

Symbolic Algorithms for Graphs and Markov Decision Processes with Fairness Objectives

Krishnendu Chatterjee¹, Monika Henzinger², Veronika Loitzenbauer³,
Simin Oraee⁴, and Viktor Toman¹

¹ IST Austria, Klosterneuburg, Austria

² University of Vienna, Austria

³ Johannes Kepler University Linz, Austria

⁴ Max Planck Institute for Software Systems, Kaiserslautern, Germany

Abstract. Given a model and a specification, the fundamental model-checking problem asks for algorithmic verification of whether the model satisfies the specification. We consider graphs and Markov decision processes (MDPs), which are fundamental models for reactive systems. One of the very basic specifications that arise in verification of reactive systems is the strong fairness (aka Streett) objective. Given different types of requests and corresponding grants, the objective requires that for each type, if the request event happens infinitely often, then the corresponding grant event must also happen infinitely often. All ω -regular objectives can be expressed as Streett objectives and hence they are canonical in verification. To handle the state-space explosion, symbolic algorithms are required that operate on a succinct implicit representation of the system rather than explicitly accessing the system. While explicit algorithms for graphs and MDPs with Streett objectives have been widely studied, there has been no improvement of the basic symbolic algorithms. The worst-case numbers of symbolic steps required for the basic symbolic algorithms are as follows: quadratic for graphs and cubic for MDPs. In this work we present the first sub-quadratic symbolic algorithm for graphs with Streett objectives, and our algorithm is sub-quadratic even for MDPs. Based on our algorithmic insights we present an implementation of the new symbolic approach and show that it improves the existing approach on several academic benchmark examples.

1 Introduction

In this work we present faster symbolic algorithms for graphs and Markov decision processes (MDPs) with strong fairness objectives. For the fundamental model-checking problem, the input consists of a *model* and a *specification*, and the algorithmic verification problem is to check whether the model *satisfies* the specification. We first describe the specific model-checking problem we consider and then our contributions.

Models: Graphs and MDPs. Two standard models for reactive systems are graphs and Markov decision processes (MDPs). Vertices of a graph represent states of a reactive system, edges represent transitions of the system, and infinite paths of the graph represent non-terminating trajectories of the reactive system. MDPs extend graphs with probabilistic transitions that represent reactive systems with uncertainty. Thus graphs and MDPs are the de-facto model of reactive systems with nondeterminism, and nondeterminism with stochastic aspects, respectively [18,3].

Specification: Strong Fairness (aka Streett) Objectives. A basic and fundamental property in the analysis of reactive systems is the *strong fairness condition*, which informally requires that if events are enabled infinitely often, then they must be executed infinitely often. More precisely, the strong fairness conditions (aka Streett objectives) consist of k types of requests and corresponding grants, and the objective requires that for each type if the request happens infinitely often, then the corresponding grant must also happen infinitely often. After safety, reachability, and liveness, the strong fairness condition is one of the most standard properties that arise in the analysis of reactive systems, and chapters of standard textbooks in verification are devoted to it (e.g., [18, Chapter 3.3], [31, Chapter 3], [2, Chapters 8, 10]). Moreover, all ω -regular objectives can be described by Streett objectives, e.g., LTL formulas and non-deterministic ω -automata can be translated to deterministic Streett automata [33] and efficient translation has been an active research area [15,22,27]. Thus Streett objectives are a canonical class of objectives that arise in verification.

Satisfaction. The basic notions of satisfaction for graphs and MDPs are as follows: For graphs the notion of satisfaction requires that there is a trajectory (infinite path) that belongs to the set of paths described by the Streett objective. For MDPs the satisfaction requires that there is a policy to resolve the nondeterminism such that the Streett objective is ensured almost-surely (with probability 1). Thus the algorithmic model-checking problem of graphs and MDPs with Streett objectives is a core problem in verification.

Explicit vs Symbolic Algorithms. The traditional algorithmic studies consider *explicit* algorithms that operate on the explicit representation of the system. In contrast, *implicit* or *symbolic* algorithms only use a set of predefined operations and do not explicitly access the system [19]. The significance of symbolic algorithms in verification is as follows: to combat the state-space explosion, large systems must be succinctly represented implicitly and then symbolic algorithms are scalable, whereas explicit algorithms do not scale as it is computationally too expensive to even explicitly construct the system.

Relevance. In this work we study symbolic algorithms for graphs and MDPs with Streett objectives. Symbolic algorithms for the analysis of graphs and MDPs are at the heart of many state-of-the-art tools such as SPIN, NuSMV for graphs [26,17] and PRISM, LiQuor, Storm for MDPs [28,16,21]. Our contributions are related to the algorithmic complexity of graphs and MDPs with Streett objectives for symbolic algorithms. We first present previous results and then our contributions.

Previous Results. The most basic algorithm for the problem for graphs is based on repeated SCC (strongly connected component) computation, and informally can be described as follows: for a given SCC, (a) if for every request type that is present in the SCC the corresponding grant type is also present in the SCC, then the SCC is identified as “good”, (b) else vertices of each request type that has no corresponding grant type in the SCC are removed, and the algorithm recursively proceeds on the remaining graph. Finally, reachability to good SCCs is computed. The current best-known symbolic algorithm for SCC computation requires $O(n)$ symbolic steps, for graphs with n vertices [24], and moreover, the algorithm is optimal [14]. For MDPs, the SCC computation has to be replaced by MEC (maximal end-component) computation, and the current best-known symbolic algorithm for MEC computation requires $O(n^2)$ symbolic steps. While there have been several explicit algorithms for graphs with Streett

objectives [25,12], MEC computation [8,9,10], and MDPs with Streett objectives [7], as well as symbolic algorithms for MDPs with Büchi objectives [11], the current best-known bounds for symbolic algorithms with Streett objectives are obtained from the basic algorithms, which are $O(n \cdot \min(n, k))$ for graphs and $O(n^2 \cdot \min(n, k))$ for MDPs, where k is the number of types of request-grant pairs.

Our Contributions. In this work our main contributions are as follows:

- We present a symbolic algorithm that requires $O(n \cdot \sqrt{m \log n})$ symbolic steps, both for graphs and MDPs, where m is the number of edges. In the case $k = O(n)$, the previous worst-case bounds are quadratic ($O(n^2)$) for graphs and cubic ($O(n^3)$) for MDPs. In contrast, we present the first sub-quadratic symbolic algorithm both for graphs as well as MDPs. Moreover, in practice, since most graphs are sparse (with $m = O(n)$), the worst-case bounds of our symbolic algorithm in these cases are $O(n \cdot \sqrt{n \log n})$. Another interesting contribution of our work is that we also present an $O(n \cdot \sqrt{m})$ symbolic steps algorithm for MEC decomposition, which is relevant for our results as well as of independent interest, as MEC decomposition is used in many other algorithmic problems related to MDPs. Our results are summarized in Table 1.
- While our main contribution is theoretical, based on the algorithmic insights we also present a new symbolic algorithm implementation for graphs and MDPs with Streett objectives. We show that the new algorithm improves (by around 30%) the basic algorithm on several academic benchmark examples from the VLTS benchmark suite [20].

Table 1: Symbolic algorithms for Streett objectives and MEC decomposition.

Problem	Symbolic Operations		Reference
	Basic Algorithm	Improved Algorithm	
Graphs with Streett	$O(n \cdot \min(n, k))$	$O(n\sqrt{m \log n})$	Theorem 2
MDPs with Streett	$O(n^2 \cdot \min(n, k))$	$O(n\sqrt{m \log n})$	Theorem 4
MEC decomposition	$O(n^2)$	$O(n\sqrt{m})$	Theorem 3

Technical Contributions. The two key technical contributions of our work are as follows:

- *Symbolic Lock Step Search:* We search for newly emerged SCCs by a local graph exploration around vertices that lost adjacent edges. In order to find small new SCCs first, all searches are conducted “in parallel”, i.e., in lock-step, and the searches stop as soon as the first one finishes successfully. This approach has successfully been used to improve explicit algorithms [25,13,9,7]. Our contribution is a non-trivial symbolic variant (Section 3) which lies at the core of the theoretical improvements.
- *Symbolic Interleaved MEC Computation:* For MDPs the identification of vertices that have to be removed can be interleaved with the computation of MECs such that in each iteration the computation of SCCs instead of MECs is sufficient to make progress [7]. We present a symbolic variant of this interleaved computation. This interleaved MEC computation is the basis for applying the lock-step search to MDPs.

2 Definitions

2.1 Basic Problem Definitions

Markov Decision Processes (MDPs) and Graphs. An MDP $P = ((V, E), (V_1, V_R), \delta)$ consists of a finite directed graph $G = (V, E)$ with a set of n vertices V and a set of m edges E , a partition of the vertices into *player 1 vertices* V_1 and *random vertices* V_R , and a probabilistic transition function δ . We call an edge (u, v) with $u \in V_1$ a *player 1 edge* and an edge (v, w) with $v \in V_R$ a *random edge*. For $v \in V$ we define $In(v) = \{w \in V \mid (w, v) \in E\}$ and $Out(v) = \{w \in V \mid (v, w) \in E\}$. The probabilistic transition function is a function from V_R to $\mathcal{D}(V)$, where $\mathcal{D}(V)$ is the set of probability distributions over V and a random edge $(v, w) \in E$ if and only if $\delta(v)[w] > 0$. Graphs are a special case of MDPs with $V_R = \emptyset$.

Plays and Strategies. A *play* or infinite path in P is an infinite sequence $\omega = \langle v_0, v_1, v_2, \dots \rangle$ such that $(v_i, v_{i+1}) \in E$ for all $i \in \mathbb{N}$; we denote by Ω the set of all plays. A *player 1 strategy* $\sigma : V^* \cdot V_1 \rightarrow V$ is a function that assigns to every finite prefix $\omega \in V^* \cdot V_1$ of a play that ends in a player 1 vertex v a successor vertex $\sigma(\omega) \in V$ such that $(v, \sigma(\omega)) \in E$; we denote by Σ the set of all player 1 strategies. A strategy is *memoryless* if we have $\sigma(\omega) = \sigma(\omega')$ for any $\omega, \omega' \in V^* \cdot V_1$ that end in the same vertex $v \in V_1$.

Objectives. An *objective* ϕ is a subset of Ω said to be winning for player 1. We say that a play $\omega \in \Omega$ *satisfies the objective* if $\omega \in \phi$. For a vertex set $T \subseteq V$ the *reachability objective* is the set of infinite paths that contain a vertex of T , i.e., $\text{Reach}(T) = \{\langle v_0, v_1, v_2, \dots \rangle \in \Omega \mid \exists j \geq 0 : v_j \in T\}$. Let $\text{Inf}(\omega)$ for $\omega \in \Omega$ denote the set of vertices that occur infinitely often in ω . Given a set TP of k pairs (L_i, U_i) of vertex sets $L_i, U_i \subseteq V$ with $1 \leq i \leq k$, the *Streott objective* is the set of infinite paths for which it holds for *each* $1 \leq i \leq k$ that whenever a vertex of L_i occurs infinitely often, then a vertex of U_i occurs infinitely often, i.e., $\text{Streott}(\text{TP}) = \{\omega \in \Omega \mid L_i \cap \text{Inf}(\omega) = \emptyset \text{ or } U_i \cap \text{Inf}(\omega) \neq \emptyset \text{ for all } 1 \leq i \leq k\}$.

Almost-Sure Winning Sets. For any measurable set of plays $A \subseteq \Omega$ we denote by $\text{Pr}_v^\sigma(A)$ the probability that a play starting at $v \in V$ belongs to A when player 1 plays strategy σ . A strategy σ is *almost-sure (a.s.) winning* from a vertex $v \in V$ for an objective ϕ if $\text{Pr}_v^\sigma(\phi) = 1$. The *almost-sure winning set* $\langle\langle 1 \rangle\rangle_{as}(P, \phi)$ of player 1 is the set of vertices for which player 1 has an almost-sure winning strategy. In graphs the existence of an almost-sure winning strategy corresponds to the existence of a play in the objective, and the set of vertices for which player 1 has an (almost-sure) winning strategy is called the *winning set* $\langle\langle 1 \rangle\rangle(P, \phi)$ of player 1.

Symbolic Encoding of MDPs. Symbolic algorithms operate on sets of vertices, which are usually described by Binary Decision Diagrams (BDDs) [29,1]. In particular Ordered Binary Decision Diagrams [6] (OBDDs) provide a canonical symbolic representation of Boolean functions. For the computation of almost-sure winning sets of MDPs it is sufficient to encode MDPs with OBDDs and one additional bit that denotes whether a vertex is in V_1 or V_R .

Symbolic Steps. One symbolic step corresponds to one primitive operation as supported by standard symbolic packages like CUDD [34]. In this paper we only allow the same basic *set-based symbolic operations* as in [32,23,5,11], namely set operations and the

following one-step symbolic operations for a set of vertices Z : (a) the one-step predecessor operator $\text{Pre}(Z) = \{v \in V \mid \text{Out}(v) \cap Z \neq \emptyset\}$; (b) the one-step successor operator $\text{Post}(Z) = \{v \in V \mid \text{In}(v) \cap Z \neq \emptyset\}$; and (c) the one-step *controllable* predecessor operator $\text{CPre}_R(Z) = \{v \in V_1 \mid \text{Out}(v) \subseteq Z\} \cup \{v \in V_R \mid \text{Out}(v) \cap Z \neq \emptyset\}$; i.e., the CPre_R operator computes all vertices such that the successor belongs to Z with positive probability. This operator can be defined using the Pre operator and basic set operations as follows: $\text{CPre}_R(Z) = \text{Pre}(Z) \setminus (V_1 \cap \text{Pre}(V \setminus Z))$. We additionally allow cardinality computation and picking an arbitrary vertex from a set as in [11].

Symbolic Model. Informally, a symbolic algorithm does not operate on explicit representation of the transition function of a graph, but instead accesses it through Pre and Post operations. For explicit algorithms, a Pre/Post operation on a set of vertices (resp., a single vertex) requires $O(m)$ (resp., the order of indegree/outdegree of the vertex) time. In contrast, for symbolic algorithms Pre/Post operations are considered unit-cost. Thus an interesting algorithmic question is whether better algorithmic bounds can be obtained considering Pre/Post as unit operations. Moreover, the basic set operations are computationally less expensive (as they encode the relationship between the state variables) compared to the Pre/Post symbolic operations (as they encode the transitions and thus the relationship between the present and the next-state variables). In all presented algorithms, the number of set operations is asymptotically at most the number of Pre/Post operations. Hence in the sequel we focus on the number of Pre/Post operations of algorithms.

Algorithmic Problem. Given an MDP P (resp. a graph G) and a set of Streett pairs TP , the problem we consider asks for a symbolic algorithm to compute the almost-sure winning set $\langle\langle 1 \rangle\rangle_{as}(P, \text{Streett}(\text{TP}))$ (resp. the winning set $\langle\langle 1 \rangle\rangle(G, \text{Streett}(\text{TP}))$), which is also called the *qualitative analysis* of MDPs (resp. graphs).

2.2 Basic Concepts related to Algorithmic Solution

Reachability. For a graph $G = (V, E)$ and a set of vertices $S \subseteq V$ the set $\text{GRAPHREACH}(G, S)$ is the set of vertices of V that *can reach* a vertex of S within G , and it can be identified with at most $|\text{GRAPHREACH}(G, S) \setminus S| + 1$ many Pre operations.

Strongly Connected Components. For a set of vertices $S \subseteq V$ we denote by $G[S] = (S, E \cap (S \times S))$ the subgraph of the graph G induced by the vertices of S . An induced subgraph $G[S]$ is strongly connected if there exists a path in $G[S]$ between every pair of vertices of S . A *strongly connected component (SCC)* of G is a set of vertices $C \subseteq V$ such that the induced subgraph $G[C]$ is strongly connected and C is a maximal set in V with this property. We call an SCC *trivial* if it only contains a single vertex and no edges; and *non-trivial* otherwise. The SCCs of G partition its vertices and can be found in $O(n)$ symbolic steps [24]. A bottom SCC C in a directed graph G is an SCC with no edges from vertices of C to vertices of $V \setminus C$, i.e., an SCC without *outgoing* edges. Analogously, a top SCC C is an SCC with no *incoming* edges from $V \setminus C$. For more intuition for bottom and top SCCs, consider the graph in which each SCC is contracted into a single vertex (ignoring edges within an SCC). In the resulting directed acyclic graph the sinks represent the bottom SCCs and the sources represent the top SCCs. Note that every graph has at least one bottom and at least one top SCC. If the graph is not

strongly connected, then there exist at least one top and at least one bottom SCC that are disjoint and thus one of them contains at most half of the vertices of G .

Random Attractors. In an MDP P the *random attractor* $Attr_R(P, W)$ of a set of vertices W is defined as $Attr_R(P, W) = \bigcup_{j \geq 0} Z_j$ where $Z_0 = W$ and $Z_{j+1} = Z_j \cup CPre_R(Z_j)$ for all $j > 0$. The attractor can be computed with at most $|Attr_R(P, W) \setminus W| + 1$ many $CPre_R$ operations.

Maximal End-Components. Let X be a vertex set without outgoing random edges, i.e., with $Out(v) \subseteq X$ for all $v \in X \cap V_R$. A sub-MDP of an MDP P induced by a vertex set $X \subseteq V$ without outgoing random edges is defined as $P[X] = ((X, E \cap (X \times X)), (V_1 \cap X, V_R \cap X), \delta)$. Note that the requirement that X has no outgoing random edges is necessary in order to use the same probabilistic transition function δ . An *end-component* (EC) of an MDP P is a set of vertices $X \subseteq V$ such that (a) X has no outgoing random edges, i.e., $P[X]$ is a valid sub-MDP, (b) the induced sub-MDP $P[X]$ is strongly connected, and (c) $P[X]$ contains at least one edge. Intuitively, an end-component is a set of vertices for which player 1 can ensure that the play stays within the set and almost-surely reaches all the vertices in the set (infinitely often). An end-component is a *maximal end-component* (MEC) if it is maximal under set inclusion. An end-component is *trivial* if it consists of a single vertex (with a self-loop), otherwise it is *non-trivial*. The *MEC decomposition* of an MDP consists of all MECs of the MDP.

Good End-Components. All algorithms for MDPs with Streett objectives are based on finding good end-components, defined below. Given the union of all good end-components, the almost-sure winning set is obtained by computing the almost-sure winning set for the reachability objective with the union of all good end-components as the target set. The correctness of this approach is shown in [7,30] (see also [3, Chap. 10.6.3]). For Streett objectives a good end-component is defined as follows. In the special case of graphs they are called good components.

Definition 1 (Good end-component). *Given an MDP P and a set $TP = \{(L_j, U_j) \mid 1 \leq j \leq k\}$ of target pairs, a good end-component is an end-component X of P such that for each $1 \leq j \leq k$ either $L_j \cap X = \emptyset$ or $U_j \cap X \neq \emptyset$. A maximal good end-component is a good end-component that is maximal with respect to set inclusion.*

Lemma 1 (Correctness of Computing Good End-Components [30, Corollary 2.6.5, Proposition 2.6.9]). *For an MDP P and a set TP of target pairs, let \mathcal{X} be the set of all maximal good end-components. Then $\llbracket 1 \rrbracket_{as}(P, Reach(\bigcup_{X \in \mathcal{X}} X))$ is equal to $\llbracket 1 \rrbracket_{as}(P, Streett(TP))$.*

Iterative Vertex Removal. All the algorithms for Streett objectives maintain vertex sets that are candidates for good end-components. For such a vertex set S we (a) refine the maintained sets according to the SCC decomposition of $P[S]$ and (b) for a set of vertices W for which we know that it cannot be contained in a good end-component, we remove its random attractor from S . The following lemma shows the correctness of these operations.

Lemma 2 (Correctness of Vertex Removal [30, Lemma 2.6.10]). *Given an MDP $P = ((V, E), (V_1, V_R), \delta)$, let X be an end-component with $X \subseteq S$ for some $S \subseteq V$. Then*

- (a) $X \subseteq C$ for one SCC C of $P[S]$ and
- (b) $X \subseteq S \setminus \text{Attr}_R(P', W)$ for each $W \subseteq V \setminus X$ and each sub-MDP P' containing X .

Let X be a good end-component. Then X is an end-component and for each index j , $X \cap U_j = \emptyset$ implies $X \cap L_j = \emptyset$. Hence we obtain the following corollary.

Corollary 1 ([30, Corollary 4.2.2]). *Given an MDP P , let X be a good end-component with $X \subseteq S$ for some $S \subseteq V$. For each i with $S \cap U_i = \emptyset$ it holds that $X \subseteq S \setminus \text{Attr}_R(P[S], L_i \cap S)$.*

For an index j with $S \cap U_j = \emptyset$ we call the vertices of $S \cap L_j$ *bad vertices*. The set of all bad vertices $\text{BAD}(S) = \bigcup_{1 \leq i \leq k} \{v \in L_i \cap S \mid U_i \cap S = \emptyset\}$ can be computed with $2k$ set operations.

3 Symbolic Divide-and-Conquer with Lock-Step Search

In this section we present a symbolic version of the lock-step search for strongly connected subgraphs [25]. This symbolic version is used in all subsequent results, i.e., the sub-quadratic symbolic algorithms for graphs and MDPs with Streett objectives, and for MEC decomposition.

Divide-and-Conquer. The common property of the algorithmic problems we consider in this work is that the goal is to identify subgraphs of the input graph $G = (V, E)$ that are strongly connected and satisfy some additional properties. The difference between the problems lies in the required additional properties. We describe and analyze the Procedure `LOCK-STEP-SEARCH` that we use in all our improved algorithms to efficiently implement a divide-and-conquer approach based on the requirement of strong connectivity, that is, we divide a subgraph $G[S]$, induced by a set of vertices S , into two parts that are not strongly connected within $G[S]$ or detect that $G[S]$ is strongly connected.

Start Vertices of Searches. The input to Procedure `LOCK-STEP-SEARCH` is a set of vertices $S \subseteq V$ and two subsets of S denoted by H_S and T_S . In the algorithms that call the procedure as a subroutine, vertices contained in H_S have lost incoming edges (i.e., they were a “head” of a lost edge) and vertices contained in T_S have lost outgoing edges (i.e., they were a “tail” of a lost edge) since the last time a superset of S was identified as being strongly connected. For each vertex h of H_S the procedure conducts a backward search (i.e., a sequence of Pre operations) within $G[S]$ to find the vertices of S that can reach h ; and analogously a forward search (i.e., a sequence of Post operations) from each vertex t of T_S is conducted.

Intuition for the Choice of Start Vertices. If the subgraph $G[S]$ is not strongly connected, then it contains at least one top SCC and at least one bottom SCC that are disjoint. Further, if for a superset $S' \supset S$ the subgraph $G[S']$ was strongly connected, then each top SCC of $G[S]$ contains a vertex that had an additional incoming edge in $G[S']$ compared to $G[S]$, and analogously each bottom SCC of $G[S]$ contains a vertex that had an additional outgoing edge. Thus by keeping track of the vertices that lost incoming or outgoing edges, the following invariant will be maintained by all our improved algorithms.

Invariant 1 (Start Vertices Sufficient). *We have $H_S, T_S \subseteq S$. Either (a) $H_S \cup T_S = \emptyset$ and $G[S]$ is strongly connected or (b) at least one vertex of each top SCC of $G[S]$ is contained in H_S and at least one vertex of each bottom SCC of $G[S]$ is contained in T_S .*

```

Procedure LOCK-STEP-SEARCH( $G, S, H_S, T_S$ )
  /* Pre and Post defined w.r.t. to  $G$  */
  1 foreach  $v \in H_S \cup T_S$  do  $C_v \leftarrow \{v\}$ 
  2 while true do
  3    $H'_S \leftarrow H_S, T'_S \leftarrow T_S$ 
  4   foreach  $h \in H_S$  do /* search for top SCC */
  5      $C'_h \leftarrow (C_h \cup \text{Pre}(C_h)) \cap S$ 
  6     if  $|C'_h \cap H'_S| > 1$  then  $H'_S \leftarrow H'_S \setminus \{h\}$ 
  7     else
  8       if  $C'_h = C_h$  then return  $(C_h, H'_S, T'_S)$ 
  9        $C_h \leftarrow C'_h$ 
 10  foreach  $t \in T_S$  do /* search for bottom SCC */
 11     $C'_t \leftarrow (C_t \cup \text{Post}(C_t)) \cap S$ 
 12    if  $|C'_t \cap T'_S| > 1$  then  $T'_S \leftarrow T'_S \setminus \{t\}$ 
 13    else
 14      if  $C'_t = C_t$  then return  $(C_t, H'_S, T'_S)$ 
 15       $C_t \leftarrow C'_t$ 
 16   $H_S \leftarrow H'_S, T_S \leftarrow T'_S$ 

```

Lock-Step Search. The searches from the vertices of $H_S \cup T_S$ are performed in *lock-step*, that is, (a) one step is performed in each of the searches before the next step of any search is done and (b) all searches stop as soon as the first of the searches finishes. This is implemented in Procedure LOCK-STEP-SEARCH as follows. A step in the search from a vertex $t \in T_S$ (and analogously for $h \in H_S$) corresponds to the execution of the iteration of the for-each loop for $t \in T_S$. In an iteration of a for-each loop we might discover that we do not need to consider this search further (see the paragraph on ensuring strong connectivity below) and update the set T_S (via T'_S) for future iterations accordingly. Otherwise the set C_t is either strictly increasing in this step of the search or the search for t terminates and we return the set of vertices in $G[S]$ that are reachable from t . So the two for-each loops over the vertices of T_S and H_S that are executed in an iteration of the while-loop perform one step of each of the searches and the while-loop stops as soon as a search stops, i.e., a return statement is executed and hence this implements properties (a) and (b) of lock-step search. Note that the while-loop terminates, i.e., a return statement is executed eventually because for all $t \in T_S$ (and resp. for all $h \in H_S$) the sets C_t are monotonically increasing over the iterations of the while-loop, we have $C_t \subseteq S$, and if some set C_t does not increase in an iteration, then it is either removed from T_S and thus not considered further or a return statement is executed. Note that when a search from a vertex $t \in T_S$ stops, it has discovered a maximal set of vertices C that can be reached from t ; and analogously for $h \in H_S$. Figure 1 shows a small intuitive example of a call to the procedure.

Comparison to Explicit Algorithm. In the *explicit* version of the algorithm [25,7] the search from vertex $t \in T_S$ performs a depth-first search that terminates exactly when every *edge* reachable from t is explored. Since any search that starts outside of a bottom SCC but reaches the bottom SCC has to explore more edges than the search started inside of the bottom SCC, the first search from a vertex of T_S that terminates has exactly

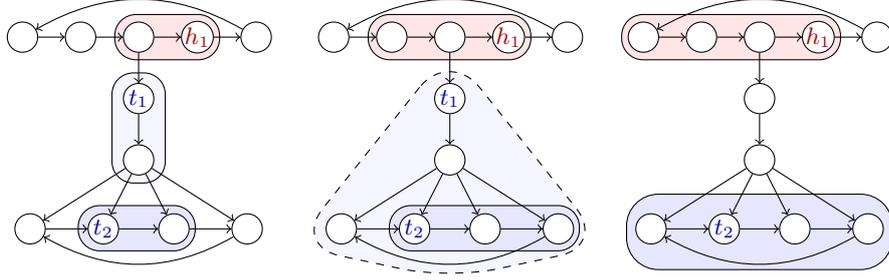


Fig. 1: An example of symbolic lock-step search showing the first three iterations of the main while-loop. Note that during the second iteration, the search started from t_1 is disregarded since it collides with t_2 . In the subsequent fourth iteration, the search started from t_2 is returned by the procedure.

explored (one of) the smallest (in the number of edges) bottom SCC(s) of $G[S]$. Thus on explicit graphs the explicit lock-step search from the vertices of $H_S \cup T_S$ finds (one of) the smallest (in the number of edges) top or bottom SCC(s) of $G[S]$ in time proportional to the number of searches times the number of edges in the identified SCC. In *symbolically* represented graphs it can happen (1) that a search started outside of a bottom (resp. top) SCC terminates earlier than the search started within the bottom (resp. top) SCC and (2) that a search started in a larger (in the number of vertices) top or bottom SCC terminates before one in a smaller top or bottom SCC. We discuss next how we address these two challenges.

Ensuring Strong Connectivity. First, we would like the set returned by Procedure LOCK-STEP-SEARCH to indeed be a top or bottom SCC of $G[S]$. For this we use the following observation for bottom SCCs that can be applied to top SCCs analogously. If a search starting from a vertex of $t_1 \in T_S$ encounters another vertex $t_2 \in T_S$, $t_1 \neq t_2$, there are two possibilities: either (1) both vertices are in the same SCC or (2) t_1 can reach t_2 but not vice versa. In Case (1) the searches from both vertices can explore all vertices in the SCC and thus it is sufficient to only search from one of them. In Case (2) the SCC of t_1 has an outgoing edge and thus cannot be a bottom SCC. Hence in both cases we can remove the vertex t_1 from the set T_S while still maintaining Invariant 1. By Invariant 1 we further have that each search from a vertex of T_S that is not in a bottom SCC encounters another vertex of T_S in its search and therefore is removed from the set T_S during Procedure LOCK-STEP-SEARCH (if no top or bottom SCC is found earlier). This ensures that the returned set is either a top or a bottom SCC.¹

Bound on Symbolic Steps. Second, observe that we can still bound the number of symbolic steps needed for the search that terminates first by the number of *vertices* in the smallest top or bottom SCC of $G[S]$, since this is an upper bound on the symbolic steps needed for the search started in this SCC. Thus provided Invariant 1, we can bound the number of symbolic steps in Procedure LOCK-STEP-SEARCH to identify a vertex set $C \subsetneq S$ such that C and $S \setminus C$ are not strongly connected in $G[S]$ by $O((|H_S| + |T_S|) \cdot \min(|C|, |S \setminus C|))$. In the algorithms that call Procedure LOCK-STEP-SEARCH we charge the number of symbolic steps in the procedure to the vertices in the

¹ To improve the practical performance, we return the updated sets H_S and T_S . By the above argument this preserves Invariant 1.

smaller set of C and $S \setminus C$; this ensures that each vertex is charged at most $O(\log n)$ times over the whole algorithm. We obtain the following result (proof in Appendix A).

Theorem 1 (Lock-Step Search). *Provided Invariant 1 holds, Procedure LOCK-STEP-SEARCH(G, S, H_S, T_S) returns a top or bottom SCC C of $G[S]$. It uses $O((|H_S| + |T_S|) \cdot \min(|C|, |S \setminus C|))$ symbolic steps if $C \neq S$ and $O((|H_S| + |T_S|) \cdot |C|)$ otherwise.*

4 Graphs with Streett Objectives

Basic Symbolic Algorithm. Recall that for a given graph (with n vertices) and a Streett objective (with k target pairs) each non-trivial strongly connected subgraph without bad vertices is a good component. The basic symbolic algorithm for graphs with Streett objectives repeatedly removes bad vertices from each SCC and then recomputes the SCCs until all good components are found. The winning set then consists of the vertices that can reach a good component. We refer to this algorithm as STRETTGRAPHBASIC. For the pseudocode and more details see Appendix B.

Proposition 1. *Algorithm STRETTGRAPHBASIC correctly computes the winning set in graphs with Streett objectives and requires $O(n \cdot \min(n, k))$ symbolic steps.*

Improved Symbolic Algorithm. In our improved symbolic algorithm we replace the recomputation of all SCCs with the search for a new top or bottom SCC with Procedure LOCK-STEP-SEARCH from vertices that have lost adjacent edges whenever there are not too many such vertices. We present the improved symbolic algorithm for graphs with Streett objectives in more detail as it also conveys important intuition for the MDP case. The pseudocode is given in Algorithm STRETTGRAPHIMPR.

Iterative Refinement of Candidate Sets. The improved algorithm maintains a set `goodC` of already identified good components that is initially empty and a set \mathcal{X} of candidates for good components that is initialized with the SCCs of the input graph G . The difference to the basic algorithm lies in the properties of the vertex sets maintained in \mathcal{X} and the way we identify sets that can be separated from each other without destroying a good component. In each iteration one vertex set S is removed from \mathcal{X} and, after the removal of bad vertices from the set, either identified as a good component or split into several candidate sets. By Lemma 2 and Corollary 1 the following invariant is maintained throughout the algorithm for the sets in `goodC` and \mathcal{X} .

Invariant 2 (Maintained Sets). *The sets in $\mathcal{X} \cup \text{goodC}$ are pairwise disjoint and for every good component C of G there exists a set $Y \supseteq C$ such that either $Y \in \mathcal{X}$ or $Y \in \text{goodC}$.*

Lost Adjacent Edges. In contrast to the basic algorithm, the subgraph induced by a set S contained in \mathcal{X} is not necessarily strongly connected. Instead, we remember vertices of S that have lost adjacent edges since the last time a superset of S was determined to induce a strongly connected subgraph; vertices that lost incoming edges are contained in H_S and vertices that lost outgoing edges are contained in T_S . In this way we maintain Invariant 1 throughout the algorithm, which enables us to use Procedure LOCK-STEP-SEARCH with the running time guarantee provided by Theorem 1.

Algorithm STREETGRAPHIMPR: Improved Alg. for Graphs with Streett Obj.

Input : graph $G = (V, E)$ and Streett pairs $TP = \{(L_i, U_i) \mid 1 \leq i \leq k\}$
Output: $\langle\langle 1 \rangle\rangle(G, \text{Streett}(TP))$

```

1  $\mathcal{X} \leftarrow \text{ALLSCCs}(G)$ ;  $\text{goodC} \leftarrow \emptyset$ 
2 foreach  $C \in \mathcal{X}$  do  $H_C \leftarrow \emptyset$ ;  $T_C \leftarrow \emptyset$ 
3 while  $\mathcal{X} \neq \emptyset$  do
4   remove some  $S \in \mathcal{X}$  from  $\mathcal{X}$ 
5    $B \leftarrow \bigcup_{1 \leq i \leq k: U_i \cap S = \emptyset} (L_i \cap S)$ 
6   while  $B \neq \emptyset$  do
7      $S \leftarrow S \setminus B$ 
8      $H_S \leftarrow (H_S \cup \text{Post}(B)) \cap S$ 
9      $T_S \leftarrow (T_S \cup \text{Pre}(B)) \cap S$ 
10     $B \leftarrow \bigcup_{1 \leq i \leq k: U_i \cap S = \emptyset} (L_i \cap S)$ 
11  if  $\text{Post}(S) \cap S \neq \emptyset$  then /*  $G[S]$  contains at least one edge */
12    if  $|H_S| + |T_S| = 0$  then  $\text{goodC} \leftarrow \text{goodC} \cup \{S\}$ 
13    else if  $|H_S| + |T_S| \geq \sqrt{m/\log n}$  then
14      delete  $H_S$  and  $T_S$ 
15       $\mathcal{C} \leftarrow \text{ALLSCCs}(G[S])$ 
16      if  $|\mathcal{C}| = 1$  then  $\text{goodC} \leftarrow \text{goodC} \cup \{S\}$ 
17      else
18        foreach  $C \in \mathcal{C}$  do  $H_C \leftarrow \emptyset$ ;  $T_C \leftarrow \emptyset$ 
19         $\mathcal{X} \leftarrow \mathcal{X} \cup \mathcal{C}$ 
20    else
21       $(C, H_S, T_S) \leftarrow \text{LOCK-STEP-SEARCH}(G, S, H_S, T_S)$ 
22      if  $C = S$  then  $\text{goodC} \leftarrow \text{goodC} \cup \{S\}$ 
23      else /* separate  $C$  and  $S \setminus C$  */
24         $S \leftarrow S \setminus C$ 
25         $H_C \leftarrow \emptyset$ ;  $T_C \leftarrow \emptyset$ 
26         $H_S \leftarrow (H_S \cup \text{Post}(C)) \cap S$ 
27         $T_S \leftarrow (T_S \cup \text{Pre}(C)) \cap S$ 
28         $\mathcal{X} \leftarrow \mathcal{X} \cup \{S\} \cup \{C\}$ 
29 return  $\text{GRAPHREACH}(G, \bigcup_{C \in \text{goodC}} C)$ 

```

Identifying SCCs. Let S be the vertex set removed from \mathcal{X} in a fixed iteration of Algorithm STREETGRAPHIMPR after the removal of bad vertices in the inner while-loop. First note that if S is strongly connected and contains at least one edge, then it is a good component. If the set S was already identified as strongly connected in a previous iteration, i.e., H_S and T_S are empty, then S is identified as a good component in line 12. If many vertices of S have lost adjacent edges since the last time a super-set of S was identified as a strongly connected subgraph, then the SCCs of $G[S]$ are determined as in the basic algorithm. To achieve the optimal asymptotic upper bound, we say that many vertices of S have lost adjacent edges when we have $|H_S| + |T_S| \geq \sqrt{m/\log n}$, while lower thresholds are used in our experimental results. Otherwise, if not too many vertices of S lost adjacent edges, then we start a symbolic *lock-step search* for top SCCs from the vertices of H_S and for bottom SCCs from the vertices of T_S using Procedure LOCK-

STEP-SEARCH. The set returned by the procedure is either a top or a bottom SCC C of $G[S]$ (Theorem 1). Therefore we can from now on consider C and $S \setminus C$ separately, maintaining Invariants 1 and 2.

Algorithm STRETTGRAPHIMPR. A succinct description of the pseudocode is as follows: Lines 1–2 initialize the set of candidates for good components with the SCCs of the input graph. In each iteration of the main while-loop one candidate is considered and the following operations are performed: (a) lines 5–10 iteratively remove all bad vertices; if afterwards the candidate is still strongly connected (and contains at least one edge), it is identified as a good component in the next step; otherwise it is partitioned into new candidates in one of the following ways: (b) if many vertices lost adjacent edges, lines 13–19 partition the candidate into its SCCs (this corresponds to an iteration of the basic algorithm); (c) otherwise, lines 20–28 use symbolic lock-step search to partition the candidate into one of its SCCs and the remaining vertices. The while-loop terminates when no candidates are left. Finally, vertices that can reach some good component are returned. We have the following result (proof in Appendix B).

Theorem 2 (Improved Algorithm for Graphs). *Algorithm* STRETTGRAPHIMPR correctly computes the winning set in graphs with StreTT objectives and requires $O(n \cdot \sqrt{m \log n})$ symbolic steps.

5 Symbolic MEC Decomposition

In this section we present a succinct description of the basic symbolic algorithm for MEC decomposition and then present the main ideas for the improved algorithm.

Basic symbolic algorithm for MEC decomposition. The basic symbolic algorithm for MEC decomposition maintains a set of identified MECs and a set of candidates for MECs, initialized with the SCCs of the MDP. Whenever a candidate is considered, either (a) it is identified as a MEC or (b) it contains vertices with outgoing random edges, which are then removed together with their random attractor from the candidate, and the SCCs of the remaining sub-MDP are added to the set of candidates. We refer to the algorithm as MECBASIC.

Proposition 2. *Algorithm* MECBASIC correctly computes the MEC decomposition of MDPs and requires $O(n^2)$ symbolic steps.

Improved symbolic algorithm for MEC decomposition. The improved symbolic algorithm for MEC decomposition uses the ideas of symbolic lock-step search presented in Section 3. Informally, when considering a candidate that lost a few edges from the remaining graph, we use the symbolic lock-step search to identify some bottom SCC. We refer to the algorithm as MECIMPR. Since all the important conceptual ideas regarding the symbolic lock-step search are described in Section 3, we relegate the technical details to Appendix C. We summarize the main result (proof in Appendix C).

Theorem 3 (Improved Algorithm for MEC). *Algorithm* MECIMPR correctly computes the MEC decomposition of MDPs and requires $O(n \cdot \sqrt{m})$ symbolic steps.

6 MDPs with Streett Objectives

Basic Symbolic Algorithm. We refer to the basic symbolic algorithm for MDPs with Streett objectives as `STREETTMDPBASIC`, which is similar to the algorithm for graphs, with SCC computation replaced by MEC computation. The pseudocode of Algorithm `STREETTMDPBASIC` together with its detailed description is presented in Appendix D.

Proposition 3. *Algorithm `STREETTMDPBASIC` correctly computes the almost-sure winning set in MDPs with Streett objectives and requires $O(n^2 \cdot \min(n, k))$ symbolic steps.*

Remark. The above bound uses the basic symbolic MEC decomposition algorithm. Using our improved symbolic MEC decomposition algorithm, the above bound could be improved to $O(n \cdot \sqrt{m} \cdot \min(n, k))$.

Improved Symbolic Algorithm. We refer to the improved symbolic algorithm for MDPs with Streett objectives as `STREETTMDPIMPR`. First we present the main ideas for the improved symbolic algorithm. Then we explain the key differences compared to the improved symbolic algorithm for graphs. A thorough description with the technical details and proofs is presented in Appendix D.

- First, we improve the algorithm by interleaving the symbolic MEC computation with the detection of bad vertices [7,30]. This allows to replace the computation of MECs in each iteration of the while-loop with the computation of SCCs and an additional random attractor computation.
 - *Intuition of interleaved computation.* Consider a candidate for a good end-component S after a random attractor to some bad vertices is removed from it. After the removal of the random attractor, the set S does not have random vertices with outgoing edges. Consider that further $\text{BAD}(S) = \emptyset$ holds. If S is strongly connected and contains an edge, then it is a good end-component. If S is not strongly connected, then $P[S]$ contains at least two SCCs and some of them might have random vertices with outgoing edges. Since end-components are strongly connected and do not have random vertices with outgoing edges, we have that (1) every good end-component is completely contained in one of the SCCs of $P[S]$ and (2) the random vertices of an SCC with outgoing edges and their random attractor do not intersect with any good end-component (see Lemma 2).
 - *Modification from basic to improved algorithm.* We use these observations to modify the basic algorithm as follows: First, for the sets that are candidates for good end-components, we do not maintain the property that they are end-components, but only that they do not have random vertices with outgoing edges (it still holds that every maximal good end-component is either already identified or contained in one of the candidate sets). Second, for a candidate set S , we repeat the removal of bad vertices until $\text{BAD}(S) = \emptyset$ holds before we continue with the next step of the algorithm. This allows us to make progress after the removal of bad vertices by computing all SCCs (instead of MECs)

- of the remaining sub-MDP. If there is only one SCC, then this is a good end-component (if it contains at least one edge). Otherwise (a) we remove from each SCC the set of random vertices with outgoing edges and their random attractor and (b) add the remaining vertices of each SCC as a new candidate set.
- Second, as for the improved symbolic algorithm for graphs, we use the symbolic lock-step search to quickly identify a top or bottom SCC every time a candidate has lost a small number of edges since the last time its superset was identified as being strongly connected. The symbolic lock-step search is described in detail in Section 3.

Using interleaved MEC computation and lock-step search leads to a similar algorithmic structure for Algorithm STRETTMDPIMPR as for our improved symbolic algorithm for graphs (Algorithm STRETTGRAPHIMPR). The key differences are as follows: First, the set of candidates for good end-components is initialized with the MECs of the input graph instead of the SCCs. Second, whenever bad vertices are removed from a candidate, also their random attractor is removed. Further, whenever a candidate is partitioned into its SCCs, for each SCC, the random attractor of the vertices with outgoing random edges is removed. Finally, whenever a candidate S is separated into C and $S \setminus C$ via symbolic lock-step search, the random attractor of the vertices with outgoing random edges is removed from C , and the random attractor of C is removed from S .

Theorem 4 (Improved Algorithm for MDPs). *Algorithm STRETTMDPIMPR correctly computes the almost-sure winning set in MDPs with Streett objectives and requires $O(n \cdot \sqrt{m} \log n)$ symbolic steps.*

7 Experiments

We present a basic prototype implementation of our algorithm and compare against the basic symbolic algorithm for graphs and MDPs with Streett objectives.

Models. We consider the academic benchmarks from the VLTS benchmark suite [20], which gives representative examples of systems with nondeterminism, and has been used in previous experimental evaluation (such as [4,11]).

Specifications. We consider random LTL formulae and use the tool Rabinizer [27] to obtain deterministic Rabin automata. Then the negations of the formulae give us Streett automata, which we consider as the specifications.

Graphs. For the models of the academic benchmarks, we first compute SCCs, as all algorithms for Streett objectives compute SCCs as a preprocessing step. For SCCs of the model benchmarks we consider products with the specification Streett automata, to obtain graphs with Streett objectives, which are the benchmark examples for our experimental evaluation. The number of transitions in the benchmarks ranges from 300K to 5Million.

MDPs. For MDPs, we consider the graphs obtained as above and consider a fraction of the vertices of the graph as random vertices, which is chosen uniformly at random. We consider 10%, 20%, and 50% of the vertices as random vertices for different experimental evaluation.

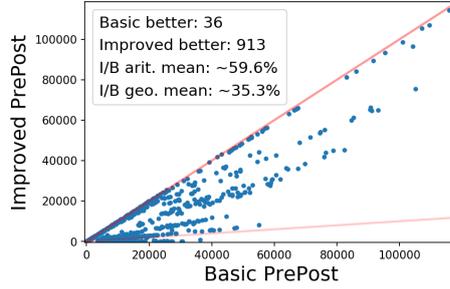
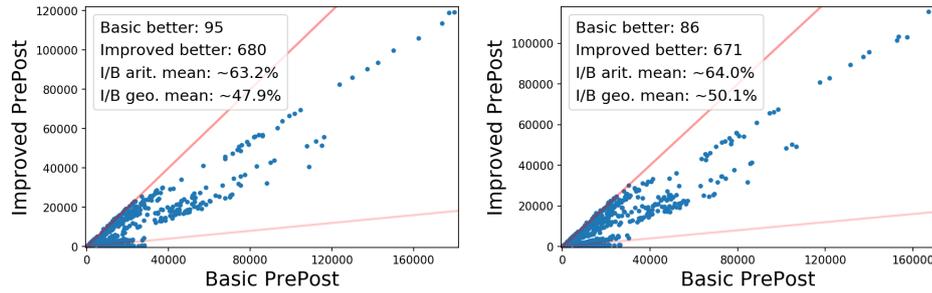
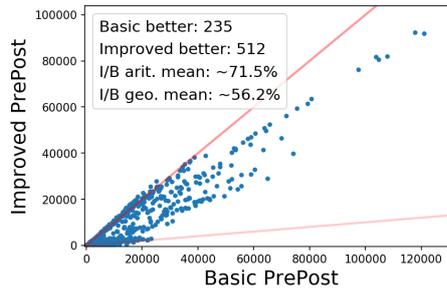


Fig. 2: Results for graphs with Streett objectives.



(a) 10% random vertices

(b) 20% random vertices



(c) 50% random vertices

Fig. 3: Results for MDPs with Streett objectives.

Experimental evaluation. In the experimental evaluation we compare the number of symbolic steps (i.e., the number of Pre/Post operations²) executed by the algorithms, the comparison of running time yields similar results and is provided in Appendix E. As the initial preprocessing step is the same for all the algorithms (computing all SCCs for graphs and all MECs for MDPs), the comparison presents the number of symbolic steps executed after the preprocessing. The experimental results for graphs are shown in Figure 2 and the experimental results for MDPs are shown in Figure 3 (in each figure the two lines represent equality and an order-of-magnitude improvement, respectively).

² Recall that the basic set operations are cheaper to compute, and asymptotically at most the number of Pre/Post operations in all the presented algorithms.

Discussion. Note that the lock-step search is the key reason for theoretical improvement, however, the improvement relies on a large number of Streett pairs. In the experimental evaluation, the LTL formulae generate Streett automata with small number of pairs, which after the product with the model accounts for an even smaller fraction of pairs as compared to the size of the state space. This has two effects:

- In the experiments the lock-step search is performed for a much smaller parameter value ($O(\log n)$) instead of the theoretically optimal bound of $\sqrt{m/\log n}$, and leads to a small improvement.
- For large graphs, since the number of pairs is small as compared to the number of states, the improvement over the basic algorithm is minimal.

In contrast to graphs, in MDPs even with small number of pairs as compared to the state-space, the interleaved MEC computation has a notable effect on practical performance, and we observe performance improvement even in large MDPs.

8 Conclusion

In this work we consider symbolic algorithms for graphs and MDPs with Streett objectives, as well as for MEC decomposition. Our algorithmic bounds match for both graphs and MDPs. In contrast, while SCCs can be computed in linearly many symbolic steps no such algorithm is known for MEC decomposition. An interesting direction of future work would be to explore further improved symbolic algorithms for MEC decomposition. Moreover, further improved symbolic algorithms for graphs and MDPs with Streett objectives is also an interesting direction of future work.

Acknowledgements. K. C. and M. H. are partially supported by the Vienna Science and Technology Fund (WWTF) grant ICT15-003. K. C. is partially supported by the Austrian Science Fund (FWF): S11407-N23 (RiSE/SHiNE), and an ERC Start Grant (279307: Graph Games). V. T. is partially supported by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 665385. V. L. is partially supported by the Austrian Science Fund (FWF): S11408-N23 (RiSE/SHiNE), the ISF grant #1278/16, and an ERC Consolidator Grant (project MPM). For M. H. and V. L. the research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement no. 340506.

References

1. Akers, S.B.: Binary decision diagrams. *IEEE Trans. Comput.* C-27(6), 509–516 (1978)
2. Alur, R., Henzinger, T.A.: *Computer-aided verification (2004)*, <http://www.cis.upenn.edu/group/cis673/>
3. Baier, C., Katoen, J.P.: *Principles of model checking*. MIT Press (2008)
4. Barnat, J., Chaloupka, J., van de Pol, J.: Distributed algorithms for SCC decomposition. *J. Log. Comput.* 21(1), 23–44 (2011)
5. Bloem, R., Gabow, H.N., Somenzi, F.: An algorithm for strongly connected component analysis in $n \log n$ symbolic steps. *Form. Methods Syst. Des.* 28(1), 37–56 (2006)

6. Bryant, R.E.: Symbolic manipulation of Boolean functions using a graphical representation. In: Conference on Design automation (DAC). pp. 688–694 (1985)
7. Chatterjee, K., Dvořák, W., Henzinger, M., Loitzenbauer, V.: Model and objective separation with conditional lower bounds: Disjunction is harder than conjunction. In: LICS. pp. 197–206 (2016)
8. Chatterjee, K., Henzinger, M.: Faster and Dynamic Algorithms For Maximal End-Component Decomposition And Related Graph Problems In Probabilistic Verification. In: SODA. pp. 1318–1336 (2011)
9. Chatterjee, K., Henzinger, M.: An $O(n^2)$ Time Algorithm for Alternating Büchi Games. In: SODA. pp. 1386–1399 (2012)
10. Chatterjee, K., Henzinger, M.: Efficient and Dynamic Algorithms for Alternating Büchi Games and Maximal End-Component Decomposition. *Journal of the ACM* 61(3), 15 (2014)
11. Chatterjee, K., Henzinger, M., Joglekar, M., Shah, N.: Symbolic algorithms for qualitative analysis of Markov decision processes with Büchi objectives. *Form. Methods Syst. Des.* 42(3), 301–327 (2013)
12. Chatterjee, K., Henzinger, M., Loitzenbauer, V.: Improved Algorithms for One-Pair and k -Pair Streett Objectives. In: LICS. pp. 269–280 (2015)
13. Chatterjee, K., Jurdziński, M., Henzinger, T.A.: Simple stochastic parity games. In: CSL. pp. 100–113 (2003)
14. Chatterjee, K., Dvořák, W., Henzinger, M., Loitzenbauer, V.: Lower bounds for symbolic computation on graphs: Strongly connected components, liveness, safety, and diameter. In: SODA. pp. 2341–2356 (2018)
15. Chatterjee, K., Gaiser, A., Kretínský, J.: Automata with generalized rabin pairs for probabilistic model checking and LTL synthesis. In: CAV. pp. 559–575 (2013)
16. Ciesinski, F., Baier, C.: LiQuor: A tool for qualitative and quantitative linear time analysis of reactive systems. In: QEST. pp. 131–132 (2006)
17. Cimatti, A., Clarke, E., Giunchiglia, F., Roveri, M.: NUSMV: a new symbolic model checker. *International Journal on Software Tools for Technology Transfer (STTT)* 2(4), 410–425 (2000)
18. Clarke, Jr., E.M., Grumberg, O., Peled, D.A.: *Model Checking*. MIT Press, Cambridge, MA, USA (1999)
19. Clarke, E., Grumberg, O., Peled, D.: *Symbolic model checking*. In: *Model Checking*. MIT Press (1999)
20. CWI/SEN2 and INRIA/VASY: The VLTS Benchmark Suite, <http://cadp.inria.fr/resources/vlts>
21. Dehnert, C., Junges, S., Katoen, J., Volk, M.: A Storm is coming: A modern probabilistic model checker. In: CAV. pp. 592–600 (2017)
22. Esparza, J., Kretínský, J.: From LTL to deterministic automata: A safriless compositional approach. In: CAV. pp. 192–208 (2014)
23. Gentilini, R., Piazza, C., Policriti, A.: Computing strongly connected components in a linear number of symbolic steps. In: SODA. pp. 573–582 (2003)
24. Gentilini, R., Piazza, C., Policriti, A.: Symbolic graphs: Linear solutions to connectivity related problems. *Algorithmica* 50(1), 120–158 (2008)
25. Henzinger, M., Telle, J.A.: Faster Algorithms for the Nonemptiness of Streett Automata and for Communication Protocol Pruning. In: SWAT. pp. 16–27 (1996)
26. Holzmann, G.J.: The model checker SPIN. *IEEE Trans. Softw. Eng.* 23(5), 279–295 (1997)
27. Komárková, Z., Kretínský, J.: Rabinizer 3: Safriless translation of LTL to small deterministic automata. In: ATVA. pp. 235–241 (2014)
28. Kwiatkowska, M.Z., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: CAV. pp. 585–591 (2011)

29. Lee, C.Y.: Representation of switching circuits by binary-decision programs. *Bell System Techn. J.* 38(4), 985–999 (1959)
30. Loitzenbauer, V.: Improved Algorithms and Conditional Lower Bounds for Problems in Formal Verification and Reactive Synthesis. Ph.D. thesis, University of Vienna (2016)
31. Manna, Z., Pnueli, A.: *Temporal Verification of Reactive Systems: Progress (Draft)* (1996)
32. Ravi, K., Bloem, R., Somenzi, F.: A comparative study of symbolic algorithms for the computation of fair cycles. In: *FMCAD*. pp. 143–160 (2000)
33. Safra, S.: On the complexity of ω -automata. In: *FOCS*. pp. 319–327 (1988)
34. Somenzi, F.: CUDD: CU decision diagram package release 3.0.0 (2015), <http://vlsi.colorado.edu/~fabio/CUDD/>

Appendix

A Details of Section 3: Symbolic Lock-Step Search

Proof (of Theorem 1).

Strong connectivity. We want to show that $C \leftarrow \text{LOCK-STEP-SEARCH}(G, S, H_S, T_S)$ is a top or bottom SCC of $G[S]$ given Invariant 1 is satisfied. By the invariant at least one vertex of each top SCC of $G[S]$ is contained in H_S and at least one vertex of each bottom SCC of $G[S]$ is contained in T_S . Suppose C is the set obtained from a search conducted by Post operations that started from within a bottom SCC \tilde{C} of $G[S]$. Since \tilde{C} is a bottom SCC and we update the search by executing Post operations (and moreover intersect with S at every update), we have $C \subseteq \tilde{C}$. Further, since \tilde{C} is an SCC, the updates with Post eventually cover all vertices of \tilde{C} , which gives us $C = \tilde{C}$. A set C_t constructed with Post operations whose start vertex t is not contained in a bottom SCC of $G[S]$ can not yield the set C since eventually it contains a bottom SCC of $G[S]$, and by Invariant 1 this SCC contains a candidate in T_S ; therefore $|C_t \cap T_S| > 1$ is satisfied at some point in the construction of C_t and then search is canceled by removing t from T_S ; note that a search starting from a bottom SCC can be canceled only if another vertex of the bottom SCC remains in T_S . By the symmetric argument for searches conducted by Pre operations that started from a vertex of a top SCC we have that the returned set C is either a top or a bottom SCC of $G[S]$.

Bound on symbolic steps. Consider (one of) the smallest top or bottom SCCs \tilde{C} of $G[S]$. Suppose w.l.o.g. that \tilde{C} is a bottom SCC. By Invariant 1 there is a search, conducted by Post operations, that starts from a vertex $t \in T_S$ within \tilde{C} and that is not canceled, and therefore this search terminates after at most $|\tilde{C}|$ many Post operations. Other searches may terminate earlier but this gives an upper bound of $O((|H_S| + |T_S|) \cdot |\tilde{C}|)$ on the number of symbolic steps until the lock-step search terminates. Finally, consider the returned set $C \leftarrow \text{LOCK-STEP-SEARCH}(G, S, H_S, T_S)$. There are two possible cases: either (i) $S = C$, which implies $C = \tilde{C}$ so the number of symbolic steps can be bounded by $O((|H_S| + |T_S|) \cdot |C|)$, or (ii) $S \neq C$. In the second case, since \tilde{C} is (some) smallest SCC, C is an SCC, and $S \setminus C$ contains at least one SCC, we have $|\tilde{C}| \leq |C|$ and $|\tilde{C}| \leq |S \setminus C|$, and hence we can bound the number of symbolic steps in this case by $O((|H_S| + |T_S|) \cdot \min(|C|, |S \setminus C|))$. \square

B Details of Section 4: Graphs with Streott Objectives

B.1 Basic Symbolic Algorithm for Graphs with Streott Objectives

The pseudocode of the basic symbolic algorithm for graphs with Streott objectives is given in Algorithm STREOTTGRAPHBASIC.

The basic symbolic algorithm for Streott objectives on graphs finds good components as follows. The algorithm maintains two sets of vertex sets: `goodC` contains identified good components and is initially empty; \mathcal{X} contains candidates for good components and is initialized with the SCCs of the input graph G . The sets in \mathcal{X} are strongly connected subgraphs of G throughout the algorithm. In each iteration of the while-loop one of

Algorithm STRETTGRAPHBASIC: Basic Algorithm for Graphs with Streett Obj.

Input : graph $G = (V, E)$ and Streett pairs $TP = \{(L_i, U_i) \mid 1 \leq i \leq k\}$
Output: $\langle\langle 1 \rangle\rangle (G, \text{Streett}(TP))$

```
1  $\mathcal{X} \leftarrow \text{ALLSCCs}(G)$ ;  $\text{goodC} \leftarrow \emptyset$ 
2 while  $\mathcal{X} \neq \emptyset$  do
3   remove some  $S \in \mathcal{X}$  from  $\mathcal{X}$ 
4    $B \leftarrow \bigcup_{1 \leq i \leq k: U_i \cap S = \emptyset} (L_i \cap S)$ 
5   if  $B \neq \emptyset$  then
6      $S \leftarrow S \setminus B$ 
7      $\mathcal{X} \leftarrow \mathcal{X} \cup \text{ALLSCCs}(G[S])$ 
8   else
9     if  $\text{Post}(S) \cap S \neq \emptyset$  then /*  $G[S]$  contains at least one edge */
10     $\text{goodC} \leftarrow \text{goodC} \cup \{S\}$ 
11 return  $\text{GRAPHREACH}(G, \bigcup_{C \in \text{goodC}} C)$ 
```

the candidate sets S maintained in \mathcal{X} is considered. If the set S does not contain bad vertices and contains at least one edge, then it is a good component and added to goodC . Otherwise, the set of bad vertices B in S is removed from S ; the subgraph induced by $S' = S \setminus B$ might not be strongly connected but every good component contained in S' must still be strongly connected, therefore the maximal strongly connected subgraphs of $G[S']$ are added to \mathcal{X} as new candidates for good components. By Lemma 2 and Corollary 1 this procedure maintains the property that every good component of G is completely contained in one of the vertex sets of goodC or \mathcal{X} . Further in each iteration either (a) vertices are removed or separated into different vertex sets or (b) a new good component is identified. Thus after at most $O(n)$ iterations the set \mathcal{X} is empty and all good components of G are contained in goodC . Furthermore, whenever bad vertices are removed from a given candidate set, the number of target pairs this candidate set intersects is reduced by one. Thus each vertex is considered in at most $O(k)$ iterations of the main while-loop. Finally, the set of vertices that can reach a good component is determined (by $O(n)$ Pre operations) and output as the winning set. Since computing SCCs can be done in $O(n)$ symbolic steps, the total number of symbolic steps of the basic algorithm is bounded by $O(n \cdot \min(n, k))$.

B.2 Improved Symbolic Algorithm for Graphs with Streett Objectives

Lemma 3 (Invariants of Improved Algorithm for Graphs). *Invariant 1 and Invariant 2 are preserved throughout Algorithm STRETTGRAPHIMPR, i.e., they hold before the first iteration, after each iteration, and after termination of the main while-loop. Further, Invariant 1 is preserved during each iteration of the main while-loop.*

Proof.

Invariant 1. Whenever a new candidate S is added as a result from ALLSCCs, it is strongly connected, and we set $H_S = T_S = \emptyset$; this in particular implies that the invariant is satisfied after the initialization of the algorithm.

By induction and Theorem 1, the invariant is satisfied whenever Procedure LOCK-STEP-SEARCH returns a candidate C and we set $H_C = T_C = \emptyset$.

Now consider an update of a candidate S where some subset B is deleted from it and assume the invariant holds before the update. In these cases we update H_S and T_S by setting $H_S \leftarrow (H_S \cup \text{Post}(B)) \cap S$ and $T_S \leftarrow (T_S \cup \text{Pre}(B)) \cap S$. This adds the vertices that remain in S and have an edge from a vertex of B to H_S and those with an edge to B to T_S . Suppose a new top (resp. bottom) SCC $\tilde{S} \subseteq S$ emerges in S by the removal of B from S . Then some vertex of \tilde{S} had an outgoing edge to B (resp. an incoming edge from B) and thus is contained in the updated set T_S (resp. H_S), maintaining the invariant. This happens whenever we remove $\text{BAD}(S)$ from S , and whenever we subtract a result from Procedure LOCK-STEP-SEARCH C from S .

Invariant 2 – Disjointness. The sets in $\mathcal{X} \cup \text{goodC}$ are pairwise disjoint at the initialization since goodC is initialized as \emptyset . Furthermore, whenever a set S is added to goodC in an iteration of the main while-loop, a superset $\tilde{S} \supseteq S$ is removed from \mathcal{X} in the same iteration of the while-loop. Therefore by induction the disjointness of the sets in $\mathcal{X} \cup \text{goodC}$ is preserved.

Invariant 2 – Containment of good components. At initialization, \mathcal{X} contains all SCCs of the input graph G . Each good component C of G is strongly connected, so there exists an SCC $Y \supseteq C$ such that $Y \in \mathcal{X}$ for each good component C .

Consider a set $S \in \mathcal{X}$ that is removed from \mathcal{X} at the beginning of an iteration of the main while-loop. Consider further a good component C of G such that $C \subseteq S$. We require that a set $Y \supseteq C$ is added to either \mathcal{X} or goodC in this iteration of the main while-loop.

First, whenever we remove $\text{BAD}(S)$ from S , by Corollary 1 we maintain the fact that $C \subseteq S$. Second, $G[S]$ contains an edge since $C \subseteq S$. Finally, one of the three cases happens:

Case (1): If $|H_S| + |T_S| = 0$, then the set $S \supseteq C$ is added to goodC .

Case (2): If $|H_S| + |T_S| \geq \sqrt{m/\log n}$, then the algorithm computes the SCCs of $G[S]$. Since $C \subseteq S$ is strongly connected, it is completely contained in some SCC Y of $G[S]$, and Y is added either to \mathcal{X} or to goodC .

Case (3): If $0 < |H_S| + |T_S| < \sqrt{m/\log n}$, then the algorithm either adds $S \supseteq C$ to goodC , or partitions S into \tilde{S} and $S \setminus \tilde{S}$. Suppose the latter case happens, then by Theorem 1 we have that \tilde{S} is an SCC of $G[S]$. Further, since $C \subseteq S$ is strongly connected, it is completely contained in some SCC of $G[S]$. Therefore either $C \subseteq \tilde{S}$ or $C \subseteq (S \setminus \tilde{S})$, and both \tilde{S} and $S \setminus \tilde{S}$ are added to \mathcal{X} .

By the above case analysis we have that a set $Y \supseteq C$ is added to either \mathcal{X} or goodC in the iteration of the main while-loop, and thus the invariant is preserved throughout the algorithm. \square

Proof (of Theorem 2).

Correctness. Whenever a candidate set S is added to goodC , it contains an edge by the check at line 11, and $\text{BAD}(S) = \emptyset$ by the check at line 6. Furthermore, (a) at line 12, S is strongly connected by Invariant 1, (b) at line 16, S is strongly connected by the result of ALLSCCS, and (c) at line 22, S is strongly connected by Theorem 1. Therefore we

have that whenever a candidate set is added to `goodC`, it is indeed a good component (soundness).

Finally, by soundness, Invariant 2, the termination of the algorithm (shown below), and the fact that $\mathcal{X} = \emptyset$ at the termination of the algorithm, we have that `goodC` contains all good components of G (completeness).

Symbolic steps analysis. By [24], the initialization with the SCCs of the input graph takes $O(n)$ symbolic steps. Furthermore, the reachability computation in the last step takes $O(n)$ Pre operations.

In each iteration of the outer while-loop, a set S is removed from \mathcal{X} and either (a) a set $S' \subseteq S$ is added to `goodC` and no set is added to \mathcal{X} or (b) at least two sets that are (proper subsets of) a partition of S are added to \mathcal{X} . Both can happen at most $O(n)$ times, thus there can be at most $O(n)$ iterations of the outer while-loop. The Pre and Post operations at lines 11, 26, and 27 can be charged to the iterations of the outer while-loop.

An iteration of the inner while-loop (lines 6-10) is executed only if some vertices B are removed from S ; the vertices of B are then not considered further. Thus there can, in total, be at most $O(n)$ Pre and Post operations over all iterations of the inner while-loop.

Note that every vertex in each of H_S and T_S can be attributed to at least one unique implicit edge deletion since we only add vertices to H_S resp. T_S that are successors resp. predecessors of vertices that were separated from S (or deleted from the maintained graph). Whenever the case $|H_S| + |T_S| \geq \sqrt{m/\log n}$ occurs, for all subsets $C \subseteq S$ that are then added to \mathcal{X} , we initialize $H_C = T_C = \emptyset$. Therefore the case $|H_S| + |T_S| \geq \sqrt{m/\log n}$ can happen at most $O(\sqrt{m \log n})$ times throughout the algorithm since there are at most m edges that can be deleted, and hence in total takes $O(n \cdot \sqrt{m \log n})$ symbolic steps.

It remains to bound the number of symbolic steps in Procedure `LOCK-STEP-SEARCH`. Let C be the set returned by the procedure; we charge the symbolic steps in this call of the procedure to the vertices of the smaller set of C and $S \setminus C$. By Theorem 1 we have either (a) $C = S$, the number of symbolic steps in this call is bounded by $O(\sqrt{m/\log n} \cdot |C|)$, and the set S is added to `goodC` or (b) $\min(|C|, |S \setminus C|) \leq |S|/2$ and the number of symbolic steps in this call is bounded by $O(\sqrt{m/\log n} \cdot \min(|C|, |S \setminus C|))$. Case (a) can happen at most once for the vertices of C , and for case (b) note that the size of a set containing a specific vertex can be halved at most $O(\log n)$ times; thus we charge each vertex at most $O(\log n)$ times. Hence we can bound the total number of symbolic steps in all calls to the procedure by $O(n \cdot \sqrt{m \log n})$. \square

C Details of Section 5: Symbolic MEC Decomposition

C.1 Basic Symbolic Algorithm for MEC decomposition

Recall that an end-component is a set of vertices that (a) has no random edges to vertices not in the set and its induced sub-MDP is (b) strongly connected and (c) contains at least one edge.

Algorithm `MECBASIC` computes all maximal end-components of a given MDP and is formulated as to highlight the similarities to the algorithms for graphs and MDPs with Streett objectives. The algorithm maintains two sets, the set `goodC` of identified

Algorithm MECBASIC: Basic Algorithm for Maximal End-Components

Input : an MDP $P = (G = (V, E), (V_1, V_R))$
Output: the set of maximal end-components of P

```
1 goodC  $\leftarrow$   $\emptyset$ 
2  $\mathcal{X} \leftarrow$  ALLSCCs( $G$ )
3 while  $\mathcal{X} \neq \emptyset$  do
4   remove some  $S \in \mathcal{X}$  from  $\mathcal{X}$ 
5    $rouT \leftarrow S \cap V_R \cap \text{Pre}(V \setminus S)$ 
6   if  $rouT \neq \emptyset$  then
7      $S \leftarrow S \setminus \text{Attr}_R(G, rouT)$ 
8      $\mathcal{X} \leftarrow \mathcal{X} \cup \text{ALLSCCs}(G[S])$ 
9   else
10    if  $\text{Post}(S) \cap S \neq \emptyset$  then /*  $G[S]$  contains at least one edge */
11    |  $\text{goodC} \leftarrow \text{goodC} \cup \{S\}$ 
12 return goodC
```

maximal end-components that is initially empty and the set \mathcal{X} of candidates for maximal end-components that is initialized with the SCCs of the MDP. In each iteration of the while-loop one set S is removed from \mathcal{X} and either (1a) identified as a maximal end-component and added to **goodC** or (1b) removed because the induced sub-MDP does not contain an edge or (2) it contains vertices with outgoing random edges. In the latter case these vertices $rouT$ are identified and their random attractor is removed from S . After this step the sub-MDP induced by the remaining vertices of S might not be strongly connected any more. Therefore the SCCs of this sub-MDP are determined and added to \mathcal{X} as new candidates for maximal end-components. Note that this maintains the invariants that (i) each set in \mathcal{X} induces a strongly connected subgraph and (ii) each end-component is a subset of one set in either **goodC** or \mathcal{X} . By (i) a set in \mathcal{X} is an end-component if it does not have outgoing random edges and the induced sub-MDP contains an edge, i.e., in particular this holds for the sets added to **goodC** (soundness). By (ii) and $\mathcal{X} = \emptyset$ at termination of the while-loop the algorithm identifies all maximal end-components of the MDP (completeness). Since both (1) and (2) can happen at most $O(n)$ times, there are $O(n)$ iterations of the while-loop. In each iteration the most expensive operations are the computation of a random attractor and of SCCs, which can both be done in $O(n)$ symbolic steps. Thus Algorithm MECBASIC correctly computes all maximal end-components of an MDP and takes $O(n^2)$ symbolic steps.

C.2 Improved Symbolic Algorithm for MEC decomposition

Informal description. We show how to determine all maximal end-components (MECs) of an MDP in $O(n\sqrt{m})$ symbolic operations. The difference to the basic algorithm lies in the way strongly connected parts of the MDP are identified after the deletion of vertices that cannot be contained in a MEC. For this the symbolic lock-step search from Section 3 is used whenever not too many edges have been deleted since the last re-computation of SCCs.

Let P be the given MDP and $G = (V, E)$ its underlying graph. The algorithm maintains two sets of vertex sets: the set `goodC` of already identified MECs that is initialized with the empty set and the set \mathcal{X} that is initialized with the SCCs of G and contains vertex sets that are candidates for MECs. The algorithm preserves the following invariant for the `goodC` and \mathcal{X} over the iterations of the while-loop and returns the set `goodC` when the set \mathcal{X} is empty after an iteration of the while-loop.

Invariant 4 (Maintained Sets). *The sets in $\mathcal{X} \cup \text{goodC}$ are pairwise disjoint and for every maximal end-component X of G there exists a set $Y \supseteq X$ such that either $Y \in \mathcal{X}$ or $Y \in \text{goodC}$.*

For each vertex set S in \mathcal{X} additionally a subset T_S of S is maintained that contains vertices that have lost outgoing edges since the last time a superset of S was identified as strongly connected. We use the following restrictions of Invariant 1 and Theorem 1 (presented in Section 3) to bottom SCCs only.

Invariant 5 (Start Vertices BSCC). *Either (a) T_S is empty and $G[S]$ is strongly connected or (b) at least one vertex of each bottom SCC of $G[S]$ is contained in T_S .*

Theorem 5 (Lock-Step Search BSCC). *Provided Invariant 5 holds, Procedure LOCK-STEP-SEARCH(G, S, \emptyset, T_S) returns a bottom SCC $C \subseteq S$ of $G[S]$ in $O(|T_S| \cdot |C|)$ symbolic steps.*

Proof. The proof of Theorem 5 is a straightforward simplification of the proof of Theorem 1 located in Appendix A. \square

Initially the sets T_S are empty. The algorithm maintains Invariant 5 for all $S \in \mathcal{X}$. This will ensure the correctness and the number of symbolic steps of Procedure LOCK-STEP-SEARCH (Section 3) as called by the algorithm.

In each iteration of the while-loop one vertex set S is removed from \mathcal{X} and processed. First the random vertices of S with edges to vertices of $V \setminus S$ are identified and their random attractor is removed from S . After this step, there are no random vertices with edges from S to $V \setminus S$. The predecessors of the removed vertices that are contained in S are added to T_S and additionally T_S is updated to only include vertices that are still in S . This preserves Invariant 5 (see also [30, Lemma 4.5.2]). The number of symbolic steps for the attractor computation can be charged to the removed vertices and is therefore bounded by $O(n)$ in total.

If afterwards $G[S]$ does not contain an edge anymore, then S is not considered further and the algorithm continues with the next iteration. Otherwise one of three cases happens.

Case (1): If T_S is empty, then by Invariant 5 $G[S]$ is strongly connected, contains at least one edge and does not contain a random vertex with edges to $V \setminus S$, i.e., S is an end-component, and by Invariant 4 it is a MEC. In this case the algorithm adds the set S to `goodC`, which preserves both invariants and can happen at most $O(n)$ times.

Case (2): If there are at least \sqrt{m} vertices in T_S , then the set T_S is deleted and as in the basic algorithm all SCCs of $G[S]$ are computed and add to \mathcal{X} as new candidates for MECs. For each of the SCCs C a set T_C is initialized with the empty set. As a vertex is

Algorithm MECIMPR: Improved Algorithm for Maximal End-Components

Input : an MDP $P = (G = (V, E), (V_1, V_R))$
Output : the set of maximal end-components of P

```
1  $\mathcal{X} \leftarrow \text{ALLSCCs}(G)$ ; goodC  $\leftarrow \emptyset$ 
2 foreach  $C \in \mathcal{X}$  do
3    $T_C \leftarrow \emptyset$ 
4 while  $\mathcal{X} \neq \emptyset$  do
5   remove some  $S \in \mathcal{X}$  from  $\mathcal{X}$ 
6    $rouT \leftarrow S \cap V_R \cap \text{Pre}(V \setminus S)$ 
7    $A \leftarrow \text{Attr}_R(G, rouT)$ 
8    $S \leftarrow S \setminus A$ 
9    $T_S \leftarrow (T_S \cup \text{Pre}(A)) \cap S$ 
10  if  $\text{Post}(S) \cap S \neq \emptyset$  then /*  $G[S]$  contains at least one edge */
11    if  $|T_S| = 0$  then
12       $\text{goodC} \leftarrow \text{goodC} \cup \{S\}$ 
13    else if  $|T_S| \geq \sqrt{m}$  then
14      delete  $T_S$ 
15       $\mathcal{C} \leftarrow \text{ALLSCCs}(G[S])$ 
16      if  $|\mathcal{C}| = 1$  then
17         $\text{goodC} \leftarrow \text{goodC} \cup \{S\}$ 
18      else
19        foreach  $C \in \mathcal{C}$  do
20           $T_C \leftarrow \emptyset$ 
21           $\mathcal{X} \leftarrow \mathcal{X} \cup \mathcal{C}$ 
22    else
23       $C \leftarrow \text{LOCK-STEP-SEARCH}(G, S, \emptyset, T_S)$ 
24      if  $\text{Post}(C) \cap C \neq \emptyset$  then /*  $G[C]$  contains at least one edge */
25         $\text{goodC} \leftarrow \text{goodC} \cup \{C\}$ 
26         $S \leftarrow S \setminus C$ 
27         $T_S \leftarrow (T_S \cup \text{Pre}(C)) \cap S$ 
28         $\mathcal{X} \leftarrow \mathcal{X} \cup \{S\}$ 
29 return goodC
```

added to a set T_S only if one of its incoming edges is removed by the algorithm, Case (2) can happen only $O(\sqrt{m})$ times over the whole algorithm. Thus the total number of symbolic steps for this case is $O(n\sqrt{m})$. Note that the Invariants 5 and 4 are preserved.

Case (3): If T_S contains less than \sqrt{m} vertices, then Procedure LOCK-STEP-SEARCH(G, S, \emptyset, T_S) is called. By Invariant 5 and Theorem 5 the procedure returns a bottom SCC C of $G[S]$ in $O(|T_S| \cdot |C|)$ many symbolic steps. Since there are no random edges between S and $V \setminus S$ in P and C has no outgoing edges in $G[S]$, we have that C is an end-component if it contains at least one edge. By Invariant 4 it is also a MEC and is correctly added to **goodC**. As the sets in **goodC** are not considered further by the algorithm, we can charge the symbolic steps of Procedure LOCK-STEP-SEARCH to the vertices of C . Thus this part takes at most $O(n\sqrt{m})$ symbolic steps over the whole algorithm. The vertices of $S \setminus C$ are added back to \mathcal{X} , which preserves Invariant 4. The predecessors of C in $S \setminus C$ are added to $T_{S \setminus C}$ and vertices of C are removed from $T_{S \setminus C}$, which preserves Invariant 5.

By the above case analysis we have that each vertex set that is added to **goodC** is indeed a MEC (soundness). By Invariant 4 and $\mathcal{X} = \emptyset$ at termination of the algorithm we further have completeness. In each iteration either S does not contain an edge and is not considered further, a set is added to **goodC** (and not contained in \mathcal{X} after that) or case (2) happens. Thus there are at most $O(n + \sqrt{m})$ iterations of the algorithm. The symbolic operations we have not yet accounted for in the analysis of the number of symbolic steps are of $O(1)$ per iteration. Hence Algorithm MECIMPR takes $O(n\sqrt{m})$ symbolic steps and correctly computes the MECs of the given MDP P .

Lemma 6 (Invariants of Improved Algorithm for MEC). *Invariant 5 and Invariant 4 are preserved throughout Algorithm MECIMPR, i.e., they hold before the first iteration, after each iteration, and after termination of the main while-loop. Further, Invariant 5 is preserved during each iteration of the main while-loop.*

Proof.

Invariant 5. The proof of maintaining Invariant 5 in Algorithm MECIMPR is a straightforward simplification of the proof of maintaining Invariant 1 in Algorithm STREETGRAPHIMPR (located in Appendix B).

Invariant 4 – Disjointness. The sets in $\mathcal{X} \cup \mathbf{goodC}$ are pairwise disjoint at the initialization since **goodC** is initialized as \emptyset . Furthermore, whenever a set S is added to **goodC** in an iteration of the main while-loop, a superset $\tilde{S} \supseteq S$ is removed from \mathcal{X} in the same iteration of the while-loop. Therefore by induction the disjointness of the sets in $\mathcal{X} \cup \mathbf{goodC}$ is preserved.

Invariant 4 – Containment of maximal end-components. At initialization, \mathcal{X} contains all SCCs of G . Each maximal end-component X of $P = (G = (V, E), (V_1, V_R), \delta)$ is strongly connected, so there exists an SCC $Y \supseteq X$ of G such that $Y \in \mathcal{X}$.

Consider a set $S \in \mathcal{X}$ that is removed from \mathcal{X} at the beginning of an iteration of the main while-loop. Consider further a maximal end-component X of P such that $X \subseteq S$. We require that a set $Y \supseteq X$ is added to either \mathcal{X} or **goodC** in this iteration of the main while-loop.

First, after we remove $\text{Attr}_R(G, S \cap V_R \cap \text{Pre}(V \setminus S))$ from S , we maintain the fact that $X \subseteq S$ by Lemma 2. Second, $G[S]$ contains an edge since $X \subseteq S$. Finally, one of the three cases happens:

Case (1): If $|T_S| = 0$, then the set $S \supseteq X$ is added to **goodC**.

Case (2): If $|T_S| \geq \sqrt{m}$, then the algorithm computes the SCCs of $G[S]$. Since $X \subseteq S$ is strongly connected, it is completely contained in some SCC Y of $G[S]$, and Y is added to \mathcal{X} .

Case (3): If $0 < |T_S| < \sqrt{m}$, then the algorithm partitions S into C and $S \setminus C$. By Theorem 5 we have that C is a (bottom) SCC of $G[S]$. Since $X \subseteq S$ is strongly connected, it is completely contained in some SCC of $G[S]$. Therefore either $X \subseteq C$ or $X \subseteq (S \setminus C)$. The set $S \setminus C$ is added to \mathcal{X} . If $X \subseteq C$, then in particular $G[C]$ contains an edge, and C is added to **goodC**.

By the above case analysis we have that a set $Y \supseteq X$ is added to either \mathcal{X} or **goodC** in the iteration of the main while-loop. \square

Proof (of Theorem 3).

Correctness. A candidate set can be added to **goodC** in three cases. When S is added to **goodC** at line 12 (resp. at line 17), then it contains an edge by the check at line 10, it is strongly connected by $|T_S| = 0$ and Invariant 5 (resp. by the result of ALLSCCs), and it has no random vertices with edges to $V \setminus S$ by the random attractor removal at lines 6-9. When C is added at line 25, then it contains an edge by the check at line 24, it is strongly connected by Theorem 5, it contains no random vertices with edges to $V \setminus S$ by the random attractor removal at lines 6-9, and it contains no random vertices with edges to $S \setminus C$ by the fact that C is a bottom SCC of $G[S]$ (see Theorem 5). Therefore we have that whenever a candidate set is added to **goodC**, it is an end-component, and by induction and Invariant 4 we have that it is a maximal end-component (soundness).

Finally, by soundness, Invariant 4, the termination of the algorithm (shown below), and the fact that $\mathcal{X} = \emptyset$ at the termination of the algorithm, we have that **goodC** contains all the maximal end-components of P (completeness).

Symbolic steps analysis. By [24], the initialization with the SCCs of a given MDP takes $O(n)$ symbolic steps.

In each iteration of the outer while-loop, a set S is removed from \mathcal{X} and (a) S is added to **goodC**, or (b) at least two sets that are (subsets of) a partition of S are added to \mathcal{X} , or (c) S is partitioned into two sets, one of them may be added to **goodC** and the other is added to \mathcal{X} . All three cases can happen at most $O(n)$ times, so there can be at most $O(n)$ iterations of the outer while-loop. The Pre and Post operations at lines 6, 9, 10, 24, and 27 can be charged to the iterations of the outer while-loop.

Each CPre_R operation executed as a part of the random attractor computation at line 7 adds at least one vertex to A , and the vertices of A are then not considered any further in the algorithm. Therefore there can, in total, be at most $O(n)$ CPre_R operations over all attractor computations at line 7.

Note that every vertex in each of T_S can be attributed to at least one unique implicit edge deletion since we only add vertices to T_S that are predecessors of the vertices that were separated from S (or deleted from the maintained graph). Whenever the case $|T_S| \geq \sqrt{m}$ occurs, for all subsets $C \subseteq S$ that are then added to \mathcal{X} , we initialize

$T_C = \emptyset$. Therefore, the case $|T_S| \geq \sqrt{m}$ can happen at most $O(\sqrt{m})$ times throughout the algorithm since there are at most m edges that can be deleted. By [24] we have a bound $O(n)$ for one iteration, so we can bound the total number of symbolic steps in all iterations of this case by $O(n \cdot \sqrt{m})$.

It remains to bound the number of symbolic steps in Procedure LOCK-STEP-SEARCH. Let C be the set returned by LOCK-STEP-SEARCH(G, S, \emptyset, T_S). By Theorem 5 and the fact that $|T_S| < \sqrt{m}$, the number of symbolic steps in this call is bounded by $O(\sqrt{m} \cdot |C|)$, and the set C is not considered further in the algorithm after this call. Hence we can bound the total number of symbolic steps in all calls of the procedure by $O(n \cdot \sqrt{m})$. \square

D Details of Section 6: MDPs with Streett Objectives

D.1 Basic Symbolic Algorithm for MDPs with Streett Objectives

Algorithm STRETTMDPBASIC: Basic Algorithm for MDPs with Streett Obj.

Input : MDP $P = ((V, E), (V_1, V_R), \delta)$ and pairs $TP = \{(L_i, U_i) \mid 1 \leq i \leq k\}$

Output: $\langle\langle 1 \rangle\rangle_{as}(P, \text{Streett}(TP))$

```

1  $\mathcal{X} \leftarrow \text{ALLMECS}(P)$ ;  $\text{goodEC} \leftarrow \emptyset$ 
2 while  $\mathcal{X} \neq \emptyset$  do
3   remove some  $S \in \mathcal{X}$  from  $\mathcal{X}$ 
4    $B \leftarrow \bigcup_{1 \leq i \leq k: U_i \cap S = \emptyset} (L_i \cap S)$ 
5   if  $B \neq \emptyset$  then
6      $S \leftarrow S \setminus \text{Attr}_R(P[S], B)$ 
7      $\mathcal{X} \leftarrow \mathcal{X} \cup \text{ALLMECS}(P[S])$ 
8   else  $\text{goodEC} \leftarrow \text{goodEC} \cup \{S\}$ 
9 return  $\langle\langle 1 \rangle\rangle_{as}(P, \text{Reach}(\bigcup_{X \in \text{goodEC}} X))$ 

```

The pseudocode of the basic symbolic algorithm for MDPs with Streett objectives is given in Algorithm STRETTMDPBASIC. The key differences compared to Algorithm STRETTGRAPHBASIC are as follows: (a) SCC computation is replaced by MEC computation; (b) along with the removal of bad vertices, their random attractor is also removed; and (c) removing the attractor ensures that the check required for trivial SCCs for graphs (line 9) is not required any further.

To compute the almost-sure winning set for MDPs with Streett objectives, we first find all (maximal) good end-components and then solve almost-sure reachability with the union of the good end-components as target set as the last step of the algorithm. This is correct by Lemma 1. Towards finding all good end-components, the algorithm maintains two sets, the set `goodEC` of identified good end-components that is initially empty and the set \mathcal{X} of end-components that are candidates for good end-components that is initialized with the MECs of the MDP. In each iteration of the while-loop one set S is removed from the set of candidates \mathcal{X} and the set of bad vertices $\text{BAD}(S)$ of S is determined. If $\text{BAD}(S)$ is empty, then S is a good end-component and added to `goodEC`. Otherwise the random attractor of $\text{BAD}(S)$ in $P[S]$ is removed from S , which by Corollary 1 does not remove any vertices that are in a good end-component. The

remaining vertices of S have no outgoing random edges and thus still induce a sub-MDP but the sub-MDP might not be strongly connected any more. Then the MECs of this sub-MDP are added to \mathcal{X} . These operations maintain the invariants that (i) each set in \mathcal{X} is an end-component and (ii) each good end-component is a subset of one set in either `goodEC` or \mathcal{X} . By (i) a set in \mathcal{X} is a (maximal) good end-component if it does not contain any bad vertices, i.e., in particular this holds for the sets added to `goodEC` (soundness). By (ii) and $\mathcal{X} = \emptyset$ at termination of the while-loop the algorithm identifies all (maximal) good end-components of the MDP (completeness). Since in each iteration of the while-loop either (1) a set is removed from \mathcal{X} and added to `goodEC` or (2) bad vertices are removed from a set and not considered further by the algorithm, there can be at most $O(n)$ iterations of the while-loop. Furthermore, whenever bad vertices are removed, then the number of target pairs a given candidate set intersects is reduced by one. Thus each vertex is considered in at most $O(k)$ iterations of the while-loop. The most expensive operation in the while-loop is the computation of the MECs. Denoting the number of symbolic steps for the MEC computation with $O(\text{MEC})$, the number of symbolic steps of Algorithm `STREETTMDPBASIC` is $O(\min(n, k) \cdot \text{MEC})$ (assuming that the number of symbolic steps for the almost-sure reachability computation is lower than that).

D.2 Improved Symbolic Algorithm for MDPs with Streett Objectives

We present the technical details regarding the improved symbolic algorithm for MDPs with Streett objectives. The main ideas of the algorithm are presented in Section 6. The pseudocode is given in Algorithm `STREETTMDPIMPR`.

The following invariant is maintained throughout Algorithm `STREETTMDPIMPR` for the sets in `goodEC` and \mathcal{X} .

Invariant 7 (Maintained Sets). *The sets in $\mathcal{X} \cup \text{goodEC}$ are pairwise disjoint and for every good end-component C of G there exists a set $Y \supseteq C$ such that either $Y \in \mathcal{X}$ or $Y \in \text{goodEC}$.*

Furthermore, the algorithm maintains the invariant that each candidate for a good end-component $S \in \mathcal{X}$ contains no random edges to vertices not in S .

Invariant 8 (No Random Outgoing Edges). *Given an MDP P and its underlying graph $G = (V, E)$, for each set $S \in \mathcal{X}$ there are no random vertices in S with edges to vertices in $V \setminus S$.*

Finally, for each candidate set $S \in \mathcal{X}$ the algorithm remembers sets H_S and T_S of vertices that have lost incoming resp. outgoing edges since the last time a superset of S was identified as being strongly connected. The algorithm maintains Invariant 1 and therefore it can use Procedure `LOCK-STEP-SEARCH` together with its correctness guarantee and bound on symbolic steps provided by Theorem 1.

Lemma 9 (Invariants of Improved Algorithm for MDPs). *Invariant 1, Invariant 7, and Invariant 8 are preserved throughout Algorithm `STREETTMDPIMPR`, i.e., they hold before the first iteration, after each iteration, and after termination of the main while-loop. Further, Invariant 1 is preserved during each iteration of the main while-loop.*

Algorithm STREETTMDPIMPR: Improved Alg. for MDPs with Streett Obj.

Input : MDP $P = ((V, E), (V_1, V_R), \delta)$ and pairs $TP = \{(L_i, U_i) \mid 1 \leq i \leq k\}$
Output : $\langle\langle 1 \rangle\rangle_{as}(P, \text{Streett}(TP))$

```

1  $\mathcal{X} \leftarrow \text{ALLMECS}(P)$ ;  $\text{goodEC} \leftarrow \emptyset$ 
2 foreach  $C \in \mathcal{X}$  do  $H_C \leftarrow \emptyset$ ;  $T_C \leftarrow \emptyset$ 
3 while  $\mathcal{X} \neq \emptyset$  do
4   remove some  $S \in \mathcal{X}$  from  $\mathcal{X}$ 
5    $B \leftarrow \bigcup_{1 \leq i \leq k: U_i \cap S = \emptyset} (L_i \cap S)$ 
6   while  $B \neq \emptyset$  do
7      $A \leftarrow \text{Attr}_R(P[S], B)$ 
8      $S \leftarrow S \setminus A$ 
9      $H_S \leftarrow (H_S \cup \text{Post}(A)) \cap S$ 
10     $T_S \leftarrow (T_S \cup \text{Pre}(A)) \cap S$ 
11     $B \leftarrow \bigcup_{1 \leq i \leq k: U_i \cap S = \emptyset} (L_i \cap S)$ 
12  if  $\text{Post}(S) \cap S \neq \emptyset$  then /*  $P[S]$  contains at least one edge */
13    if  $|H_S| + |T_S| = 0$  then  $\text{goodEC} \leftarrow \text{goodEC} \cup \{S\}$ 
14    else if  $|H_S| + |T_S| \geq \sqrt{m/\log n}$  then
15      delete  $H_S$  and  $T_S$ 
16       $\mathcal{C} \leftarrow \text{ALLSCCS}(P[S])$ 
17      if  $|\mathcal{C}| = 1$  then  $\text{goodEC} \leftarrow \text{goodEC} \cup \{S\}$ 
18      else
19        foreach  $C \in \mathcal{C}$  do
20           $\text{rout} \leftarrow C \cap V_R \cap \text{Pre}(S \setminus C)$ 
21           $A \leftarrow \text{Attr}_R(P[C], \text{rout})$ 
22           $C \leftarrow C \setminus A$ 
23           $H_C \leftarrow \text{Post}(A) \cap C$ 
24           $T_C \leftarrow \text{Pre}(A) \cap C$ 
25           $\mathcal{X} \leftarrow \mathcal{X} \cup \{C\}$ 
26      else
27         $(C, H_S, T_S) \leftarrow \text{LOCK-STEP-SEARCH}(G, S, H_S, T_S)$ 
28        if  $C = S$  then  $\text{goodEC} \leftarrow \text{goodEC} \cup \{S\}$ 
29        else /* separate  $C$  and  $S \setminus C$  */
30           $\text{rout}_C \leftarrow C \cap V_R \cap \text{Pre}(S \setminus C)$  /* empty if  $C$  bottom SCC */
31           $A_C \leftarrow \text{Attr}_R(P[C], \text{rout}_C)$  /*  $= \text{Attr}_R(P[S], S \setminus C) \cap C$  */
32           $A_S \leftarrow \text{Attr}_R(P[S], C)$ 
33           $C \leftarrow C \setminus A_C$ 
34           $S \leftarrow S \setminus A_S$ 
35           $H_C \leftarrow \text{Post}(A_C) \cap C$ 
36           $T_C \leftarrow \text{Pre}(A_C) \cap C$ 
37           $H_S \leftarrow (H_S \cup \text{Post}(A_S)) \cap S$ 
38           $T_S \leftarrow (T_S \cup \text{Pre}(A_S)) \cap S$ 
39           $\mathcal{X} \leftarrow \mathcal{X} \cup \{S\} \cup \{C\}$ 
38 return  $\langle\langle 1 \rangle\rangle_{as}(P, \text{Reach}(\bigcup_{C \in \text{goodEC}} C))$ 

```

Proof.

Invariant 1. The proof is a minor extension of the maintenance proof for Algorithm STREETGRAPHIMPR that is given in Appendix B. In terms of strong connectivity of a candidate S and the maintenance of the sets H_S and T_S , the only difference to the graph case is that after an SCC C is computed by ALLSCCs or Procedure LOCK-STEP-SEARCH, another subset of vertices A (vertices with outgoing random edges and their random attractor) is removed from C . In this case the invariant is maintained by initializing H_C resp. T_C with the vertices of $C \setminus A$ with edges from resp. to vertices of A , i.e., $H_C \leftarrow \text{Post}(A) \cap C$ and $T_C \leftarrow \text{Pre}(A) \cap C$.

Invariant 7 – Disjointness. The sets in $\mathcal{X} \cup \text{goodEC}$ are pairwise disjoint at the initialization since goodEC is initialized as \emptyset . Furthermore, whenever a set S is added to goodEC in an iteration of the main while-loop, a superset $\tilde{S} \supseteq S$ is removed from \mathcal{X} in the same iteration of the while-loop. Therefore by induction the disjointness of the sets in $\mathcal{X} \cup \text{goodEC}$ is preserved.

Invariant 7 – Containment of good end-components. At initialization, \mathcal{X} contains all MECs of the input MDP $P = (G = (V, E), (V_1, V_R), \delta)$. Each good end-component C of P is an end-component, so there exists a MEC $Y \supseteq C$ such that $Y \in \mathcal{X}$ for each good end-component C .

Consider a set $S \in \mathcal{X}$ that is removed from \mathcal{X} at the beginning of an iteration of the main while-loop. Consider further a good end-component C of P such that $C \subseteq S$. We require that a set $Y \supseteq C$ is added to either \mathcal{X} or goodEC in this iteration of the main while-loop.

First, whenever we remove $\text{Attr}_R(P[S], \text{BAD}(S))$ from S , by Corollary 1, we maintain the fact that $C \subseteq S$. Second, $P[S]$ contains an edge since $C \subseteq S$. Finally, one of the three cases happens:

Case (1): If $|H_S| + |T_S| = 0$, then the set $S \supseteq C$ is added to goodEC .

Case (2): If $|H_S| + |T_S| \geq \sqrt{m/\log n}$, then the algorithm computes the SCCs of $P[S]$. If S itself is the (sole) SCC of $P[S]$, then it is added to goodEC . Otherwise, since $C \subseteq S$ is strongly connected, it is completely contained in some SCC Y of $P[S]$. Furthermore, since C has no outgoing random edges, by Lemma 2 it is contained in Y even after we remove $\text{Attr}_R(P[Y], Y \cap V_R \cap \text{Pre}(S \setminus Y))$ from it. Finally, Y is added to \mathcal{X} .

Case (3): If $0 < |H_S| + |T_S| < \sqrt{m/\log n}$, then the algorithm either adds $S \supseteq C$ to goodEC , or partitions S into \tilde{S} and $S \setminus \tilde{S}$. Suppose the latter case happens, then by Theorem 1 we have that \tilde{S} is an SCC of $P[S]$. Further, since $C \subseteq S$ is strongly connected, it is completely contained in some SCC of $P[S]$. Therefore either $C \subseteq \tilde{S}$ or $C \subseteq (S \setminus \tilde{S})$. If $C \subseteq \tilde{S}$, then by Lemma 2 after the removal of $\text{Attr}_R(P[\tilde{S}], \tilde{S} \cap V_R \cap \text{Pre}(S \setminus \tilde{S}))$ from \tilde{S} we maintain that $C \subseteq \tilde{S}$. If $C \subseteq (S \setminus \tilde{S})$, then by Lemma 2 after the removal of $\text{Attr}_R(P[S], \tilde{S})$ from $(S \setminus \tilde{S})$ we maintain that $C \subseteq (S \setminus \tilde{S})$. Finally, both \tilde{S} and $S \setminus \tilde{S}$ are added to \mathcal{X} .

By the above case analysis we have that a set $Y \supseteq C$ is added to either \mathcal{X} or goodEC in the iteration of the main while-loop.

Invariant 8. Given an MDP, the set \mathcal{X} is initialized with the MECs of the MDP, and by definition they have no random outgoing edges. Therefore the invariant holds before the first iteration of the main while-loop.

Consider a candidate set $S \in \mathcal{X}$ in a given iteration of the main while-loop. By the induction hypothesis, S has no random vertices with edges to $V \setminus S$. First, some bad vertices can be iteratively removed from S . At each such removal, the random attractor to these vertices is removed from S as well. After the removal, by the definition of a random attractor, S has no random outgoing edges to the attractor, and therefore by induction has no random outgoing edges to $V \setminus S$. Second, S may be partitioned into at least two proper subsets. Then for each such subset C , the random attractor to random vertices in C with edges to $S \setminus C$ is removed from C . By induction and the definition of a random attractor, after the removal C contains no random outgoing edges to $V \setminus C$ and adding it to \mathcal{X} preserves the invariant. \square

Proof (of Theorem 4).

Correctness. Whenever a candidate set S is added to `goodEC`, it contains an edge by the check at line 12, $\text{BAD}(S) = \emptyset$ by the check at line 6, and it has no outgoing random edges by Invariant 8 and the random attractor removal at line 8. Furthermore, (a) at line 13, S is strongly connected by Invariant 1, (b) at line 17, S is strongly connected by the result of `ALLSCCS`, and (c) at line 28, S is strongly connected by Theorem 1. Therefore we have that whenever a candidate set is added to `goodEC`, it is indeed a good end-component (soundness).

Finally, by soundness, Invariant 7, the termination of the algorithm (shown below), and the fact that $\mathcal{X} = \emptyset$ at the termination of the algorithm, we have that `goodEC` contains all good end-components of G (completeness).

Symbolic steps analysis. When using our improved symbolic algorithm for MEC decomposition, the initialization takes $O(n \cdot \sqrt{m})$ symbolic steps by Theorem 3.

In each iteration of the outer while-loop, a set S is removed from \mathcal{X} and either (a) a set $S' \subseteq S$ is added to `goodEC` and no set is added to \mathcal{X} or (b) at least two sets that are (subsets of) a partition of S are added to \mathcal{X} . Both can happen at most $O(n)$ times, thus there can be at most $O(n)$ iterations of the outer while-loop. The Pre and Post operations at lines 12, 30, 33, 34, 35, and 36 can be charged to the iterations of the outer while-loop.

An iteration of the inner while-loop (line 6) is executed only if some vertices B are removed from S ; the vertices of B are then not considered further. Thus there can, in total, be at most $O(n)$ Post operations at line 9 and Pre operations at line 10 over all iterations of the inner while-loop.

Similarly, each `CPreR` operation executed as a part of a random attractor computation adds at least one vertex to the attractor, and the vertices of the attractor are then not considered any further in the algorithm. Therefore there can, in total, be at most $O(n)$ `CPreR` operations over all attractor computations at lines 7, 21, 30, and 30.

Note that every vertex in each of H_S and T_S can be attributed to at least one unique implicit edge deletion since we only add vertices to H_S resp. T_S that are successors resp. predecessors of vertices that were separated from S (or deleted from the maintained graph). Whenever the case $|H_S| + |T_S| \geq \sqrt{m/\log n}$ occurs, for all subsets $C \subseteq S$ that are then added to \mathcal{X} , we initialize $H_C = T_C = \emptyset$. Therefore, the case $|H_S| + |T_S| \geq \sqrt{m/\log n}$ can happen at most $O(\sqrt{m \log n})$ times throughout the algorithm since there are at most m edges that can be deleted. In one iteration of this case, the number

of symbolic steps executed by ALLSCCs together with symbolic steps executed at lines 20, 23, and 24, is bounded by $O(n)$ [24].

It remains to bound the number of symbolic steps in Procedure LOCK-STEP-SEARCH. Let C be the set returned by the procedure; we charge the symbolic steps in this call of the procedure to the vertices of the smaller set of C and $S \setminus C$. By Theorem 1 we have either (a) $C = S$, the number of symbolic steps in this call is bounded by $O(\sqrt{m/\log n} \cdot |C|)$, and the set S is added to `goodEC` or (b) $\min(|C|, |S \setminus C|) \leq |S|/2$ and the number of symbolic steps in this call is bounded by $O(\sqrt{m/\log n} \cdot \min(|C|, |S \setminus C|))$. Case (a) can happen at most once for the vertices of C , and for case (b) note that the size of a set containing a specific vertex can be halved at most $O(\log n)$ times; thus we charge each vertex at most $O(\log n)$ times. Hence we can bound the total number of symbolic steps in all calls to the procedure by $O(n \cdot \sqrt{m \log n})$. \square

E Details of Section 7: Experiments

We present the results of the experimental evaluation when comparing based on the time. In all the figures, both axes plot the amount of seconds spent on the execution. Similar to the case of symbolic steps, we begin the measurement after the initial preprocessing step (computing all SCCs for graphs and all MECs for MDPs) is finished. The results for graphs are shown in Figure 4 and the results for MDPs are shown in Figure 5.

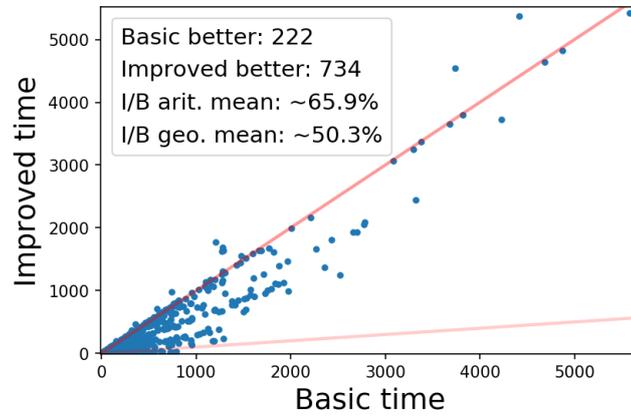
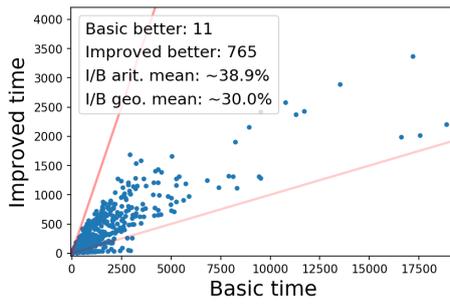
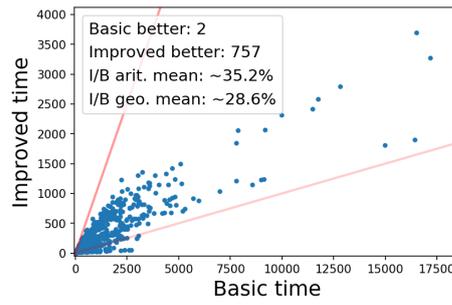


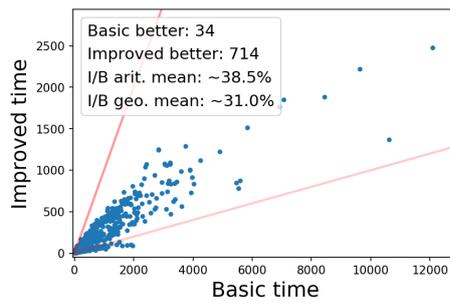
Fig. 4: Comparison of time for graphs with Streett objectives.



(a) 10% random vertices



(b) 20% random vertices



(c) 50% random vertices

Fig. 5: Comparison of time for MDPs with Streett objectives.