Symbolic algorithms for qualitative analysis of Markov decision processes with Büchi objectives

**Krishnendu Chatterjee · Monika Henzinger ·
Manas Joglekar · Nisarg Shah**

Abstract We consider Markov decision processes (MDPs) with Büchi (liveness) objectives. We consider the problem of computing the set of *almost-sure* winning states from where the objective can be ensured with probability 1. Our contributions are as follows: First, we present the first subquadratic symbolic algorithm to compute the almost-sure winning set for MDPs with Büchi objectives; our algorithm takes $O(n \cdot \sqrt{m})$ symbolic steps as compared to the previous known algorithm that takes $O(n^2)$ symbolic steps, where n is the number of states and m is the number of edges of the MDP. In practice MDPs have constant out-degree, and then our symbolic algorithm takes $O(n \cdot \sqrt{n})$ symbolic steps, as compared to the previous known $O(n^2)$ symbolic steps algorithm. Second, we present a new algorithm, namely *win-lose* algorithm, with the following two properties: (a) the algorithm iteratively computes subsets of the almost-sure winning set and its complement, as compared to all previous algorithms that discover the almost-sure winning set upon termination; and (b) requires $O(n \cdot \sqrt{K})$ symbolic steps, where K is the maximal number of edges of strongly connected components (scc's) of the MDP. The win-lose algorithm requires symbolic computation of scc's. Third, we improve the algorithm for symbolic scc computation; the previous known algorithm takes linear symbolic steps, and our new algorithm improves

K. Chatterjee (✉)
IST Austria, Klosterneuburg, Austria
e-mail: krishnendu.chatterjee@ist.ac.at

M. Henzinger
University of Vienna, Vienna, Austria

M. Joglekar
Stanford University, Palo Alto, USA

N. Shah
Carnegie Mellon University, Pittsburgh, USA

the constants associated with the linear number of steps. In the worst case the previous known algorithm takes $5 \cdot n$ symbolic steps, whereas our new algorithm takes $4 \cdot n$ symbolic steps.

1 Introduction

Markov decision processes   The standard model of systems in verification of probabilistic systems is *Markov decision processes (MDPs)* that exhibit both probabilistic and nondeterministic behavior [13]. MDPs have been used to model and solve control problems for stochastic systems [11]: there, nondeterminism represents the freedom of the controller to choose a control action, while the probabilistic component of the behavior describes the system response to control actions. MDPs have also been adopted as models for concurrent probabilistic systems [7], probabilistic systems operating in open environments [19], and under-specified probabilistic systems [2]. A *specification* describes the set of desired behaviors of the system, which in the verification and control of stochastic systems is typically an ω-regular set of paths. The class of ω-regular languages extends classical regular languages to infinite strings, and provides a robust specification language to express all commonly used specifications, such as safety, liveness, fairness, etc. [24]. Parity objectives are a canonical way to define such ω-regular specifications. Thus MDPs with parity objectives provide the theoretical framework to study problems such as the verification and control of stochastic systems.

Qualitative and quantitative analysis   The analysis of MDPs with parity objectives can be classified into qualitative and quantitative analysis. Given an MDP with parity objective, the *qualitative analysis* asks for the computation of the set of states from where the parity objective can be ensured with probability 1 (almost-sure winning). The more general *quantitative analysis* asks for the computation of the maximal probability at each state with which the controller can satisfy the parity objective.

Importance of qualitative analysis   The qualitative analysis of MDPs is an important problem in verification that is of interest irrespective of the quantitative analysis problem. There are many applications where we need to know whether the correct behavior arises with probability 1. For instance, when analyzing a randomized embedded scheduler, we are interested in whether every thread progresses with probability 1 [9]. Even in settings where it suffices to satisfy certain specifications with probability $p < 1$, the correct choice of p is a challenging problem, due to the simplifications introduced during modeling. For example, in the analysis of randomized distributed algorithms it is quite common to require correctness with probability 1 (see, e.g., [16, 17, 23]). Furthermore, in contrast to quantitative analysis, qualitative analysis is robust to numerical perturbations and modeling errors in the transition probabilities, and consequently the algorithms for qualitative analysis are combinatorial. Finally, for MDPs with parity objectives, the best known algorithms and all algorithms used in practice first perform the qualitative analysis, and then perform a quantitative analysis on the result of the qualitative analysis [6–8]. Thus qualitative analysis of MDPs with parity objectives is one of the most fundamental and core problems in verification of probabilistic systems. One of the key challenges in probabilistic verification is to obtain efficient and symbolic algorithms for qualitative analysis of MDPs with parity objectives, as symbolic algorithms allow to handle MDPs with a large state space.

Previous results   The qualitative analysis of MDPs with parity objectives is achieved by iteratively applying solutions of the qualitative analysis of MDPs with Büchi objectives [6–8]. The qualitative analysis of an MDP with a parity objective with d priorities can be achieved by $O(d)$ calls to an algorithm for qualitative analysis of MDPs with Büchi objectives, and hence we focus on the qualitative analysis of MDPs with Büchi objectives. The classical algorithm for qualitative analysis of MDPs with Büchi objectives works in $O(n \cdot m)$ time, where n is the number of states, and m is the number of edges of the MDP [7, 8]. The classical algorithm can be implemented symbolically, and it takes at most $O(n^2)$ symbolic steps. An improved algorithm for the problem was given in [5] that works in $O(m \cdot \sqrt{m})$ time. The algorithm of [5] crucially depends on maintaining the same number of edges in certain forward searches. Thus the algorithm needs to explore edges of the graph explicitly and is inherently non-symbolic. A recent $O(m \cdot n^{2/3})$ time algorithm for the problem was given in [4]; however the algorithm requires the dynamic-tree data structure of Sleator-Tarjan [20], and such data structures cannot be implemented symbolically. In the literature, there is no symbolic subquadratic algorithm for qualitative analysis of MDPs with Büchi objectives.

Our contribution   In this work our main contributions are as follows.

1. We present a new and simpler subquadratic algorithm for qualitative analysis of MDPs with Büchi objectives that runs in $O(m \cdot \sqrt{m})$ time, and show that the algorithm can be implemented symbolically. The symbolic algorithm takes at most $O(n \cdot \sqrt{m})$ symbolic steps, and thus we obtain the first symbolic subquadratic algorithm. In practice, MDPs often have constant out-degree: for example, see [10] for MDPs with large state space but constant number of actions, or [11, 18] for examples from inventory management where MDPs have constant number of actions (the number of actions corresponds to the out-degree of MDPs). For MDPs with constant out-degree our new symbolic algorithm takes $O(n \cdot \sqrt{n})$ symbolic steps, as compared to $O(n^2)$ symbolic steps of the previous best known algorithm.

2. All previous algorithms for the qualitative analysis of MDPs with Büchi objectives iteratively discover states that are guaranteed to be not almost-sure winning, and only when the algorithm terminates the almost-sure winning set is discovered. We present a new algorithm (namely *win-lose* algorithm) that iteratively discovers both states in the almost-sure winning set and its complement. Thus if the problem is to decide whether a given state s is almost-sure winning, and the state s is almost-sure winning, then the win-lose algorithm can stop at an intermediate iteration unlike all the previous algorithms. Our algorithm works in time $O(\sqrt{K_E} \cdot m)$ time, where K_E is the maximal number of edges of any scc of the MDP (in this paper we write scc for maximal strongly connected component). We also show that the win-lose algorithm can be implemented symbolically, and it takes at most $O(\sqrt{K_E} \cdot n)$ symbolic steps.

3. Our win-lose algorithm requires to compute the scc decomposition of a graph in $O(n)$ symbolic steps. The scc decomposition problem is one of the most fundamental problem in the algorithmic study of graph problems. The symbolic scc decomposition problem has many other applications in verification: for example, checking emptiness of ω-automata, and bad-cycle detection problems in model checking, see [3] for other applications. An $O(n \cdot \log n)$ symbolic step algorithm for scc decomposition was presented in [3], and the algorithm was improved in [12]. The algorithm of [12] is a linear symbolic step scc decomposition algorithm that requires at most $\min\{5 \cdot n, 5 \cdot D \cdot N + N\}$ symbolic steps, where D is the diameter of the graph, and N is the number of scc's of the graph. We present an improved version of the symbolic scc decomposition algorithm. Our algorithm improves the constants of the number of the linear symbolic steps. Our algorithm requires

at most $\min\{3 \cdot n + N, 5 \cdot D^* + N\}$ symbolic steps, where D^* is the sum of the diameters of the scc's of the graph. Thus, in the worst case, the algorithm of [12] requires $5 \cdot n$ symbolic steps, whereas our algorithm requires $4 \cdot n$ symbolic steps. Moreover, the number of symbolic steps of our algorithm is always bounded by the number of symbolic steps of the algorithm of [12] (i.e. our algorithm is never worse).

Our experimental results show that our new algorithms perform better than the previous known algorithms both for qualitative analysis of MDPs with Büchi objectives and symbolic scc computation.

2 Definitions

Markov decision processes (MDPs)   A *Markov decision process (MDP)* $G = ((S, E), (S_1, S_P), \delta)$ consists of a directed graph (S, E), a partition (S_1, S_P) of the *finite* set S of states, and a probabilistic transition function $\delta : S_P \to \mathcal{D}(S)$, where $\mathcal{D}(S)$ denotes the set of probability distributions over the state space S. The states in S_1 are the *player*-1 states, where player 1 decides the successor state, and the states in S_P are the *probabilistic (or random)* states, where the successor state is chosen according to the probabilistic transition function δ. We assume that for $s \in S_P$ and $t \in S$, we have $(s, t) \in E$ iff $\delta(s)(t) > 0$, and we often write $\delta(s, t)$ for $\delta(s)(t)$. For a state $s \in S$, we write $E(s)$ to denote the set $\{t \in S \mid (s, t) \in E\}$ of possible successors. For technical convenience we assume that every state in the graph (S, E) has at least one outgoing edge, i.e., $E(s) \neq \emptyset$ for all $s \in S$. We will denote by $n = |S|$ and $m = |E|$ the size of the state space and the number of transitions (or edges), respectively.

Plays and strategies   An infinite path, or a *play*, of an MDP G is an infinite sequence $\omega = \langle s_0, s_1, s_2, \ldots \rangle$ of states such that $(s_k, s_{k+1}) \in E$ for all $k \in \mathbb{N}$. We write Ω for the set of all plays, and for a state $s \in S$, we write $\Omega_s \subseteq \Omega$ for the set of plays that start from the state s. A *strategy* for player 1 is a function $\sigma : S^* \cdot S_1 \to \mathcal{D}(S)$ that chooses the probability distribution over the successor states for all finite sequences $\mathbf{w} \in S^* \cdot S_1$ of states ending in a player-1 state (the sequence represents a prefix of a play). A strategy must respect the edge relation: for all $\mathbf{w} \in S^*$ and $s \in S_1$, if $\sigma(\mathbf{w} \cdot s)(t) > 0$, then $t \in E(s)$. A strategy is *deterministic (pure)* if it chooses a unique successor for all histories (rather than a probability distribution), otherwise it is *randomized*. Player 1 follows the strategy σ if in each player-1 move, given that the current history of the play is $\mathbf{w} \in S^* \cdot S_1$, she chooses the next state according to $\sigma(\mathbf{w})$. We denote by Σ the set of all strategies for player 1. A *memoryless* player-1 strategy does not depend on the history of the play but only on the current state; i.e., for all $\mathbf{w}, \mathbf{w}' \in S^*$ and for all $s \in S_1$ we have $\sigma(\mathbf{w} \cdot s) = \sigma(\mathbf{w}' \cdot s)$. A memoryless strategy can be represented as a function $\sigma : S_1 \to \mathcal{D}(S)$, and a pure memoryless strategy can be represented as $\sigma : S_1 \to S$.

Once a starting state $s \in S$ and a strategy $\sigma \in \Sigma$ is fixed, the outcome of the MDP is a random walk ω_s^σ for which the probabilities of events are uniquely defined, where an *event* $\mathcal{A} \subseteq \Omega$ is a measurable set of plays. For a state $s \in S$ and an event $\mathcal{A} \subseteq \Omega$, we write $\mathrm{Pr}_s^\sigma(\mathcal{A})$ for the probability that a play belongs to $\mathcal{A}$ if the game starts from the state s and player 1 follows the strategy σ.

Objectives   We specify *objectives* for player 1 by providing a set of *winning* plays $\Phi \subseteq \Omega$. We say that a play ω *satisfies* the objective Φ if $\omega \in \Phi$. We consider ω-*regular objectives* [24], specified as parity conditions. We also consider the special case of Büchi objectives.

- *Büchi objectives.* Let T be a set of target states. For a play $\omega = \langle s_0, s_1, \ldots \rangle \in \Omega$, we define $\mathrm{Inf}(\omega) = \{s \in S \mid s_k = s \text{ for infinitely many } k\}$ to be the set of states that occur infinitely often in ω. The Büchi objectives require that some state of T be visited infinitely often, and defines the set of winning plays $\mathrm{Büchi}(T) = \{\omega \in \Omega \mid \mathrm{Inf}(\omega) \cap T \neq \emptyset\}$.
- *Parity objectives.* For $c, d \in \mathbb{N}$, we write $[c..d] = \{c, c+1, \ldots, d\}$. Let $p : S \to [0..d]$ be a function that assigns a *priority* $p(s)$ to every state $s \in S$, where $d \in \mathbb{N}$. The *parity objective* is defined as $\mathrm{Parity}(p) = \{\omega \in \Omega \mid \min(p(\mathrm{Inf}(\omega))) \text{ is even}\}$. In other words, the parity objective requires that the minimum priority visited infinitely often is even. In the sequel we will use Φ to denote parity objectives.

Qualitative analysis: almost-sure winning   Given a player-1 objective Φ, a strategy $\sigma \in \Sigma$ is *almost-sure winning* for player 1 from the state s if $\mathrm{Pr}_s^\sigma(\Phi) = 1$. The *almost-sure winning set* $\langle\langle 1 \rangle\rangle_{almost}(\Phi)$ for player 1 is the set of states from which player 1 has an almost-sure winning strategy. The qualitative analysis of MDPs correspond to the computation of the almost-sure winning set for a given objective Φ. It follows from the results of [7, 8] that for all MDPs and all reachability and parity objectives, if there is an almost-sure winning strategy, then there is a pure memoryless almost-sure winning strategy. The qualitative analysis of MDPs with parity objectives is achieved by iteratively applying the solutions of qualitative analysis of MDPs with Büchi objectives [6, 8], and hence in this work we will focus on qualitative analysis for Büchi objectives.

Theorem 1 [7, 8] *For all MDPs G, and all reachability and parity objectives Φ, there exists a pure memoryless strategy σ_* such that for all $s \in \langle\langle 1 \rangle\rangle_{almost}(\Phi)$ we have $\mathrm{Pr}_s^{\sigma_*}(\Phi) = 1$.*

Scc and bottom scc   Given a graph $G = (S, E)$, a set C of states is a strongly connected component if for all $s, t \in C$ there is a path from s to t going through states in C. We will denote by scc a maximal strongly connected component. An scc C is a bottom scc if for all $s \in C$ all out-going edges are in C, i.e., $E(s) \subseteq C$.

Markov chains, closed recurrent sets   A Markov chain is a special case of MDP with $S_1 = \emptyset$, and hence for simplicity a Markov chain is a tuple $((S, E), \delta)$ with a probabilistic transition function $\delta : S \to \mathcal{D}(S)$, and $(s, t) \in E$ iff $\delta(s, t) > 0$. A *closed recurrent* set C of a Markov chain is a bottom scc in the graph (S, E). Let $\mathcal{C} = \bigcup_{C \text{ is closed recurrent}} C$. It follows from the results on Markov chains [15] that for all $s \in S$, the set $\mathcal{C}$ is reached with probability 1 in finite time, and for all C such that C is closed recurrent, for all $s \in C$ and for all $t \in C$, if the starting state is s, then the state t is visited infinitely often with probability 1.

Markov chain from an MDP and memoryless strategy   Given an MDP $G = ((S, E), (S_1, S_P), \delta)$ and a memoryless strategy $\sigma_* : S_1 \to \mathcal{D}(S)$ we obtain a Markov chain $G' = ((S, E'), \delta')$ as follows: $E' = E \cap (S_P \times S) \cup \{(s, t) \mid s \in S_1, \sigma_*(s)(t) > 0\}$; and $\delta'(s, t) = \delta(s, t)$ for $s \in S_P$, and $\delta'(s, t) = \sigma(s)(t)$ for $s \in S_1$ and $t \in E(s)$. We will denote by G_{σ_*} the Markov chain obtained from an MDP G by fixing a memoryless strategy σ_* in the MDP.

Symbolic encoding of an MDP   All algorithms of the paper will only depend on the graph (S, E) of the MDP and the partition (S_1, S_P), and not on the probabilistic transition function δ. Thus the symbolic encoding of an MDP is obtained as the standard encoding of a transition system (with an OBDD [22]), with one additional bit, and the bit denotes whether a state belongs to S_1 or S_P. Also note that if the state bits already encode whether a state belongs to S_1 or S_P, then the additional bit is not required.

Symbolic step   To define the symbolic complexity of an algorithm an important concept to clarify is the notion of one symbolic step. In this work we adopt the following convention: one symbolic step corresponds to one primitive operation, where a primitive operation is one that is supported by the standard symbolic package like CuDD [22]. For example, the one-step predecessor and successor operators, obtaining an OBDD for a cube (a path from root to a leaf node with constant 1) of an OBDD, etc. are all supported as primitive operations in CuDD [22]. If the number of state variables to encode the OBDD is ℓ, then the symbolic steps that operate on the transition relation (such as predecessor and successor operations) consider OBDDs with $2 \cdot \ell$ variables, as compared to other operations like union, intersection, obtaining cubes, etc., that operate on OBDDs with ℓ variables. Thus the operations on transitions relations are quadratic as compared to other operations and we will focus on the number of symbolic steps that are on transition relations. Our algorithms will require cardinality computation of a set, and we will argue how to use primitive OBDD operations on OBDDs with ℓ variables to achieve the symbolic cardinality computation.

3 Symbolic algorithms for Büchi objectives

In this section we will present a new improved algorithm for the qualitative analysis of MDPs with Büchi objectives, and then present a symbolic implementation of the algorithm. Thus we obtain the first symbolic subquadratic algorithm for the problem. We start with the notion of *attractors* that is crucial for our algorithm.

Random and player 1 attractor   Given an MDP G, let $U \subseteq S$ be a subset of states. The *random attractor* $Attr_R(U)$ is defined inductively as follows: $X_0 = U$, and for $i \geq 0$, let $X_{i+1} = X_i \cup \{s \in S_1 \mid E(s) \subseteq X_i\} \cup \{s \in S_P \mid E(s) \cap X_i \neq \emptyset\}$. In other words, X_{i+1} consists of (a) states in X_i, (b) player-1 states whose all successors are in X_i and (c) random states that have at least one edge to X_i. Then $Attr_R(U) = \bigcup_{i \geq 0} X_i$. The definition of *player-1 attractor* $Attr_1(U)$ is analogous and is obtained by exchanging the role of random states and player 1 states in the above definition.

Property of attractors   Given an MDP G, and set U of states, let $A = Attr_R(U)$. Then from A player 1 cannot ensure to avoid U, in other words, for all states in A and for all player 1 strategies, the set U is reached with positive probability. For $A = Attr_1(U)$ there is a player 1 memoryless strategy to ensure that the set U is reached with certainty. The computation of random and player 1 attractors is the computation of alternating reachability and can be achieved in $O(m)$ time [14], and can be achieved in $O(n)$ symbolic steps.

3.1 A new subquadratic algorithm

The classical algorithm for computing the almost-sure winning set in MDPs with Büchi objectives has $O(n \cdot m)$ running time, and the symbolic implementation of the algorithm takes at most $O(n^2)$ symbolic steps. A subquadratic algorithm, with $O(m \cdot \sqrt{m})$ running time, for the problem was presented in [5]. The algorithm of [5] uses a mix of backward exploration and forward exploration. Every forward exploration step consists of executing a set of DFSs (depth first searches) simultaneously for a specified number of edges, and must maintain the exploration of the same number of edges in each of the DFSs. The algorithm thus depends crucially on maintaining the number of edges traversed explicitly, and hence the algorithm has no symbolic implementation. In this section we present a new subquadratic

algorithm to compute $\langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$. The algorithm is simpler as compared to the algorithm of [5] and we will show that our new algorithm can be implemented symbolically. Our new algorithm has some similar ideas as the algorithm of [5] in mixing backward and forward exploration, but the key difference is that the new algorithm never stops the forward exploration after a certain number of edges, and hence need not maintain the traversed edges explicitly. Thus the new algorithm is simpler, and our correctness and running time analysis proofs are different. We show that our new algorithm works in $O(m \cdot \sqrt{m})$ time, and requires at most $O(n \cdot \sqrt{m})$ symbolic steps.

Improved algorithm for almost-sure Büchi Our algorithm iteratively removes states from the graph, until the almost-sure winning set is computed. At iteration i, we denote the remaining subgraph as (S_i, E_i), where S_i is the set of remaining states, E_i is the set of remaining edges, and the set of remaining target states is T_i (i.e., $T_i = S_i \cap T$). The set of states removed will be denoted by Z_i, i.e., $S_i = S \setminus Z_i$. The algorithm will ensure that (a) $Z_i \subseteq S \setminus \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$; and (b) for all $s \in S_i \cap S_P$ we have $E(s) \cap Z_i = \emptyset$. In every iteration the algorithm identifies a set Q_i of states such that there is no path from Q_i to the set T_i. Hence clearly $Q_i \subseteq S \setminus \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$. By the random attractor property from $Attr_R(Q_i)$ the set Q_i is reached with positive probability against any strategy for player 1. The algorithm maintains the set L_{i+1} of states that were removed from the graph since (and including) the last iteration of Case 1, and the set J_{i+1} of states that lost an edge to states removed from the graph since the last iteration of Case 1. Initially $L_0 := J_0 := \emptyset$, $Z_0 := \emptyset$, and let $i := 0$ and we describe the iteration i of our algorithm. We call our algorithm IMPRALGO (Improved Algorithm) and the pseudocode is given as Algorithm 1.

1. *Case 1.* If $((|J_i| > \sqrt{m})$ or $i = 0)$, then
   (a) Let Y_i be the set of states that can reach the current target set T_i (this can be computed in $O(m)$ time by a graph reachability algorithm).
   (b) Let $Q_i := S_i \setminus Y_i$, i.e., there is no path from Q_i to T_i.
   (c) $Z_{i+1} := Z_i \cup Attr_R(Q_i)$. The set $Attr_R(Q_i)$ is removed from the graph.
   (d) The set L_{i+1} is the set of states removed from the graph in this iteration (i.e., $L_{i+1} := Attr_R(Q_i))$ and J_{i+1} be the set of states in the remaining graph with an edge to L_{i+1}.
   (e) If Q_i is empty, the algorithm stops, otherwise $i := i + 1$ and go to the next iteration.
2. *Case 2.* Else ($|J_i| \leq \sqrt{m}$ and $i > 0$), then
   (a) We do a lock-step search from every state s in J_i as follows: we do a DFS from s and (a) if the DFS tree reaches a state in T_i, then we stop the DFS search from s; and (b) if the DFS is completed without reaching a state in T_i, then we stop the entire lock-step search, and all states in the DFS tree are identified as Q_i. The set $Attr_R(Q_i)$ is removed from the graph and $Z_{i+1} := Z_i \cup Attr_R(Q_i)$. If DFS searches from all states s in J_i reach the set T_i, then the algorithm stops.
   (b) The set L_{i+1} is the set of states removed from the graph since the last iteration of Case 1 (i.e., $L_{i+1} := L_i \cup Attr_R(Q_i)$, where Q_i is the DFS tree that stopped without reaching T_i in the previous step of this iteration) and J_{i+1} be the set of states in the remaining graph with an edge to L_{i+1}, i.e., $J_{i+1} := (J_i \setminus Attr_R(Q_i)) \cup X_i$, where X_i is the subset of states of S_i with an edge to $Attr_R(Q_i)$.
   (c) $i := i + 1$ and go to the next iteration.

Correctness and running time analysis We first prove the correctness of the algorithm.

Lemma 1 *Algorithm* IMPRALGO *correctly computes the set* $\langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$.

Algorithm 1 ImprAlgo

Input: An MDP $G = ((S, E), (S_1, S_P), \delta)$ with Büchi set T.
Output: $\langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$, i.e., the almost-sure winning set for player 1.
1. $i := 0$; $S_0 := S$; $E_0 := E$; $T_0 := T$;
2. $L_0 := Z_0 := J_0 := \emptyset$;
3. **if** ($|J_i| > \sqrt{m}$ or $i = 0$) **then**
 3.1. $Y_i := \textbf{Reach}(T_i, (S_i, E_i))$; (i.e., compute the set Y_i that can reach T_i in the graph (S_i, E_i))
 3.2. $Q_i := S_i \setminus Y_i$;
 3.3. **if** ($Q_i = \emptyset$) **then goto** line 6;
 3.4. **else goto** line 5;
4. **else** (i.e., $J_i \leq \sqrt{m}$ and $i > 0$)
 4.1. **for each** $s \in J_i$
  4.1.1. $DFS_{i,s} := s$; (initializing DFS-trees)
 4.2. **for each** $s \in J_i$
  4.2.1. Do 1 step of DFS from $DFS_{i,s}$, unless it has encountered a state from T_i
  4.2.2. If DFS encounters a state from T_i, mark that DFS as stopped
  4.2.3. **if** DFS completes without meeting T_i **then**
   4.2.3.1. $Q_i := DFS_{i,s}$;
   4.2.3.2. **goto** line 5;
  4.2.4. **if** all DFSs meet T_i **then**
   4.2.4.1. **goto** line 6;
5. Removal of attractor of Q_i in the following steps
 5.1 $Z_{i+1} := Z_i \cup Attr_R(Q_i, (S_i, E_i), (S_1 \cap S_i, S_P \cap S_i))$;
 5.2. $S_{i+1} := S_i \setminus Z_{i+1}$; $E_{i+1} := E_i \cap S_{i+1} \times S_{i+1}$;
 5.4. **if** the last goto call from step 3.4 (i.e. Case 1 is executed) **then**
  5.4.1 $L_{i+1} := Attr_R(Q_i, (S_i, E_i), (S_1 \cap S_i, S_P \cap S_i))$;
 5.5. **else** $L_{i+1} := L_i \cup Attr_R(Q_i, (S_i, E_i), (S_1 \cap S_i, S_P \cap S_i))$;
 5.6 $J_{i+1} := E^{-1}(L_{i+1}) \cap S_{i+1}$;
 5.7. $i := i + 1$;
 5.8. **goto** line 3;
6. **return** $S \setminus Z_i$;

Proof We consider an iteration i of the algorithm. Recall that in this iteration Y_i is the set of states that can reach T_i and Q_i is the set of states with no path to T_i. Thus the algorithm ensures that in every iteration i, for the set of states Q_i identified by the algorithm there is no path to the set T_i, and hence from Q_i the set T_i cannot be reached with positive probability. Clearly, from Q_i the set T_i cannot be reached with probability 1. Since from $Attr_R(Q_i)$ the set Q_i is reached with positive probability against all strategies for player 1, it follows that from $Attr_R(Q_i)$ the set T_i cannot be ensured to be reached with probability 1. Thus for the set Z_i of removed states we have $Z_i \subseteq S \setminus \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$. It follows that all the states removed by the algorithm over all iterations are not part of the almost-sure winning set.

To complete the correctness argument we show that when the algorithm stops, the remaining set is $\langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$. When the algorithm stops, let S_* be the set of remaining states and T_* be the set of remaining target states. It follows from above that $S \setminus S_* \subseteq S \setminus \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$ and to complete the proof we show $S_* \subseteq \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$. The following assertions hold: (a) for all $s \in S_* \cap S_P$ we have $E(s) \subseteq S_*$, and (b) for all states $s \in S_*$ there is a path to the set T_*. We prove (a) as follows: whenever the algorithm removes a set Z_i, it is a random attractor, and thus if a state $s \in S_* \cap S_P$ has an edge (s, t) with $t \in S \setminus S_*$, then s would have been included in $S \setminus S_*$, and thus (a) follows. We prove (b) as follows: (i) If the algorithm stops in Case 1, then $Q_i = \emptyset$, and it follows that every state in S_* can reach T_*. (ii) We now consider the case when the algorithm stops in Case 2: In this case every state in J_i has a path to $T_i = T_*$, this is because if there is a state s in J_i with no path to T_i, then the DFS tree from s would have been identified as Q_i in step 2 (a) and the algorithm would not have stopped. It follows that there is no bottom scc in the graph induced

by S_* that does not intersect T_*: because if there is a bottom scc that does not contain a state from J_i and also does not contain a target state, then it would have been identified in the last iteration of Case 1. Since every state in S_* has an out-going edge, it follows every state in S_* has a path to T_*. Hence (b) follows. Consider a shortest path (or the BFS tree) from all states in S_* to T_*, and for a state $s \in S_* \cap S_1$, let s' be the successor for the shortest path, and we consider the pure memoryless strategy σ_* that chooses the shortest path successor for all states $s \in (S_* \setminus T_*) \cap S_1$, and in states in $T_* \cap S_1$ choose any successor in S_*. Let $\ell = |S_*|$ and let α be the minimum of the positive transition probability of the MDP. For all states $s \in S_*$, the probability that T_* is reached within ℓ steps is at least α^ℓ, and it follows that the probability that T_* is not reached within $k \cdot \ell$ steps is at most $(1 - \alpha^\ell)^k$, and this goes to 0 as k goes to ∞. It follows that for all $s \in S_*$ the pure memoryless strategy σ_* ensures that T_* is reached with probability 1. Moreover, the strategy ensures that S_* is never left, and hence it follows that T_* is visited infinitely often with probability 1. It follows that $S_* \subseteq \langle\!\langle 1 \rangle\!\rangle_{almost}(\text{Büchi}(T_*)) \subseteq \langle\!\langle 1 \rangle\!\rangle_{almost}(\text{Büchi}(T))$ and hence the correctness follows. $\qquad\square$

We now analyze the running time of the algorithm.

Lemma 2 *Given an MDP G with m edges, Algorithm* IMPRALGO *takes $O(m \cdot \sqrt{m})$ time.*

Proof The total work of the algorithm, when Case 1 is executed, over all iterations is at most $O(\sqrt{m} \cdot m)$: this follows because between two iterations of Case 1 at least $\sqrt{m}$ edges must have been removed from the graph (since $|J_i| > \sqrt{m}$ everytime Case 1 is executed other than the case when $i = 0$), and hence Case 1 can be executed at most $m/\sqrt{m} = \sqrt{m}$ times. Since each iteration can be achieved in $O(m)$ time, the $O(m \cdot \sqrt{m})$ bound for Case 1 follows. We now show that the total work of the algorithm, when Case 2 is executed, over all iterations is at most $O(\sqrt{m} \cdot m)$. The argument is as follows: consider an iteration i such that Case 2 is executed. Then we have $|J_i| \leq \sqrt{m}$. Let Q_i be the DFS tree in iteration i while executing Case 2, and let $E(Q_i) = \bigcup_{s \in Q_i} E(s)$. The lock-step search ensures that the number of edges explored in this iteration is at most $|J_i| \cdot |E(Q_i)| \leq \sqrt{m} \times |E(Q_i)|$. Since Q_i is removed from the graph we *charge* the work of $\sqrt{m} \cdot |E(Q_i)|$ to edges in $E(Q_i)$, charging work $\sqrt{m}$ to each edge. Since there are at most m edges, the total charge of the work over all iterations when Case 2 is executed is at most $O(m \cdot \sqrt{m})$. Note that if instead of $\sqrt{m}$ we would have used a bound k in distinguishing Case 1 and Case 2, we would have achieved a running time bound of $O(m^2/k + m \cdot k)$, which is optimized by $k = \sqrt{m}$. Our desired result follows. $\quad\square$

This gives us the following result.

Theorem 2 *Given an MDP G and a set T of target states, the algorithm* IMPRALGO *correctly computes the set $\langle\!\langle 1 \rangle\!\rangle_{almost}(\text{Büchi}(T))$ in time $O(m \cdot \sqrt{m})$.*

3.2 Symbolic implementation of IMPRALGO

In this subsection we will a present symbolic implementation of each of the steps of algorithm IMPRALGO. The symbolic algorithm depends on the following symbolic operations that can be easily achieved with an OBDD implementation. For a set $X \subseteq S$ of states, let

$$\text{Pre}(X) = \big\{ s \in S \mid E(s) \cap X \neq \emptyset \big\};$$

$$\text{Post}(X) = \bigg\{ t \in S \mid t \in \bigcup_{s \in X} E(s) \bigg\};$$

$$\mathsf{CPre}(X) = \big\{s \in S_P \mid E(s) \cap X \neq \emptyset\big\} \cup \big\{s \in S_1 \mid E(s) \subseteq X\big\}.$$

In other words, $\mathsf{Pre}(X)$ is the predecessors of states in X; $\mathsf{Post}(X)$ is the successors of states in X; and $\mathsf{CPre}(X)$ is the set of states Y such that for every random state in Y there is a successor in X, and for every player 1 state in Y all successors are in Y. All these operations are supported as primitive operations by the CuDD package.

We now present a symbolic version of IMPRALGO. For the symbolic version the basic steps are as follows: (i) Case 1 of the algorithm is same as Case 1 of IMPRALGO, and (ii) Case 2 is similar to Case 2 of IMPRALGO, and the only change in Case 2 is instead of lock-step search exploring the same number of edges, we have lock-step search that executes the same number of symbolic steps. The details of the symbolic implementation are as follows, and we will refer to the algorithm as SYMBIMPRALGO.

1. *Case 1.* In Case 1(a) we need to compute reachability to a target set T. The symbolic implementation is standard and done as follows: $X_0 := T$ and $X_{i+1} := X_i \cup \mathsf{Pre}(X_i)$ until $X_{i+1} = X_i$. The computation of the random attractor is also standard and is achieved as above replacing Pre by CPre. It follows that every iteration of Case 1 can be achieved in $O(n)$ symbolic steps.
2. *Case 2.* For analysis of Case 2 we present a symbolic implementation of the lock-step forward search. The lock-step ensures that each search executes the same number of symbolic steps. The implementation of the forward search from a state s in iteration i is achieved as follows: $P_0 := \{s\}$ and $P_{j+1} := P_j \cup \mathsf{Post}(P_j)$ unless $P_{j+1} = P_j$ or $P_j \cap T_i \neq \emptyset$. If $P_j \cap T_i \neq \emptyset$, then the forward search is stopped from s. If $P_{j+1} = P_j$ and $P_j \cap T_i = \emptyset$, then we have identified that there is no path from states in P_j to T_i.
3. *Symbolic computation of cardinality of sets.* The other key operation required by the algorithm is determining whether the size of set J_i is at least $\sqrt{m}$ or not. Below we describe the details of this symbolic operation.

Symbolic computation of cardinality   Given a symbolic description of a set X and a number k, our goal is to determine whether $|X| \leq k$. A naive way is to check for each state, whether it belongs to X. But this takes time proportional to the size of state space and also is not symbolic. We require a procedure that uses the structure of a BDD and directly finds the states which this BDD represents. It should also take into account that if more than k states are already found, then no more computation is required. We present the following procedure to accomplish the same. A *cube* of a BDD is a path from root node to leaf node where the leaf node is the constant 1 (i.e. true). Thus, each cube represents a set of states present in the BDD which are exactly the states found by doing every possible assignment of the variables not occurring in the cube. For an explicit implementation: consider a procedure that uses Cudd_ForEachCube (from CUDD package, see [22] for symbolic implementation) to iterate over the cubes of a given OBDD in the same manner the successor function works on a binary tree. If ℓ is the number of variables not occurring in a particular cube, we get 2^ℓ states from that cube which are part of the OBDD. We keep on summing up all such states until they exceed k. If it does exceed, we stop and say that $|X| > k$. Else we terminate when we have exhausted all cubes and we get $|X| \leq k$. Thus we require $\min(k, |BDD(X)|)$ symbolic steps, where $BDD(X)$ is the size of the OBDD of X. We also note, that this method operates on OBDDs that represent set of states, and these OBDDs only use $\log(n)$ variables compared to $2 \cdot \log(n)$ variables used by OBDDs representing transitions (edge relation). Hence, the operations mentioned are cheaper as compared to Pre and Post computations.

Correctness and runtime analysis   The correctness of SYMBIMPRALGO is established following the correctness arguments for algorithm IMPRALGO. We now analyze the worst case number of symbolic steps. The total number of symbolic steps executed by Case 1 over all iterations is $O(n \cdot \sqrt{m})$ since between two executions of Case 1 at least $\sqrt{m}$ edges are removed, and every execution is achieved in $O(n)$ symbolic steps. The work done for the symbolic cardinality computation is charged to the edges already removed from the graph, and hence the total number of symbolic steps over all iterations for the size computations is $O(m)$. We now show that the total number of symbolic steps executed over all iterations of Case 2 is $O(n \cdot \sqrt{m})$. The analysis is achieved as follows. Consider an iteration i of Case 2, and let the number of states removed in the iteration be n_i. Then the number of symbolic steps executed in this iteration for each of the forward search is at most n_i, and since $|J_i| \leq \sqrt{m}$, it follows that the number of symbolic steps executed is at most $n_i \cdot \sqrt{m}$. Since we remove n_i states, we *charge* each state removed from the graph with $\sqrt{m}$ symbolic steps for the total $n_i \cdot \sqrt{m}$ symbolic steps. Since there are at most n states, the total charge of symbolic steps over all iterations is $O(n \cdot \sqrt{m})$. Thus it follows that we have a symbolic algorithm to compute the almost-sure winning set for MDPs with Büchi objectives in $O(n \cdot \sqrt{m})$ symbolic steps.

Theorem 3 *Given an MDP G and a set T of target states, the symbolic algorithm* SYMBIMPRALGO *correctly computes* $\langle\!\langle 1 \rangle\!\rangle_{almost}(\text{Büchi}(T))$ *in* $O(n \cdot \sqrt{m})$ *symbolic steps.*

Remark 1 In many practical cases, MDPs have constant out-degree and hence we obtain a symbolic algorithm that works in $O(n \cdot \sqrt{n})$ symbolic steps, as compared to the previous known (symbolic implementation of the classical) algorithm that requires $\Omega(n^2)$ symbolic steps.

Remark 2 Note that in our algorithm we used $\sqrt{m}$ to distinguish between Case 1 and Case 2 to obtain the optimal time complexity. However, our algorithm could also be parametrized with a parameter k to distinguish between Case 1 and Case 2, and then the number of symbolic steps required is $O(\frac{n \cdot m}{k} + n \cdot k)$. For example, if $m = O(n)$, by choosing $k = \log n$, we obtain a symbolic algorithm that requires $O(\frac{n^2}{\log n})$ symbolic steps as compared to the $O(n^2)$ symbolic steps of the previous known algorithms. In other words our algorithm can be easily parametrized to provide a trade-off between the number of forward searches and speed up in the number of symbolic steps.

3.3 Optimized SYMBIMPRALGO

In the worst case, the SYMBIMPRALGO algorithm takes $O(n \cdot \sqrt{m})$ steps. However it is easy to construct a family of MDPs with n states and $O(n)$ edges, where the classical algorithm takes $O(n)$ symbolic steps, whereas SYMBIMPRALGO requires $\Omega(n \cdot \sqrt{n})$ symbolic steps. One approach to obtain an algorithm that takes at most $O(n \cdot \sqrt{n})$ symbolic steps and no more than linearly many symbolic steps of the classical algorithm is to dovetail (or run in lock-step) the classical algorithm and SYMBIMPRALGO, and stop when either of them stops. This approach will take time at least twice the minimum running time of the classical algorithm and SYMBIMPRALGO. We show that a much smarter dovetailing is possible (at the level of each iteration). We now present the smart dovetailing algorithm, and we call the algorithm SMDVSYMBIMPRALGO. The basic change is in Case 2 of SYMBIMPRALGO. We now describe the changes in Case 2:

– At the beginning of an execution of Case 2 at iteration i such that the last execution was Case 1, we initialize a set U_i to T_i. Every time a post computation ($\mathsf{Post}(P_j)$) is done, we update U_i by $U_{i+1} := U_i \cup \mathsf{Pre}(U_i)$ (this is the backward exploration step of the classical algorithm and it is dovetailed with the forward exploration step in every iteration). For the forward exploration step, we continue the computation of P_j unless $P_{j+1} = P_j$ or $P_j \cap U_i \neq \emptyset$ (i.e., SYMBIMPRALGO checked the emptiness of intersection with T_i, whereas in SMDVSYMBIMPRALGO the emptiness of the intersection is checked with U_i). If $U_{i+1} = U_i$ (i.e., a fixpoint is reached), then $S_i \setminus U_i$ and its random attractor is removed from the graph.

Correctness and symbolic steps analysis   We present the correctness and number of symbolic steps required analysis for the algorithm SMDVSYMBIMPRALGO. The correctness analysis is same as IMPRALGO and the only change is as follows (we describe iteration i): (a) if in Case 2 we obtain a set $P_j = P_{j+1}$ and its intersection with U_i is empty, then there is no path from P_j to U_i and since $T_i \subseteq U_i$, it follows that there is no path from P_j to U_i; (b) if $P_j \cap U_i \neq \emptyset$, then since U_i is obtained as the backward exploration from T_i, every state in U_i has a path to T_i, and it follows that there is a path from the starting state of P_j to U_i and hence to T_i; and (c) if $U_i = \mathsf{Pre}(U_i)$, then U_i is the set of states that can reach T_i and all the other states can be removed. Thus the correctness follows similar to the arguments for IMPRALGO. The key idea of the running time analysis is as follows:

1. Case 1 of the algorithm is same to Case 1 of SYMBIMPRALGO, and in Case 2 the algorithm also runs like SYMBIMPRALGO, but for every symbolic step (Post computation) of SYMBIMPRALGO, there is an additional (Pre) computation. Hence the total number of symbolic steps of SMDVSYMBIMPRALGO is at most twice the number of symbolic steps of SYMBIMPRALGO. However, the optimized step of maintaining the set U_i which includes T_i may allow to stop several of the forward exploration as they may intersect with U_i earlier than intersection with T_i.
2. Case 1 of the algorithm is same as in Case 1 of the classical algorithm. In Case 2 of the algorithm the backward exploration step is the same as the classical algorithm, and (i) for every Pre computation, there is an additional Post computation and (ii) for every check whether $U_i = \mathsf{Pre}(U_i)$, there is a check whether $P_j = P_{j+1}$ or $P_j \cap U_i \neq \emptyset$. It follows that the total number of symbolic steps of Case 1 and Case 2 over all iterations is at most twice the number of symbolic steps of the classical algorithm. The cardinality computation takes additional $O(m)$ symbolic steps over all iterations.

Hence we obtain the following result.

Theorem 4 *Given an MDP G and a set T of target states, the symbolic algorithm* SMDVSYMBIMPRALGO *correctly computes* $\langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$ *and requires at most*

$$\min\{2 \cdot \mathsf{SymbStep}(\text{SYMBIMPRALGO}), 2 \cdot \mathsf{SymbStep}(\text{CLASSICAL}) + O(m)\}$$

symbolic steps, where $\mathsf{SymbStep}$ *is the number of symbolic steps of an algorithm.*

Observe that it is possible that the number of symbolic steps and running time of SMDVSYMBIMPRALGO is smaller than both SYMBIMPRALGO and CLASSICAL (in contrast to a simple dovetailing of SYMBIMPRALGO and CLASSICAL, where the running time and symbolic steps is twice that of the minimum). It is straightforward to construct a family of examples where SMDVSYMBIMPRALGO takes linear ($O(n)$) symbolic steps, however both CLASSICAL and SYMBIMPRALGO take at least $O(n \cdot \sqrt{n})$ symbolic steps.

Algorithm 2 WinLose

Input: An MDP $G = ((S, E), (S_1, S_P), \delta)$ with Büchi set T.
Output: $\langle\!\langle 1 \rangle\!\rangle_{almost}(\text{Büchi}(T))$, i.e., the almost-sure winning set for player 1.
1. $W := W_1 := W_2 := \emptyset$
2. **while**($W \neq S$) **do**
   2.1. $SCCS :=$**SCC-Decomposition**$(S \setminus W)$ (i.e. scc decomposition of the graph induced by $S \setminus W$)
   2.2. **for each** C in $SCCS$
      2.2.1. **if** $(E(C) \subset C \cup W)$ **then** (checks if C is a bottom scc in graph induced by $S \setminus W$)
         2.2.1.1. **if** $C \cap T \neq \emptyset$ or $E(C) \cap W_1 \neq \emptyset$ **then**
            2.2.1.1.1. $W_1 := W_1 \cup C$
         2.2.1.2. **else**
            2.2.1.2.1. $W_2 := W_2 \cup C$
   2.3. $W_1 := Attr_1(W_1, (S, E), (S_1, S_P))$
   2.4. $W_2 := Attr_R(W_2, (S, E), (S_1, S_P))$
   2.5. $W := W_1 \cup W_2$
3. **return** W_1

---

4 The win-lose algorithm

All the algorithms known for computing the almost-sure winning set (including the algorithms presented in the previous section) iteratively compute the set of states from where it is guaranteed that there is no almost-sure winning strategy for the player. The almost-sure winning set is discovered only when the algorithm stops. In this section, first we will present an algorithm that iteratively computes two sets W_1 and W_2, where W_1 is a subset of the almost-sure winning set, and W_2 is a subset of the complement of the almost-sure winning set. The algorithm has $O(K \cdot m)$ running time, where K is the size of the maximal strongly connected component (scc) of the graph of the MDP. We will first present the basic version of the algorithm, and then present an improved version of the algorithm, using the techniques to obtain IMPRALGO from the classical algorithm, and finally present the symbolic implementation of the new algorithm.

4.1 The basic win-lose algorithm

The basic steps of the new algorithm are as follows. The algorithm maintains W_1 and W_2, that are guaranteed to be subsets of the almost-sure winning set and its complement respectively. Initially $W_1 = \emptyset$ and $W_2 = \emptyset$. We also maintain that $W_1 = Attr_1(W_1)$ and $W_2 = Attr_R(W_2)$. We denote by W the union of W_1 and W_2. We describe an iteration of the algorithm and we will refer to the algorithm as the WINLOSE algorithm (pseudocode is given as Algorithm 2).

1. *Step 1.* Compute the scc decomposition of the remaining graph of the MDP, i.e., scc decomposition of the MDP graph induced by $S \setminus W$.
2. *Step 2.* For every bottom scc C in the remaining graph: if $C \cap \text{Pre}(W_1) \neq \emptyset$ or $C \cap T \neq \emptyset$, then $W_1 = Attr_1(W_1 \cup C)$; else $W_2 = Attr_R(W_2 \cup C)$, and the states in W_1 and W_2 are removed from the graph.

The stopping criterion is as follows: the algorithm stops when $W = S$. Observe that in each iteration, a set C of states is included in either W_1 or W_2, and hence W grows in each iteration. Observe that our algorithm has the flavor of iterative scc decomposition algorithm for computing maximal end-component decomposition of MDPs.

Correctness of the algorithm Note that in Step 2 we ensure that $Attr_1(W_1) = W_1$ and $Attr_R(W_2) = W_2$, and hence in the remaining graph there is no state of player 1 with an edge to W_1 and no random state with an edge to W_2. We show by induction that after every iteration $W_1 \subseteq \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$ and $W_2 \subseteq S \setminus \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$. The base case (with $W_1 = W_2 = \emptyset$) follows trivially. We prove the inductive case considering the following two cases.

1. Consider a bottom scc C in the remaining graph such that $C \cap \text{Pre}(W_1) \neq \emptyset$ or $C \cap T \neq \emptyset$. Consider the randomized memoryless strategy σ for the player that plays all edges in C uniformly at random, i.e., for $s \in C$ we have $\sigma(s)(t) = \frac{1}{|E(s) \cap C|}$ for $t \in E(s) \cap C$. If $C \cap \text{Pre}(W_1) \neq \emptyset$, then the strategy ensures that W_1 is reached with probability 1, since $W_1 \subseteq \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$ by inductive hypothesis it follows $C \subseteq \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$. Hence $Attr_1(W_1 \cup C) \subseteq \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$. If $C \cap T \neq \emptyset$, then since there is no edge from random states to W_2, it follows that under the randomized memoryless strategy σ, the set C is a closed recurrent set of the resulting Markov chain, and hence every state is visited infinitely often with probability 1. Since $C \cap T \neq \emptyset$, it follows that $C \subseteq \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$, and hence $Attr_1(W_1 \cup C) \subseteq \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$.

2. Consider a bottom scc C in the remaining graph such that $C \cap \text{Pre}(W_1) = \emptyset$ and $C \cap T = \emptyset$. Then consider any strategy for player 1: (a) If a play starting from a state in C stays in the remaining graph, then since C is a bottom scc, it follows that the play stays in C with probability 1. Since $C \cap T = \emptyset$ it follows that T is never visited. (b) If a play leaves C (note that C is a bottom scc of the remaining graph and not the original graph, and hence a play may leave C), then since $C \cap \text{Pre}(W_1) = \emptyset$, it follows that the play reaches W_2, and by hypothesis $W_2 \subseteq S \setminus \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$. In either case it follows that $C \subseteq S \setminus \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$. It follows that $Attr_R(W_2 \cup C) \subseteq S \setminus \langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$.

The correctness of the algorithm follows as when the algorithm stops we have $W_1 \cup W_2 = S$.

Running time analysis In each iteration of the algorithm at least one state is removed from the graph, and every iteration takes at most $O(m)$ time: in every iteration, the scc decomposition of step 1 and the attractor computation in step 2 can be achieved in $O(m)$ time. Hence the naive running of the algorithm is $O(n \cdot m)$. The desired $O(K \cdot m)$ bound is achieved by considering the standard technique of running the algorithm on the scc decomposition of the MDP. In other words, we first compute the scc of the graph of the MDP, and then proceed bottom up computing the partition W_1 and W_2 for an scc C once the partition is computed for all states below the scc. Observe that the above correctness arguments are still valid. The running time analysis is as follows: let ℓ be the number of scc's of the graph, and let n_i and m_i be the number of states and edges of the i-th scc. Let $K = \max\{n_i \mid 1 \leq i \leq \ell\}$. Our algorithm runs in time $O(m) + \sum_{i=1}^{\ell} O(n_i \cdot m_i) \leq O(m) + \sum_{i=1}^{\ell} O(K \cdot m_i) = O(K \cdot m)$.

Theorem 5 *Given an MDP with a Büchi objective, the* WINLOSE *algorithm iteratively computes the subsets of the almost-sure winning set and its complement, and in the end correctly computes the set* $\langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$ *and the algorithm runs in time* $O(K_S \cdot m)$, *where* K_S *is the maximum number of states in an scc of the graph of the MDP.*

4.2 Improved WINLOSE algorithm and symbolic implementation

Improved WINLOSE *algorithm* The improved version of the WINLOSE algorithm performs a forward exploration to obtain a bottom scc like Case 2 of IMPRALGO. At iteration i, we denote the remaining subgraph as (S_i, E_i), where S_i is the set of remaining states,

and E_i is the set of remaining edges. The set of states removed will be denoted by Z_i, i.e., $S_i = S \setminus Z_i$, and Z_i is the union of W_1 and W_2. In every iteration the algorithm identifies a set C_i of states such that C_i is a bottom scc in the remaining graph, and then it follows the steps of the WINLOSE algorithm. We will consider two cases. The algorithm maintains the set L_{i+1} of states that were removed from the graph since (and including) the last iteration of Case 1, and the set J_{i+1} of states that lost an edge to states removed from the graph since the last iteration of Case 1. Initially $J_0 := L_0 := Z_0 := W_1 := W_2 := \emptyset$, and let $i := 0$, and we describe the iteration i of our algorithm. We call our algorithm IMPRWINLOSE (pseudocode is given as Algorithm 3).

1. *Case 1.* If $((|J_i| > \sqrt{m})$ or $i = 0)$, then
   (a) Compute the scc decomposition of the remaining graph.
   (b) For each bottom scc C_i, if $C_i \cap T \neq \emptyset$ or $C_i \cap \mathsf{Pre}(W_1) \neq \emptyset$, then $W_1 := Attr_1(W_1 \cup C_i)$, else $W_2 := Attr_R(W_2 \cup C_i)$.
   (c) $Z_{i+1} := W_1 \cup W_2$. The set $Z_{i+1} \setminus Z_i$ is removed from the graph.
   (d) The set L_{i+1} is the set of states removed from the graph in this iteration and J_{i+1} be the set of states in the remaining graph with an edge to L_{i+1}.
   (e) If Z_i is S, the algorithm stops, otherwise $i := i + 1$ and go to the next iteration.
2. *Case 2.* Else $((|J_i| \leq \sqrt{m})$ and $i > 0)$, then
   (a) Consider the set J_i to be the set of vertices in the graph that lost an edge to the states removed since the last iteration that executed Case 1.
   (b) We do a lock-step search from every state s in J_i as follows: we do a DFS from s, until the DFS stops. Once the DFS stops we have identified a bottom scc C_i.
   (c) If $C_i \cap T \neq \emptyset$ or $C_i \cap \mathsf{Pre}(W_1) \neq \emptyset$, then $W_1 := Attr_1(W_1 \cup C_i)$, else $W_2 := Attr_R(W_2 \cup C_i)$.
   (d) $Z_{i+1} := W_1 \cup W_2$. The set $Z_{i+1} \setminus Z_i$ is removed from the graph.
   (e) The set L_{i+1} is the set of states removed from the graph since the last iteration of Case 1 and J_{i+1} be the set of states in the remaining graph with an edge to L_{i+1}.
   (f) If $Z_i = S$, the algorithm stops, otherwise $i := i + 1$ and go to the next iteration.

Correctness and running time   The correctness proof of IMPRWINLOSE is similar as the correctness argument of WINLOSE algorithm. One additional care requires to be taken for Case 2: we need to show that when we terminate the lockstep DFS search in Case 2, then we obtain a bottom scc. First, we observe that in iteration i, when Case 2 is executed, each bottom scc must contain a state from J_i, since it was not a bottom scc in the last execution of Case 1. Second, among all the lockstep DFSs, the first one that terminates must be a bottom scc because the DFS search from a state of J_i that does not belong to a bottom scc explores states of bottom scc's below it. Since Case 2 stops when the first DFS terminates we obtain a bottom scc. The rest of the correctness proof is identical as the proof for the WINLOSE algorithm. The running time analysis of the algorithm is similar to IMPRALGO algorithm, and this shows the algorithm runs in $O(m \cdot \sqrt{m})$ time. Applying the IMPRWINLOSE algorithm bottom up on the scc decomposition of the MDP gives us a running time of $O(m \cdot \sqrt{K_E})$, where K_E is the maximum number of edges of an scc of the MDP.

Theorem 6 *Given an MDP with a Büchi objective, the* IMPRWINLOSE *algorithm iteratively computes the subsets of the almost-sure winning set and its complement, and in the end correctly computes the set* $\langle\!\langle 1 \rangle\!\rangle_{almost}(\text{Büchi}(T))$. *The algorithm* IMPRWINLOSE *runs in time* $O(\sqrt{K_E} \cdot m)$, *where* K_E *is the maximum number of edges in an scc of the graph of the MDP.*

Algorithm 3 ImprWinLose

Input: An MDP $G = ((S, E), (S_1, S_P), \delta)$ with Büchi set T.
Output: $\langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$, i.e., the almost-sure winning set for player 1.
1. $i := 0; S_0 := S; E_0 := E; T_0 := T;$
2. $W_1 := W_2 := L_0 := Z_0 := J_0 := \emptyset;$
3. **if** ($|J_i| > \sqrt{m}$ or $i = 0$) **then**
   3.1. $SCCS := \textbf{SCC-Decomposition}(S_i)$ (scc decomposition of graph induced by S_i)
   3.2. **for each** C in $SCCS$
      3.2.1. **if** ($E_i(C) \subset C$) **then** (checks if C is a bottom scc in graph induced by S_i)
         3.2.1.1. **if** $C \cap T \neq \emptyset$ or $E(C) \cap W_1 \neq \emptyset$ **then**
            3.2.1.1.1. $W_1 := W_1 \cup C$
         3.2.1.2. **else**
            3.2.1.2.1. $W_2 := W_2 \cup C$
   3.3. **goto** line 5
4. **else** (i.e., $J_i \leq \sqrt{m}$ and $i > 0$)
   4.1. **for each** $s \in J_i$
      4.1.1. $DFS_{i,s} := s$ (initializing DFS-trees)
   4.2. **for each** $s \in J_i$
      4.2.1. Do 1 step of DFS from $DFS_{i,s}$
      4.2.2. **if** DFS completes **then**
         4.2.2.1. $C := DFS_{i,s}$
         4.2.2.2. **if** $C \cap T \neq \emptyset$ or $E(C) \cap W_1 \neq \emptyset$ **then**
            4.2.2.2.1. $W_1 := W_1 \cup C$
         4.2.2.3. **else**
            4.2.2.3.1. $W_2 := W_2 \cup C$
         4.2.2.4. **goto** line 5
5. Removal of W_1 and W_2 states in the following steps
   5.1. $W_1 := Attr_1(W_1, (S_i, E_i), (S_1 \cap S_i, S_P \cap S_i))$
   5.2. $W_2 := Attr_R(W_2, (S_i, E_i), (S_1 \cap S_i, S_P \cap S_i))$
   5.3. $Z_{i+1} := Z_i \cup W_1 \cup W_2$
   5.4. $S_{i+1} := S_i \setminus Z_{i+1}; E_{i+1} := E_i \cap S_{i+1} \times S_{i+1}$
   5.5. **if** the last **goto** call was from line 3.3 **then**
      5.5.1. $L_{i+1} := Z_{i+1} \setminus Z_i$
   5.6. **else**
      5.6.1. $L_{i+1} := L_i \cup (Z_{i+1} \setminus Z_i)$
   5.7. $J_{i+1} := E^{-1}(L_{i+1}) \cap S_{i+1}$
   5.8. **if** $Z_{i+1} = S$ **then**
      5.8.1. goto line 6
   5.9. $i := i + 1;$ **goto** line 3
6. **return** W_1

Symbolic implementation The symbolic implementation of IMPRWINLOSE algorithm is obtained in a similar fashion as SYMBIMPRALGO was obtained from IMPRALGO. The only additional step required is the symbolic scc computation. It follows from the results of [12] that scc decomposition can be computed in $O(n)$ symbolic steps. In the following section we will present an improved symbolic scc computation algorithm. The correctness proof of SYMBIMPRWINLOSE is similar to IMPRWINLOSE algorithm. For the correctness of the SYMBIMPRWINLOSE algorithm we again need to take care that when we terminate in Case 2, then we have identified a bottom scc. Note that for symbolic step forward search we cannot guarantee that the forward search that stops first gives a bottom scc. For Case 2 of the SYMBIMPRWINLOSE we do in lockstep both symbolic forward and backward searches, stop when both the searches stop and gives the same result. Thus we ensure when we terminate an iteration of Case 2 we obtain a bottom scc. The correctness then follows from the correctness arguments of WINLOSE and IMPRWINLOSE. The symbolic steps required analysis is same as for SYMBIMPRALGO.

Corollary 1 *Given an MDP with a Büchi objective, the symbolic* IMPRWINLOSE *algorithm (*SYMBIMPRWINLOSE*) iteratively computes the subsets of the almost-sure winning set and its complement, and in the end correctly computes the set* $\langle\langle 1 \rangle\rangle_{almost}(\text{Büchi}(T))$. *The algorithm* SYMBIMPRWINLOSE *requires* $O(\sqrt{K_E} \cdot n)$ *symbolic steps, where* K_E *is the maximum number of edges in an scc of the graph of the MDP.*

Remark 3 It is clear from the complexity of the WINLOSE and IMPRWINLOSE algorithms that they would perform better for MDPs where the graph has many small scc's, rather than few large ones.

5 Improved symbolic SCC algorithm

A symbolic algorithm to compute the scc decomposition of a graph in $O(n \cdot \log n)$ symbolic steps was presented in [3]. The algorithm of [3] was based on forward and backward searches. The algorithm of [12] improved the algorithm of [3] to obtain an algorithm for scc decomposition that takes at most linear number of symbolic steps. In this section we present an improved version of the algorithm of [12] that improves the constants of the number of linear symbolic steps required. In Sect. 5.1 we present the improved algorithm and correctness, and some further technical details are presented in Sect. 8.1 of Appendix.

5.1 Improved algorithm and correctness

We first describe the main ideas of the algorithm of [12] and then present our improved algorithm. The algorithm of [12] improves the algorithm of [3] by maintaining the right order for forward sets. The notion of *spine-sets* and *skeleton of a forward set* was designed for this purpose.

Spine-sets and skeleton of a forward set Let $G = (S, E)$ be a directed graph. Consider a finite path $\tau = (s_0, s_1, \ldots, s_\ell)$, such that for all $0 \le i \le \ell - 1$ we have $(s_i, s_{i+1}) \in E$. The path is *chordless* if for all $0 \le i < j \le \ell$ such that $j - i > 1$, there is no edge from s_i to s_j. Let $U \subseteq S$. The pair (U, s) is a *spine-set* of G iff G contains a chordless path whose set of states is U that ends in s. For a state s, let $\mathsf{FW}(s)$ denote the set of states that is reachable from s (i.e., reachable by a forward search from s). The set (U, t) is a *skeleton of* $\mathsf{FW}(s)$ iff t is a state in $\mathsf{FW}(s)$ whose distance from s is maximum and U is the set of states on a shortest path from s to t. The following lemma was shown in [12] establishing relation of skeleton of forward set and spine-set.

Lemma 3 [12] *Let* $G = (S, E)$ *be a directed graph, and let* $\mathsf{FW}(s)$ *be the forward set of* $s \in S$. *The following assertions hold*: (1) *If* (U, t) *is a skeleton of the forward-set* $\mathsf{FW}(s)$, *then* $U \subseteq \mathsf{FW}(s)$. (2) *If* (U, t) *is a skeleton of* $\mathsf{FW}(s)$, *then* (U, t) *is a spine-set in* G.

The intuitive idea of the algorithm The algorithm of [12] is a recursive algorithm, and in every recursive call the scc of a state s is determined by computing $\mathsf{FW}(s)$, and then identifying the set of states in $\mathsf{FW}(s)$ having a path to s. The choice of the state to be processed next is guided by the implicit inverse order associated with a possible spine-set. This is achieved as follows: whenever a forward-set $\mathsf{FW}(s)$ is computed, a skeleton of such a forward set is also computed. The order induced by the skeleton is then used for the subsequent computations. Thus the symbolic steps performed to compute $\mathsf{FW}(s)$ are distributed over the scc

computation of the states belonging to a skeleton of FW(s). The key to establish the linear complexity of symbolic steps is the amortized analysis. We now present the main procedure SCCFIND and the main sub-procedure SKELFWD of the algorithm from [12].

Procedures SCCFIND *and* SKELFWD   The main procedure of the algorithm is SCCFIND that calls SKELFWD as a sub-procedure. The input to SCCFIND is a graph (S, E) and (A, B), where either $(A, B) = (\emptyset, \emptyset)$ or $(A, B) = (U, \{s\})$, where (U, s) is a spine-set. If S is $\emptyset$, then the algorithm stops. Else, (a) if (A, B) is $(\emptyset, \emptyset)$, then the procedure picks an arbitrary s from S and proceeds; (b) otherwise, the sub-procedure SKELFWD is invoked to compute the forward set of s together with the skeleton (U', s') of such a forward set. The SCCFIND procedure has the following local variables: FWSet, NewSet, NewState and SCC. The variable FWSet that maintains the forward set, whereas NewSet and NewState maintain U' and $\{s'\}$, respectively. The variable SCC is initialized to s, and then augmented with the scc containing s. The partition of the scc's is updated and finally the procedure is recursively called over:

1. the subgraph of (S, E) is induced by $S \setminus$ FWSet and the spine-set of such a subgraph obtained from $(U, \{t\})$ by subtracting SCC;
2. the subgraph of (S, E) induced by FWSet $\setminus$ SCC and the spine-set of such a subgraph obtained from (NewSet, NewState) by subtracting SCC.

The SKELFWD procedure takes as input a graph (S, E) and a state s, first it computes the forward set FW(s), and second it computes the skeleton of the forward set. The forward set is computed by symbolic breadth first search, and the skeleton is computed with a stack. The detailed pseudocodes are in the following subsection. We will refer to this algorithm of [12] as SYMBOLICSCC. The following result was established in [12]: for the proof of the constant 5, refer to the appendix of [12] and the last sentence explicitly claims that every state is charged at most 5 symbolic steps.

Theorem 7 [12] *Let $G = (S, E)$ be a directed graph. The algorithm* SYMBOLICSCC *correctly computes the scc decomposition of G in* $\min\{5 \cdot |S|, 5 \cdot D(G) \cdot N(G) + N(G)\}$ *symbolic steps, where $D(G)$ is the diameter of G, and $N(G)$ is the number of scc's in G.*

Improved symbolic algorithm   We now present our improved symbolic scc algorithm and refer to the algorithm as IMPROVEDSYMBOLICSCC. Our algorithm mainly modifies the sub-procedure SKELFWD. The improved version of SKELFWD procedure takes an additional input argument Q, and returns an additional output argument that is stored as a set P by the calling SCCFIND procedure. The calling function passes the set U as Q. The way the output P is computed is as follows: at the end of the forward search we have the following assignment: $P :=$ FWSet $\cap\, Q$. After the forward search, the skeleton of the forward set is computed with the help of a stack. The elements of the stacks are sets of states stored in the forward search. The spine set computation is similar to SKELFWD, the difference is that when elements are popped of the stack, we check if there is a non-empty intersection with P, if so, we break the loop and return. Moreover, for the backward searches in SCCFIND we initialize SCC by P rather than s. We refer to the new sub-procedure as IMPROVEDSKELFWD (detailed pseudocode in the following subsection).

Correctness   Since s is the last element of the spine set U, and P is the intersection of a forward search from s with U, it means that all elements of P are both reachable from s (since P is a subset of FW(s)) and can reach s (since P is a subset of U). It follows that P

is a subset of the scc containing s. The part of the spine-set computation that was skipped, would have consisted of states that are reachable from s, and can reach P, and thus belong to the scc containing s. Since the scc containing s is to be removed after the call, not computing the spine-set beyond P does not change the future function calls, i.e., the value of U', since the omitted parts of NewSet are in the scc containing s. The modification of starting the backward search from P does not change the result, since P will anyway be included in the backward search. So the IMPROVEDSYMBOLICSCC algorithm gives the same result as SYMBOLICSCC, and the correctness follows from Theorem 7.

Symbolic steps analysis   We present two upper bounds on the number of symbolic steps of the algorithm. Intuitively following are the symbolic operations that need to be accounted for: (1) when a state is included in a spine set for the first time in IMPROVEDSKELFWD sub-procedure which has two parts: the first part is the forward search and the second part is computing the skeleton of the forward set; (2) when a state is already in a spine set and is found in forward search of IMPROVEDSKELFWD and (3) the backward search for determining the scc. We now present the number of symbolic steps analysis for IMPROVEDSYMBOLICSCC.

1. There are two parts of IMPROVEDSKELFWD, (i) a forward search and (ii) a backward search for skeleton computation of the forward set. For the backward search, we show that the number of steps performed equals the size of NewSet computed. One key idea of the analysis is the proof where we show that a state becomes part of spine-set at most once, as compared to the algorithm of [12] where a state can be part of spine-set at most twice. Because, when it is already part of a spine-set, it will be included in P and we stop the computation of spine-set when an element of P gets included. We now split the analysis in two cases: (a) states that are included in spine-set, and (b) states that are not included in spine-set.
   (a) We charge one symbolic step for the backward search of IMPROVEDSKELFWD (spine-set computation) to each element when it first gets inserted in a spine-set. For the forward search, we see that the number of steps performed is the size of spine-set that would have been computed if we did not stop the skeleton computation. But by stopping it, we are only omitting states that are part of the scc. Hence we charge one symbolic step to each state getting inserted into spine-set for the first time and each state of the scc. Thus, a state getting inserted in a spine-set is charged two symbolic steps (for forward and backward search) of IMPROVEDSKELFWD the first time it is inserted.
   (b) A state not inserted in any spine-set is charged one symbolic step for backward search which determines the scc.
   Along with the above symbolic steps, one step is charged to each state for the forward search in IMPROVEDSKELFWD at the time its scc is being detected. Hence each state gets charged at most three symbolic steps. Besides, for computing NewState, one symbolic step is required per scc found. Thus the total number of symbolic steps is bounded by $3 \cdot |S| + N(G)$, where $N(G)$ is the number of scc's of G.
2. Let D^* be the sum of diameters of the scc's in a G. Consider a scc with diameter d. In any scc the spine-set is a shortest path, and hence the size of the spine-set is bounded by d. Thus the three symbolic steps charged to states in spine-set contribute to at most $3 \cdot d$ symbolic steps for the scc. Moreover, the number of iterations of forward search of IMPROVEDSKELFWD charged to states belonging to the scc being computed are at most d. The number of iterations of the backward search to compute the scc is also at most d. Hence, the two symbolic steps charged to states not in any spine-set also contribute at most $2 \cdot d$ symbolic steps for the scc. Finally, computation of NewSet takes one

Table 1 The average symbolic steps required by symbolic algorithms for MDPs with Büchi objectives

Number of states	Classical	SYMBIMPRALGO	SMDvSYMBIMPRALGO	SYMBIMPRWINLOSE
5000	30731	3478	3898	3573
10000	103977	6622	7490	6815
20000	306015	12010	13212	13687

Table 2 The average running time required in sec by symbolic algorithms for MDPs with Büchi objectives

Number of states	Classical	SYMBIMPRALGO	SMDvSYMBIMPRALGO	SYMBIMPRWINLOSE
5000	78.8	9.7	10.2	10.8
10000	563.7	40.3	43.0	46.1
20000	3974.4	186.4	192.3	217.4

symbolic step per scc. Hence we have $5 \cdot d + 1$ symbolic steps for a scc with diameter d. We thus obtain an upper bound of $5D^* + N(G)$ symbolic steps.

It is straightforward to argue that the number of symbolic steps of IMPROVEDSCCFIND is at most the number of symbolic steps of SCCFIND. The detailed pseudocode and technical details of the running time analysis is presented in Appendix.

Theorem 8 *Let $G = (S, E)$ be a directed graph. The algorithm* IMPROVEDSYMBOLICSCC *correctly computes the scc decomposition of G in* $\min\{3 \cdot |S| + N(G), 5 \cdot D^*(G) + N(G)\}$ *symbolic steps, where $D^*(G)$ is the sum of diameters of the scc's of G, and $N(G)$ is the number of scc's in G.*

Remark 4 Observe that in the worst case SCCFIND takes $5 \cdot n$ symbolic steps, whereas IMPROVEDSCCFIND takes at most $4 \cdot n$ symbolic steps. Thus our algorithm improves the constant of the number of linear symbolic steps required for symbolic scc decomposition.

6 Experimental results

In this section we present our experimental results. We first present the results for symbolic algorithms for MDPs with Büchi objectives and then for symbolic scc decomposition.

Symbolic algorithm for MDPs with Büchi objectives   We implemented all the symbolic algorithms (including the classical one) and ran the algorithms on randomly generated graphs. If we consider arbitrarily randomly generated graphs, then in most cases it gives rise to trivial MDPs. Hence we generated more structured MDP graphs. First we generated a large number of MDPs and as a first step chose the MDP graphs where all the algorithms required large number of symbolic steps, and then generated large number of MDP graphs randomly by small perturbations of the graphs chosen in the first step. Our results of average symbolic steps required are shown in Table 1 and show that the new algorithms perform significantly better than the classical algorithm. The running time comparison is given in Table 2.

Table 3 The average symbolic steps required for scc computation

Number of states	Algorithm from [12]	Our algorithm	Percentage improvement
10000	1043	878	15.83
25000	2649	2264	14.53
50000	6299	5394	14.36

Table 4 The average running time required in sec for scc computation

Number of states	Algorithm from [12]	Our algorithm	Percentage improvement
10000	7.7	6.3	17.53
25000	48.3	40.0	16.98
50000	180.8	152.5	15.67

Symbolic scc computation We implemented the symbolic scc decomposition algorithm from [12] and our new symbolic algorithm. A comparative study of the algorithm of [12] and the algorithm of [3] was done in [21], and it was found that the performances were comparable. Hence we only perform the comparison of the algorithm of [12] and our new algorithm. We ran the algorithms on randomly generated graphs. Again arbitrarily randomly generated graphs in many cases gives rise to graphs that are mostly disconnected or completely connected. Hence we generated random graphs by first constructing a topologically sorted order of the scc's and then adding edges randomly respecting the topologically sorted order. Our results of average symbolic steps are shown in Table 3 and shows that our new algorithm performs better (around 15 % improvement over the algorithm of [12]). The running time comparison is shown in Table 4.

In all cases, our implementations were the basic implementation of the algorithms, and more optimized implementations would lead to improved performance results. The source codes, sample examples for the experimental results, and other details of the implementation are available at http://www.cs.cmu.edu/~nkshah/SymbolicMDP.

Symbolic scc for academic benchmarks We also considered the symbolic scc computation on the academic benchmark graphs: we considered the VLTS benchmarks (obtained from http://www.inrialpes.fr/vasy/cadp/resources/benchmark_bcg.html) which are also used for scc computation algorithms comparison in [1]. Our results are shown in Table 5. If the input graph has one scc, then the algorithm from [12] and our improved algorithm is the same. In the experimental results in all cases where the number of scc's is one there is no improvement (as the algorithms coincide), and in one other degenerate case where all scc's are self-loops there was no improvement. Other than the cases where there were no improvements, in one out of twenty-three cases there was less than 1 % improvement, in five cases between 1.5–3 % improvement, in four cases between 3–5 % improvement, in seven cases between 5–10 % improvement, and in six cases more than 10 % improvement.

7 Conclusion

In this work we considered a core problem of probabilistic verification which is to compute the set of almost-sure winning states in MDPs with Büchi objectives. We presented the first

Table 5 Symbolic steps for scc decomposition for VLTS benchmarks

Name	#states	#transitions	#SCC	MaxSCC	Algo. from [12]	Our algo.	% Impr
vasy_0_1	289	1224	49	16	547	492	10.05
vasy_1_4	1183	4464	25	319	644	535	16.93
vasy_5_9	5486	9676	2525	450	11873	11407	3.92
vasy_8_24	8879	24411	2197	2184	11898	10940	8.05
vasy_8_38	8921	38424	7412	6	28785	28187	2.08
vasy_10_56	10849	56156	2	10848	136	105	22.79
vasy_18_73	18746	73043	33	18714	333	282	15.32
vasy_25_25	25217	25216	25217	1	126083	126083	0.00
vasy_40_60	40006	60007	1	40006	120017	120017	0.00
vasy_52_318	52268	318126	11828	20368	48288	46786	3.11
vasy_65_2621	65537	2621480	2	65536	106	77	27.36
vasy_66_1302	66929	1302664	66929	1	292459	277794	5.01
vasy_69_520	69754	520633	52138	8	271454	253484	6.62
vasy_83_325	83436	325584	253	75078	1861	1801	3.22
vasy_116_368	116456	368569	116456	1	543946	513510	5.60
vasy_164_1619	164865	1619204	31965	4096	117632	116906	0.62
vasy_166_651	166464	651168	484	149769	3591	3512	2.20
vasy_386_1171	386496	1171872	8089	86733	34595	33959	1.84
vasy_574_13561	574057	13561040	460	573248	1473	1448	1.70
vasy_6020_19353	6020550	19353474	2041	6013920	20543	18557	9.67
cwi_1_2	1952	2387	1	1952	104	104	0.00
cwi_3_14	3996	14552	3996	1	18478	17503	5.28
cwi_142_925	142472	925429	142472	1	617584	591791	4.18
cwi_214_684	214202	684419	74016	86560	483052	441562	8.59
cwi_371_641	371804	641565	2	371803	678	473	30.24
cwi_2165_8723	2165446	8723465	47926	423505	191327	186847	2.34
cwi_7838_59101	7838608	59101007	1	7838608	822	822	0.00

symbolic sub-quadratic algorithm for the problem, and also a new symbolic sub-quadratic algorithm (IMPRWINLOSE algorithm). As compared to all previous algorithms which identify the almost-sure winning states upon termination, the IMPRWINLOSE algorithm can potentially discover almost-sure winning states in intermediate steps as well. Finally we considered another core graph theoretic problem in verification which is the symbolic scc decomposition problem. We presented an improved algorithm for the problem. The previous best known algorithm for the problem required $5 \cdot n$ symbolic steps in the worst case and our new algorithm takes at most $4 \cdot n$ symbolic steps, where n is the number of states of the graph. Our basic implementation shows that our new algorithms perform favorably over the old algorithms. Optimized implementations of the new algorithms and detailed experimental studies would be an interesting direction for future work.

Appendix

8.1 Technical details of improved symbolic scc algorithm

The pseudocode of SCCFIND is formally given as Algorithm 4. The correctness analysis and the analysis of the number of symbolic steps is given in [12]. The pseudocode of IMPROVEDSCCFIND is formally given as Algorithm 5. The main changes of IMPROVED-SCCFIND from SCCFIND are as follows: (1) instead of SKELFWD the algorithm IMPROVEDSCCFIND calls procedure IMPROVEDSKELFWD that returns an additional set P and IMPROVEDSKELFWD is invoked with an additional argument that is U; (2) in line 4 of IMPROVEDSCCFIND the set SCC is initialized to P instead of s. The main difference of IMPROVEDSKELFWD from SKELFWD is as follows: (1) the set P is computed in line 4 of IMPROVEDSKELFWD as FWSet $\cap$ Q, where Q is the set passed by IMPROVEDSCCFIND as the argument; and (2) in the while loop it is checked if the element popped intersects with P and if yes, then the procedure breaks the while loop. The correctness argument from the correctness of SCCFIND is already shown in Sect. 5.1.

Symbolic steps analysis We now present the detailed symbolic steps analysis of the algorithm. As noted in Sect. 3.2, common symbolic operations on a set of states are Pre, Post and CPre. We note that these operations involve symbolic sets of $2 \cdot \log(n)$ variables, as compared to symbolic sets of $\log(n)$ variables required for operations such as union, intersection and set difference. Thus only Pre, Post and CPre are counted as symbolic steps, as done in [12]. The total number of other symbolic operations is also $O(|S|)$. We note that only lines 5 and 10 of IMPROVEDSCCFIND and lines 3.3 and 7.3 of IMPROVEDSKELFWD involve Pre and Post operations.

In the following, we charge the costs of these lines to states in order to achieve the $3 \cdot |S| + N(G)$ bound for symbolic steps. We define subspine-set as NewSet returned by IMPROVEDSKELFWD and show the following result.

Lemma 4 *For any spine-set U and its end vertex u, T is a subspine-set iff $U \setminus T \subseteq \mathsf{SCC}(u)$.*

Proof Note that while constructing a subspine-set T, we stop the construction when we find any state $v \in$ FWSet $\cap$ U from the spine set. Now clearly since $v \in U$, there is a path from v to u. Also, since we found this state in $FW(u)$, there is a path from u to v. Hence, $v \in \mathsf{SCC}(u)$. Also, each state that we are omitting by stopping construction of T has the property that there is a path from u to that state and a path from that state to v. This implies that all the states we are omitting in construction of T are in $\mathsf{SCC}(u)$. □

Note that since we pass NewSet \ SCC in the subsequent call to IMPROVEDSCCFIND, it will actually be a spine set for the reduced problem. In the following lemma we show that any state can be part of subspine-set at most once, as compared to twice in the SCCFIND procedure in [12]. This lemma is one of the key points that lead to the improved analysis of symbolic steps required.

Lemma 5 *Any state v can be part of subspine-set at most once.*

Proof In [12], the authors show that any state v can be included in spine sets at most twice in SKELFWD. The second time the state v is included is in line 6 of SKELFWD when the SCC(v) of the state is to be found. In contrast, IMPROVEDSKELFWD checks intersection of

---

Algorithm 4 SCCFIND

    Input: $(S, E, \langle U, s \rangle)$, i.e., a graph (S, E) with spine set (U, s).
    Output: SCCPartition i.e. the set of SCCs of the graph (S, E)
    Initialize SCCPartition $:= \emptyset$; FWSet $:= \emptyset$;
  1. **if** $(S = \emptyset)$ **then**
    1.1 **return**;
  2. **if** $(U = \emptyset)$ **then**
    2.1 $s := \text{pick}(S)$
  3. $\langle$FWSet, NewSet, NewState$\rangle := \text{SKELFWD}(S, E, s)$
  4. SCC $= s$
  5. **while** $(((\text{Pre}(\text{SCC}) \cap \text{FWSet}) \setminus \text{SCC}) \neq \emptyset)$ **do**
    5.1 SCC $:= \text{SCC} \cup (\text{Pre}(\text{SCC}) \cap \text{FWSet})$
  6. SCCPartition $:= \text{SCCPartition} \cup \{\text{SCC}\}$
(Recursive call on $S \setminus \text{FWSet}$)
  7. $S' := S \setminus \text{FWSet}$
  8. $E' := E \cap (S' \times S')$
  9. $U' := U \setminus \text{SCC}$
  10. $s' := \text{Pre}(\text{SCC} \cap U) \cap (S \setminus \text{SCC})$
  11. SCCPartition $:= \text{SCCPartition} \cup \text{SCCFIND}(S', E', \langle U', s' \rangle)$
(Recursive call on FWSet $\setminus$ SCC)
  12. $S' := \text{FWSet} \setminus \text{SCC}$
  13. $E' := E \cap (S' \times S')$
  14. $U' := \text{NewSet} \setminus \text{SCC}$
  15. $s' := \text{NewState} \setminus \text{SCC}$
  16. SCCPartition $:= \text{SCCPartition} \cup \text{SCCFIND}(S', E', \langle U', s' \rangle)$
  17. Return SCCPartition

Procedure SKELFWD
    Input: (S, E, s), i.e., a graph (S, E) with a state $s \in S$.
    Output: $\langle$FWSet, NewSet, NewState$\rangle$,
       i.e. forward set FWSet, new spine-set NewSet and NewState $\in$ NewSet
  1. Let stack be an empty stack of sets of nodes
  2. $L := s$
  3. **while** $(L \neq \emptyset)$ **do**
    3.1 Push(stack, L)
    3.2 FWSet $:= \text{FWSet} \cup L$
    3.3 $L := \text{Post}(L) \setminus \text{FWSet}$
  4. $L := \text{Pop}(\text{stack})$
  5. NewSet $:= \text{NewState} := \text{pick}(L)$
  6. **while** $(\text{stack} \neq \emptyset)$ **do**
    6.1 $L := \text{Pop}(\text{stack})$
    6.2 NewSet $:= \text{NewSet} \cup \text{pick}(\text{Pre}(\text{NewSet}) \cap L)$
7. **return** $\langle$FWSet, NewSet, NewState$\rangle$

---

the subspine-set being constructed with the set P that contains the states of SCC(v) which are already in a subspine-set. When this happens, it stops the construction of the subspine-set. Now if v is already included in the subspine-set, then it will be part of P and would not be included in subspine-set again. Hence, v can be part of subspine-set at most once. $\qquad \square$

Lemma 6 *States added in* SCC *by iteration of line* 5 *of* IMPROVEDSCCFIND *are exactly the states which are not part of any subspine-set.*

Proof We see that in line 5 of IMPROVEDSCCFIND, we start from SCC $= P$ and then we find the SCC by backward search. Also, P has all the states from SCC which are part of any subspine-set. Hence, the extra states that are added in SCC are states which are never included in a subspine-set. $\qquad \square$

Algorithm 5 IMPROVEDSCCFIND

Input: $(S, E, \langle U, s \rangle)$, i.e., a graph (S, E) with spine set (U, s).
Output: SCCPartition, the set of SCCs of the graph (S, E)
Initialize SCCPartition $:= \emptyset$; FWSet $:= \emptyset$;
1. **if** $(S = \emptyset)$ **then**
    1.1 **return**;
2. **if** $(U = \emptyset)$ **then**
    2.1 $s := \text{pick}(S)$
3. $\langle$FWSet, NewSet, NewState, $P\rangle := $ IMPROVEDSKELFWD(S, E, U, s)
4. SCC $= P$
5. **while** $(((\text{Pre}(\text{SCC}) \cap \text{FWSet}) \setminus \text{SCC}) \neq \emptyset)$ **do**
    5.1 SCC $:= \text{SCC} \cup (\text{Pre}(\text{SCC}) \cap \text{FWSet})$
6. SCCPartition $:= \text{SCCPartition} \cup \{\text{SCC}\}$
(Recursive call on $S \setminus \text{FWSet}$)
  7. $S' := S \setminus \text{FWSet}$
  8. $E' := E \cap (S' \times S')$
  9. $U' := U \setminus \text{SCC}$
  10. $s' := \text{Pre}(\text{SCC} \cap U) \cap (S \setminus \text{SCC})$
  11. SCCPartition $:= \text{SCCPartition} \cup $ IMPROVEDSCCFIND$(S', E', \langle U', s' \rangle)$
(Recursive call on FWSet $\setminus$ SCC)
  12. $S' := \text{FWSet} \setminus \text{SCC}$
  13. $E' := E \cap (S' \times S')$
  14. $U' := \text{NewSet} \setminus \text{SCC}$
  15. $s' := \text{NewState} \setminus \text{SCC}$
  16. SCCPartition $:= \text{SCCPartition} \cup $ IMPROVEDSCCFIND$(S', E', \langle U', s' \rangle)$

Procedure IMPROVEDSKELFWD
Input: (S, E, Q, s), i.e., a graph (S, E) with a set Q and a state $s \in S$.
Output: $\langle$FWSet, NewSet, NewState, $P\rangle$
1. Let stack be an empty stack of sets of nodes
2. $L := s$
3. **while** $(L \neq \emptyset)$ **do**
    3.1 Push(stack, L)
    3.2 FWSet $:= \text{FWSet} \cup L$
    3.3 $L := \text{Post}(L) \setminus \text{FWSet}$
4. $P := \text{FWSet} \cap Q$
5. $L := \text{Pop}(\text{stack})$
6. NewSet $:= $ NewState $:= \text{pick}(L)$
7. **while** $(\text{stack} \neq \emptyset)$ **do**
    7.1 $L := \text{Pop}(\text{stack})$
    7.2 **if** $(L \cap P \neq \emptyset)$ **then**
        7.2.1 **break while loop**
    7.3 **else** NewSet $:= \text{NewSet} \cup \text{pick}(\text{Pre}(\text{NewSet}) \cap L)$
8. **return** $\langle$FWSet, NewSet, NewState, $P\rangle$

Charging symbolic steps to states   We now consider three cases to charge symbolic steps to states and scc's.

1. *Charging states included in subspine-set.* First, we see that the number of times the loop of line 3 in IMPROVEDSKELFWD is executed is equal to the size of the spine set that SKELFWD would have computed. Using Lemma 4, we can charge one symbolic step to each state of the subspine-set and each state of the SCC that is being computed. Now, the number of times line 7.3 of IMPROVEDSKELFWD is executed equals the size of subspine-set that is computed. Hence, we charge one symbolic step to each state of subspine-set for this line.

    Now we summarize the symbolic steps charged to each state which is part of some subspine-set. First time when a state gets into a subspine-set, it is charged two steps,

one for line 3.3 and one for line 7.3 of IMPROVEDSKELFWD. If its SCC is not found in the same call to IMPROVEDSCCFIND, then it comes into action once again when its SCC is being found. By Lemma 5, it is never again included in a subspine set. Hence in this call to IMPROVEDSKELFWD, it is only charged one symbolic step for line 3.3 and none for line 7.3 as line 7.3 is charged to states that become part of the newly constructed subspine-set. Also because of Lemma 6, since this state is in a subspine-set, it is not charged anything for line 5 of IMPROVEDSCCFIND. Hence, a state that occurs in any subspine-set is charged at most three symbolic steps.

2. *Charging states not included in subspine-set.* For line 5 of IMPROVEDSCCFIND, the number of times it is executed is the number of states that are added to SCC after initialization to SCC = P. Using Lemma 6, we charge one symbolic step to each state of this SCC that is never a part of any subspine-set. Also, we might have charged one symbolic step to such a state for line 3.3 of IMPROVEDSKELFWD when we called it. Hence, each such state is charged at most two symbolic steps.

3. *Charging SCCs.* For line 10 of IMPROVEDSCCFIND, we see that it is executed only once in a call to IMPROVEDSCCFIND that computes a SCC. Hence, the total number of times line 10 is executed equals $N(G)$, the number of SCCs of the graph. Hence, we charge each SCC one symbolic step for this line.

The above argument shows that the number of symbolic steps that the algorithm IMPROVEDSCCFIND requires is at most $3 \cdot |S| + N(G)$. This completes the formal proof of Theorem 8.

We now present an example that presents a family of graphs with $k \cdot n$ states, where the SCCFIND algorithm takes almost $5 \cdot k \cdot n$ symbolic steps, whereas the IMPROVEDSCCFIND algorithm takes at most $3 \cdot k \cdot n + n$ symbolic steps.

Example 1 Let $k, n \in \mathbb{N}$. Consider a graph with $k \cdot n$ states such that the states are numbered from 1 to $k \cdot n$. The edges are as follows: (1) for all states $1 \leq i \leq k \cdot n - 1$, there is an edge $(i, i + 1)$ (i.e., the states are all in a line); and (2) for all $1 \leq i \leq n$, there is an edge from state $k \cdot i$ to state $(i - 1) \cdot k + 1$. We will show that the SCCFIND algorithm requires roughly $(5 \cdot k - 1) \cdot n$ symbolic steps on this graph. Note the number of scc's in this graph is n, and hence by Theorem 8 the IMPROVEDSCCFIND algorithm takes at most $3 \cdot k \cdot n + n$ symbolic steps.

We now analyze the symbolic steps required by the SCCFIND algorithm, and our analysis is in two steps.

1. *Step 1.* In the beginning, starting from the state 1 of the graph, the algorithm performs a forward and backward search to find the spine set. This will have a cost of two symbolic steps per state (one symbolic step while going forward, and one symbolic step while going backward), except for the first vertex which gets charged only one symbolic step (it does not get charged while going backwards). This gives a cost of $2 \cdot k \cdot n - 1$ symbolic steps. After this, the first scc will be found with an additional cost of only k, and the first discovered scc consists of states $\{1, 2, \ldots, k\}$. Hence the total symbolic steps required is $2 \cdot k \cdot n - 1 + k$.

2. *Step 2.* After Step 1, the algorithm will start finding scc's from the end of the spine set. So consider the set of last k states from $k \cdot (n - 1) + 1$ to $n \cdot k$. The algorithm will pick the last state $n \cdot k$, find a spine set, consisting of the last k states (states $k \cdot (n - 1) + 1$ to $n \cdot k$). This will have a cost of $2 \cdot k - 1$ symbolic steps (k symbolic steps for the forward search, and $k - 1$ symbolic steps for the backward search). After this, the algorithm will find the scc containing the last state $n \cdot k$ (i.e., the scc that consists of states $\{k \cdot (n - 1) + 1, k \cdot (n - 1) + 2, \ldots, n \cdot k\}$), and this takes k symbolic steps as there are k vertices in the

scc. Now Step 2 is repeated with state $k \cdot (n - 1)$, and then repeated for state $k \cdot (n - 2)$ and so on. So the cost for every scc, except the very first one, is $2k - 1 + k = 3k - 1$. Since there are n scc's, the total number of symbolic steps required for Step 2 is at least $(3 \cdot k - 1) \cdot (n - 1)$.

Hence the total symbolic steps required for the algorithm is at least

$$2 \cdot k \cdot n - 1 + k + (3 \cdot k - 1) \cdot (n - 1) = (5 \cdot k - 1) \cdot n - 2 \cdot k.$$

Note that with $k = n$, SCCFIND takes at least $5 \cdot n^2 - 3 \cdot n$ symbolic steps, whereas the IMPROVEDSCCFIND takes at most $3 \cdot n^2 + n$ symbolic steps.

References

1. Barnat J, Chaloupka J, van de Pol J (2011) Distributed algorithms for SCC decomposition. J Log Comput 21(1):23–44
2. Bianco A, de Alfaro L (1995) Model checking of probabilistic and nondeterministic systems. In: FSTTCS 95. LNCS, vol 1026. Springer, Berlin, pp 499–513
3. Bloem R, Gabow HN, Somenzi F (2000) An algorithm for strongly connected component analysis in log symbolic steps. In: FMCAD, pp 37–54
4. Chatterjee K, Henzinger M (2011) Faster and dynamic algorithms for maximal end-component decomposition and related graph problems in probabilistic verification. In: SODA, pp 1318–1336
5. Chatterjee K, Jurdziński M, Henzinger TA (2003) Simple stochastic parity games. In: CSL'03. LNCS, vol 2803. Springer, Berlin, pp 100–113
6. Chatterjee K, Jurdziński M, Henzinger TA (2004) Quantitative stochastic parity games. In: SODA'04. SIAM, Philadelphia, pp 121–130
7. Courcoubetis C, Yannakakis M (1995) The complexity of probabilistic verification. J ACM 42(4):857–907
8. de Alfaro L (1997) Formal verification of probabilistic systems. PhD thesis, Stanford University
9. de Alfaro L, Faella M, Majumdar R, Raman V (2005) Code-aware resource management. In: EM-SOFT 05. ACM, New York
10. de Alfaro L, Roy P (2007) Magnifying-lens abstraction for Markov decision processes. In: CAV, pp 325–338
11. Filar J, Vrieze K (1997) Competitive Markov decision processes. Springer, Berlin
12. Gentilini R, Piazza C, Policriti A (2003) Computing strongly connected components in a linear number of symbolic steps. In: SODA, pp 573–582
13. Howard H (1960) Dynamic programming and Markov processes. MIT Press, Cambridge
14. Immerman N (1981) Number of quantifiers is better than number of tape cells. J Comput Syst Sci 22:384–406
15. Kemeny JG, Snell JL, Knapp AW (1966) Denumerable Markov chains. Van Nostrand, Princeton
16. Kwiatkowska M, Norman G, Parker D (2000) Verifying randomized distributed algorithms with prism. In: Workshop on advances in verification (WAVE'00)
17. Pogosyants A, Segala R, Lynch N (2000) Verification of the randomized consensus algorithm of Aspnes and Herlihy: a case study. Distrib Comput 13(3):155–186
18. Puterman ML (1994) Markov decision processes. Wiley, New York
19. Segala R (1995) Modeling and verification of randomized distributed real-time systems. PhD thesis, MIT Technical Report MIT/LCS/TR-676
20. Sleator DD, Tarjan RE (1983) A data structure for dynamic trees. J Comput Syst Sci 26(3):362–391
21. Somenzi F. Personal communication
22. Somenzi F (1998) Colorado university decision diagram package. http://vlsi.colorado.edu/pub/
23. Stoelinga MIA (2002) Fun with FireWire: experiments with verifying the IEEE1394 root contention protocol. In: Formal aspects of computing
24. Thomas W (1997) Languages, automata, and logic. In: Rozenberg G, Salomaa A (eds) Beyond words. Handbook of formal languages, vol 3. Springer, Berlin, pp 389–455. Chap 7