Faster Algorithms for Quantitative Analysis of MCs and MDPs with Small Treewidth*

Ali Asadi[1], Krishnendu Chatterjee[2], Amir Kafshdar Goharshady[2], Kiarash Mohammadi[3], and Andreas Pavlogiannis[4]

[1] Sharif University of Technology, Tehran, Iran, aasadi@ce.sharif.edu
[2] IST Austria, Klosterneuburg, Austria, {kchatterjee, goharshady}@ist.ac.at
[3] Ferdowsi University of Mashhad, Mashhad, Iran, kiarash.km@gmail.com
[4] Aarhus University, Aarhus, Denmark, pavlogiannis@cs.au.dk

Abstract. Discrete-time Markov Chains (MCs) and Markov Decision Processes (MDPs) are two standard formalisms in system analysis. Their main associated *quantitative* objectives are hitting probabilities, discounted sum, and mean payoff. Although there are many techniques for computing these objectives in general MCs/MDPs, they have not been thoroughly studied in terms of parameterized algorithms, particularly when treewidth is used as the parameter. This is in sharp contrast to *qualitative* objectives for MCs, MDPs and graph games, for which treewidth-based algorithms yield significant complexity improvements. In this work, we show that treewidth can also be used to obtain faster algorithms for the quantitative problems. For an MC with n states and m transitions, we show that each of the classical quantitative objectives can be computed in $O((n+m)\cdot t^2)$ time, given a tree decomposition of the MC with width t. Our results also imply a bound of $O(\kappa\cdot(n+m)\cdot t^2)$ for each objective on MDPs, where κ is the number of strategy-iteration refinements required for the given input and objective. Finally, we make an experimental evaluation of our new algorithms on low-treewidth MCs and MDPs obtained from the DaCapo benchmark suite. Our experiments show that on low-treewidth MCs and MDPs, our algorithms outperform existing well-established methods by one or more orders of magnitude.

Keywords: Markov Chains · Markov Decision Processes · Treewidth

1 Introduction

MCs. Perhaps the most standard formalism for modeling randomness in discrete-time systems is that of discrete-time Markov Chains (MCs). MCs have immense

*A longer version is available at [1]. The research was partly supported by Austrian Science Fund (FWF) Grant No. NFN S11407-N23 (RiSE/SHiNE), Vienna Science and Technology Fund (WWTF) Project ICT15-003, the Facebook PhD Fellowship Program, and DOC Fellowship No. 24956 of the Austrian Academy of Sciences (ÖAW).

applications in verification, and are used to express randomness both in the system and in the environment [11]. Besides the theoretical appeal, the analysis of MCs is also a core component in several model checkers [19,30].

MDPs. When the system exhibits both stochastic and non-deterministic behavior, the standard model of MCs is lifted to Markov Decision Processes (MDPs). For example, MDPs are used to model stochastic controllers, where the non-determinism models freedom of the controller and randomness models the behavior of the system. MDPs are also a topic of active study in verification [14].

Quantitative Analysis. Three of the most standard analysis objectives for MCs are the following: The *hitting probabilities* objective takes as input a set of target vertices $\mathfrak{T}$ of the MC, and asks to compute for each vertex u, the probability that a random walk from u eventually hits $\mathfrak{T}$. The *discounted sum* objective takes as input a discount factor $\lambda \in (0, 1)$ and a reward function R that assigns a reward to each edge. The task is to compute for each vertex u the expected reward of a random walk starting from u, where the value of the walk is the sum of the rewards along its edges, discounted by the factor λ at each step*. Finally, the *mean payoff* objective is similar to discounted sum, except that the value of a walk is the long-run average of the rewards along its edges. In MDPs, the analyses ask for a strategy that maximizes the respective quantity.

Analysis Algorithms. Given the importance of quantitative objectives for MCs and MDPs, there have been various techniques for solving them efficiently. For MCs, the hitting probabilities and discounted sum objectives reduce to solving a system of linear equations [32]. For MDPs, all three objectives reduce to solving a linear program [32]. Besides the LP formulation, two popular approaches for solving quantitative objectives on MDPs are value iteration [3] and strategy iteration [28]. Value iteration is the most commonly used method in verification and operates by computing optimal policies for successive finite horizons. However, this process leads only to approximations of the optimal values, and for some objectives no stopping criterion for the optimal strategy is known [2]. In cases where such criteria are known (e.g. [35]), the number of iterations necessary before the numbers can be rounded to provide an optimal solution can be extremely high [10]. Nevertheless, value iteration has proved to be very successful in practice and is included in many probabilistic model checkers, such as [30,19]. On the other hand, strategy iteration lies on the observation that given a fixed strategy, the MDP reduces to an MC, and hence one can compute the value of each vertex using existing techniques on MCs. Then, the strategy can be refined to a new strategy that improves the value of each vertex. The running time of strategy iteration can be written as $O(\kappa \cdot f)$, where κ is the number of strategy refinements and f is the time for evaluating the strategy. Although κ can be exponentially large [20], it behaves as a small constant in practice, which makes strategy iteration work well in practice [29].

Treewidth. *Treewidth* is a well-studied graph parameter. Many classes of graphs which arise in practice have constant treewidth. An example is that Control

*The undiscounted sum objective is obtained by letting $\lambda = 1$ and our algorithms for discounted sum can be slightly modified to handle this case, too.

Flow Graphs (CFGs) of `goto`-free programs in many programming languages have constant treewidth [37,26,17]. Treewidth has important algorithmic implications, as many graph problems admit (more) efficient solutions on graphs of low treewidth [15,37,23,24]. In program analysis, treewidth has been exploited to develop improvements for register allocation [37], algebraic-path analysis [13], dataflow analysis [16,8], data-dependence analysis [7], and model checking [33,22].

Our Contributions. The contributions of this work are as follows:

1. *Theoretical Contributions.* Our general theoretical result is a linear-time algorithm for solving systems of linear equations whose primal graph has low treewidth. Given a linear system S of m equations over n unknowns, and a tree decomposition of the primal graph of S that has width t, our algorithm solves S in time $O((n+m)\cdot t^2)$. Given an MC M of treewidth t and a corresponding tree decomposition, our algorithm directly implies similar running times for the hitting probabilities and discounted sum objectives for M. In addition, we develop an algorithm that solves the mean-payoff objective for M in time $O((n+m)\cdot t^2)$. Our results on MCs also imply upper-bounds for the running time of strategy iteration on low-treewidth MDPs. Given an MDP P with treewidth t and a quantitative objective, our results imply that P can be solved in time $O(\kappa\cdot(n+m)\cdot t^2)$, where κ is the number of iterations until strategy iteration stabilizes for the respective input and objective.
2. *Practical Contributions.* We develop two practical algorithms for solving the hitting probabilities and discounted sum objectives on low-treewidth MCs. Although these algorithms have the same worst-case complexity of $O((n+m)\cdot t^2)$ as our general solution, they avoid its most practically time-consuming step, i.e. applying the Gram-Schmidt process, and replace it with simple changes to the MC. We report on an implementation of these algorithms and their performance in solving MCs and MDPs with low treewidth.

The existing works closest to this paper are [12,23]. The work of [12] considers the maximal end-component decomposition and the almost-sure reachability set computation in low-treewidth MDPs. These are both *qualitative* objectives, and thus very different from the *quantitative* objectives we consider here, which cannot be solved by [12]. Specifically, the main problem solved by [12] is almost-sure reachability, i.e. reachability with probability 1, which is a very special qualitative case of computing hitting probabilities. The work of [23] develops an algorithm for solving linear systems of low treewidth. Considering the computational complexity when applied to MCs/MDPs of treewidth t, the algorithms we develop in this work are a factor t faster compared to [23]. On the practical side, the algorithms in [23] have more complicated intermediate steps, which we expect will lead to large constant factors in the runtime of their implementations.

2 Preliminaries

Discrete Probability Distributions. Given a finite set X, a probability distribution over X is a function $d : X \to [0,1]$ such that $\sum_{x\in X} d(x) = 1$. We denote the set of all probability distributions over X by $\mathcal{D}(X)$.

Markov Chains (MCs). A *Markov chain* $C = (V, E, \delta)$ consists of a finite directed graph (V, E) and a probabilistic transition function $\delta : V \to \mathcal{D}(V)$, such that for any pair u, v of vertices, we have $\delta(u)(v) > 0$ only if $(u, v) \in E$. In an MC C, we start a random walk from a vertex $v_0 \in V$ and at each step, being in a vertex v, we probabilistically choose one of the successors of v and go there. The probability with which a successor w is chosen is given by $\delta(v)(w)$. Let $O \subseteq V^\omega$ be a measurable set of infinite paths on V, we use the notation $Pr_{v_0}(O)$ to denote the probability that our infinite random walk starting from v_0 is a member of O.

Markov Decision Processes (MDPs). A *Markov decision process* is a tuple $P = (V, E, V_1, V_P, \delta)$ which consists of a finite directed graph (V, E), a partitioning of V into two sets V_1 and V_P, and a probabilistic transition function $\delta : V_P \to \mathcal{D}(V)$, such that for any $(u, v) \in V_P \times V$, we have $\delta(u)(v) > 0$ only if $(u, v) \in E$. We assume that all vertices of an MDP have at least one outgoing edge. Intuitively, an MDP is a one-player game with two types of vertices: those controlled by Player 1, i.e. V_1, and those that behave probabilistically, i.e. V_P.

Strategies. In an MDP P, a *strategy* is a function $\sigma : V_1 \to V$, such that for every $v \in V_1$ we have $(v, \sigma(v)) \in E$. Informally, a strategy is a recipe for Player 1 that tells her which successor to choose based on the current state[†]. Given an MDP P with a strategy σ, we start a random walk from a vertex $v_0 \in V$ and at each step, being in a vertex v, choose the successor as follows: (i) if $v \in V_1$, then we go to $\sigma(v)$, and (ii) if $v \in V_P$ we act as in the case of MCs, i.e. we go to each successor w with probability $\delta(v)(w)$. As before, given a measurable set $O \subseteq V^\omega$ of infinite paths on V, we define $Pr^\sigma_{v_0}(O)$ as the probability that our infinite random walk becomes a member of O. Note that an MDP with a fixed strategy σ is basically an MC, in which for every $v \in V_1$ we have $\delta(v)(\sigma(v)) = 1$.

Hitting Probabilities [32]***.*** Let $C = (V, E, \delta)$ be an MC and $\mathfrak{T} \subseteq V$ a designated set of *target* vertices. We define $\mathit{Hit}(\mathfrak{T}) \subseteq V^\omega$ as the set of all infinite sequences of vertices that intersect $\mathfrak{T}$. The *Hitting probability* $\mathit{HitPr}(u, \mathfrak{T})$ is defined as $Pr_u(\mathit{Hit}(\mathfrak{T}))$. In other words, $\mathit{HitPr}(u, \mathfrak{T})$ is the probability of eventually reaching $\mathfrak{T}$, assuming that we start our random walk at u. In case of MDPs, we assume that the player aims to maximize the hitting probability by choosing the best possible strategy. Therefore, we define $\mathit{HitPr}(u, \mathfrak{T})$ as $\max_\sigma Pr^\sigma_u(\mathit{Hit}(\mathfrak{T}))$.

Discounted Sums of Rewards [34]***.*** Let $C = (V, E, \delta)$ be an MC and $R : E \to \mathbb{R}$ a *reward function* that assigns a real value to each edge. Also, let $\lambda \in (0, 1)$ be a *discount factor*. Given an infinite path $\pi = v_0, v_1, \dots$ over (V, E), we define the total reward $R(\pi)$ of π as

$$\sum_{i=0}^{\infty} \lambda^i \cdot R(v_i, v_{i+1}) = R(v_0, v_1) + \lambda \cdot R(v_1, v_2) + \lambda^2 \cdot R(v_2, v_3) + \dots.$$

For $u \in V$ we define $\mathit{ExpDisSum}(u)$ as the expected value of the reward of our random walk if we begin it at u, i.e. $\mathit{ExpDisSum}(u) := \mathbb{E}_u[R(\pi)]$. As in the previous case, when considering MDPs, we assume that the player aims to max-

[†]We only consider pure memoryless strategies because they are sufficient for our use-cases, i.e. there always exists an optimal strategy that is pure and memoryless [29].

imize the discounted sum, hence given an MDP $P = (V, E, V_1, V_P, \delta)$, a reward function R and a discount factor λ, we define $\mathit{ExpDisSum}(u) := \max_\sigma \mathbb{E}_u^\sigma[R(\pi)]$.
Mean Payoff [34,29]**.** Let C be an MC and R a reward function. Given an infinite path $\pi = v_0, v_1, \ldots$ over C, we define the n-step average reward of π as

$$R(\pi[0..n]) := \frac{1}{n} \sum_{i=1}^{n} R(v_{i-1}, v_i).$$

Given a start vertex $u \in V$, the expected *long-time average* or *mean payoff* value from u is defined as $\mathit{ExpMP}(u) := \lim_{n\to\infty} \mathbb{E}_u[R(\pi[0..n])]$. In other words, $\mathit{ExpMP}(u)$ captures the expected reward per step in a random walk starting at u. For an MDP P, we define $\mathit{ExpMP}(u) := \max_\sigma \lim_{n\to\infty} \mathbb{E}_u^\sigma[R(\pi[0..n])]$. The limits in the former definitions are guaranteed to exist [34,29].
Problems. We consider the following problems for both MCs and MDPs:

- Given a target set $\mathfrak{T}$ compute $\mathit{HitPr}(u, \mathfrak{T})$ for every vertex u.
- Given a reward function R and a discount factor λ compute $\mathit{ExpDisSum}(u)$ for every vertex u.
- Given a reward function R, compute $\mathit{ExpMP}(u)$ for every vertex u.

Solving MCs [32]**.** A classical approach to the above problems for MCs is to reduce them to solving systems of linear equations. In case of hitting probabilities, we define one variable x_u for each vertex u, whose value in the solution to the system would be equal to $\mathit{HitPr}(u, \mathfrak{T})$. The system is constructed as follows:

- We add the equation $x_\mathfrak{t} = 1$ for every $\mathfrak{t} \in \mathfrak{T}$, and
- For every vertex $u \notin \mathfrak{T}$ with successors $u_1, \ldots, u_k$, we add the equation $x_u = \sum_{i=1}^{k} \delta(u)(u_i) \cdot x_{u_i}$.

If every vertex can reach a target, then it is well-known that the resulting system has a unique solution in which the value assigned to each x_u is equal to $\mathit{HitPr}(u, \mathfrak{T})$. A similar approach can be used in the case of discounted sums. We define one variable y_u per vertex u and if the successors of u are $u_1, \ldots, u_k$, then we add the equation $y_u = \sum_{i=1}^{k} \delta(u)(u_i) \cdot (R(u, u_i) + \lambda \cdot y_{u_i})$.
Primal Graphs. Let S be a system of linear equations with m equations and n unknowns (variables). The primal graph $G(S)$ of S is an undirected graph with n vertices, each corresponding to one unknown in S, in which there is an edge between two unknowns x and y iff there exists an equation in S that contains both x and y with non-zero coefficients.
Solving MDPs. There are two classical approaches to solving the above problems for MDPs. One is to reduce the problem to Linear Programming (LP) in a manner similar to the reduction from MC to linear systems [21]. The other approach is to use dynamic programming [3]. We consider a widely-used variety of dynamic programming, called *strategy iteration* or *policy iteration* [28].
Strategy Iteration (SI) [3]**.** In SI we start with an arbitrary initial strategy σ_0 and attempt to find a better strategy in each step. Formally, assume that our strategy after i iterations is σ_i. Then, we compute $\mathit{val}_i(u) = \mathit{HitPr}^{\sigma_i}(u, \mathfrak{T})$ for every vertex u. This is equivalent to computing hitting probabilities in the MC that is obtained by considering our MDP together with the strategy σ_i. We use the values $\mathit{val}_i(u)$ to obtain a better strategy σ_{i+1} as follows: for every vertex $v \in V_1$ with successors $v_1, v_2, \ldots, v_k$, we set $\sigma_{i+1}(v) = \arg\max_{v_j} \mathit{val}_i(v_j)$.

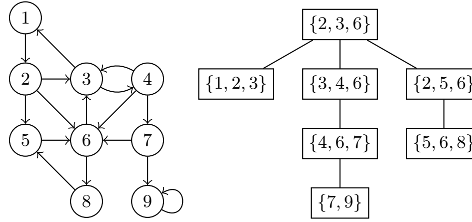

Fig. 1. A graph G (left) and a tree decomposition of G with width 2 (right).

In case of discounted sum, we let $val_i(u) = ExpDisSum^{\sigma_i}(u)$ and $\sigma_{i+1}(v) = \arg\max_{v_j} R(v, v_j) + \lambda \cdot val_i(v_j)$. We repeat these steps until we reach a point where our strategy converges. It is well-known that strategy iteration always converges to the optimal strategy σ_κ, and at that point the values val_κ will be the desired hitting probabilities/discounted sums [28,21]. Given that SI solves the classic problems above on MDPs by several calls to a procedure for solving the same problems on MCs, our runtime improvements for MCs are naturally extended to MDPs. So, in the sequel we focus on MCs.

Tree Decompositions [36]***.*** Given a directed or undirected graph $G = (V, E)$, a *tree decomposition* of G is a tree (T, E_T) such that:

- Each vertex $b \in T$ of the tree is associated with a subset $V_b \subseteq V$ of vertices of the graph. For clarity, we reserve the word "vertex" for vertices of G and use the word "bag" to refer to vertices of T. Also, we define $E_b := \{(u, v) \in E \mid u, v \in V_b\}$.
- Each vertex appears in at least one bag, i.e. $\bigcup_{b \in T} V_b = V$.
- Each edge appears in at least one bag, i.e. $\bigcup_{b \in T} E_b = E$.
- Each vertex appears in a connected subtree of T. In other words, for all $b, b', b'' \in T$, if b'' is in the unique path between b and b', then $V_b \cap V_{b'} \subseteq V_{b''}$.

Treewidth [36]***.*** The *width* of a tree decomposition is the size of its largest bag minus one, i.e. $w(T) = \max_{b \in T} |V_b| - 1$. A tree decomposition of G is called *optimal* if it has the smallest possible width. The *treewidth* $tw(G)$ of G is defined as the width of its optimal tree decomposition(s).

Computing Treewidth and Tree Decompositions. The problem of computing the treewidth of a graph is solvable in linear time when parameterized by the treewidth itself [6]. The algorithm in [6] also finds an optimal tree decomposition in linear time. Moreover, [37] proves control-flow graphs of structured programs have constant treewidth and provides a linear-time algorithm for producing the tree decomposition by a single parse of the program.

3 Algorithms for MCs with Constant Treewidth

We now consider quantitative problems on MCs. As mentioned before, our improvements carry over to MDPs using SI. We build on classical state-elimination

algorithms such as those used in [18,27]. The main novelty of our approach is that we use the tree decompositions to obtain a suitable *order* for eliminating vertices. This specific ordering significantly reduces the runtime complexity of classical state-elimination algorithms from cubic to linear.

3.1 A Simple Algorithm for Computing Hitting Probabilities

We begin by looking into the problem of computing hitting probabilities for general MCs without exploiting the treewidth. Without loss of generality, we can assume that our target set contains a single vertex. Otherwise, we add a new vertex $\mathfrak{t}$ and add edges with probability 1 from every target vertex to $\mathfrak{t}$. This will keep the hitting probabilities intact. Consider our MC $C = (V, E, \delta)$ and our target vertex $\mathfrak{t} \in V$. If there is only one vertex in the MC then there is not much to solve. We just return that $\mathit{HitPr}(\mathfrak{t}, \mathfrak{t}) = 1$. Otherwise, we take an arbitrary vertex $u \neq \mathfrak{t}$ and try to remove it from the MC to obtain a smaller MC that can in turn be solved using the same method. We should do this in a manner that does not change $\mathit{HitPr}(v, \mathfrak{t})$ for any vertex $v \neq u$. Figure 2 shows how to remove a vertex u from C in order to obtain a smaller MC $\overline{C} = (V \setminus \{u\}, \overline{E}, \overline{\delta})^{\ddagger}$. Basically, we remove u and all of its edges, and instead add new edges from every predecessor u' to every successor u''. We also update the transition function δ by setting $\overline{\delta}(u')(u'') = \delta(u')(u'') + \delta(u')(u) \cdot \delta(u)(u'')$. It is easy to verify that for every $v \neq u$, we have $\overline{\mathit{HitPr}}(v, \mathfrak{t}) = \mathit{HitPr}(v, \mathfrak{t})$. Hence, we can compute hitting probabilities for every vertex $v \neq u$ in $\overline{C}$ instead of C. Finally, if $u_1, u_2, \ldots, u_k$ are the successors of u in C, we know that $\mathit{HitPr}(u, \mathfrak{t}) = \sum_{i=1}^{k} \delta(u)(u_i) \cdot \mathit{HitPr}(u_i, \mathfrak{t}) = \sum_{i=1}^{k} \delta(u)(u_i) \cdot \overline{\mathit{HitPr}}(u_i, \mathfrak{t})$. Hence, we can easily compute the hitting probability for u using this formula. A pseudocode of is available in [1].

A special case arises when there is a self-loop transition from u to u. If $\delta(u)(u) = 1$, i.e. u is an absorbing trap, then we can simply remove u, noting that $\mathit{HitPr}(u, \mathfrak{t}) = 0$. On the other hand if $0 < \delta(u)(u) < 1$, then we should distribute $\delta(u)(u)$ proportionately among the other successors of u because staying for a finite number of steps in the same vertex u does not change the hitting property of a path, and the probability of staying at u forever is 0.

Removing each vertex can take at most $O(n^2)$ time, given that it has $O(n)$ predecessors and successors. We should remove $n - 1$ vertices, leading to a total runtime of $O(n^3)$, which is worse than the reduction to system of linear equations and then applying Gaussian elimination. However, the runtime can be significantly improved if we remove vertices in an order that guarantees every vertex has a low degree upon removal.

3.2 Computing Hitting Probabilities in Constant Treewidth

The main idea behind our algorithm is simple: we take the algorithm from the previous section and use tree decompositions to obtain an ordering for the removal of vertices. Given that we can choose any bag in T as the root, without

‡We always use $\overline{C}$ to denote an MC that is obtained from C by removing one vertex. We apply this rule across our notation, e.g. $\overline{\delta}$ is the respective transition function.

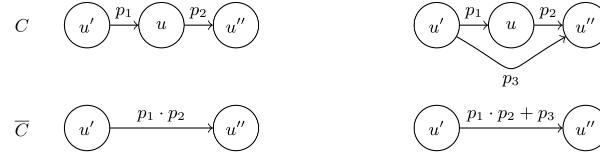

Fig. 2. Removing a vertex u. The vertex u' is a predecessor of u and u'' is one of its successors. The left side shows the changes when there is no edge from u' to u'' and the right side shows the other case, where $(u', u'') \in E$. Edge labels are δ values.

loss of generality, we assume that the target vertex $\mathfrak{t}$ is in the root bag[§]. We base our approach on the following lemmas:

Lemma 1. *Let $l \in T$ be a* leaf *bag of the tree decomposition (T, E_T) of our MC C, and let $\bar{l}$ be the parent of l. If $V_l \subseteq V_{\bar{l}}$, then $(T \setminus \{l\}, E_T \setminus \{(\bar{l}, l)\})$ is also a valid tree decomposition for C.*

Proof. We just need to check that all the required properties of a tree decomposition hold after removal of l. Given that $V_l \subseteq V_{\bar{l}}$, any vertex that appears in l is also in $\bar{l}$ and hence removal of l does not cause any vertex to be unrepresented in the tree decomposition. The same applies to edges. Moreover, removing a *leaf* bag cannot disconnect the previously-connected set of bags containing a vertex.

Lemma 2. *Let $l \in T$ be a bag of the tree decomposition (T, E_T) and assume that the vertex $u \in V$ only appears in V_l, i.e. it does not appear in the vertex set of any other bag. Then, u has at most $|V_l|$ predecessors/successors in C.*

Proof. If u' is a predecessor/successor of u, then there is an edge between them. By definition, a tree decomposition should cover every edge. Hence, there should be a bag b such that $u, u' \in V_b$. By assumption, u only appears in V_l. Hence, every predecessor/successor u' must also appear in V_l.

The Algorithm. The above lemmas provide a convenient order for removing vertices. At each step, we choose an arbitrary *leaf* bag l. If there is a vertex u that appears *only* in V_l, then we eliminate u as in Figure 2. In this case, Lemma 2 guarantees that u has $O(t)$ predecessors and successors. Otherwise, $V_l \subseteq V_{\bar{l}}$ (recall that each vertex appears in a connected subtree) and we can remove l from our tree decomposition according to Lemma 1. See [1] for a pseudocode.
Example. Consider the graph and tree decomposition in Figure 1 with an arbitrary transition probability function δ and target vertex $\mathfrak{t} = 6$. On this example, our algorithm would first choose an arbitrary leaf bag, say $\{7, 9\}$ and then realize that 9 has only appeared in this bag. Hence it removes vertex 9 from the MC

[§] If $|\mathfrak{T}| \geq 2$, we use the same technique as in the previous section to have only one target $\mathfrak{t}$. To keep the tree decomposition valid, we add $\mathfrak{t}$ to every bag.

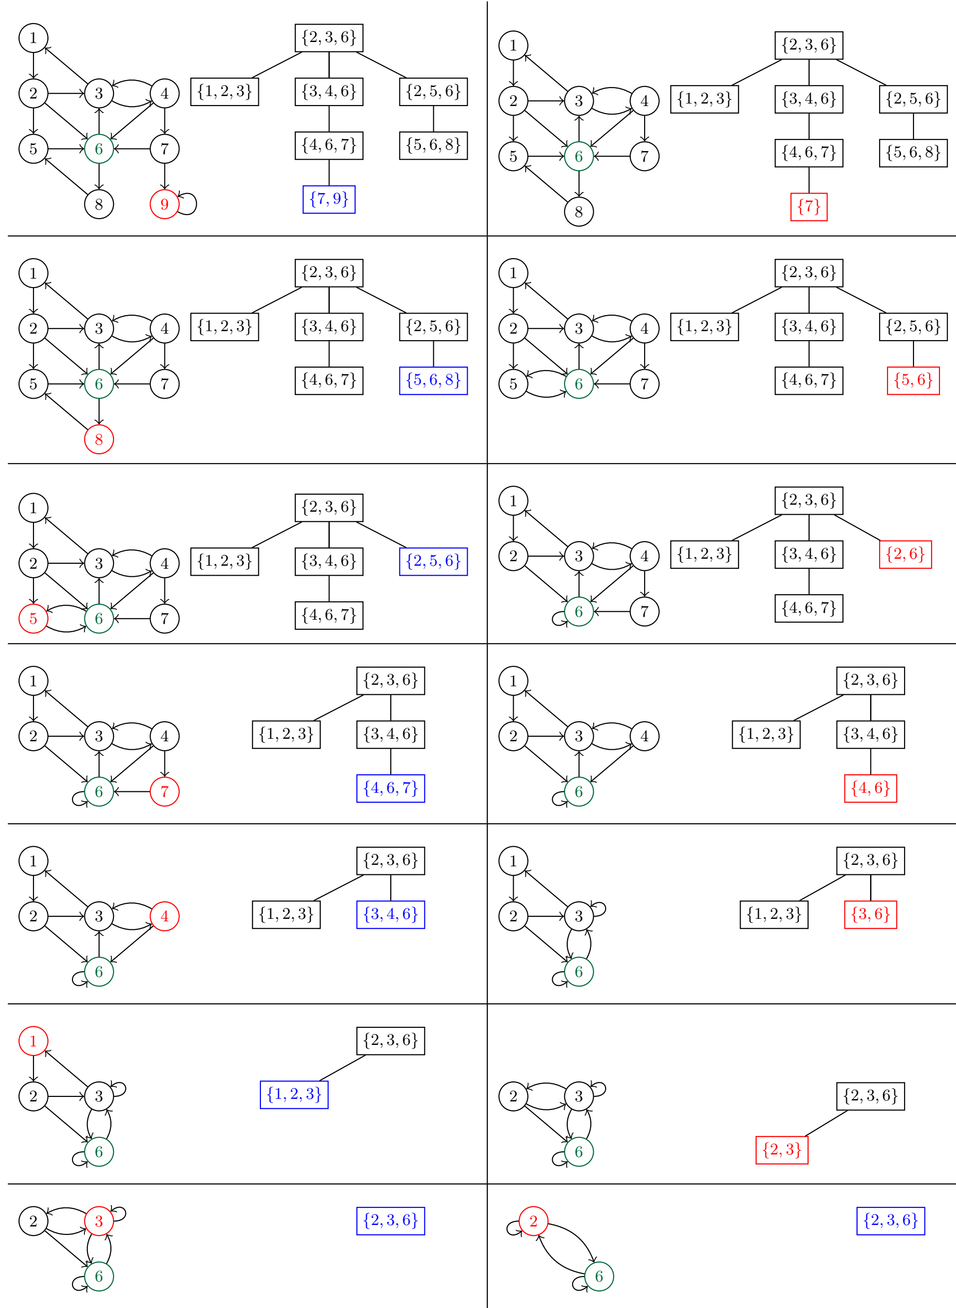

Fig. 3. The Steps Taken by our Algorithm on the Graph and Tree Decomposition in Figure 1. The target vertex $\mathfrak{t} = 6$ is shown in green. At each step the vertex/bag that is being removed is shown in red. An active bag whose vertices, but not itself, are considered for removal is shown in blue. After removing vertex 2, the graph has only one vertex and the base case of the algorithm is run.

using the same procedure as in the previous section. In the next iteration, it chooses the bag $\{7\}$ and realizes that the set of vertices in this bag is a subset of vertices that appear in its parent. Hence, it removes this unnecessary bag. The algorithm continues similarly, until only the target vertex 6 remains, at which point the problem is trivial. Figure 3 shows all the steps of our algorithm. Note that because the width of our tree decomposition is 2, at each step when we are removing a vertex u, it has at most 3 neighbors (counting itself).

Note that throughout this algorithm the tree decomposition remains valid, because we are only adding edges between vertices that are already in the same leaf bag l. Given that we remove at most $O(n)$ bags and $n-1$ vertices and that removing each vertex takes only $O(t^2)$, the total runtime is $O(n \cdot t^2)$.

Theorem 1. *Given an MC with n vertices and treewidth t and an optimal tree decomposition of the MC, our algorithm computes hitting probabilities from every vertex to a designated target set in $O(n \cdot t^2)$.*

3.3 Computing Expected Discounted Sums in Constant Treewidth

We use a similar approach for handling the discounted sum problem. The only difference is in how a vertex is removed. Given an MC $C = (V, E, \delta)$, a tree decomposition (T, E_T) of C, a reward function $R : E \to \mathbb{R}$ and a discount factor $\lambda \in (0, 1)$, we first add a new vertex called $\hat{\mathbf{1}}$ to the MC. The vertex $\hat{\mathbf{1}}$ is disjoint from all other vertices and only has a single self-loop with probability 1 and reward $1 - \lambda$. In other words, we define $\delta(\hat{\mathbf{1}})(\hat{\mathbf{1}}) = 1$ and $R(\hat{\mathbf{1}}, \hat{\mathbf{1}}) = 1 - \lambda$. We also add $\hat{\mathbf{1}}$ to the vertex set of every bag. The reason behind this gadget is that we have $\mathit{ExpDisSum}(\hat{\mathbf{1}}) = (1 - \lambda) \cdot (1 + \lambda + \lambda^2 + \ldots) = 1$.

In our algorithm, the requirement that for all u, v we should have $0 \leq \delta(u)(v) \leq 1$ is unnecessary and becomes untenable, too. Therefore, we allow $\delta(u)(v)$ to have any real value, and use the linear system interpretation of C as in Section 2, i.e. instead of considering C as an MC, we consider it to be a representation of the linear system S_C defined as follows:

- For every vertex $u \in V$, the system S_C contains one unknown y_u, and
- For every vertex $u \in V$, whose successors are $u_1, u_2, \ldots, u_k$, the system S_C contains an equation $\mathfrak{e}_u := y_u = \sum_{i=1}^{k} \delta(u)(u_i) \cdot (R(u, u_i) + \lambda \cdot y_{u_i})$.

As mentioned in Section 2, in the solution to S_C, the value assigned to the unknown y_u is equal to $\mathit{ExpDisSum}(u)$ in the MC C. However, the definition above does not depend on the fact that C is an MC and can also be applied if δ has arbitrary real values.

Now suppose that we want to remove a vertex $u \neq \hat{\mathbf{1}}$ with successors $u_1, \ldots, u_k$ from C. This is equivalent to removing y_u from S_C without changing the values of other unknowns in the solution. Given that we have $y_u = \sum_{i=1}^{k} \delta(u)(u_i) \cdot (R(u, u_i) + \lambda \cdot y_{u_i})$, we can simply replace every occurrence of y_u in other equations with the right-hand-side expression of this equation. If $u' \neq u$ is a predecessor of u, then we have $y_{u'} = A + \delta(u')(u) \cdot (R(u', u) + \lambda \cdot y_u)$, where A is an expression that depends on other successors of u'. We can rewrite this equation as $y_{u'} = A + \delta(u')(u) \cdot R(u', u) + \sum_{i=1}^{k} \delta(u')(u) \cdot \delta(u)(u_i) \cdot \lambda \cdot (R(u, u_i) + \lambda \cdot y_{u_i})$.

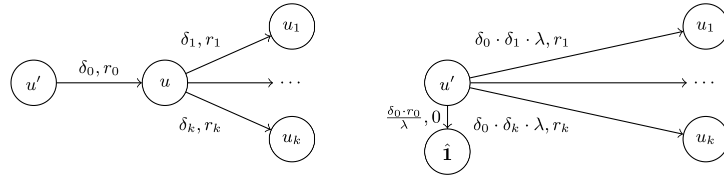

Fig. 4. Removing u from C (left) to obtain $\overline{C}$ (right). The vertex u' is a predecessor of u and $u_1, \ldots, u_k$ are its successors. Each edge is labelled with its δ and R values.

This is equivalent to obtaining a new $\overline{C}$ from C by removing the vertex u and adding the following edges from every predecessor u' of u:

- An edge $(u', \hat{\mathbf{1}})$, such that $R(u', \hat{\mathbf{1}}) = 0$ and $\delta(u')(1) = \frac{1}{\lambda} \cdot (\delta(u')(u) \cdot R(u', u))$,
- An edge (u', u_i) to every successor u_i of u, such that $R(u', u_i) = R(u, u_i)$ and $\delta(u')(u_i) = \delta(u')(u) \cdot \delta(u)(u_i) \cdot \lambda$.

This construction is shown in Figure 4. Using this construction, the value of y_v remains the same in solutions of S_C and $S_{\overline{C}}$. There are a few special cases, e.g. when the graph has parallel edges or self-loops. See [1] for details.

As in the previous section, we can solve the problem on the smaller $\overline{C}$ and then use the equation $\mathfrak{e}_u$ to compute the value of y_u in the solution to S_C. This algorithm's runtime can be analyzed exactly as before. We have to remove n vertices and each removal takes $O(n^2)$ for a total runtime of $O(n^3)$. To obtain a better algorithm that exploits tree decompositions, we can use the exact same removal order as in the previous section, leading to the same runtime, i.e. $O(n \cdot t^2)$. Note that we have added $\hat{\mathbf{1}}$ to the associated vertex set of every bag, so the tree decomposition always remains valid throughout our algorithm.

Theorem 2. *Given an MC with n vertices and treewidth t and an optimal tree decomposition of the MC, the algorithm described in this section computes expected discounted sums from every vertex of the MC in $O(n \cdot t^2)$.*

3.4 Systems of Equations with Constant-Treewidth Primal Graphs

The ideas used in the previous section can be extended to obtain faster algorithms for solving any linear system whose primal graph has a small treewidth. However, new subtleties arise, given that general linear systems might have no solution or infinitely many solutions. In contrast, the systems S_C discussed in the previous section were guaranteed to have a unique solution. We consider a system S of m linear equations over n real unknowns as input, and assume that its primal graph $G(S)$ has treewidth t. Our algorithm for solving S is similar to our previous algorithms, and is actually what most students are taught in junior high school. We take an arbitrary unknown x and choose an arbitrary equation $\mathfrak{e}$ in which x appears with a non-zero coefficient. We then rewrite $\mathfrak{e}$ as $x = R_x$,

where R_x is a linear expression based on other unknowns. Finally, we replace every occurrence of x in other equations with R_x and solve the resulting smaller system $\overline{S}$. If $\overline{S}$ has no solutions or inifinitely many solutions, then so does S. Otherwise, we evaluate R_x in the solution of $\overline{S}$ to get the solution value for x. Using this algorithm, we have to remove $O(n)$ unknowns. When removing x, we might have to replace an expression of size $O(n)$, i.e. R_x, in $O(m)$ potential other equations where x has appeared. Hence, the overall runtime is $O(n^2 \cdot m)$.

Given a tree decomposition (T, E_T) of the primal graph $G(S)$, we choose the unknows in the usual order, i.e. we always choose an unknown x that appears only in a leaf bag. If x does not appear in any equations, then we can simply remove it and then S is satisfiable iff $\overline{S}$ is satisfiable. Moreover, if S is satisfiable, then it has infinitely many solutions, given that x is not restricted. Otherwise, there is an equation $\mathfrak{e}$ in which x appears with non-zero coefficient, and hence we can rewrite this equation as $x = R_x$. Note that x has $O(t)$ neighbors in $G(S)$, given that it only appears in a leaf bag and all of its neighbors should also appear in the same bag, hence the length of R_x is $O(t)$, too. The problem is that x might have appeared in any of the other $O(m)$ equations. Hence, replacing it with R_x in every equation will lead to a runtime of $O(m \cdot t)$. We repeat this for every unknown, so our total runtime is $O(n \cdot m \cdot t)$, which is not linear.

The crucial observation is that while x might have appeared in as many as m equations, not all of them are linearly independent. Let $\mathfrak{E}_x$ be the set of equations containing x and l be the leaf bag in which x appears and assume that $V_l = \{x, y_1, \ldots, y_{k-1}\}$. Then the only unknowns that can appear together with x in an equation are $y_1, \ldots, y_{k-1}$. In other words, all equations in $\mathfrak{E}_x$ are over V_l. Hence, we can apply the Gram-Schmidt process on $\mathfrak{E}_x$ to remove the unnecessary equations and only keep at most k equations that form an orthogonal basis (or alternatively realize that the system is unsatisfiable). Given that we are operating in dimension $k = O(t)$, this will take $O(t^2 \cdot |\mathfrak{E}_x|)$ time. See [1] for a pseudocode. As in previous algorithms, our approach always keeps the tree decomposition valid. Moreover, as argued above, its runtime is $O((n+m) \cdot t^2)$.

Theorem 3. *Given a system of m linear equations over n unknowns, its primal graph, and a tree decomposition of the primal graph with width t, our algorithm solves the system in time $O((n+m) \cdot t^2)$.*

The algorithm can easily be extended to find a basis for the solution set. Moreover, it can also be combined with the algorithms in the previous sections to solve the mean-payoff objective, hence we have:

Theorem 4 (Proof and Details in [1]). *Given an MC with n vertices and treewidth t and an optimal tree decomposition, expected mean payoffs from every vertex can be computed in $O(n \cdot t^2)$.*

4 Experimental Results

We now report on a C/C++ implementation of our algorithms and provide a performance comparison with previous approaches. See [1] for details.

Fig. 5. Experimental Results for Computing Hitting Probabilities in MCs.

Compared Approaches. We consider the hitting probability and discounted sum problems for MCs and MDPs. In the case of MCs, we directly use our algorithms from Section 3.2 and Section 3.3. For MDPs, we use strategy iteration, where we use the above algorithms for the strategy evaluation step in each iteration. We compare our approach with the following alternatives:

- *Classical Approaches.* In case of MCs, we compare against an implementation of Gaussian elimination (Gauss). For MDPs, we consider our own implementation of value iteration (VI) and strategy iteration (SI).
- *Numerical and Industrial Optimizers.* We use Matlab [31] and Gurobi [25] to solve systems of linear equalities corresponding to MCs. For MDPs, we use Matlab [31], Gurobi [25] and lpsolve [4] to handle the corresponding LPs.
- *Probabilistic Model Checkers.* The model checkers Storm [19] and Prism [30] have standard procedures for computing hitting probabilities, but not for discounted sums. Hence, we compare with them on hitting probabilities only.

Benchmarks. We used CFGs of the 40 Java programs from the DaCapo suite [5] as our benchmarks. They have between 33 and 103918 vertices and transitions. To obtain MDPs, we randomly (with probability 1/2) turned each vertex into a Player 1 or a probabilistic one. We assigned random probabilities to each outgoing edge of a probabilistic vertex. To obtain MCs, we did the same, except that all vertices are probabilistic. For the hitting probabilities problem, we chose one random vertex from each connected component of the control flow graphs as a target. In case of discounted sum, we uniformly chose a discount factor between 0 and 1 for each instance, and also assigned random integral rewards between -1000 to 1000 to each edge. Finally, we used JTDec [9] to compute tree decompositions. In each case the width of the obtained decomposition was no more than 9. The time and memory used for computing a tree decomposition

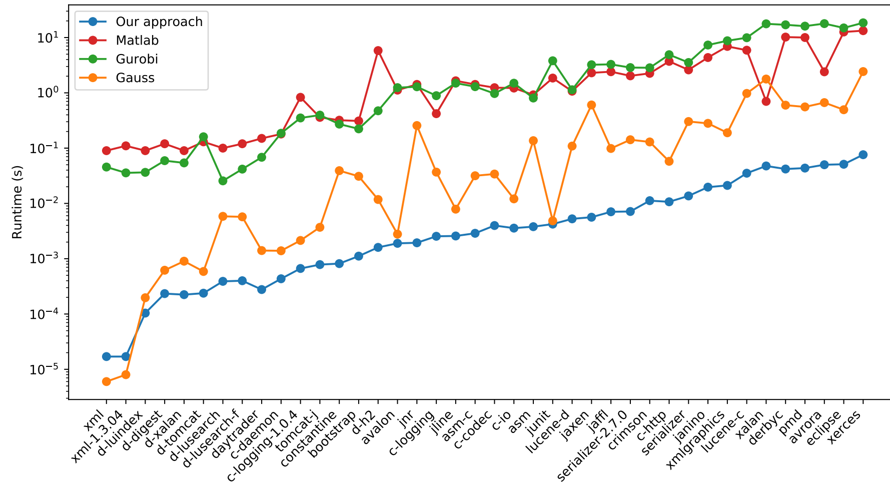

Fig. 6. Experimental Results for Computing Expected Discounted Sums in MCs.

are negligible, given that it is obtained by a single pass over the program. See [1] for details of the benchmarks and the motivation for using them.

Results. The runtimes for computing the values of hitting probabilities and discounted sums are shown in Figures 5–8. The benchmarks are on the x-axes and ordered by their size. Note that the y-axes are in *logarithmic scale*. For example, Figure 5 shows results for computing hitting probabilities in MCs, where Prism is the slowest tool by far, while our approach comfortably beats every other method. The gap is more apparent in MDPs (Figures 7–8). Overall, we see that our new algorithms consistently outperform both classical approaches like VI and SI, and highly optimized solvers and model checkers like Gurobi, Prism and Storm, by one or more orders of magnitude. Hence, the theoretical improvements are also realized in practice. See [1] for raw numbers. It is also noteworthy that the outputs of the different approaches agreed with each other within an error range of 10^{-5}. Finally, note that our approach is only applicable to instances with small treewidth, such as CFGs of structured programs. For MCs/MDPs with arbitrary treewidth, the problem of computing an optimal tree decomposition is NP-hard [6].

References

1. Asadi, A., et al.: Faster algorithms for quantitative analysis of Markov chains and Markov decision processes with small treewidth. arXiv preprint:2004.08828 (2020)
2. Ashok, P., Chatterjee, K., Daca, P., Křetínský, J., Meggendorfer, T.: Value iteration for long-run average reward in Markov decision processes. In: CAV (2017)
3. Bellman, R.: A Markovian decision process. Journal of Mathematics and Mechanics (1957)
4. Berkelaar, M., Eikland, K., Notebaert, P.: lpsolve Linear Programming system

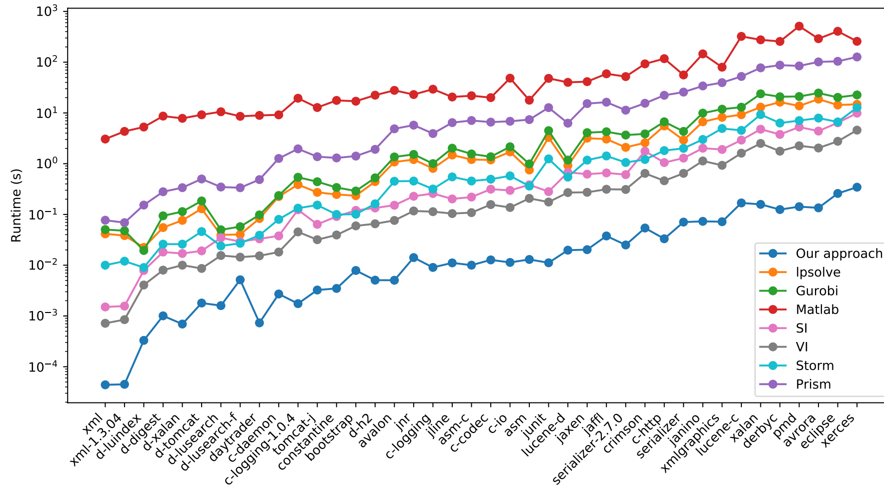

Fig. 7. Experimental Results for Computing Hitting Probabilities in MDPs.

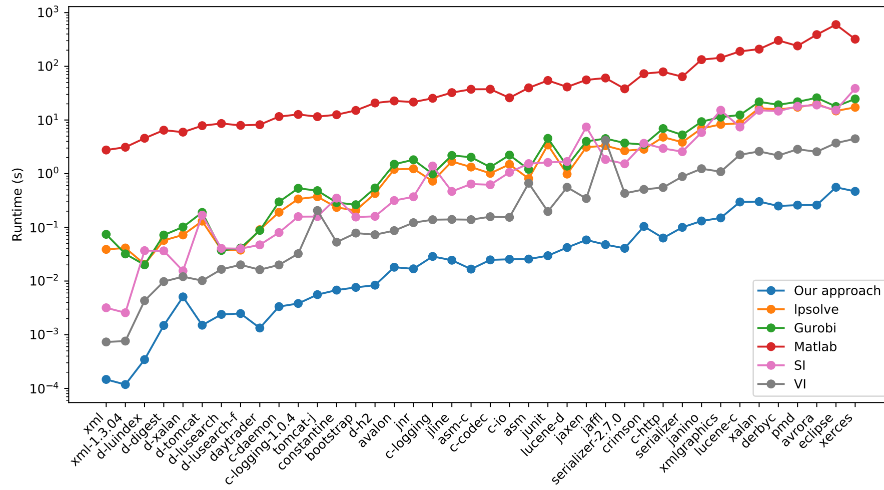

Fig. 8. Experimental Results for Computing Expected Discounted Sums in MDPs.

5. Blackburn, S.M., et al.: The DaCapo benchmarks: Java benchmarking development and analysis. In: OOPSLA (2006)
6. Bodlaender, H.L.: A linear-time algorithm for finding tree-decompositions of small treewidth. SIAM Journal on Computing **25**(6) (1996)
7. Chatterjee, K., Choudhary, B., Pavlogiannis, A.: Optimal dyck reachability for data-dependence and alias analysis. In: POPL (2017)
8. Chatterjee, K., Goharshady, A.K., Ibsen-Jensen, R., Pavlogiannis, A.: Optimal and perfectly parallel algorithms for on-demand data-flow analysis. In: ESOP (2020)

9. Chatterjee, K., Goharshady, A.K., Pavlogiannis, A.: JTDec: A tool for tree decompositions in soot. In: ATVA (2017)
10. Chatterjee, K., Henzinger, T.A.: Value iteration. In: Model Checking (2008)
11. Chatterjee, K., Henzinger, T.A., Jobstmann, B., Singh, R.: Measuring and synthesizing systems in probabilistic environments. In: CAV (2010)
12. Chatterjee, K., Łącki, J.: Faster algorithms for Markov decision processes with low treewidth. In: CAV (2013)
13. Chatterjee, K., et al.: Algorithms for algebraic path properties in concurrent systems of constant treewidth components. TOPLAS **40**(3) (2018)
14. Chatterjee, K., et al.: Symbolic algorithms for graphs and Markov decision processes with fairness objectives. In: CAV (2018)
15. Chatterjee, K., et al.: Efficient parameterized algorithms for data packing. In: POPL (2019)
16. Chatterjee, K., et al.: Faster algorithms for dynamic algebraic queries in basic RSMs with constant treewidth. TOPLAS (2019)
17. Chatterjee, K., et al.: The treewidth of smart contracts. In: SAC (2019)
18. Daws, C.: Symbolic and parametric model checking of discrete-time Markov chains. In: ICTAC (2004)
19. Dehnert, C., et al.: A storm is coming: A modern probabilistic model checker. In: CAV (2017)
20. Fearnley, J.: Exponential lower bounds for policy iteration. In: ICALP (2010)
21. Feinberg, E.A.: Handbook of Markov decision processes. Springer (2012)
22. Ferrara, A., Pan, G., Vardi, M.Y.: Treewidth in verification: Local vs. global. In: LPAR (2005)
23. Fomin, F.V., et al.: Fully polynomial-time parameterized computations for graphs and matrices of low treewidth. TALG **14**(3) (2018)
24. Goharshady, A.K., Mohammadi, F.: An efficient algorithm for computing network reliability in small treewidth. Reliability Engineering & System Safety **193** (2020)
25. Gurobi Optimization, L.: Gurobi optimizer (2019), `http://www.gurobi.com`
26. Gustedt, J., et al.: The treewidth of Java programs. In: ALENEX (2002)
27. Hahn, E.M., Hermanns, H., Wachter, B., Zhang, L.: Param: A model checker for parametric markov models. In: CAV (2010)
28. Howard, R.A.: Dynamic programming and Markov processes. (1960)
29. Křetínský, J., Meggendorfer, T.: Efficient strategy iteration for mean payoff in Markov decision processes. In: ATVA (2017)
30. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: CAV (2011)
31. MATLAB: The MathWorks Inc. (2018)
32. Norris, J.R.: Markov chains. Cambridge University Press (1998)
33. Obdržálek, J.: Fast mu-calculus model checking when tree-width is bounded. In: CAV (2003)
34. Puterman, M.L.: Markov Decision Processes. Wiley (2014)
35. Quatmann, T., Katoen, J.P.: Sound value iteration. In: CAV (2018)
36. Robertson, N., Seymour, P.D.: Graph minors. iii. planar tree-width. Journal of Combinatorial Theory, Series B **36**(1) (1984)
37. Thorup, M.: All structured programs have small tree width and good register allocation. Information and Computation **142**(2) (1998)