdot \llbracket t = s \rrbracket$ for some $x \in \mathbb{R}_{>0}$. We have $\text{vt}_s(\pi) = \frac{1}{x} \cdot \text{tr}_{\text{rew}}(\pi)$ holds for all paths $\pi \in \text{Paths}^{\mathcal{D}}$. Thus, we obtain that

$$\begin{aligned} \mathbb{E}[\text{vt}_s] &= \int_{\text{Paths}^{\mathcal{D}}} \text{vt}_s \, d\text{Pr} && \text{by Definition 4} \\ &= \int_{\text{Paths}^{\mathcal{D}}} \frac{1}{x} \cdot \text{tr}_{\text{rew}} \, d\text{Pr} \\ &= \frac{1}{x} \cdot \mathbb{E}[\text{tr}_{\text{rew}}]. \end{aligned}$$

□

Theorem 2. *For all states $s \in S$, it holds that $\mathbb{E}^{\mathcal{C}}[\text{vt}_s] = \frac{1}{r(s)} \cdot \mathbb{E}^{\text{emb}(\mathcal{C})}[\text{vt}_s]$.*

Proof. Let $s \in S$ be an arbitrary but fixed state in the CTMC \mathcal{C} . Consider the reward structure rew with $\text{rew}(t) = \llbracket s = t \rrbracket$ for all $t \in S$. It follows that for all paths $\pi \in \text{Paths}^{\mathcal{C}}$, the visiting time $\text{vt}_s(\pi)$ is equal to the total reward $\text{tr}_{\text{rew}}(\pi)$ of π . Thus,

$$\begin{aligned} \mathbb{E}^{\mathcal{C}}[\text{vt}_s] &= \int_{\text{Paths}^{\mathcal{C}}} \text{vt}_s \, d\text{Pr}^{\mathcal{C}} \\ &= \int_{\text{Paths}^{\mathcal{C}}} \text{tr}_{\text{rew}} \, d\text{Pr}^{\mathcal{C}} \\ &= \mathbb{E}^{\mathcal{C}}[\text{tr}_{\text{rew}}]. \end{aligned}$$

Consider the reward structure for the DTMC $\text{emb}(\mathcal{C})$ with $\text{rew}^{\text{emb}(\mathcal{C})}(t) = \frac{\text{rew}(t)}{r(t)}$ for all $t \in S$, or equivalently

$$\text{rew}^{\text{emb}(\mathcal{C})}(t) = \frac{\llbracket s = t \rrbracket}{r(t)} = \frac{\llbracket s = t \rrbracket}{r(s)}.$$

It holds that $\mathbb{E}^{\mathcal{C}}[\text{tr}_{\text{rew}}] = \mathbb{E}^{\text{emb}(\mathcal{C})}[\text{tr}_{\text{rew}^{\text{emb}(\mathcal{C})}}]$ (see, e.g., [42, Page 31]). Consequently,

$$\begin{aligned} \mathbb{E}^{\mathcal{C}}[\text{vt}_s] &= \mathbb{E}^{\mathcal{C}}[\text{tr}_{\text{rew}}] \\ &= \mathbb{E}^{\text{emb}(\mathcal{C})}[\text{tr}_{\text{rew}^{\text{emb}(\mathcal{C})}}] \\ &= \frac{1}{r(s)} \cdot \mathbb{E}^{\text{emb}(\mathcal{C})}[\text{vt}_s] && \text{by Lemma 1.} \end{aligned}$$

□

A.2 Proofs of Section 4 (Accurately Computing EVTs)

Proofs of Section 4.1 (Value Iteration)

Theorem 3. *The EVT_s-operator from Definition 5 has the following properties:*

- (i) $(\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}$ is the unique fixed point of Φ .
- (ii) For all $\mathbf{x}^{(0)} \in \mathbb{R}^{|S_{\text{tr}}|}$ we have $\lim_{k \rightarrow \infty} \Phi^{(k)}(\mathbf{x}^{(0)}) = (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}$.

Proof. To prove claim (i), we write the linear equation system from Theorem 1 in matrix notation:

$$\tau_{\text{init}} + \mathbf{Q}^T \cdot \mathbf{x} = \mathbf{x} \quad \Leftrightarrow \quad \Phi(\mathbf{x}) = \mathbf{x}.$$

By Theorem 1, the vector $(\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}$ is the unique solution \mathbf{x} of the equation system above. Thus, the unique fixed point of Φ is $(\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}$.

To show claim (ii), we prove that Φ is a contraction mapping on a suitable normed vector space and apply Banach's fixed point theorem (see, e.g., [50, Theorem 6.2.3]). First, we describe the normed vector space. The matrix \mathbf{Q} satisfies $\lim_{k \rightarrow \infty} \mathbf{Q}^k = \mathbf{0}$ and thus the spectral radius $\rho(\mathbf{Q})$ is strictly smaller than 1 (see, e.g., [34, Theorem 5.6.12]). Thus, there is a matrix norm $\|\cdot\|_{\mathbf{Q}}$ satisfying $\|\mathbf{Q}\|_{\mathbf{Q}} = \rho(\mathbf{Q}) < 1$ (see, e.g., [34, Lemma 5.6.10]). It follows that there is a vector norm $\|\cdot\|_{\mathbf{Q}}$ on $\mathbb{R}^{|S_{\text{tr}}|}$ such that for all $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$ we have $\|\mathbf{Q} \cdot \mathbf{x}\|_{\mathbf{Q}} \leq \|\mathbf{Q}\|_{\mathbf{Q}} \cdot \|\mathbf{x}\|_{\mathbf{Q}}$ (see, e.g., [34, Theorem 5.7.13]). As Cauchy sequences are the only sequences that converge in a finite-dimensional normed vector space, we have that $(\mathbb{R}^{|S_{\text{tr}}|}, \|\cdot\|_{\mathbf{Q}})$ is a complete normed vector space (see, e.g., [34, Theorem 5.4.10]). The contraction property of Φ follows straightforwardly from the definition of $\|\cdot\|_{\mathbf{Q}}$: For any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^{|S_{\text{tr}}|}$ we have

$$\begin{aligned} \|\Phi(\mathbf{x}) - \Phi(\mathbf{y})\|_{\mathbf{Q}} &= \|\mathbf{Q}(\mathbf{x} - \mathbf{y})\|_{\mathbf{Q}} \\ &\leq \underbrace{\|\mathbf{Q}\|_{\mathbf{Q}}}_{<1} \cdot \|\mathbf{x} - \mathbf{y}\|_{\mathbf{Q}}. \end{aligned} \quad [34, \text{Theorem 5.7.13}]$$

We have shown that Φ is a contraction mapping on $(\mathbb{R}^{|S_{\text{tr}}|}, \|\cdot\|_{\mathbf{Q}})$ and thus Banach's fixed point theorem (see, e.g., [50, Theorem 6.2.3]) implies that $(\Phi^{(k)}(\mathbf{x}))_{k \in \mathbb{N}}$ converges to the unique fixed point $(\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}$ of Φ with respect to the norm $\|\cdot\|_{\mathbf{Q}}$. Lastly, $\lim_{k \rightarrow \infty} \Phi^{(k)}(\mathbf{x})$ holds with respect to any norm on $\mathbb{R}^{|S_{\text{tr}}|}$ (see, e.g., [34, Corollary 5.4.6.]), which concludes the proof. \square

Theorem 9. *For input DTMC \mathcal{D} , stopping criterion $\text{crit} \in \{\text{abs}, \text{rel}\}$, and any termination threshold $\epsilon > 0$, Algorithm 1 terminates.*

Proof. We first proof termination for the absolute termination criterion. The sequence $(\mathbf{x}^{(k)})_{k \in \mathbb{N}}$ converges to $(\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}$ in $\mathbb{R}^{|S_{\text{tr}}|}$, by Theorem 3. Thus, for $\epsilon' = \frac{\epsilon}{2} > 0$ there exists $N \in \mathbb{N}$ such that for all $k > N$

$$\max_{s \in S_{\text{tr}}} \left\{ \left| \mathbf{x}^{(k)}(s) - \mathbb{E}[\mathbf{vt}_s] \right| \right\} < \epsilon'.$$

As the inequality holds for all $k > N$ we conclude that

$$\max_{s \in S_{\text{tr}}} \left\{ \left| \mathbf{x}^{(k+1)}(s) - \mathbb{E}[\mathbf{vt}_s] \right| \right\} < \epsilon'.$$

By the triangular inequality and homogeneity of the max norm, we have that

$$\begin{aligned} \max_{s \in S_{\text{tr}}} \left\{ \left| \mathbf{x}^{(k)}(s) - \mathbf{x}^{(k+1)}(s) \right| \right\} &\leq \max_{s \in S_{\text{tr}}} \left\{ \left| \mathbf{x}^{(k)}(s) - \mathbb{E}[\mathbf{vt}_s] \right| \right\} \\ &\quad + \max_{s \in S_{\text{tr}}} \left\{ \left| \mathbf{x}^{(k+1)}(s) - \mathbb{E}[\mathbf{vt}_s] \right| \right\} \end{aligned}$$

$$\leq 2 \cdot \epsilon' = \epsilon.$$

Thus, the absolute stopping criterion is satisfied after a finite number of iterations.

Next, we show that the relative stopping criterion is also satisfied for all $\epsilon > 0$ within a finite number of iterations. Recall that S_{tr} only contains reachable states, i.e., $(\mathbb{E}[\mathbf{v}\mathbf{t}_s])_{s \in S_{\text{tr}}} > 0$. Let $s \in S_{\text{tr}}$. Since $(\mathbf{x}^{(k)}(s))_{k \in \mathbb{N}}$ converges to $\mathbb{E}[\mathbf{v}\mathbf{t}_s]$ in $\mathbb{R}^{|S_{\text{tr}}|}$, for $\epsilon' = \frac{1}{4} \cdot \mathbb{E}[\mathbf{v}\mathbf{t}_s] \cdot \epsilon > 0$ there exists $N \in \mathbb{N}$ such that for all $k > N$,

$$\left| \mathbf{x}^{(k)}(s) - \mathbb{E}[\mathbf{v}\mathbf{t}_s] \right| < \epsilon'.$$

Again, by triangular inequality and homogeneity, we obtain that for all $k > N$

$$\left| \mathbf{x}^{(k)}(s) - \mathbf{x}^{(k+1)}(s) \right| < 2 \cdot \epsilon' = \frac{1}{2} \cdot \mathbb{E}[\mathbf{v}\mathbf{t}_s] \cdot \epsilon,$$

or, equivalently

$$2 \cdot \left| \frac{\mathbf{x}^{(k)}(s) - \mathbf{x}^{(k+1)}(s)}{\mathbb{E}[\mathbf{v}\mathbf{t}_s]} \right| < \epsilon. \quad (1)$$

Next, let $\epsilon_s = \frac{1}{2} \cdot \mathbb{E}[\mathbf{v}\mathbf{t}_s] > 0$. By convergence of $(\mathbf{x}^{(k)}(s))_{k \in \mathbb{N}}$, there exists $N_s \in \mathbb{N}$ such that for all $k > N_s$,

$$\left| \mathbf{x}^{(k)}(s) - \mathbb{E}[\mathbf{v}\mathbf{t}_s] \right| < \epsilon_s = \frac{1}{2} \cdot \mathbb{E}[\mathbf{v}\mathbf{t}_s]. \quad (2)$$

Thus,

$$\begin{aligned} \mathbf{x}^{(k)}(s) &= \left| \mathbb{E}[\mathbf{v}\mathbf{t}_s] - (\mathbb{E}[\mathbf{v}\mathbf{t}_s] - \mathbf{x}^{(k)}(s)) \right| && \text{since } \mathbf{x}^{(k)}(s) \geq 0 \\ &\geq \left| \mathbb{E}[\mathbf{v}\mathbf{t}_s] \right| - \left| \mathbb{E}[\mathbf{v}\mathbf{t}_s] - \mathbf{x}^{(k)}(s) \right| \\ &> \mathbb{E}[\mathbf{v}\mathbf{t}_s] - \frac{1}{2} \cdot \mathbb{E}[\mathbf{v}\mathbf{t}_s] && \text{by Equation (2)} \\ &= \frac{1}{2} \cdot \mathbb{E}[\mathbf{v}\mathbf{t}_s]. \end{aligned}$$

Consequently, for all $k > \max\{N, (N_s)_{s \in S_{\text{tr}}}\}$ it holds that $\mathbf{x}^{(k+1)}(s) \geq \frac{1}{2} \cdot \mathbb{E}[\mathbf{v}\mathbf{t}_s] > 0$ and by using Equation (1) we obtain

$$\left| \frac{\mathbf{x}^{(k)}(s) - \mathbf{x}^{(k+1)}(s)}{\mathbf{x}^{(k+1)}(s)} \right| \leq 2 \cdot \left| \frac{\mathbf{x}^{(k)}(s) - \mathbf{x}^{(k+1)}(s)}{\mathbb{E}[\mathbf{v}\mathbf{t}_s]} \right| < \epsilon.$$

By considering all states in S_{tr} , there exists $N' \in \mathbb{N}$ such that for all $k > N'$

$$\max_{s \in S_{\text{tr}}} \left\{ \left| \frac{\mathbf{x}^{(k)}(s) - \mathbf{x}^{(k+1)}(s)}{\mathbf{x}^{(k+1)}(s)} \right| \right\} < \epsilon.$$

Thus, the relative stopping criterion is satisfied after a finite number of iterations. \square

Proofs of Section 4.2 (Interval Iteration)

Lemma 2. Let $\mathbf{u}, \mathbf{l} \in \mathbb{R}^{|S_{\text{tr}}|}$ with $\mathbf{l} \leq (\mathbb{E}[\mathbf{v}\mathbf{t}_s])_{s \in S_{\text{tr}}} \leq \mathbf{u}$. Then,

- (i) $\Phi_{\max}^{(k)}(\mathbf{l}) \leq \Phi_{\max}^{(k+1)}(\mathbf{l})$ and $\Phi_{\min}^{(k)}(\mathbf{u}) \geq \Phi_{\min}^{(k+1)}(\mathbf{u})$ for all $k \in \mathbb{N}$.
- (ii) $\Phi^{(k)}(\mathbf{l}) \leq \Phi_{\max}^{(k)}(\mathbf{l}) \leq (\mathbb{E}[\mathbf{v}\mathbf{t}_s])_{s \in S_{\text{tr}}} \leq \Phi_{\min}^{(k)}(\mathbf{u}) \leq \Phi^{(k)}(\mathbf{u})$ for all $k \in \mathbb{N}$.
- (iii) $\lim_{k \rightarrow \infty} \Phi_{\max}^{(k)}(\mathbf{l}) = \lim_{k \rightarrow \infty} \Phi_{\min}^{(k)}(\mathbf{u}) = (\mathbb{E}[\mathbf{v}\mathbf{t}_s])_{s \in S_{\text{tr}}}$.

Proof. The proof is analog to the proof of [7, Lemma 3.3]. Claim (i) follows directly from the definition of the operators Φ_{max} and Φ_{min} .

To show claim (ii), we first prove that $\Phi^{(k)}(\mathbf{l}) \leq \Phi_{max}^{(k)}(\mathbf{l})$ holds for all $k \in \mathbb{N}$, by induction over $k \in \mathbb{N}$. For $k = 0$, the claim holds trivially. For the induction step, assume $\Phi^{(k)}(\mathbf{l}) \leq \Phi_{max}^{(k)}(\mathbf{l})$ holds for an arbitrary but fixed $k > 0$. By the monotonicity of Φ follows that

$$\begin{aligned} \Phi(\Phi^{(k)}(\mathbf{l})) &\leq \Phi(\Phi_{max}^{(k)}(\mathbf{l})) && \text{using } \Phi^{(k)}(\mathbf{l}) \leq \Phi_{max}^{(k)}(\mathbf{l}) \\ &\leq \max \left\{ \Phi_{max}^{(k)}(\mathbf{l}), \Phi(\Phi_{max}^{(k)}(\mathbf{l})) \right\} \\ &= \Phi_{max}^{(k+1)}(\mathbf{l}). \end{aligned}$$

Secondly, we prove $\Phi_{max}^{(k)}(\mathbf{l}) \leq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}}$ by induction over $k \in \mathbb{N}$. For the induction base $\Phi_{max}^{(0)}(\mathbf{l}) = \mathbf{l} \leq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}}$ holds by assumption. For the induction step, assume $\Phi_{max}^{(k)}(\mathbf{l}) \leq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}}$. Then, by the monotonicity of Φ ,

$$\Phi(\Phi_{max}^{(k)}(\mathbf{l})) \leq \Phi((\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}}).$$

Since $(\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}}$ is the unique fixed point of Φ , we get

$$\Phi_{max}^{(k+1)}(\mathbf{l}) = \max \left\{ \Phi_{max}^{(k)}(\mathbf{l}), \Phi(\Phi_{max}^{(k)}(\mathbf{l})) \right\} \leq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}}.$$

The reasoning for $\Phi_{min}^{(k)}(\mathbf{u}) \leq \Phi^{(k)}(\mathbf{u})$ works analogously.

In order to prove (iii), exploit that $\Phi^{(k)}(\mathbf{l}) \leq \Phi_{max}^{(k)}(\mathbf{l}) \leq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}}$ holds for all $k \in \mathbb{N}$ by claim (ii). Additionally, $(\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}} = \lim_{k \rightarrow \infty} \Phi^{(k)}(\mathbf{l})$ holds by Theorem 3. Claim (iii) follows as $(\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}} = \lim_{k \rightarrow \infty} \Phi^{(k)}(\mathbf{l}) \leq \lim_{k \rightarrow \infty} \Phi_{max}^{(k)}(\mathbf{l}) \leq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}}$. Again, the reasoning for the sequence of upper bounds is analogous. \square

Theorem 4 (Correctness of Algorithm 2). *Given an input DTMC \mathcal{D} , initial vectors $\mathbf{l}^{(0)}, \mathbf{u}^{(0)}$ with $\mathbf{l}^{(0)} \leq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}} \leq \mathbf{u}^{(0)}$, $\text{crit} \in \{\text{abs}, \text{rel}\}$, and a threshold $\epsilon > 0$, Algorithm 2 terminates and returns a vector $\mathbf{x}^{res} \in \mathbb{R}^{|S_{tr}|}$ satisfying $\text{diff}^{\text{crit}}(\mathbf{x}^{res}, (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}}) \leq \epsilon$.*

Proof. First, we show that the algorithm terminates after a finite number of iterations for any $\epsilon > 0$ using the absolute and relative stopping criterion. From Lemma 2 (iii), it follows that the sequences $(\mathbf{l}^{(k)})_{k \in \mathbb{N}}$ and $(\mathbf{u}^{(k)})_{k \in \mathbb{N}}$ converge to $(\mathbb{E}[\mathbf{vt}_s])_{s \in S_{tr}}$. Thus, we have that for all $\epsilon > 0$ there exist $N_l, N_u \in \mathbb{N}$ such that for all $k_l > N_l$ and $k_u > N_u$, it holds that

$$\max_{s \in S_{tr}} \left\{ |\mathbf{l}^{(k)}(s) - \mathbb{E}[\mathbf{vt}_s]| \right\} < \epsilon \quad \text{and} \quad \max_{s \in S_{tr}} \left\{ |\mathbf{u}^{(k)}(s) - \mathbb{E}[\mathbf{vt}_s]| \right\} < \epsilon.$$

Thus, for all $k > \max\{N_l, N_u\}$ we have

$$\max_{s \in S_{tr}} \left\{ |\mathbf{l}^{(k)}(s) - \mathbf{u}^{(k)}(s)| \right\} < 2 \cdot \epsilon. \quad (3)$$

Using Equation (3), we can conclude that for all $k > \max\{N_l, N_u\}$ Thus the absolute stopping criterion is satisfied after finitely many iterations.

Second, we consider the relative stopping criterion for any $\epsilon > 0$. Let $s \in S_{tr}$. Since S_{tr} only contains reachable states, we have that $\mathbb{E}[\mathbf{vt}_s] > 0$. By a similar

reasoning as in the proof of Theorem 9, we obtain that there exists $N_s \in \mathbb{N}$ such that for all $k > N_s$

$$|\mathbf{l}^{(k)}(s)| > \frac{1}{2} \cdot \mathbb{E}[\mathbf{v}\mathbf{t}_s] > 0. \quad (4)$$

Analogously to the proof of Theorem 9, it follows that for all $s \in S_{\text{tr}}$ and $\epsilon > 0$, there exists $N \in \mathbb{N}$ such that for all $k > N$,

$$\left| \frac{\mathbf{l}^{(k)}(s) - \mathbf{u}^{(k)}(s)}{\mathbf{l}^{(k)}(s)} \right| < 2 \cdot \epsilon.$$

By considering all transient states, we obtain that there exists $N' \in \mathbb{N}_{>0}$ such for all $k > N'$,

$$\max_{s \in S_{\text{tr}}} \left\{ \left| \frac{\mathbf{l}^{(k)}(s) - \mathbf{u}^{(k)}(s)}{\mathbf{l}^{(k)}(s)} \right| \right\} < 2 \cdot \epsilon.$$

Thus, the relative stopping criterion is satisfied and the algorithm terminates after a finite number of iterations.

Next, we prove ϵ -soundness of Algorithm 2 with respect to the absolute and relative difference. The algorithm returns $\mathbf{x}^{res} = \frac{1}{2}(\mathbf{l} + \mathbf{u})$, with $\mathbf{l}, \mathbf{u} \in \mathbb{R}^{|S_{\text{tr}}|}$ such that $\text{diff}^{crit}(\mathbf{u}, \mathbf{l}) \leq 2 \cdot \epsilon$ for $crit \in \{abs, rel\}$. Let $\mathbf{x}^{true} = (\mathbb{E}[\mathbf{v}\mathbf{t}_s])_{s \in S_{\text{tr}}}$. First, we consider the absolute difference diff^{abs} . For all $s \in S_{\text{tr}}$, we have

$$\begin{aligned} |\mathbf{x}^{res}(s) - \mathbf{x}^{true}(s)| &= \left| \frac{1}{2} \cdot (\mathbf{l}(s) + \mathbf{u}(s)) - \mathbf{x}^{true}(s) \right| \\ &= \frac{1}{2} \cdot |\mathbf{l}(s) + \mathbf{u}(s) - 2 \cdot \mathbf{x}^{true}(s)| \\ &\leq \frac{1}{2} \cdot |\mathbf{l}(s) + \mathbf{u}(s) - 2 \cdot \mathbf{l}(s)| \quad \mathbf{0} \leq \mathbf{l}(s) \leq \mathbf{x}^{true}(s) \leq \mathbf{u}(s) \\ &= \frac{1}{2} \cdot \underbrace{|\mathbf{u}(s) - \mathbf{l}(s)|}_{\leq 2 \cdot \epsilon} \\ &\leq \epsilon \quad \text{termination criterion.} \end{aligned}$$

This implies that $\text{diff}^{abs}(\mathbf{x}^{res}, \mathbf{x}^{true}) = \max_{s \in S_{\text{tr}}} \{|\mathbf{x}^{res}(s) - \mathbf{x}^{true}(s)|\} \leq \epsilon$. Second, we prove that the returned result is ϵ -sound if the relative difference is used. Let $s \in S_{\text{tr}}$ be a transient state. The restriction of S_{tr} to reachable states guarantees that $\mathbf{x}^{true} = (\mathbb{E}[\mathbf{v}\mathbf{t}_s])_{s \in S_{\text{tr}}} > \mathbf{0}$. This yields $\mathbf{u}_s \neq 0$ by Lemma 2 (ii). Thus, due to the termination criterion, it holds that $\mathbf{l}(s) \neq 0$. Otherwise, we would have that $\text{diff}^{rel}(\mathbf{u}, \mathbf{l}) = \infty \not\leq \epsilon$. As a result, we obtain

$$\begin{aligned} \left| \frac{\mathbf{x}^{res}(s) - \mathbf{x}^{true}(s)}{\mathbf{x}^{true}(s)} \right| &= \left| \frac{\frac{1}{2} \cdot (\mathbf{l}(s) + \mathbf{u}(s)) - \mathbf{x}^{true}(s)}{\mathbf{x}^{true}(s)} \right| \\ &= \left| \frac{\mathbf{l}(s) + \mathbf{u}(s) - 2 \cdot \mathbf{x}^{true}(s)}{2 \cdot \mathbf{x}^{true}(s)} \right| \\ &\leq \left| \frac{\mathbf{l}(s) + \mathbf{u}(s) - 2 \cdot \mathbf{l}(s)}{2 \cdot \mathbf{l}(s)} \right| \quad \mathbf{0} < \mathbf{l}(s) \leq \mathbf{x}^{true}(s) \leq \mathbf{u}(s) \end{aligned}$$

$$= \frac{1}{2} \cdot \underbrace{\left| \frac{\mathbf{u}(s) - \mathbf{l}(s)}{\mathbf{l}(s)} \right|}_{\leq 2 \cdot \epsilon} \leq \epsilon \quad \text{termination criterion.}$$

Consequently,

$$\text{diff}^{rel}(\mathbf{x}^{res}, \mathbf{x}^{true}) \stackrel{\mathbf{x}^{true} > \mathbf{0}}{=} \max_{s \in S_{tr}} \left\{ \left| \frac{\mathbf{x}^{res}(s) - \mathbf{x}^{true}(s)}{\mathbf{x}^{true}(s)} \right| \right\} \leq \epsilon.$$

□

Proofs of Section 4.3 (Topological Algorithm) To show Lemma 3, we provide an auxiliary lemma which is a generalization of Theorem 6, originally presented by Kemeny and Snell in [38, Theorem 3.3.7]: The probability $\Pr((S \setminus C) \mathbf{U} \{s\})$ of reaching a transient state s in an SCC C , and only visiting states in $S \setminus C$ prior to reaching s , depends on the EVTs of the predecessors of s that lie outside the SCC C :

Lemma 5. *Let $s \in S_{tr}$ be a transient state in \mathcal{D} and let C be the unique SCC of s in \mathcal{D} . Then,*

$$\Pr((S \setminus C) \mathbf{U} \{s\}) = \iota_{\text{init}}(s) + \sum_{\substack{s' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(s', s) \cdot \mathbb{E}[\mathbf{vt}_{s'}].$$

Proof. The proof is similar to [38, Theorem 3.3.7], but lifted to transient states.

Consider a transient state $s \in S_{tr}$ of \mathcal{D} that is contained in SCC C . Let $S_? = \{s' \in C' \mid C' \in \text{SCC}^{\mathcal{D}} \wedge C' \hookrightarrow C\}$. Since each state in C can reach s , $S_?$ contains exactly those states that can reach s and lie outside C . Thus for all states $s' \in S \setminus (S_? \cup \{s\})$, we have $\Pr_{s'}((S \setminus C) \mathbf{U} \{s\}) = 0$.

If $S_? = \emptyset$, we conclude that $\mathbf{P}(s', s) = 0$ for all states $s' \in S \setminus C$. We directly obtain from [6, Theorem 10.15] that

$$\begin{aligned} \Pr((S \setminus C) \mathbf{U} \{s\}) &= \iota_{\text{init}}(s) \\ &= \iota_{\text{init}}(s) + \sum_{\substack{s' \in C' \\ C' \hookrightarrow C}} \underbrace{\mathbf{P}(t, s)}_{=0} \cdot \mathbb{E}[\mathbf{vt}_{s'}]. \end{aligned}$$

We now assume $S_? \neq \emptyset$. By definition of $S_?$, we have that a state in $S_?$ can only be reached by other states in $S_?$, i.e., $\Pr_t(\diamond \{s\}_?) = 0$ holds for all $t \notin S_?$. Consequently, it is sufficient to consider only states in $S_?$ for calculating the probability to reach a state in $S_?$. In particular, for the matrix $\mathbf{A} = (\mathbf{P}(s', t'))_{s', t' \in S_?}$ containing the transition probabilities within $S_?$, it holds that $\mathbf{A}^k(s', t')$ equals $\mathbf{P}^k(s', t')$, i.e., the probability to be in state t' at the k^{th} step from s' for all $s', t' \in S_?$ and $k \in \mathbb{N}$ by [6, Remark 10.22].

Since s is a transient state and all states s' in $S_?$ can reach s , we conclude that $S_?$ does not contain recurrent states. Thus, $\lim_{k \rightarrow \infty} \mathbf{A}^k = \mathbf{0}$ by [6, Theorem 10.27]. From [38, Theorem 3.2.4], it follows that $(\mathbf{I} - \mathbf{A})$ is invertible, where $(\mathbf{I} - \mathbf{A})^{-1} = (\mathbb{E}_{t'}[\mathbf{vt}_{s'}])_{s', t' \in S_?}$.

[6, Theorem 10.15] implies that the vector $(\Pr_{s'}((S \setminus C) \mathbf{U}\{s\}))_{s' \in S_\gamma}$ coincides with the vector $(\mathbf{I} - \mathbf{A})^{-1} \cdot (\mathbf{P}(t', s))_{t' \in S_\gamma}$. Thus, by taking the initial distribution into account, we obtain

$$\begin{aligned} \Pr((S \setminus C) \mathbf{U}\{s\}) &= \iota_{\text{init}}(s) + \sum_{s' \in S_\gamma} \iota_{\text{init}}(s') \cdot \Pr_{s'}((S \setminus C) \mathbf{U}\{s\}) \\ &= \iota_{\text{init}}(s) + ((\iota_{\text{init}}(s'))_{s' \in S_\gamma})^T \cdot (\mathbf{I} - \mathbf{A})^{-1} \cdot (\mathbf{P}(s', s))_{s' \in S_\gamma} \\ &= \iota_{\text{init}}(s) + ((\mathbf{I} - \mathbf{A}^T)^{-1} \cdot (\iota_{\text{init}}(s'))_{t' \in S_\gamma})^T \cdot (\mathbf{P}(s', s))_{s' \in S_\gamma} \\ &= \iota_{\text{init}}(s) + (\mathbb{E}[\mathbf{v}t_{s'}])_{s' \in S_\gamma} \cdot (\mathbf{P}(s', s))_{s' \in S_\gamma}. \end{aligned}$$

□

The following definition was omitted in the main text:

Definition 7 (SCC restriction). For $C \in \text{SCC}^\mathcal{D}$ and $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$, the parametric SCC restriction $\mathcal{D}|_C[\mathbf{x}]$ of \mathcal{D} to C is defined as $\mathcal{D}|_C[\mathbf{x}] = (S^{\mathcal{D}|_C}, \mathbf{P}^{\mathcal{D}|_C}, \iota_{\text{init}}^{\mathcal{D}|_C[\mathbf{x}]})$, where

- $S^{\mathcal{D}|_C} = C$ if C is a BSCC, and $S^{\mathcal{D}|_C} = C \uplus \{s_{\text{rec}}\}$ otherwise,
- $\mathbf{P}^{\mathcal{D}|_C} : S^{\mathcal{D}|_C} \times S^{\mathcal{D}|_C} \rightarrow [0, 1]$, where

$$\mathbf{P}^{\mathcal{D}|_C}(s, t) = \begin{cases} \mathbf{P}^\mathcal{D}(s, t) & \text{if } s, t \in C, \\ \sum_{s' \in S \setminus C} \mathbf{P}^\mathcal{D}(s, s'), & \text{if } s \in C \text{ and } t = s_{\text{rec}}, \\ 0 & \text{if } s = s_{\text{rec}} \text{ and } t \in C, \\ 1 & \text{if } s = t = s_{\text{rec}}, \end{cases}$$

- $\iota_{\text{init}}^{\mathcal{D}|_C[\mathbf{x}]} : S^{\mathcal{D}|_C} \rightarrow \mathbb{R}_{\geq 0}$, where

$$\iota_{\text{init}}^{\mathcal{D}|_C[\mathbf{x}]}(s) = \begin{cases} \iota_{\text{init}}^\mathcal{D}(s) + \sum_{s' \in S \setminus C} \mathbf{P}^\mathcal{D}(s', s) \cdot \mathbf{x}(s') & \text{if } s \in C, \\ 1 - \sum_{t \in C} \iota_{\text{init}}^{\mathcal{D}|_C[\mathbf{x}]}(t) & \text{if } s = s_{\text{rec}}. \end{cases}$$

Further, we define the DTMC $\mathcal{D}|_C = (S^{\mathcal{D}|_C}, \mathbf{P}^{\mathcal{D}|_C}, \iota_{\text{init}}^{\mathcal{D}|_C})$ as the parametric SCC restriction $\mathcal{D}|_C[\mathbf{x}]$ where \mathbf{x} is chosen such that $\iota_{\text{init}}^{\mathcal{D}|_C[\mathbf{x}]} = \iota_{\text{init}}^{\mathcal{D}|_C}$ with

$$\iota_{\text{init}}^{\mathcal{D}|_C}(s) = \begin{cases} \Pr^\mathcal{D}((S \setminus C) \mathbf{U}\{s\}) & \text{if } s \in C, \\ 1 - \sum_{t \in C} \Pr^\mathcal{D}((S \setminus C) \mathbf{U}\{t\}) & \text{if } s = s_{\text{rec}}. \end{cases}$$

Lemma 3. For non-bottom SCC C and $s \in C$ we have $\mathbb{E}^\mathcal{D}[\mathbf{v}t_s] = \mathbb{E}^{\mathcal{D}|_C}[\mathbf{v}t_s]$.

Proof. By Theorem 3 (i), the vector $(\mathbb{E}^\mathcal{D}[\mathbf{v}t_s])_{s \in S_{\text{tr}}}$ is the unique fixed point of the EVTs operator Φ . In particular, for all $s \in S_{\text{tr}}$,

$$\mathbb{E}^\mathcal{D}[\mathbf{v}t_s] = \iota_{\text{init}}^\mathcal{D}(s) + \sum_{t \in S_{\text{tr}}} \mathbf{P}^\mathcal{D}(t, s) \cdot \mathbb{E}^\mathcal{D}[\mathbf{v}t_t]. \quad (5)$$

Let C be a non-bottom SCC and let $\mathcal{D}|_C = (S^{\mathcal{D}|_C}, \mathbf{P}^{\mathcal{D}|_C}, \iota_{\text{init}}^{\mathcal{D}|_C})$ be the restriction of \mathcal{D} to C . By Equation (5) and Lemma 5, we obtain for all states $s \in C$,

$$\begin{aligned} \mathbb{E}^{\mathcal{D}}[\mathbf{v}_s] &= \underbrace{\iota_{\text{init}}^{\mathcal{D}}(s) + \sum_{t \in S \setminus C} \mathbf{P}^{\mathcal{D}}(t, s) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}_t]}_{\Pr^{\mathcal{D}}((S \setminus C) \mathbf{u}_{\{s\}})} + \sum_{t \in C} \mathbf{P}^{\mathcal{D}}(t, s) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}_t] \\ &= \iota_{\text{init}}^{\mathcal{D}|_C}(s) + \sum_{t \in C} \mathbf{P}^{\mathcal{D}}(t, s) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}_t] \quad \text{by Definition 7.} \end{aligned}$$

Thus, $(\mathbb{E}^{\mathcal{D}}[\mathbf{v}_s])_{s \in C}$ is a fixed point of the function $\Phi^{\mathcal{D}|_C} : \mathbb{R}^{|C|} \rightarrow \mathbb{R}^{|C|}$ with

$$\Phi^{\mathcal{D}|_C}(\mathbf{x}) = \left(\iota_{\text{init}}^{\mathcal{D}|_C}(s) + \sum_{t \in C} \mathbf{P}^{\mathcal{D}|_C}(t, s) \cdot \mathbf{x}(t) \right)_{s \in C}.$$

The function $\Phi^{\mathcal{D}|_C}$ coincides with the EVTs-operator tailored to $\mathcal{D}|_C$ and from Theorem 3, it follows that $(\mathbb{E}^{\mathcal{D}|_C}[\mathbf{v}_s])_{s \in C}$ is the *unique* fixed point of $\Phi^{\mathcal{D}|_C}$. Consequently, $\mathbb{E}^{\mathcal{D}}[\mathbf{v}_s] = \mathbb{E}^{\mathcal{D}|_C}[\mathbf{v}_s]$. \square

Lemma 4. *Let $\epsilon \in [0, 1)$ and let $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$ such that for every non-bottom SCC C , $(\mathbf{x}(s))_{s \in C}$ satisfies $\text{diff}^{\text{rel}}((\mathbf{x}(s))_{s \in C}, (\mathbb{E}^{\mathcal{D}|_C}[\mathbf{x}][\mathbf{v}_s])_{s \in C}) \leq \epsilon$. Then*

$$\text{diff}^{\text{rel}}(\mathbf{x}, (\mathbb{E}^{\mathcal{D}}[\mathbf{v}_s])_{s \in S_{\text{tr}}}) \leq (1 + \epsilon)^{L+1} - 1,$$

where $L = \max_{\kappa \in \text{Chains}_{\text{tr}}^{\mathcal{D}}|\kappa|}$ is the largest length of a chain of non-bottom SCCs.

Proof. Let $\epsilon \in (0, 1)$ and $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$ as specified in Lemma 4. First, we show that for all non-bottom SCCs C and $s \in C$ we have

$$\mathbf{x}(s) \leq \mathbb{E}^{\mathcal{D}}[\mathbf{v}_s] \cdot (1 + \epsilon)^{k+1},$$

where $k \in \mathbb{N}$ is the length of the longest SCC chain $C_0 \hookrightarrow \dots \hookrightarrow C_k \in \text{Chains}_{\text{tr}}^{\mathcal{D}}$ with $C_k = C$. We prove the claim by induction over the chain length $k \in \mathbb{N}$.

Consider an SCC C with $k = 0$. Analogously to the proof of Lemma 6, we obtain that $\mathcal{D}|_C[\mathbf{x}]$ is equivalent to $\mathcal{D}|_C$. Consequently, for every $s \in C$,

$$\begin{aligned} \mathbf{x}(s) &\leq \mathbb{E}^{\mathcal{D}|_C}[\mathbf{x}][\mathbf{v}_s] \cdot (1 + \epsilon) \\ &= \mathbb{E}^{\mathcal{D}}[\mathbf{v}_s] \cdot (1 + \epsilon). \end{aligned}$$

Next, let $k \in \mathbb{N}_{>0}$. Let C be a non-bottom SCC such that the longest SCC chain ending in C has length k . Consider a state $s \in C$. By the same reasoning as in Equation (12) in the proof of Lemma 6, we obtain

$$\mathbb{E}^{\mathcal{D}|_C}[\mathbf{x}][\mathbf{v}_s] = \sum_{t \in C} \sum_{i=0}^{\infty} \mathbf{P}^i(t, s) \cdot \left(\iota_{\text{init}}^{\mathcal{D}}(t) + \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \mathbf{x}(t') \right).$$

By the induction hypothesis, we have $\mathbf{x}(t') \leq \mathbb{E}^{\mathcal{D}}[\mathbf{v}_{t'}] \cdot (1 + \epsilon)^k$ for all states $t' \in C'$ in an SCC $C' \neq C$ with $C' \hookrightarrow C$. This yields

$$\mathbb{E}^{\mathcal{D}|_C}[\mathbf{x}][\mathbf{v}_s] \leq \sum_{t \in C} \sum_{i=0}^{\infty} \mathbf{P}^i(t, s) \cdot \left(\iota_{\text{init}}^{\mathcal{D}}(t) + \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}_{t'}] \cdot (1 + \epsilon)^k \right).$$

We have $(\mathbb{E}_t^{\mathcal{D}|C}[\mathbf{v}\mathbf{t}_s])_{t,s \in C} = (\mathbf{I} - (\mathbf{Q}^{\mathcal{D}|C})^T)^{-1}$ by [38, Theorem 3.2.4]. The inverse if $\mathbf{I} - (\mathbf{Q}^{\mathcal{D}|C})^T$ is given by the converging Neumann series $(\mathbf{I} - (\mathbf{Q}^{\mathcal{D}|C})^T)^{-1} = \sum_{i=0}^{\infty} (\mathbf{Q}^{\mathcal{D}|C})^i$ (see [48]). Thus, we obtain $\mathbb{E}_t^{\mathcal{D}|C}[\mathbf{v}\mathbf{t}_s] = \sum_{i=0}^{\infty} \mathbf{P}^i(t, s)$. Continuing the reasoning above and applying [7, Lemma 3.5.], results in

$$\begin{aligned} \mathbb{E}^{\mathcal{D}|C}[\mathbf{x}][\mathbf{v}\mathbf{t}_s] &\leq \sum_{t \in C} \mathbb{E}_t^{\mathcal{D}|C}[\mathbf{v}\mathbf{t}_s] \cdot \left(\iota_{\text{init}}^{\mathcal{D}}(t) + \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_{t'}] \cdot (1 + \epsilon)^k \right) \\ &= \sum_{t \in C} \frac{\Pr_t^{\mathcal{D}|C}(\diamond\{s\})}{1 - \Pr_s^{\mathcal{D}|C}(\bigcirc\diamond\{s\})} \\ &\quad \cdot \left(\iota_{\text{init}}^{\mathcal{D}}(t) + \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_{t'}] \cdot (1 + \epsilon)^k \right) \end{aligned}$$

Since ϵ is larger than 0, we infer $(1 + \epsilon)^k \geq 1$. Together with Lemma 5, we have

$$\mathbb{E}^{\mathcal{D}|C}[\mathbf{x}][\mathbf{v}\mathbf{t}_s] \leq \sum_{t \in C} \frac{\Pr_t^{\mathcal{D}|C}(\diamond\{s\})}{1 - \Pr_s^{\mathcal{D}|C}(\bigcirc\diamond\{s\})} \quad (6)$$

$$\cdot \left(\iota_{\text{init}}^{\mathcal{D}}(t) + \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_{t'}] \right) \cdot (1 + \epsilon)^k \quad (7)$$

$$= \sum_{t \in C} \frac{\Pr_t^{\mathcal{D}|C}(\diamond\{s\})}{1 - \Pr_s^{\mathcal{D}|C}(\bigcirc\diamond\{s\})} \cdot \Pr^{\mathcal{D}}\left((S \setminus C) \mathbf{U}\{t\}\right) \cdot (1 + \epsilon)^k$$

By applying the definition of the initial distribution $\iota_{\text{init}}^{\mathcal{D}|C}$ of the SCC restriction $\mathcal{D}|_C$ (see Definition 7), we obtain that

$$\begin{aligned} \mathbb{E}^{\mathcal{D}|C}[\mathbf{x}][\mathbf{v}\mathbf{t}_s] &\leq \sum_{t \in C} \frac{\Pr_t^{\mathcal{D}|C}(\diamond\{s\})}{1 - \Pr_s^{\mathcal{D}|C}(\bigcirc\diamond\{s\})} \cdot \iota_{\text{init}}^{\mathcal{D}|C}(t) \cdot (1 + \epsilon)^k \\ &= \frac{\Pr^{\mathcal{D}|C}(\diamond\{s\})}{1 - \Pr_s^{\mathcal{D}|C}(\bigcirc\diamond\{s\})} \cdot (1 + \epsilon)^k \quad \text{by [6, Theorem 10.15]} \\ &= \mathbb{E}^{\mathcal{D}|C}[\mathbf{v}\mathbf{t}_s] \cdot (1 + \epsilon)^k \quad \text{by [7, Lemma 3.5.]} \end{aligned}$$

Since the relative difference between $(\mathbf{x}(s))_{s \in C}$ and $(\mathbb{E}^{\mathcal{D}|C}[\mathbf{x}][\mathbf{v}\mathbf{t}_s])_{s \in C}$ is bounded by ϵ , we have

$$\begin{aligned} \mathbf{x}(s) &\leq \mathbb{E}^{\mathcal{D}|C}[\mathbf{x}][\mathbf{v}\mathbf{t}_s] \cdot (1 + \epsilon) \\ &\leq \mathbb{E}^{\mathcal{D}|C}[\mathbf{v}\mathbf{t}_s] \cdot (1 + \epsilon)^k \cdot (1 + \epsilon) \\ &\leq \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s] \cdot (1 + \epsilon)^{k+1} \quad \text{by Lemma 3,} \end{aligned}$$

which proves that the claim holds for all $k \in \mathbb{N}$.

Next, let $s \in S_{\text{tr}}$ and let C be the SCC containing s . L is the length of the longest non-bottom SCC chain in \mathcal{D} . In particular, this implies that the longest SCC chain ending in C has length at most L . Thus,

$$\mathbf{x}(s) \leq \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s] \cdot (1 + \epsilon)^{L+1} = \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s] + \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s] \cdot ((1 + \epsilon)^{L+1} - 1)$$

$$\begin{aligned}
&\Leftrightarrow \mathbf{x}(s) - \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s] \leq \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s] \cdot ((1 + \epsilon)^{L+1} - 1) \\
&\Leftrightarrow \frac{\mathbf{x}(s) - \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]}{\mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]} \leq (1 + \epsilon)^{L+1} - 1 \quad (\text{since } \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s] > 0.)
\end{aligned}$$

We have that $(1 + \epsilon)^{L+1} - 1 > 0$ because $\epsilon > 0$. Thus, if $x(s) > \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]$,

$$\left| \frac{\mathbf{x}(s) - \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]}{\mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]} \right| \leq (1 + \epsilon)^{L+1} - 1, \quad (8)$$

which yields an upper bound on the relative difference in case $x(s) > \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]$.

Next, consider $\mathbf{x}(s) < \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]$. We derive a lower bound of \mathbf{x} . The reasoning works analogously as above, except that we exploit the fact that $(1 - \epsilon) < 1$ holds by choice of $\epsilon < 1$ to establish the analog to Equation (6). As a result, for all states $s \in S_{\text{tr}}$,

$$\begin{aligned}
&\mathbf{x}(s) \geq \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s] \cdot (1 - \epsilon)^{L+1} = \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s] + \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s] \cdot ((1 - \epsilon)^{L+1} - 1) \\
&\Leftrightarrow \mathbf{x}(s) - \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s] \geq \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s] \cdot ((1 - \epsilon)^{L+1} - 1) \\
&\Leftrightarrow \frac{\mathbf{x}(s) - \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]}{\mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]} \geq (1 - \epsilon)^{L+1} - 1 \quad (\text{since } \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s] > 0.)
\end{aligned}$$

We have $(1 - \epsilon)^{L+1} - 1 < 0$ because $\epsilon < 1$. For $\mathbf{x}(s) < \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]$, we conclude

$$\begin{aligned}
\left| \frac{\mathbf{x}(s) - \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]}{\mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]} \right| &\leq -((1 - \epsilon)^{L+1} - 1) \\
&= 1 - (1 - \epsilon)^{L+1}.
\end{aligned} \quad (9)$$

By combining Equations (8) and (9), we infer that

$$\left| \frac{\mathbf{x}(s) - \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]}{\mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]} \right| \leq \begin{cases} (1 + \epsilon)^{L+1} - 1 & \text{if } \mathbf{x}(s) > \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s], \\ 1 - (1 - \epsilon)^{L+1} & \text{if } \mathbf{x}(s) < \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s], \\ 0 & \text{otherwise.} \end{cases}$$

This implies that

$$\begin{aligned}
\text{diff}^{\text{rel}}(\mathbf{x}, (\mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s])_{s \in S_{\text{tr}}}) &= \max_{s \in S_{\text{tr}}} \left\{ \left| \frac{\mathbf{x}(s) - \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]}{\mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_s]} \right| \right\} \\
&\leq \max \{ 1 - (1 - \epsilon)^{L+1}, (1 + \epsilon)^{L+1} - 1 \}.
\end{aligned}$$

Lastly, we show that $1 - (1 - \epsilon)^{L+1} \leq (1 + \epsilon)^{L+1} - 1$ holds for all $\epsilon \in (0, 1)$ and $L \in \mathbb{N}$ by induction on L .

For $L = 0$ we have

$$1 - (1 - \epsilon)^1 = \epsilon = (1 + \epsilon)^1 - 1.$$

By assuming that the claim holds for fixed $L \in \mathbb{N}$, we get

$$\begin{aligned}
1 - (1 - \epsilon)^{L+2} &= (1 - (1 - \epsilon)^{L+1}) \cdot (1 - \epsilon) + \epsilon \\
&\leq ((1 + \epsilon)^{L+1} - 1) \cdot (1 - \epsilon) + \epsilon \quad \text{by the induction hypothesis} \\
&\leq ((1 + \epsilon)^{L+1} - 1) \cdot (1 + \epsilon) + \epsilon \\
&= (1 + \epsilon)^{L+2} - (1 + \epsilon) + \epsilon \\
&= (1 + \epsilon)^{L+2} - 1,
\end{aligned}$$

which concludes the proof of Lemma 4. \square

Theorem 5 (Correctness of Algorithm 3). *Algorithm 3 returns a vector $\mathbf{x}^{res} \in \mathbb{R}^{|S_{tr}|}$ such that $\text{diff}^{rel}(\mathbf{x}^{res}, (\mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{tr}}) \leq \epsilon$.*

Proof. First, termination of the algorithm is obvious assuming that the procedure used to approximate the EVTs (Line 5) terminates.

Second, we show ϵ -soundness regarding the relative difference. We have $\sigma = \sqrt[L+1]{1+\epsilon} - 1$, where L is the length of the longest SCC chain in $\text{Chains}_{tr}\mathcal{D}$. By assumption, ϵ is in $(0, 1)$ which guarantees that $\sigma \in (0, 1)$. Consequently, we can apply Lemma 4, which yields

$$\text{diff}^{rel}(\mathbf{x}^{res}, (\mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{tr}}) \leq (1 + \sigma)^{L+1} - 1. \quad (10)$$

By substituting $\sigma = \sqrt[L+1]{1+\epsilon} - 1$, we obtain that

$$\begin{aligned} \text{diff}^{rel}(\mathbf{x}^{res}, (\mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{tr}}) &\leq (1 + \sigma)^{L+1} - 1 \\ &= (1 + \sqrt[L+1]{1+\epsilon} - 1)^{L+1} - 1 \\ &= \epsilon, \end{aligned}$$

which shows that the vector \mathbf{x}^{res} is an ϵ -sound approximation of $(\mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{tr}}$. \square

A.3 Proofs of Section 5 (Stationary Distributions)

We establish some basics concerning the propagation of relative errors for (non-negative) values. Claims made about the error propagations in Section 5 follow straightforwardly. Let $x \in \mathbb{R}_{\geq 0}$ be an approximation of $x^* \in \mathbb{R}_{\geq 0} \setminus \{0\}$ with relative precision $\epsilon \in (0, 1)$, i.e., $\text{diff}^{rel}(x, x^*) = |\delta_x| \leq \epsilon$, where $\delta_x := \frac{x-x^*}{x^*}$. It holds that $x = x^* \cdot (1 + \delta_x)$. Let y, y^* , and δ_y be similar.

We first argue that the sum $x + y$ is an approximation of $x^* + y^*$ with relative precision ϵ . Using $-\epsilon \leq \delta_x, \delta_y \leq \epsilon$ and $x + y = x^* \cdot (1 + \delta_x) + y^* \cdot (1 + \delta_y)$ we get

$$(x^* + y^*) \cdot (1 - \epsilon) \leq x + y \leq (x^* + y^*) \cdot (1 + \epsilon).$$

It follows that there is a $\delta^+ \in [-\epsilon, \epsilon]$ with $x + y = (x^* + y^*) \cdot (1 + \delta^+)$, i.e.,

$$\text{diff}^{rel}(x + y, x^* + y^*) = |\delta^+| \leq \epsilon.$$

Next, we consider multiplication. We have

$$x \cdot y = x^* \cdot (1 + \delta_x) \cdot y^* \cdot (1 + \delta_y) = x^* \cdot y^* \cdot (1 + \delta_x + \delta_y + \delta_x \cdot \delta_y).$$

With $\delta^\bullet := \delta_x + \delta_y + \delta_x \cdot \delta_y$ it follows that

$$\text{diff}^{rel}(x \cdot y, x^* \cdot y^*) = |\delta^\bullet| \leq 2\epsilon + \epsilon^2.$$

For division, we observe that

$$\frac{x}{y} = \frac{x^*}{y^*} \cdot \frac{1 + \delta_x}{1 + \delta_y} = \frac{x^*}{y^*} \cdot \frac{1 + \delta_x + \delta_y - \delta_y}{1 + \delta_y} = \frac{x^*}{y^*} \cdot \left(1 + \frac{\delta_x - \delta_y}{1 + \delta_y}\right).$$

With $\delta' := \frac{\delta_x - \delta_y}{1 + \delta_y}$ we get

$$\text{diff}^{rel}\left(\frac{x}{y}, \frac{x^*}{y^*}\right) = |\delta'| \leq \frac{2\epsilon}{1 - \epsilon}.$$

Theorem 7. *The stationary distribution $\theta^{\mathcal{B}}$ of an irreducible DTMC \mathcal{B} is given by $\theta^{\mathcal{B}}(s) = \frac{\mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_s]}{\sum_{t \in S} \mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_t]}$. Further, if $\text{diff}^{\text{rel}}(\mathbf{x}, (\mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_s])_{s \in S}) \leq \epsilon$ for $\mathbf{x} \in \mathbb{R}^{|S|}$ and $\epsilon \in (0, 1)$, then $\text{diff}^{\text{rel}}\left(\frac{\mathbf{x}}{\sum_{s \in S} \mathbf{x}(s)}, \theta^{\mathcal{B}}\right) \leq \frac{2\epsilon}{1-\epsilon}$.*

Proof. The equation system for computing the EVTs in $\mathcal{B}^{\hat{r}^{\hat{v}}}$ is:

$$\begin{aligned} \text{For all } s \in S \setminus \{v\} : \mathbf{x}(s) &= \sum_{r \in S} \mathbf{P}(r, s) \cdot \mathbf{x}(r) \\ \mathbf{x}(v) &= 1. \end{aligned}$$

Since $\{\hat{v}\}$ is the only BSCC and thus is reached with probability 1, the unique solution $(\mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_s])_{s \in S}$ to this equation system satisfies

$$\begin{aligned} \mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_v] &= 1 \\ &= \Pr^{\mathcal{B}^{\hat{r}^{\hat{v}}}}(\diamond\{\hat{v}\}) \\ &= \sum_{r \in S} \hat{\mathbf{P}}(r, \hat{v}) \cdot \mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_r] && \text{by Theorem 6} \\ &= \sum_{r \in S} \mathbf{P}(r, v) \cdot \mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_r]. \end{aligned}$$

We have thus shown that the equality

$$\mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_s] = \sum_{r \in S} \mathbf{P}(r, s) \cdot \mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_r]$$

holds for *all* $s \in S$, i.e., the expected visiting times $(\mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_s])_{s \in S}$ constitute a left eigenvector of the probability transition matrix \mathbf{P} of the original irreducible DTMC \mathcal{B} . Normalising the values so that they sum up to 1 yields:

$$\frac{\mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_s]}{\sum_{t \in S} \mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_t]} = \sum_{r \in S} \mathbf{P}(r, s) \cdot \frac{\mathbf{x}(r)}{\sum_{t \in S} \mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_t]}.$$

Since \mathcal{B} is irreducible, [41, Thm. 4.18] implies that the steady state distribution $\theta^{\mathcal{B}}$ is uniquely defined by $\theta^{\mathcal{B}}(s) = \frac{\mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_s]}{\sum_{t \in S} \mathbb{E}^{\mathcal{B}^{\hat{r}^{\hat{v}}}}[\mathbf{v}_t]}$.

The second claim of the proof follows directly from the propagation of relative errors from the beginning of this section. \square

A.4 Proofs of Section 6 (Conditional Expected Rewards)

Theorem 8. *Let $\text{rew} : S \rightarrow \mathbb{Q}_{\geq 0}$ with $\text{rew}(S_{\text{re}}) = \{0\}$. Then the equation system*

$$\forall s \in S_{\text{tr}} : \quad \mathbf{y}(s) = \text{rew}(s) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}_s] + \sum_{t \in S_{\text{tr}}} \mathbf{P}(t, s) \cdot \mathbf{y}(t)$$

has a unique solution $(\mathbf{y}(s))_{s \in S_{\text{tr}}}$ and for all absorbing $r \in S$, $\Pr^{\mathcal{D}}(\diamond\{r\}) > 0$,

$$\mathbb{E}^{\mathcal{D}}[\text{tr}_{\text{rew}} | \diamond\{r\}] = \frac{\sum_{t \in S_{\text{tr}}} \mathbf{P}(t, r) \cdot \mathbf{y}(t)}{\ell_{\text{init}}(r) + \sum_{t \in S_{\text{tr}}} \mathbf{P}(t, r) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}_t]}.$$

Proof. Throughout the proof we use matrix-vector notation for the sake of brevity and assume that $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}})$ is a DTMC with $\mathbf{P} = \begin{pmatrix} \mathbf{Q} & \mathbf{R} \\ \mathbf{0} & \mathbf{I} \end{pmatrix}$, i.e., \mathbf{Q} is the transition matrix from transient to transient states, \mathbf{R} is the transition matrix from transient to recurrent states. Furthermore, we assume that all recurrent states are absorbing (this is why there is an identity matrix in the lower right corner of \mathbf{P}).

First assume that $\text{rew}(S_{\text{tr}}) = 1$, i.e., all transient states have unit reward. Let $r \in S_{\text{re}}$. We can rewrite the integral in the definition of $\mathbb{E}^{\mathcal{D}}[\text{tr}_{\text{rew}} | \diamond\{r\}]$ as follows:

$$\begin{aligned} & \int_{\diamond\{r\}} \text{tr}_{\text{rew}} d\text{Pr}^{\mathcal{D}} \\ &= \sum_{k \geq 0} \int_{S \setminus \{r\} \mathbf{U}^{=k}\{r\}} \text{tr}_{\text{rew}} d\text{Pr}^{\mathcal{D}} \quad (\text{because } \diamond\{r\} = \bigsqcup_{k \geq 0} S \setminus \{r\} \mathbf{U}^{=k}\{r\}) \\ &= \sum_{k \geq 0} k \cdot \text{Pr}^{\mathcal{D}}(S \setminus \{r\} \mathbf{U}^{=k}\{r\}) \quad (\text{by the assumptions about } \text{rew}) . \end{aligned}$$

Recall that the transition matrix of \mathcal{D} is $\mathbf{P} = \begin{pmatrix} \mathbf{Q} & \mathbf{R} \\ \mathbf{0} & \mathbf{I} \end{pmatrix}$. Define $\tilde{\mathbf{P}} = \begin{pmatrix} \mathbf{Q} & \mathbf{R} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}$ and note that for all inegers $k \geq 1$ we have $\tilde{\mathbf{P}}^k = \begin{pmatrix} \mathbf{Q}^k & \mathbf{Q}^{k-1}\mathbf{R} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}$. Further, it can be shown that

$$\text{Pr}^{\mathcal{D}}(S \setminus \{r\} \mathbf{U}^{=k}\{r\}) = (\iota_{\text{init}} \cdot \tilde{\mathbf{P}}^k)_r .$$

Assume that ι_{init} is the row vector $\iota_{\text{init}} = (\tau_{\text{init}} \ \rho_{\text{init}})$ (i.e., τ_{init} and ρ_{init} are the initial probabilities of the transient and the recurrent states, respectively). Then

$$\begin{aligned} & \int_{\diamond\{r\}} \text{tr}_{\text{rew}} d\text{Pr}^{\mathcal{D}} \\ &= \sum_{k \geq 0} k \cdot (\iota_{\text{init}} \cdot \tilde{\mathbf{P}}^k)_r \quad (\text{see above}) \\ &= \sum_{k \geq 1} k \cdot (\tau_{\text{init}} \cdot \mathbf{Q}^{k-1} \cdot \mathbf{R})_r \quad (\text{equation for } \tilde{\mathbf{P}}^k \text{ above}) \\ &= \sum_{k \geq 1} k \cdot \sum_{t \in S_{\text{tr}}} (\tau_{\text{init}} \cdot \mathbf{Q}^{k-1})_t \cdot \mathbf{R}_{t,r} \quad (\text{rewriting}) \\ &= \sum_{t \in S_{\text{tr}}} \left(\sum_{k \geq 1} k \cdot (\tau_{\text{init}} \cdot \mathbf{Q}^{k-1})_t \right) \cdot \mathbf{R}_{t,r} \quad (\text{rewriting}) \quad (\star) \end{aligned}$$

Now let $t \in S_{\text{tr}}$. The next goal is to derive an equation system for the terms $\sum_{k \geq 1} k \cdot (\tau_{\text{init}} \cdot \mathbf{Q}^{k-1})_t$:

$$\begin{aligned} & \sum_{k \geq 1} k \cdot (\tau_{\text{init}} \cdot \mathbf{Q}^{k-1})_t \\ &= \sum_{k \geq 1} (\tau_{\text{init}} \cdot \mathbf{Q}^{k-1})_t + \sum_{k \geq 1} (k-1) \cdot (\tau_{\text{init}} \cdot \mathbf{Q}^{k-1})_t \quad (\text{rewriting}) \\ &= \sum_{k \geq 0} (\tau_{\text{init}} \cdot \mathbf{Q}^k)_t + \sum_{k \geq 0} k \cdot (\tau_{\text{init}} \cdot \mathbf{Q}^k)_t \quad (\text{index shifts}) \end{aligned}$$

$$\begin{aligned}
&= \left(\tau_{\text{init}} \cdot \sum_{k \geq 0} \mathbf{Q}^k \right)_t + \sum_{k \geq 0} k \cdot \sum_{t' \in S_{\text{tr}}} (\tau_{\text{init}} \cdot \mathbf{Q}^{k-1})_{t'} \cdot \mathbf{Q}_{t',t} \quad (\text{rewriting}) \\
&= \mathbb{E}^{\mathcal{D}}[\mathbf{v}t_t] + \sum_{t' \in S_{\text{tr}}} \left(\sum_{k \geq 1} k \cdot (\tau_{\text{init}} \cdot \mathbf{Q}^{k-1})_{t'} \right) \cdot \mathbf{Q}_{t',t} \quad (\text{by Theorem 1; rewriting})
\end{aligned}$$

In matrix-vector notation, the above equations can be written as

$$\begin{aligned}
\mathbf{y} &= \mathbb{E}^{\mathcal{D}}[\mathbf{v}t] + \mathbf{y} \cdot \mathbf{Q} \\
\iff \mathbf{y} \cdot (\mathbf{I} - \mathbf{Q}) &= \mathbb{E}^{\mathcal{D}}[\mathbf{v}t]
\end{aligned}$$

where $\mathbf{y} = \left(\sum_{k \geq 1} k \cdot (\tau_{\text{init}} \cdot \mathbf{Q}^{k-1})_t \right)_{t \in S_{\text{tr}}}$ and $\mathbb{E}^{\mathcal{D}}[\mathbf{v}t] = (\mathbb{E}^{\mathcal{D}}[\mathbf{v}t_t])_{t \in S_{\text{tr}}}$ are interpreted as row vectors. Since $\mathbf{I} - \mathbf{Q}$ is invertible [38], \mathbf{y} is the unique solution of this equation system. Theorem 8 with $\text{rew}(S_{\text{tr}}) = \{1\}$ now follows with (\star) (numerator) and Theorem 6 (denominator).

Now we discuss the case where the reward function rew assigns arbitrary non-negative rational rewards to the transient states S_{tr} .

First assume that rew takes positive integer values for the transient states, i.e., $\text{rew}(S_{\text{tr}}) \subseteq \mathbb{N}_{\geq 1}$. We proceed by induction on $\max_{s \in S_{\text{tr}}} \text{rew}(s)$. The base case where all states have reward 1 was already treated above. For the inductive case let $\max_{s \in S_{\text{tr}}} \text{rew}(s) = n > 1$ and consider all $s \in S_{\text{tr}}$ with $\text{rew}(s) = n$. We construct new a DTMC \mathcal{D}' with additional states s' for each such s and reward function rew' such that $\text{rew}'(s') = 1$, $\text{rew}'(s) = n - 1$ by redirecting all incoming transitions of s to s' and adding a single transition with probability 1 from s' to s . Intuitively, we simulate the reward n of s by inserting a unique predecessor with reward 1 right before s , and decreasing the reward of s to $n - 1$. It is clear that conditional expected rewards in \mathcal{D} and \mathcal{D}' are the same. Now we can apply the induction hypothesis to \mathcal{D}' (the induction hypothesis is just the statement of Theorem 8), i.e.,

$$\begin{aligned}
\mathbf{y}(s) &= (n - 1) \cdot \mathbb{E}^{\mathcal{D}'}[\mathbf{v}t_s] + \mathbf{y}(s') \\
\mathbf{y}(s') &= \mathbb{E}^{\mathcal{D}'}[\mathbf{v}t_{s'}] + \sum_{t \in S_{\text{tr}}} \mathbf{P}(t, s') \cdot \mathbf{y}(t)
\end{aligned}$$

But since s is a direct successor of s' , $\mathbb{E}^{\mathcal{D}'}[\mathbf{v}t_{s'}] = \mathbb{E}^{\mathcal{D}'}[\mathbf{v}t_s]$. Combining the above two equations yields $\mathbf{y}(s) = n \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}t_s] + \sum_{t \in S_{\text{tr}}} \mathbf{P}(t, s) \cdot \mathbf{y}(t)$ as desired.

Now assume that $\text{rew}(S_{\text{tr}}) \subseteq \mathbb{N}$ (including 0). We can make a similar construction as above for transient states $s \in S_{\text{tr}}$ with $\text{rew}(s) = 0$. We “eliminate” (see, e.g., [5]) each such s in \mathcal{D} to obtain a modified DTMC \mathcal{D}' with fewer states where each state has reward in $\mathbb{N}_{\geq 1}$. This is achieved by first scaling the outgoing transition of s by the factor $\frac{1}{1 - \mathbf{P}(s, s)}$, and then “gluing” each pair of incoming and outgoing transition of s together by multiplying their probabilities. The initial probability mass of $t_{\text{init}}(s)$ is distributed to the successors of s . Then s is not reachable anymore and can be removed from \mathcal{D}' . However, it can be shown that the effect of this elimination procedure on the transitions of \mathcal{D}' is really equivalent to simply adding the equation

$$\mathbf{y}(s) = \sum_{t \in S_{\text{tr}}} \mathbf{P}(t, s) \cdot \mathbf{y}(t) = 0 \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}t_s] + \sum_{t \in S_{\text{tr}}} \mathbf{P}(t, s) \cdot \mathbf{y}(t)$$

to the system; indeed, eliminating the variable $\mathbf{y}(s)$ with Gaussian elimination is the same as altering the transitions as described above [5].

Finally, if $\text{rew}(S_{\text{tr}}) \subseteq \mathbb{Q}_{\geq 0}$, we can consider $\text{rew}' : S \rightarrow \mathbb{N}$ where $\text{rew}'(s) = M \cdot \text{rew}(s)$ for all $s \in S$, M the greatest common integer multiple of the denominators of the rational numbers occurring in rew . Let the vector \mathbf{y}' be the unique solution of

$$\mathbf{y}' = (\text{rew}'(s) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{\text{tr}}} + \mathbf{y}' \cdot \mathbf{Q} ,$$

i.e., $\mathbf{y}' = (\text{rew}'(s) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{\text{tr}}} \cdot (\mathbf{I} - \mathbf{Q})^{-1}$. Then

$$\frac{1}{M} \mathbf{y}' = (\text{rew}(s) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{\text{tr}}} \cdot (\mathbf{I} - \mathbf{Q})^{-1} = \mathbf{y} . \quad (\dagger)$$


Now

$$\begin{aligned} & \mathbb{E}^{\mathcal{D}}[\text{tr}_{\text{rew}} | \diamond\{r\}] \\ &= \frac{1}{M} \cdot \mathbb{E}^{\mathcal{D}}[\text{tr}_{\text{rew}'} | \diamond\{r\}] && \text{(linearity of expected value)} \\ &= \frac{1}{M} \cdot \frac{\sum_{t \in S_{\text{tr}}} \mathbf{P}(t, r) \cdot \mathbf{y}'(t)}{\iota_{\text{init}}(r) + \sum_{t \in S_{\text{tr}}} \mathbf{P}(t, r) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{vt}_t]} && \text{(Theorem 8 for integer rewards)} \\ &= \frac{\sum_{t \in S_{\text{tr}}} \mathbf{P}(t, r) \cdot \mathbf{y}(t)}{\iota_{\text{init}}(r) + \sum_{t \in S_{\text{tr}}} \mathbf{P}(t, r) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{vt}_t]} && \text{by } (\dagger) \end{aligned}$$

which concludes the proof. \square

Remark 3. As an (immediate) consequence of Theorem 8, if $\mathbf{l} \leq \mathbb{E}^{\mathcal{D}}[\mathbf{vt}] \leq \mathbf{u}$, and $\mathbf{y}_{\mathbf{l}}$ and $\mathbf{y}_{\mathbf{u}}$ are the unique solutions of $(\mathbf{I} - \mathbf{Q}^T) \cdot \mathbf{y}_{\mathbf{l}} = \mathbf{l}$ and $(\mathbf{I} - \mathbf{Q}^T) \cdot \mathbf{y}_{\mathbf{u}} = \mathbf{u}$, respectively, then

$$\frac{(\mathbf{R}^T \cdot \mathbf{y}_{\mathbf{l}})_r}{\iota_{\text{init}}(r) + (\mathbf{R}^T \cdot \mathbf{u})_r} \leq \mathbb{E}^{\mathcal{D}}[\text{tr}_{\text{rew}} | \diamond\{r\}] \leq \frac{(\mathbf{R}^T \cdot \mathbf{y}_{\mathbf{u}})_r}{\iota_{\text{init}}(r) + (\mathbf{R}^T \cdot \mathbf{l})_r} .$$

Fig. 8: A DTMC for which VI does not provide ϵ -sound results.

B Value Iteration is not ϵ -Sound for EVTs

We show that the absolute or relative difference between the returned result and the true EVTs cannot be bounded by a predefined $\epsilon > 0$, in general.

Example 7. Assume that we use the standard VI algorithm with $\mathbf{x}^{(0)} = \mathbf{0}$ and $\epsilon = 0.1$ to compute the EVTs for the states s_1 and s_2 of the DTMC depicted in Figure 8, where $p \in [0, 1)$. It can be verified that the EVT of state s_1 equals 1, while the EVT of state s_2 equals $0.1 \cdot \frac{1}{1-p}$. The algorithm computes $\mathbf{x}^{(1)} = (1, 0)^T$ and $\mathbf{x}^{(2)} = (1, 0.1)^T$ in the first and second iteration. Thus, the absolute difference between the vectors $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ equals 0.1 and the algorithm terminates. Hence, the absolute difference between the returned vector and the true result equals

$$\text{diff}^{abs} \left(\begin{pmatrix} 1 \\ 0.1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0.1 \cdot \frac{1}{1-p} \end{pmatrix} \right) = \left| 0.1 - 0.1 \cdot \frac{1}{1-p} \right|.$$

In particular, the difference cannot be bounded by any $\epsilon > 0$ since $|\frac{1}{1-p}|$ can increase arbitrarily as p is getting closer to 1.

Using the relative termination criterion, the algorithm stops after at most 11 iterations for any $p < 1$. The reason for this is that $\mathbf{x}^{(10)}(s_2) = 0.1 \cdot \sum_{k=0}^8 p^k$ and $\mathbf{x}^{(11)}(s_2) = 0.1 \cdot \sum_{k=0}^9 p^k$. Thus,

$$\begin{aligned} \text{diff}^{rel}(\mathbf{x}^{10}, \mathbf{x}^{11}) &= \frac{0.1 \cdot \sum_{k=0}^8 p^k - 0.1 \cdot \sum_{k=0}^9 p^k}{0.1 \cdot \sum_{k=0}^9 p^k} \\ &= \frac{p^9}{\sum_{i=0}^9 p^i} \\ &\stackrel{p < 1}{\leq} \frac{p^9}{10 \cdot p^9} = 0.1. \end{aligned}$$

By a similar argument as above, we obtain that this criterion cannot guarantee ϵ -soundness: The true EVTs for s_2 can become arbitrarily high by setting p closer to 1. However, the returned value for s_2 equals $\mathbf{x}^{(11)}(s) = 0.1 \cdot \sum_{k=0}^9 p^k$, which is smaller than 1 for any $p < 1$. In this case, it can be verified that the relative difference is p^{10} , which implies that we can find $p < 1$ for every $\epsilon \in (0, 1)$ such that the relative difference is greater than ϵ .

C Gauss-Seidel Extensions

C.1 Gauss-Seidel Value Iteration

We lift an established optimization of the classic VI approach, called Gauss-Seidel VI [50], to EVTs. The VI algorithm, as described in Algorithm 1, uses the vector $\mathbf{x}^{(k-1)}$ of the previous iteration to perform the update of the current approximation $\mathbf{x}^{(k)} = \Phi(\mathbf{x}^{(k-1)})$. The idea of the Gauss-Seidel optimization is to compute the value of a state $s_i \in S_{\text{tr}}$ by considering the already updated values for states $s_j \in S_{\text{tr}}$ with $i < j$.

We consider an arbitrary but fixed linear order relation \preceq on the set of transient states and assume that the states are indexed such that $s_i \preceq s_j$ if and only if $i \leq j$ holds for all $s_i, s_j \in S_{\text{tr}}$.

Definition 8. The Gauss-Seidel EVTs-operator $\Gamma: \mathbb{R}^{|S_{\text{tr}}|} \rightarrow \mathbb{R}^{|S_{\text{tr}}|}$ is defined as $\Gamma(\mathbf{x}) = (\gamma_{s_i}(\mathbf{x}))_{s_i \in S_{\text{tr}}}$ for all $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$, where

$$\gamma_{s_i}(\mathbf{x}) = \iota_{\text{init}}(s_i) + \sum_{j=1}^{i-1} \mathbf{P}(s_j, s_i) \cdot \gamma_{s_j}(\mathbf{x}) + \sum_{j=i}^{|S_{\text{tr}}|} \mathbf{P}(s_j, s_i) \cdot \mathbf{x}(s_j).$$

The Gauss-Seidel variant of Algorithm 1 is obtained by replacing Φ with the Gauss-Seidel EVTs-operator Γ in Line 2. Then, given the input vector $\mathbf{x}^{(0)} \in \mathbb{R}^{|S_{\text{tr}}|}$, the algorithm computes the vector $\Gamma^{(k)}(\mathbf{x}^{(0)})$ for increasing $k \in \mathbb{N}$. The convergence results from standard VI also apply to the Gauss-Seidel variant:

Theorem 10. It holds that

- (i) $(\mathbb{E}[\mathbf{v}_t]_s)_{s \in S_{\text{tr}}}$ is the unique fixed point of Γ in $\mathbb{R}^{|S_{\text{tr}}|}$ and
- (ii) for all $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$, we have $\lim_{k \rightarrow \infty} \Gamma^{(k)}(\mathbf{x}) = (\mathbb{E}[\mathbf{v}_t]_s)_{s \in S_{\text{tr}}}$.

Proof. The proof of Theorem 10 is based on the proof provided by Puterman to establish the classical Gauss-Seidel VI approach [50, Theorem 6.3.4]. To be more precise, the statements (i) and (ii) can be shown by formulating Gauss-Seidel VI in terms of a regular splitting (see e.g., [57, Chapter 6.3]) The linear system $(\mathbf{I} - \mathbf{Q}^T) \cdot \mathbf{x} = \tau_{\text{init}}$ from Theorem 1 can be solved by splitting $(\mathbf{I} - \mathbf{Q}^T)$ into $(\mathbf{I} - \mathbf{Q}^T) = \mathbf{L} - \mathbf{U}$. In particular, the splitting is represented by $\mathbf{L} = (\mathbf{I} - \mathbf{Q}_L^T)$ and $\mathbf{U} = (\mathbf{Q}_U^T)$, where \mathbf{Q}_L^T and \mathbf{Q}_U^T are the strictly lower triangular and upper triangular part of \mathbf{Q}^T , respectively. Thus, for $n = |S_{\text{tr}}|$, we have

$$\mathbf{Q}_L^T = \begin{pmatrix} 0 & 0 & \cdots & 0 \\ \mathbf{Q}_{21}^T & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{Q}_{n1}^T & \mathbf{Q}_{n2}^T & \cdots & 0 \end{pmatrix} \quad \text{and} \quad \mathbf{Q}_U^T = \begin{pmatrix} \mathbf{Q}_{11}^T & \mathbf{Q}_{12}^T & \cdots & \mathbf{Q}_{1n}^T \\ 0 & \mathbf{Q}_{22}^T & \cdots & \mathbf{Q}_{2n}^T \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{Q}_{nn}^T \end{pmatrix}.$$

From the proof of Theorem 3, we recall that $\lim_{k \rightarrow \infty} (\mathbf{Q}^k)^k = \mathbf{0}$. It follows that $\lim_{k \rightarrow \infty} (\mathbf{Q}_L^T)^k = \mathbf{0}$ and thus, the spectral radius $\rho(\mathbf{Q}_L^T)$ is strictly smaller than 1. Consequently, the inverse of \mathbf{L} is well-defined by the Neumann series $\mathbf{L}^{-1} = (\mathbf{I} - \mathbf{Q}_L^T)^{-1} = \sum_{k=0}^{\infty} (\mathbf{Q}_L^T)^k$ (see [48]). In particular, we have that $\mathbf{Q}_L^T \geq \mathbf{0}$,

which implies that the infinite sum converges to a non-negative matrix. Since $\mathbf{U} = \mathbf{Q}_U^T$ is also non-negative, we can conclude that the splitting above is *regular*. Since $(\mathbf{I} - \mathbf{Q}^T)$ is invertible, we obtain that the iterative scheme

$$\mathbf{x}^{(k)} = \begin{cases} \mathbf{x} & \text{if } k = 0, \\ \mathbf{L}^{-1} \mathbf{U} \cdot \mathbf{x}^{(k-1)} + \mathbf{L}^{-1} \cdot \tau_{\text{init}} & \text{otherwise} \end{cases}$$

defines a sequence $(\mathbf{x}^{(k)})_{k \in \mathbb{N}}$ that converges to the unique solution of the linear system $(\mathbf{I} - \mathbf{Q}^T) \cdot \mathbf{x} = \tau_{\text{init}}$ for any starting vector $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$. In fact, the unique solution is equal to the vector $(\mathbb{E}[\mathbf{v}_s])_{s \in S_{\text{tr}}}$ by Theorem 1.

It remains to show that the sequence $(\mathbf{x}^{(k)})_{k \in \mathbb{N}}$ defined by the iterative scheme indeed coincides with the sequence $(\Gamma^{(k)}(\mathbf{x}))_{k \in \mathbb{N}}$ generated by iteratively applying the operator Γ to the vector \mathbf{x} . The proof is conducted by induction over $k \in \mathbb{N}$: For $k = 0$, the claim trivially holds. Assuming $\Gamma^{(k)}(\mathbf{x}) = \mathbf{x}^{(k)}$ holds for a fixed $k > 0$, we obtain $\Gamma^{(k+1)}(\mathbf{x}) = \Gamma(\mathbf{x}^{(k)})$. In particular, for all $s_i \in \{1, \dots, S_{\text{tr}}\}$, we obtain

$$\begin{aligned} \gamma_{s_i}(\mathbf{x}^{(k)}) &= \iota_{\text{init}}(s_i) + \sum_{j=1}^{i-1} \mathbf{P}(s_j, s_i) \cdot \gamma_{s_j}(\mathbf{x}^{(k)}(s_j)) + \sum_{j=i}^{|S_{\text{tr}}|} \mathbf{P}(s_j, s_i) \cdot \mathbf{x}^{(k)}(s_j) \\ &= \tau_{\text{init}}(s_i) + \sum_{j=1}^{i-1} \underbrace{\mathbf{Q}^T(s_i, s_j)}_{=\mathbf{L}_{ij}} \cdot \gamma_{s_j}(\mathbf{x}^{(k)}(s_j)) + \sum_{j=i}^{|S_{\text{tr}}|} \underbrace{\mathbf{Q}^T(s_i, s_j)}_{=\mathbf{U}_{ij}} \cdot \mathbf{x}^{(k)}(s_j) \\ &= \tau_{\text{init}}(s_i) - \sum_{j=1}^{i-1} \mathbf{L}_{ij} \cdot \gamma_{s_j}(\mathbf{x}^{(k)}(s_j)) + \sum_{j=i}^{|S_{\text{tr}}|} \mathbf{U}_{ij} \cdot \mathbf{x}^{(k)}(s_j). \end{aligned}$$

This can equivalently be written as

$$\left(\gamma_{s_i}(\mathbf{x}^{(k)}) \right)_{s_i \in S_{\text{tr}}} = \mathbf{L}^{-1} \mathbf{U} \cdot \Gamma^{(k)}(\mathbf{x}) + \mathbf{L}^{-1} \cdot \tau_{\text{init}},$$

which concludes the proof of Theorem 10. \square

One advantage of this modification is that the updates are performed in-place. As a result, the implementation can be realized using only one storage vector. In contrast, the standard VI algorithm requires two vectors: one for the values of the previous iteration and another one for the current iteration. Moreover, compared to standard VI, the Gauss-Seidel version can lead to faster convergence because more recent values are used for the update of the approximations. However, the Gauss-Seidel version of the VI algorithm also uses the absolute or relative stopping criterion. Thus, applying Gauss-Seidel VI for obtaining the EVTs of the DTMC considered in Example 7 can also yield inaccurate results without any guarantees on the accuracy.

C.2 Gauss-Seidel Interval Iteration

Analogous to the Gauss-Seidel version of VI, we can apply the Gauss-Seidel optimization to the II algorithm, where the update of the upper and lower bounds of EVTs of a state $s_i \in S_{\text{tr}}$ are performed based on the already updated

values for the states s_1, \dots, s_{i-1} . Again, we assume that the states are indexed such that $s_i \preceq s_j$ holds for all $s_i, s_j \in S_{\text{tr}}$, if and only if $i \leq j$ for an arbitrary but fixed linear order relation \preceq on the set of transient states.

Definition 9. The Gauss-Seidel Max EVTs-operator $\Gamma_{\max}: \mathbb{R}^{|S_{\text{tr}}|} \rightarrow \mathbb{R}^{|S_{\text{tr}}|}$ is defined as $\Gamma_{\max}(\mathbf{x}) = \max \left\{ \mathbf{x}, (\gamma_{s_i}^{\max}(\mathbf{x}))_{s_i \in S_{\text{tr}}} \right\}$ for all $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$, where

$$\gamma_{s_i}^{\max}(\mathbf{x}) = \iota_{\text{init}}(s_i) + \sum_{j=1}^{i-1} \mathbf{P}(s_j, s_i) \cdot (\Gamma_{\max}(\mathbf{x}))(s_j) + \sum_{j=i}^n \mathbf{P}(s_j, s_i) \cdot \mathbf{x}(s_j),$$

and the Gauss-Seidel Min EVTs-operator $\Gamma_{\min}: \mathbb{R}^{|S_{\text{tr}}|} \rightarrow \mathbb{R}^{|S_{\text{tr}}|}$ is given by $\Gamma_{\min}(\mathbf{x}) = \min \left\{ \mathbf{x}, (\gamma_{s_i}^{\min}(\mathbf{x}))_{s_i \in S_{\text{tr}}} \right\}$, where

$$\gamma_{s_i}^{\min}(\mathbf{x}) = \iota_{\text{init}}(s_i) + \sum_{j=1}^{i-1} \mathbf{P}(s_j, s_i) \cdot (\Gamma_{\min}(\mathbf{x}))(s_j) + \sum_{j=i}^n \mathbf{P}(s_j, s_i) \cdot \mathbf{x}(s_j).$$

The Gauss-Seidel II algorithm corresponds to Algorithm 2, where the operators Γ_{\max} and Γ_{\min} are used instead of Φ_{\max} and Φ_{\min} . The correctness results stated in Theorem 4 can be directly lifted to the Gauss-Seidel II variant of Algorithm 2 by using the following analog to Lemma 2.

Theorem 11. Let $\mathbf{u}, \mathbf{l} \in \mathbb{R}^{|S_{\text{tr}}|}$ such that $\mathbf{l} \leq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}} \leq \mathbf{u}$. For the operators Γ_{\max} and Γ_{\min} it holds that

- (i) $\Gamma_{\max}^{(k)}(\mathbf{l}) \leq \Gamma_{\max}^{(k+1)}(\mathbf{l})$ and $\Gamma_{\min}^{(k)}(\mathbf{u}) \geq \Gamma_{\min}^{(k+1)}(\mathbf{u})$ for all $k \in \mathbb{N}$,
- (ii) $\Phi^{(k)}(\mathbf{l}) \leq \Gamma_{\max}^{(k)}(\mathbf{l}) \leq (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}} \leq \Gamma_{\min}^{(k)}(\mathbf{u}) \leq \Phi^{(k)}(\mathbf{u})$ for all $k \in \mathbb{N}$, and
- (iii) $\lim_{k \rightarrow \infty} \Gamma_{\max}^{(k)}(\mathbf{l}) = (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}$ and $\lim_{k \rightarrow \infty} \Gamma_{\min}^{(k)}(\mathbf{u}) = (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}$.

Proof. Using the monotonicity of Γ the proof works analogously to the proof of Lemma 2, where the occurrences of Φ_{\max} and Φ_{\min} are replaced by Γ_{\max} and Γ_{\min} , respectively. \square

D Absolute ϵ -Sound Topological Computation

The total error can exceed ϵ if an imprecise method is used in Line 5 of Algorithm 3 even if it guarantees ϵ -sound approximations. In the following, we calculate an upper bound on the accumulated error with respect to the absolute difference.

To study the extent of the error, we consider three factors that have an impact on the error accumulation; the incoming transitions from different SCCs, the recurrence probability, and the length of the SCC chains. Below, we discuss the intuition behind each cause. For simplicity, we assume that the parametric SCC restriction $\mathcal{D}|_C[\mathbf{x}]$ is a well-defined DTMC for each SCC C .

- *Incoming transitions*: As elaborated above, the SCC restriction $\mathcal{D}|_C[\mathbf{x}]$ for the non-bottom SCC C is based on approximations $\mathbf{x}(t)$ of the EVTs of states t in $\mathcal{D}|_{C'}[\mathbf{x}]$ for SCCs $C' \hookrightarrow C$. We assume the vector \mathbf{x} satisfies $\mathbf{x}(t) = \mathbb{E}^{\mathcal{D}}[\mathbf{v}_t] + \epsilon_t$ for each $t \in C'$, $C' \hookrightarrow C$, and $|\epsilon_t| \leq \epsilon$. For a state $s \in C$, the initial distribution of $\mathcal{D}|_C[\mathbf{x}]$ (see Definition 7) is specified by

$$\iota_{\text{init}}^{\mathcal{D}|_C[\mathbf{x}]}(s) = \underbrace{\iota_{\text{init}}^{\mathcal{D}}(s) + \sum_{t \in S_{\text{tr}} \setminus C} \mathbf{P}^{\mathcal{D}}(t, s) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}_t]}_{=\iota_{\text{init}}^{\mathcal{D}|_C}(s)} + \sum_{t \in S_{\text{tr}} \setminus C} \mathbf{P}^{\mathcal{D}}(t, s) \cdot \epsilon_t.$$

Hence, the error in the initial distribution is given by $\sum_{t \in S_{\text{tr}} \setminus C} \mathbf{P}^{\mathcal{D}}(t, s) \epsilon_t$, which can generally exceed ϵ .

- *Recurrence probability*: To study the effect of recurrence probability, by [7, Lemma 3.5.] which states that EVTs in $\mathcal{D}|_C[\mathbf{x}]$ equal

$$\mathbb{E}^{\mathcal{D}|_C[\mathbf{x}]}[\mathbf{v}_s] = \Pr^{\mathcal{D}|_C[\mathbf{x}]}(\diamond\{s\}) \cdot \frac{1}{1 - \Pr_s^{\mathcal{D}|_C[\mathbf{x}]}(\bigcirc\Diamond\{s\})}.$$

Assuming the initial distribution of $\mathcal{D}|_C[\mathbf{x}]$ is erroneous yields that the reachability probability $\Pr^{\mathcal{D}|_C[\mathbf{x}]}(\diamond\{s\})$ in the parametric SCC restriction $\mathcal{D}|_C[\mathbf{x}]$ differs from the corresponding probability $\Pr^{\mathcal{D}|_C}(\diamond\{s\})$ in the exact SCC restriction $\mathcal{D}|_C$. It follows that the error in the reachability probability $\Pr^{\mathcal{D}|_C[\mathbf{x}]}(\diamond\{s\})$ is weighted by $\frac{1}{1 - \Pr_s^{\mathcal{D}|_C[\mathbf{x}]}(\bigcirc\Diamond\{s\})}$. Thus, the difference between the EVTs in $\mathcal{D}|_C[\mathbf{x}]$ and the EVTs of the DTMC \mathcal{D} increases as the recurrence probability tends to 1.

- *Chain length*: The analysis of an SCC C in Algorithm 3 is based on a potentially erroneous model $\mathcal{D}|_C[\mathbf{x}]$. Consequently, the EVTs in $\mathcal{D}|_C[\mathbf{x}]$ might differ from the EVTs in the original model \mathcal{D} as described above. Moreover, the vector $(\hat{\mathbf{x}}(s))_{s \in C}$ that is computed in Line 5 of Algorithm 3 is ϵ -sound with respect to the EVTs in $\mathcal{D}|_C[\mathbf{x}]$ and the approximated EVT $\mathbf{x}(s)$ of the state $s \in C$ can differ by an additional error ϵ from the — already erroneous — EVT of s in $\mathcal{D}|_C[\mathbf{x}]$. Thus, the error is propagated further since the values $\mathbf{x}(s)$ are considered in the analysis of the successor SCCs of C . In particular, the error accumulates up to L times if the longest chain in \mathcal{D} that ends in a non-bottom SCC has length L .

Lemma 6. *Let $\epsilon > 0$ and let $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$ be a vector such that for every non-bottom SCC C , the vector $(\mathbf{x}(s))_{s \in C}$ is an ϵ -sound approximation of the EVTs*

$(\mathbb{E}^{\mathcal{D}|_C[\mathbf{x}]}[\mathbf{vt}_s])_{s \in C}$ with respect to the absolute difference, i.e.,

$$\text{diff}^{abs} \left((\mathbf{x}(s))_{s \in C}, (\mathbb{E}^{\mathcal{D}|_C[\mathbf{x}]}[\mathbf{vt}_s])_{s \in C} \right) \leq \epsilon.$$

Then, the EVTs in \mathcal{D} satisfy

$$\text{diff}^{abs} \left(\mathbf{x}, (\mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{\text{tr}}} \right) \leq \epsilon \cdot \left(1 + \frac{1}{q} \cdot \sum_{i=1}^L T^i \right),$$

where

- T is the maximal number of incoming transitions into an SCC, i.e.,

$$T = \max_{C \in \text{SCC}_{\text{tr}}^{\mathcal{D}}} \left\{ \left| \{(s, t) \mid s \in S \setminus C, t \in C \text{ and } \mathbf{P}^{\mathcal{D}}(s, t) > 0\} \right| \right\},$$
- For each $s \in S_{\text{tr}}$, q under-approximates the probability of starting in s and never revisit s from the next step on: $0 < q \leq \min_{s \in S_{\text{tr}}} \left\{ 1 - \Pr_s^{\mathcal{D}}(\bigcirc \diamond \{s\}) \right\}$,
- L is the length of the longest chain ending in a non-bottom SCC, i.e.,

$$L = \max_{\kappa \in \text{Chains}_{\text{tr}} \mathcal{D}} \{ |\kappa| \}.$$

Proof. Consider $\epsilon > 0$ and $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$ as described above. Let C be a non-bottom SCC and let $\mathcal{D}|_C = (S^{\mathcal{D}|_C}, \mathbf{P}^{\mathcal{D}|_C}, \iota_{\text{init}}^{\mathcal{D}|_C})$ be the corresponding SCC restriction. We show $\mathbf{x}(s) \leq \mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s] + \epsilon \cdot \left(1 + \frac{1}{q} \cdot \sum_{i=1}^L T^i \right)$ holds for all $s \in S_{\text{tr}}$. Proving that $\mathbf{x}(s) \geq \mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s] - \epsilon \cdot \left(1 + \frac{1}{q} \cdot \sum_{i=1}^L T^i \right)$ works analogously.

We start by showing that for all $s \in C$, it holds that $\mathbf{x}(s) \leq \mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s] + \text{err}(s)$, where

$$\text{err}(s) = \epsilon \cdot \left(1 + \frac{1}{1 - \Pr_s^{\mathcal{D}|_C}(\bigcirc \diamond \{s\})} \cdot \sum_{i=1}^k T^i \right) \quad (11)$$

and k is the length of the longest SCC chain $C_0 \hookrightarrow \dots \hookrightarrow C_k \in \text{Chains}_{\text{tr}} \mathcal{D}$ such that $C_k = C$. The proof is by induction over the chain length $k \in \mathbb{N}$.

If $k = 0$, we know that C is not reachable via other SCCs. Thus, for the SCC restriction $\mathcal{D}|_C = (S^C, \mathbf{P}^C, \iota_{\text{init}}^C)$ and the parametric SCC restriction $\mathcal{D}|_C[\mathbf{x}] = (S^{\mathcal{D}|_C}, \mathbf{P}^{\mathcal{D}|_C}, \iota_{\text{init}}^{\mathcal{D}|_C[\mathbf{x}]})$ we have for $s \in C$

$$\iota_{\text{init}}^{\mathcal{D}|_C}(s) = \iota_{\text{init}}^{\mathcal{D}}(s) + \sum_{\substack{t \in C' \neq C \\ C' \hookrightarrow C}} \mathbf{P}(t, s) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s] = \iota_{\text{init}}^{\mathcal{D}}(s) = \iota_{\text{init}}^{\mathcal{D}|_C[\mathbf{x}]}(s).$$

We conclude that $\mathcal{D}|_C[\mathbf{x}]$ is equal to the DTMC $\mathcal{D}|_C$. Hence, for every $s \in C$, we have

$$\begin{aligned} \mathbf{x}(s) &\leq \mathbb{E}^{\mathcal{D}|_C[\mathbf{x}]}[\mathbf{vt}_s] + \epsilon && \text{by assumption on } \mathbf{x} \\ &= \mathbb{E}^{\mathcal{D}|_C}[\mathbf{vt}_s] + \epsilon \\ &= \mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s] + \text{err}(s) && \text{by Lemma 3.} \end{aligned}$$

Let C be a non-bottom SCC, where $k > 0$ is the length of the longest SCC chain ending in C . We assume that the claim holds for all SCCs that are reachable by a chain that has at most length $k - 1$.

Furthermore, let $\mathcal{D}|_C[\mathbf{x}]$ be the parametric restriction to C . By Remark 2, we have

$$(\mathbb{E}^{\mathcal{D}|_C[\mathbf{x}]}[\mathbf{v}\mathbf{t}_s])_{s \in C} = (\mathbf{I} - (\mathbf{Q}^{\mathcal{D}|_C})^T)^{-1} \cdot \tau_{\text{init}}^{\mathcal{D}|_C[\mathbf{x}]}.$$

We know that the inverse of $\mathbf{I} - (\mathbf{Q}^{\mathcal{D}|_C})^T$ is given by the following converging Neumann series (see [48]): $(\mathbf{I} - (\mathbf{Q}^{\mathcal{D}|_C})^T)^{-1} = \sum_{i=0}^{\infty} (\mathbf{Q}^{\mathcal{D}|_C})^i$. Thus, for a single state $s \in C$, we have

$$\begin{aligned} \mathbb{E}^{\mathcal{D}|_C[\mathbf{x}]}[\mathbf{v}\mathbf{t}_s] &= \sum_{t \in C} \sum_{i=0}^{\infty} \mathbf{P}^i(t, s) \cdot \iota_{\text{init}}^{\mathcal{D}|_C[\mathbf{x}]}(t) \\ &= \sum_{t \in C} \sum_{i=0}^{\infty} \mathbf{P}^i(t, s) \cdot \left(\iota_{\text{init}}^{\mathcal{D}}(t) + \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \mathbf{x}(t') \right). \end{aligned} \quad (12)$$

If $C' \hookrightarrow C$ holds, we conclude that the longest chain of non-bottom SCCs $\kappa \in \text{Chains}_{\text{tr}} \mathcal{D}$ that ends in C' satisfies $|\kappa| < k$. Hence, by the induction hypothesis, it holds that $\mathbf{x}(t') \leq \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_{t'}] + \text{err}(t')$ for every $t' \in C'$. Consequently, we obtain that

$$\begin{aligned} \mathbb{E}^{\mathcal{D}|_C[\mathbf{x}]}[\mathbf{v}\mathbf{t}_s] &\leq \sum_{t \in C} \sum_{i=0}^{\infty} \mathbf{P}^i(t, s) \cdot \left(\iota_{\text{init}}^{\mathcal{D}}(t) + \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot (\mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_{t'}] + \text{err}(t')) \right) \\ &= \sum_{t \in C} \sum_{i=0}^{\infty} \mathbf{P}^i(t, s) \cdot \left(\iota_{\text{init}}^{\mathcal{D}}(t) + \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \mathbb{E}^{\mathcal{D}}[\mathbf{v}\mathbf{t}_{t'}] \right. \\ &\quad \left. + \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \text{err}(t') \right) \\ &= \sum_{t \in C} \sum_{i=0}^{\infty} \mathbf{P}^i(t, s) \cdot \left(\Pr^{\mathcal{D}}((S \setminus C) \mathbf{U} \{t\}) + \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \text{err}(t') \right) \\ &\quad \text{by Lemma 5} \\ &= \sum_{t \in C} \sum_{i=0}^{\infty} \mathbf{P}^i(t, s) \cdot \left(\iota_{\text{init}}^{\mathcal{D}|_C}(t) + \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \text{err}(t') \right) \\ &\quad \text{by Definition 7.} \end{aligned}$$

From Theorem 1 it follows that

$$(\mathbb{E}^{\mathcal{D}|_C}[\mathbf{v}\mathbf{t}_s])_{s \in C} = (\mathbf{I} - (\mathbf{Q}^{\mathcal{D}|_C})^T)^{-1} \cdot \iota_{\text{init}}^{\mathcal{D}|_C}.$$

Moreover, the inverse of $(\mathbf{I} - (\mathbf{Q}^{\mathcal{D}|_C})^T)$ is given by the following converging Neumann series (see [48]):

$$(\mathbf{I} - (\mathbf{Q}^{\mathcal{D}|_C})^T)^{-1} = \sum_{i=0}^{\infty} (\mathbf{Q}^{\mathcal{D}|_C})^i$$

It follows that $\mathbb{E}^{\mathcal{D}|_C}[\mathbf{v}_s] = \sum_{t \in C} \sum_{i=0}^{\infty} \mathbf{P}^i(t, s) \cdot \ell_{\text{init}}^{\mathcal{D}|_C}(t)$. By continuing the reasoning above, we obtain that

$$\begin{aligned} \mathbb{E}^{\mathcal{D}|_C[\mathbf{x}]}[\mathbf{v}_s] &\leq \left(\sum_{t \in C} \sum_{i=0}^{\infty} \mathbf{P}^i(t, s) \cdot \ell_{\text{init}}^{\mathcal{D}|_C}(t) \right) \\ &\quad + \left(\sum_{t \in C} \sum_{i=0}^{\infty} \mathbf{P}^i(t, s) \sum_{t' \in C'} \mathbf{P}(t', t) \cdot \text{err}(t') \right) \\ &\leq \mathbb{E}^{\mathcal{D}|_C}[\mathbf{v}_s] + \sum_{t \in C} \sum_{i=0}^{\infty} \mathbf{P}^i(t, s) \cdot \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \text{err}(t'). \end{aligned}$$

[38, Theorem 3.2.4] implies that the EVT of s in the DTMC $\mathcal{D}|_C$ starting from a state t is equal to $\mathbb{E}_t^{\mathcal{D}|_C}[\mathbf{v}_s] = \sum_{i=0}^{\infty} \mathbf{P}^i(t, s)$. Thus,

$$\begin{aligned} \mathbb{E}^{\mathcal{D}|_C[\mathbf{x}]}[\mathbf{v}_s] &\leq \mathbb{E}^{\mathcal{D}|_C}[\mathbf{v}_s] + \sum_{t \in C} \mathbb{E}_t^{\mathcal{D}|_C}[\mathbf{v}_s] \cdot \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \text{err}(t') \\ &= \mathbb{E}^{\mathcal{D}|_C}[\mathbf{v}_s] + \sum_{t \in C} \frac{\Pr_t^{\mathcal{D}|_C}(\diamond\{s\})}{1 - \Pr_s^{\mathcal{D}|_C}(\bigcirc\Diamond\{s\})} \cdot \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \text{err}(t') \quad \text{by [7, Lemma 3.5.]} \\ &\leq \mathbb{E}^{\mathcal{D}|_C}[\mathbf{v}_s] + \frac{1}{1 - \Pr_s^{\mathcal{D}|_C}(\bigcirc\Diamond\{s\})} \cdot \sum_{t \in C} \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \mathbf{P}(t', t) \cdot \text{err}(t'). \quad (13) \end{aligned}$$

For each SCC chain $\kappa \in \text{Chains}_{\text{tr}} \mathcal{D}$ that ends in an SCC C' with $C' \hookrightarrow C$ we have that $|\kappa| \leq k - 1$. Consequently, for each $t' \in C'$, we obtain that

$$\text{err}(t') \leq \epsilon \cdot \left(1 + \frac{1}{1 - \Pr_{t'}^{\mathcal{D}|_C}(\bigcirc\Diamond\{t'\})} \cdot \sum_{i=1}^{k-1} T^i \right). \quad (14)$$

Moreover, for all $t \in C$, it holds that $\mathbf{P}(t', t)$ is at most the probability of leaving the SCC C' . In particular, $\mathbf{P}(t', t) \leq 1 - \Pr_{t'}^{\mathcal{D}|_C}(\bigcirc\Diamond\{t'\})$, which implies that

$$\frac{\mathbf{P}(t', t)}{1 - \Pr_{t'}^{\mathcal{D}|_C}(\bigcirc\Diamond\{t'\})} \leq \llbracket \mathbf{P}(t', t) > 0 \rrbracket. \quad (15)$$

By combining Equations (14) and (15), we obtain that

$$\begin{aligned} \mathbf{P}(t', t) \cdot \text{err}(t') &\leq \mathbf{P}(t', t) \cdot \epsilon \cdot \left(1 + \frac{1}{1 - \Pr_{t'}^{\mathcal{D}|_C}(\bigcirc\Diamond\{t'\})} \cdot \sum_{i=1}^{k-1} T^i \right) \\ &= \epsilon \cdot \left(\mathbf{P}(t', t) + \frac{\mathbf{P}(t', t)}{1 - \Pr_{t'}^{\mathcal{D}|_C}(\bigcirc\Diamond\{t'\})} \cdot \sum_{i=1}^{k-1} T^i \right) \\ &\leq \llbracket \mathbf{P}(t', t) > 0 \rrbracket \cdot \epsilon \cdot \left(1 + \sum_{i=1}^{k-1} T^i \right) \end{aligned}$$

Using this inequality in Equation (13) yields

$$\begin{aligned} \mathbb{E}^{\mathcal{D}|C}[\mathbf{x}][\mathbf{v}_s] &\leq \mathbb{E}^{\mathcal{D}|C}[\mathbf{v}_s] + \frac{1}{1 - \Pr_s^{\mathcal{D}|C}(\bigcirc\Diamond\{s\})} \\ &\quad \cdot \sum_{t \in C} \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \llbracket \mathbf{P}(t', t) > 0 \rrbracket \cdot \epsilon \cdot \left(1 + \sum_{i=1}^{k-1} T^i\right) \end{aligned}$$

The sum $\sum_{t \in C} \sum_{\substack{t' \in C' \\ C' \hookrightarrow C}} \llbracket \mathbf{P}(t', t) > 0 \rrbracket$ is equal to the number of incoming transitions to C . Since T is an upper bound for this number, we conclude that

$$\begin{aligned} \mathbb{E}^{\mathcal{D}|C}[\mathbf{x}][\mathbf{v}_s] &\leq \mathbb{E}^{\mathcal{D}|C}[\mathbf{v}_s] + \frac{1}{1 - \Pr_s^{\mathcal{D}|C}(\bigcirc\Diamond\{s\})} \cdot T \cdot \epsilon \cdot \left(1 + \sum_{i=1}^{k-1} T^i\right) \\ &= \mathbb{E}^{\mathcal{D}|C}[\mathbf{v}_s] + \frac{1}{1 - \Pr_s^{\mathcal{D}|C}(\bigcirc\Diamond\{s\})} \cdot \epsilon \cdot \sum_{i=1}^k T^i. \end{aligned}$$

By assumption, the vector $(\mathbf{x}(s))_{s \in C}$ is ϵ -sound with respect to the absolute difference from the vector $(\mathbb{E}^{\mathcal{D}|C}[\mathbf{x}][\mathbf{v}_s])_{s \in C}$. Consequently, we have that

$$\begin{aligned} \mathbf{x}(s) &\leq \mathbb{E}^{\mathcal{D}|C}[\mathbf{x}][\mathbf{v}_s] + \epsilon \\ &\leq \mathbb{E}^{\mathcal{D}|C}[\mathbf{v}_s] + \frac{1}{1 - \Pr_s^{\mathcal{D}|C}(\bigcirc\Diamond\{s\})} \cdot \epsilon \cdot \sum_{i=1}^k T^i + \epsilon \\ &\leq \mathbb{E}^{\mathcal{D}|C}[\mathbf{v}_s] + \epsilon \cdot \left(1 + \frac{1}{1 - \Pr_s^{\mathcal{D}|C}(\bigcirc\Diamond\{s\})} \cdot \sum_{i=1}^k T^i\right) \\ &\leq \underbrace{\mathbb{E}^{\mathcal{D}}[\mathbf{v}_s] + \epsilon \cdot \left(1 + \frac{1}{1 - \Pr_s^{\mathcal{D}|C}(\bigcirc\Diamond\{s\})} \cdot \sum_{i=1}^k T^i\right)}_{=err(s)} \quad \text{by Lemma 3.} \end{aligned}$$

Thus, we have shown that the claim holds for all $k \in \mathbb{N}$.

Next, let $s \in S_{\text{tr}}$ and C be the SCC C of s . Additionally, assume that k is the length of the longest SCC chain $\kappa \in \text{Chains}_{\text{tr}}\mathcal{D}$ ending in C . The length of the longest SCC chain containing only transient SCCs in \mathcal{D} is equal to L and by the choice of q , we have that $\frac{1}{1 - \Pr_s^{\mathcal{D}|C}(\bigcirc\Diamond\{s\})} \leq \frac{1}{q}$. Thus,

$$\begin{aligned} \mathbf{x}(s) &\leq \mathbb{E}^{\mathcal{D}}[\mathbf{v}_s] + \epsilon \cdot \left(1 + \frac{1}{1 - \Pr_s^{\mathcal{D}|C}(\bigcirc\Diamond\{s\})} \cdot \sum_{i=1}^k T^i\right) \\ &\leq \mathbb{E}^{\mathcal{D}}[\mathbf{v}_s] + \epsilon \cdot \left(1 + \frac{1}{q} \cdot \sum_{i=1}^L T^i\right), \end{aligned}$$

which concludes the proof of Lemma 6. \square

Using the upper bounds on the absolute differences provided in Lemma 6, we derive an algorithm for the topological computation of ϵ -sound results. The value q satisfying $0 < q \leq \min_{s \in S_{\text{tr}}} \left\{ (1 - \Pr_s^{\mathcal{D}}(\bigcirc\Diamond\{s\})) \right\}$, can be computed

Algorithm 4: Sound topological computation of EVTs diff^{abs} .

Input: DTMC \mathcal{D} , $\epsilon > 0$ for $\text{crit} = \text{abs}$ and $\epsilon \in (0, 1)$ for $\text{crit} = \text{rel}$

Output: $\mathbf{x} \in \mathbb{R}^{|S_{\text{tr}}|}$ with $\text{diff}^{\text{crit}}(\mathbf{x}, (\mathbb{E}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}) \leq \epsilon$

```

1  $\{C_1 \dots C_n\} \leftarrow \text{SCC}_{\text{tr}}^{\mathcal{D}}$  // obtain the non-bottom SCCs,  $C_i \hookrightarrow C_j$  implies  $i < j$ 
2  $\mathbf{x} \leftarrow \mathbf{0}$ 
3 compute  $q$  with  $0 < q \leq \min_{s \in S_{\text{tr}}} \{(1 - \Pr_s^{\mathcal{D}}(\bigcirc \Diamond \{s\}))\}$  // e.g., [7, Section 3.3]
4  $T \leftarrow \max_{C \in \text{SCC}_{\text{tr}}^{\mathcal{D}}} \{|\{(s, t) \mid s \in S \setminus C, t \in C \text{ and } \mathbf{P}(s, t) > 0\}|\}$ 
5  $L \leftarrow \max_{\kappa \in \text{Chains}_{\text{tr}} \mathcal{D}} \{|\kappa|\}$ 
6  $\sigma \leftarrow \frac{\epsilon}{1 + \frac{1}{q} \cdot \sum_{i=1}^L T^i}$  // compute the absolute threshold
7 for  $k \leftarrow 1$  to  $n$  do
8   compute  $\hat{\mathbf{x}} \in \mathbb{R}^{|C|}$  with  $\text{diff}^{\text{crit}}(\hat{\mathbf{x}}, (\mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in C_k}) \leq \sigma$  // using e.g., II
9   for  $s \in C_k$  do  $\mathbf{x}(s) \leftarrow \hat{\mathbf{x}}(s)$ 
10 return  $\mathbf{x}$ 

```

automatically, for example, by exploiting the graph-based methods proposed by Baier et al. [7, Section 3.3] that bound the recurrence probabilities $\Pr_s^{\mathcal{D}}(\bigcirc \Diamond \{s\})$ from above.

The absolute ϵ -sound topological computation of EVTs is described in Algorithm 4. The procedure works analogously to the relative ϵ -sound topological computation of EVTs described in Section 4.3. The difference is that the choice of σ is based on the bounds from Lemma 6. As a result, the termination threshold σ ensures that the *absolute* difference between the result \mathbf{x} and the EVTs $(\mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}$ does not exceed ϵ .

Theorem 12. For the input DTMC \mathcal{D} , and any threshold $\epsilon > 0$ Algorithm 4 terminates and returns a vector $\mathbf{x}^{\text{res}} \in \mathbb{R}^{|S_{\text{tr}}|}$ with $\text{diff}^{abs}(\mathbf{x}^{\text{res}}, (\mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}) \leq \epsilon$.

Proof. After each non-bottom SCC has been considered, the algorithm terminates. We prove ϵ -soundness w.r.t. the absolute difference. Consider σ after the execution of Line 6 in Algorithm 4, i.e., $\sigma = \frac{\epsilon}{1 + \frac{1}{q} \cdot \sum_{i=1}^L T^i}$ where

- q satisfies $0 < q \leq \min_{s \in S_{\text{tr}}} \{(1 - \Pr_s^{\mathcal{D}}(\bigcirc \Diamond \{s\}))\}$
- $L = \max_{\kappa \in \text{Chains}_{\text{tr}} \mathcal{D}} \{|\kappa|\}$, and
- $T = \max_{C \in \text{SCC}_{\text{tr}}^{\mathcal{D}}} \{|\{(t', t) \mid t' \in S \setminus C, t \in C \text{ and } \mathbf{P}(t', t) > 0\}|\}$.

By the definition of q and Lemma 6, we obtain that the returned vector $\mathbf{x}^{\text{res}} \in S_{\text{tr}}$ satisfies

$$\text{diff}^{abs}(\mathbf{x}^{\text{res}}, (\mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}) \leq \sigma \cdot \left(1 + \sum_{i=1}^L T^i \cdot \frac{1}{q}\right). \quad (16)$$

Hence, we conclude

$$\text{diff}^{abs}(\mathbf{x}^{\text{res}}, (\mathbb{E}^{\mathcal{D}}[\mathbf{vt}_s])_{s \in S_{\text{tr}}}) \leq \sigma \cdot \left(1 + \sum_{i=1}^L T^i \cdot \frac{1}{q}\right)$$

$$\begin{aligned} &= \frac{\epsilon}{1 + \sum_{i=1}^L T^i \cdot \frac{1}{q}} \cdot \left(1 + \sum_{i=1}^L T^i \cdot \frac{1}{q} \right) \\ &= \epsilon. \end{aligned}$$

□

E Additional Experimental Results

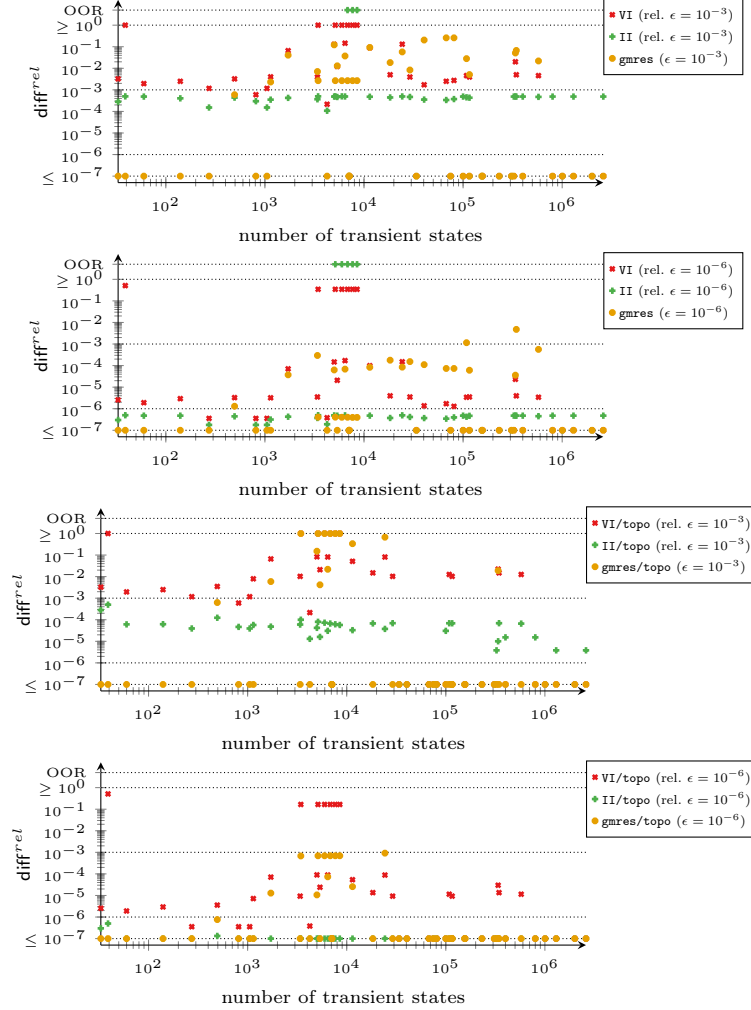


Fig. 9: Four scatter plots for the comparison of the relative difference (y -axis) between the true EVTs and results obtained by the different methods for each instance (x -axis, represented by the number of transient states). The methods include GMRES, VI and II using the relative termination criterion with a threshold of $\epsilon = 10^{-3}$ or $\epsilon = 10^{-6}$. The 4 dotted lines indicate, from bottom to top, if the result either yields a relative difference that is less or equal to the threshold 10^{-7} , exactly equal to a precision of 10^{-6} or 10^{-3} , or greater or equal to 1. The topmost line indicates that no result was obtained, for example, due to timeouts.




Fig. 10: Scatter plots depicting the runtime for EVTs computation (in seconds). The first two plots compare the (topological) II approach using the absolute stopping criterion and II using the relative stopping criterion. The rightmost plot compares LU factorization using floating point arithmetic and exact arithmetic.




Fig. 11: Scatter plots for the comparison of the relative difference (y -axis) between the true stationary distribution and results obtained by the different methods for each instance (x -axis, represented by the number of transient states). LU factorization, GMRES and II use the relative termination criterion with a threshold of $\epsilon = 10^{-3}$ or $\epsilon = 10^{-6}$.

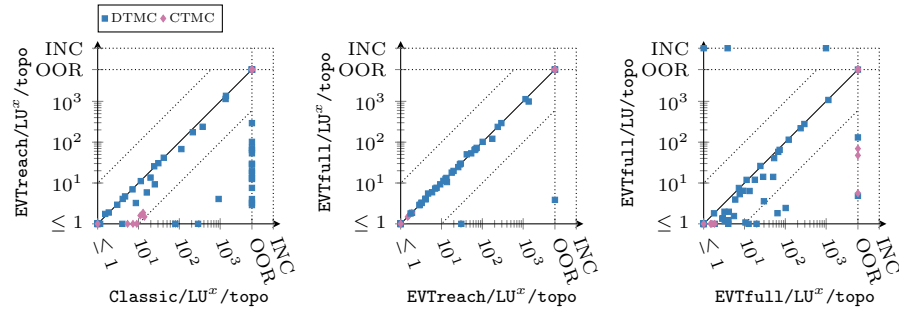


Fig. 12: Scatter plots depicting the runtime for stationary distribution computation (in seconds). The first two plots compare the **Classic**, **EVTreach** and the **EVTfull** approach using LU^x . II using the relative stopping criterion. The rightmost plot compares LU factorization using floating point arithmetic and exact arithmetic.