

Witnessing Subsystems for Probabilistic Systems with Low Tree Width

Simon Jantsch Jakob Piribauer Christel Baier

Technische Universität Dresden*

{simon.jantsch, jakob.piribauer, christel.baier}@tu-dresden.de

A standard way of justifying that a certain probabilistic property holds in a system is to provide a witnessing subsystem (also called critical subsystem) for the property. Computing *minimal* witnessing subsystems is NP-hard already for acyclic Markov chains, but can be done in polynomial time for Markov chains whose underlying graph is a tree. This paper considers the problem for probabilistic systems that are *similar* to trees or paths. It introduces the parameters *directed tree-partition width* (dtpw) and *directed path-partition width* (dppw) and shows that computing minimal witnesses remains NP-hard for Markov chains with bounded dppw (and hence also for Markov chains with bounded dtpw). By observing that graphs of bounded dtpw have bounded width with respect to all known tree similarity measures for directed graphs, the hardness result carries over to these other tree similarity measures. Technically, the reduction proceeds via the conceptually simpler *matrix-pair chain problem*, which is introduced and shown to be NP-complete for nonnegative matrices of fixed dimension. Furthermore, an algorithm which aims to utilise a given directed tree partition of the system to compute a minimal witnessing subsystem is described. It enumerates partial subsystems for the blocks of the partition along the tree order, and keeps only necessary ones. A preliminary experimental analysis shows that it outperforms other approaches on certain benchmarks which have directed tree partitions of small width.

1 Introduction

The ability to justify and explain a verification result is an important feature in the context of verification. For example, classical model checking algorithms for linear temporal logic (LTL) return a *counterexample* if the system does not satisfy the property. In this case a counterexample is typically an ultimately periodic trace of the system which does not satisfy the formula. *What* constitutes a valid counterexample is dependent on both the kind of system and the kind of specification that is considered [13].

For *probabilistic systems*, which can be modeled using *discrete-time Markov chains* (DTMC) or the more general *Markov decision processes* (MDP), a single trace of the system is usually not enough to justify that a given (probabilistic) property holds [1]. A typical class of properties in the probabilistic setting are *probabilistic reachability constraints*, where one asks whether the (maximal or minimal, for MDPs) probability to reach a set of goal states satisfies a threshold condition. The motivation for considering reachability (apart from the fact that it is a fundamental property) is that one can treat ω -regular properties using methods for reachability over the product of the system with an automaton for the property [36]. To justify that the probability to reach the goal is higher than some threshold in a DTMC one can return a *set of traces* of the system whose probability exceeds the threshold [21]. Another notion of counterexamples for probabilistic reachability constraints are *witnessing subsystems* (also called

*This work was supported by DFG grant 389792660 as part of TRR 248 – CPEC, see <https://perspicuous-computing.science>, the Cluster of Excellence EXC 2050/1 (CeTI, project ID 390696704, as part of Germany’s Excellence Strategy), DFG-projects BA-1679/11-1, BA-1679/12-1 and the Research Training Group QuantLA (GRK 1763).

critical subsystems) [16, 35, 36]. The idea is to justify a lower bound on the reachability probability by providing a *subsystem* which *by itself* already exceeds the threshold. Apart from providing explanations to a human working with the system model, probabilistic explanations have been used in automated analysis frameworks such as *counterexample-guided abstraction refinement* [22] or *counterexample-guided inductive synthesis* [2, 8]. In all these applications it is important to find *small* explanations (all paths of the system or the entire system as subsystem are trivial explanations in case the property holds).

Computing *minimal* witnessing subsystems in terms of number of states is computationally difficult. The corresponding decision problem, henceforth called the *witness problem*, is NP-complete already for acyclic DTMCs [16]. Known algorithms rely on *mixed-integer linear programming* (MILP) [16, 35, 36] or *vertex enumeration* [16]. On the other hand, the problem is in P for DTMCs whose underlying graph is a tree [16]. This leads to the natural question of whether efficient algorithms exist for systems whose underlying graph is *similar* to a tree. A parameter that measures this is the *tree width* of a graph, which has been studied extensively in graph theory [6]. Several NP-hard problems for graphs, for example the 3-coloring problem, are in P for graphs with bounded tree width [6]. Faster algorithms for standard problems in probabilistic model checking were proposed for systems of small tree width [3, 11]. Algorithms for non-probabilistic quantitative verification problems on models with low tree width were considered in [10]. One motivation for studying such systems is that *control-flow graphs* of programs in languages such as JAVA and C, under certain syntactic restrictions, are known to have bounded tree width [19, 34]. A stronger notion than tree width is *path width* [31], which intuitively measures how similar a graph is to a path and has been applied in fields such as graph drawing [14] and natural language processing [28].

The standard notion of tree width is defined for undirected graphs. Related notions have been considered for directed graphs, although here the theory is not as mature and as of now there is no standard notion [26, 30, 32]. A stronger notion than tree width for undirected graphs is that of *tree-partition width* [20, 33, 37], which requires a partition of the graph whose induced quotient structure is a tree. To the best of our knowledge, *tree-partition width* has not been studied for directed graphs so far.

Contributions.

- The paper introduces a tree-similarity measure called directed tree-partition width (dtpw), which can be seen as the directed analogue to the tree-partition width for undirected graphs that is known from the literature. We show that deciding whether there exists a directed tree partition with width at most k is NP-complete (Section 3).
- The second main contribution is NP-completeness for the witness problem in Markov chains with bounded dtpw (Section 4). This implies that the problem is NP-hard for bounded-width Markov chains with respect to the following tree-similarity measures for directed graphs: directed tree width (from [26]), bounded D-width (from [32]) and bounded undirected tree width. The reduction proceeds via the conceptually simpler d -dimensional matrix-pair chain problem, which we introduce and show to be NP-complete (Section 4.1) for fixed d , even for nonnegative matrices.
- We describe a dedicated algorithm that computes minimal witnesses by proceeding bottom-up along a given directed tree partition. It enumerates partial subsystems for the blocks of the partition, but keeps only necessary ones (Section 5). On certain instances that do have a good tree-decomposition, our prototype implementation significantly outperforms the standard MILP approach (Section 5.3).

2 Preliminaries

Graphs, partitions and quotients. A set of subsets $\{S_1, \dots, S_n\}$ of a given set S is a *partition* of S if $S_i \cap S_j = \emptyset$ for all $1 \leq i < j \leq n$ and $\bigcup_{1 \leq i \leq n} S_i = S$. The elements of a partition are called *blocks*. Given a

graph $G = (V, E)$ and a partition $\mathcal{P} = \{V_1, \dots, V_n\}$ of V , the *quotient* of G under \mathcal{P} is the graph $(\mathcal{P}, E_{\mathcal{P}})$, where $(V_i, V_j) \in E_{\mathcal{P}}$ iff there is some $(s, t) \in E$ such that $s \in V_i$ and $t \in V_j$. A directed graph is a *tree* if its underlying undirected graph is a tree and every node has in-degree at most one. A directed graph is a *path* if it is a tree and all vertices have outdegree at most one.

Probabilistic systems. A *discrete-time Markov chain* (DTMC) \mathcal{M} is a tuple (S, P, ι) where S is a set of states, $P : S \times S \rightarrow \mathbb{Q}_{\geq 0}$ is the *probabilistic transition matrix*, which needs to satisfy $\sum_{s' \in S} P(s, s') \leq 1$ for all $s \in S$, and where the *initial distribution* ι satisfies $\sum_{s \in S} \iota(s) \leq 1$.¹ A state s with $\sum_{s' \in S} P(s, s') = 0$ is called a trap state. A path is a finite or infinite sequence $s_0 s_1 s_2 \dots \in S^{\leq \omega}$ such that $\iota(s_0) > 0$ and for all $i \geq 0$: $P(s_i, s_{i+1}) > 0$. A *maximal path* is a path that is infinite or that ends in a trap state. The set of maximal paths of \mathcal{M} is called $\text{Paths}(\mathcal{M})$ and carries a (sub-)probability measure whose associated σ -algebra is generated by the cylinder sets $\text{Cyl}(\tau) = \{\pi \in \text{Paths}(\mathcal{M}) \mid \tau \text{ is a prefix of } \pi\}$, which have probability $\text{Cyl}(s_0 \dots s_n) = \iota(s_0) \cdot \prod_{0 \leq i < n} P(s_i, s_{i+1})$. The probability of a measurable set $\Pi \subseteq \text{Paths}(\mathcal{M})$ is denoted by $\text{Pr}_{\mathcal{M}}(\Pi)$. A *Markov decision process* (MDP) \mathcal{M} is a tuple $(S, \text{Act}, P, \iota)$ where S is a set of states, $\text{Act} : S \rightarrow 2^A$ is a function that assigns to each state a finite set of *actions* from the set A , $P : S \times A \times S \rightarrow \mathbb{Q}_{\geq 0}$ satisfies $\sum_{s' \in S} P(s, \alpha, s') \leq 1$ for all (s, α) with $s \in S$ and $\alpha \in \text{Act}(s)$, and ι satisfies $\sum_{s \in S} \iota(s) \leq 1$. The paths of \mathcal{M} are finite or infinite sequences $s_0 \alpha_0 s_1 \alpha_1 \dots \in (S \times A)^{\leq \omega}$ such that $\iota(s_0) > 0$ and for all $i \geq 0$: $P(s_i, \alpha_i, s_{i+1}) > 0$ and $\alpha_i \in \text{Act}(s_i)$. The set $\text{Paths}_{\text{fin}}(\mathcal{M})$ denotes the finite paths ending in a state. The set $\text{Paths}(\mathcal{M})$ is the set of maximal paths, which are infinite or end in a state s with $\sum_{s' \in S} P(s, \alpha, s') = 0$ for all actions $\alpha \in \text{Act}(s)$. A *scheduler* of \mathcal{M} is a function $\mathfrak{S} : \text{Paths}_{\text{fin}}(\mathcal{M}) \rightarrow A$ satisfying $\mathfrak{S}(s_0 \alpha_0 \dots s_n) \in \text{Act}(s_n)$. Every scheduler \mathfrak{S} of \mathcal{M} induces a (possibly infinite) Markov chain $\mathcal{M}_{\mathfrak{S}}$ and thereby a (sub-)probability measure on $\text{Paths}(\mathcal{M})$. The maximal, respectively minimal, probabilities of some path property $\Pi \subseteq \text{Paths}(\mathcal{M})$ are defined as $\mathbf{Pr}_{\mathcal{M}}^{\max}(\Pi) = \sup_{\mathfrak{S}} \text{Pr}_{\mathcal{M}_{\mathfrak{S}}}(\Pi)$ and $\mathbf{Pr}_{\mathcal{M}}^{\min}(\Pi) = \inf_{\mathfrak{S}} \text{Pr}_{\mathcal{M}_{\mathfrak{S}}}(\Pi)$, where \mathfrak{S} ranges over all schedulers of \mathcal{M} . For ω -regular properties this notation is justified as the supremum, respectively infimum, is attained by some scheduler. For more details see [4, Chapter 10]. The underlying graph of an MDP $\mathcal{M} = (S, \text{Act}, P, \iota)$ has vertices S and edges: $\{(s, s') \in S \times S \mid \text{there exists } \alpha \in \text{Act} \text{ such that } P(s, \alpha, s') > 0\}$. We denote by $\mathcal{M}(s)$ the MDP one gets by replacing the initial distribution in \mathcal{M} by the dirac distribution on $s \in S$.

Witnessing subsystems. Let $\mathcal{M} = (S, \text{Act}, P, \iota)$ be an MDP. A *subsystem* of \mathcal{M} is an MDP $\mathcal{M}' = (S', \text{Act}', P', \iota')$ where $S' \subseteq S$ and for all (s, α) with $s \in S'$ and $\alpha \in \text{Act}(s)$ and $s' \in S'$: $P'(s, \alpha, s') \in \{P(s, \alpha, s'), 0\}$. Similarly, we require for all $s \in S$: $\iota'(s) \in \{\iota(s), 0\}$. In words, edges of the subsystem either retain the probability of the original system, or have probability zero. This again results in a sub-stochastic MDP. The subsystem *induced* by a set of states $S' \subseteq S$ is defined as $\mathcal{M}_{S'} = (S', \text{Act}', P', \iota')$ where $P'(s, \alpha, s') = P(s, \alpha, s')$ if $s, s' \in S'$, and otherwise $P'(s, \alpha, s') = 0$, and similarly for ι' . Given an MDP $\mathcal{M}' = (S', \text{Act}', P', \iota')$, we assume that there is a set of trap states Goal . By $\diamond \text{Goal}$, we denote the set of paths that contain a state $t \in \text{Goal}$. Any subsystem \mathcal{M}' of \mathcal{M} satisfies $\mathbf{Pr}_{\mathcal{M}'}^{\max}(\diamond \text{Goal}) \leq \mathbf{Pr}_{\mathcal{M}}^{\max}(\diamond \text{Goal})$ and $\mathbf{Pr}_{\mathcal{M}'}^{\min}(\diamond \text{Goal}) \leq \mathbf{Pr}_{\mathcal{M}}^{\min}(\diamond \text{Goal})$ as the probability to reach Goal cannot increase under any scheduler by setting transition probabilities to 0. The subsystem \mathcal{M}' is a *witness* for $\mathbf{Pr}_{\mathcal{M}}^*(\diamond \text{Goal}) \geq \lambda$ if $\mathbf{Pr}_{\mathcal{M}'}^*(\diamond \text{Goal}) \geq \lambda$, where $*$ $\in \{\min, \max\}$ and $\lambda \in [0, 1]$. The definition for DTMCs is analogous.

¹We require the transition matrix and initial distribution to be *sub-stochastic* here. To obtain a *stochastic* transition matrix and initial distribution one can add a state *fail* together with edges to *fail* carrying the missing probability. Using sub-stochastic DTMCs technically simplifies the treatment of subsystems and allows to ignore the *fail* state in the underlying graph.

3 Directed tree- and path-partition width

We propose a natural extension of tree-partition width [37] to directed graphs. In what follows, let $G = (V, E)$ be a fixed finite directed graph.

Definition 1 (Directed tree partition). *A partition $\mathcal{P} = \{V_1, \dots, V_n\}$ of V is a directed tree partition of G if the quotient of G under \mathcal{P} is a tree. We denote by $DTP(G)$ the set of directed tree partitions of G .*

Definition 2 (Directed tree-partition width (dtpw)). *The directed tree-partition width of graph G is:*

$$\text{dtpw}(G) := \min_{\mathcal{P} \in DTP(G)} \max_{S \in \mathcal{P}} |S|$$

Replacing *tree* by *path* in definitions 1 and 2 leads to the notions of *directed path partition* and *directed path-partition width* (dppw). Any strongly connected component of a graph needs to be included in a single block of the partition, which distinguishes this notion from other notions of tree width for directed graphs. In particular, it is not the case that the standard tree width of an undirected graph G_u equals the directed tree-partition width of the directed graph one gets by including both edges (u, v) and (v, u) whenever u and v are connected in G_u .

Proposition 3. *If a class \mathcal{C} of graphs has bounded directed tree-partition width, then \mathcal{C} has bounded directed tree width (from [26]), bounded D -width (from [32]) and bounded undirected tree width.*

Deciding whether a directed tree partition with small width exists is NP-complete. The reduction goes from the *oneway bisection problem* [15] in directed graphs, which asks whether there exists a partition of the given graph into two equally-sized vertex sets V_0, V_1 such that all edges go from V_0 to V_1 .

Proposition 4. *The two problems (1) decide $\text{dppw}(G) \leq k$ and (2) decide $\text{dtpw}(G) \leq k$, given a directed graph G and $k \in \mathbb{N}$, are both NP-complete.*

4 Hardness of the witness problem for DTMCs with low directed path-partition width

The witness problem for probabilistic reachability in DTMCs is defined as follows.

Definition 5 (Witness problem). *The witness problem takes as input a DTMC \mathcal{M} , $k \in \mathbb{N}$ and $\lambda \in \mathbb{Q}$, and asks whether there exists a witnessing subsystem for $\Pr_{\mathcal{M}}(\diamond \text{Goal}) \geq \lambda$ with at most k states.*

The problem is NP-complete for acyclic DTMCs, but in P for DTMCs whose underlying graph is a tree [16]. In this section we prove NP-hardness of the witness problem for DTMCs with directed path-partition width at most 6 (and hence also for DTMCs with directed tree-partition at most 6). The idea is to relate different subsets of the blocks of a directed path partition to matrices that describe how the reachability probability is passed on through the system. This leads us to a natural intermediate problem:

Definition 6 (d -dimensional matrix-pair chain problem). *The d -dimensional matrix-pair chain problem (d -MCP) takes as input a sequence $(M_0^1, M_1^1), \dots, (M_0^n, M_1^n)$, where $M_i^j \in \mathbb{Q}^{d \times d}$, a starting vector $\iota \in \mathbb{Q}^{1 \times d}$, final vector $f \in \mathbb{Q}^{d \times 1}$, and $\lambda \in \mathbb{Q}$ (with all numbers encoded in binary) and asks whether there exists a tuple $(\sigma_1, \dots, \sigma_n) \in \{0, 1\}^n$ such that:*

$$\iota \cdot M_{\sigma_1}^1 \cdots M_{\sigma_n}^n \cdot f \geq \lambda$$

The nonnegative variant of the problem restricts all input numbers to be nonnegative.

In the following we will show that the 2-MCP is NP-hard and can be reduced to the *nonnegative* 3-MCP. The reason that we are interested in the nonnegative variant is that we would like to reduce it to the witness problem, where one can naturally encode (sub)stochastic matrices, but it is unclear how to deal with negative values.

The chain of polynomial reductions that the argument uses is the following:



4.1 Hardness of the matrix-pair chain problem

To show NP-hardness of the 2-MCP we reduce from the *partition* problem, which is among Karp's 21 NP-complete problems [27]. Given a finite set $S = \{s_1, \dots, s_n\} \subseteq \mathbb{Z}$, whose elements are encoded in binary, it asks whether there exists $W \subseteq S$ such that $\sum W = \sum(S \setminus W)$, where $\sum X = \sum_{x \in X} x$. The main idea of the polynomial reduction is to relate each s_i to a pair of matrices M_0^i, M_1^i where M_0^i realizes a clockwise rotation by an angle which corresponds to the value of s_i , and M_1^i realizes the counter-clockwise rotation by the same angle. Then for all $\sigma_1, \dots, \sigma_n \in \{0, 1\}^n$ we have that $W = \{s_i \mid \sigma_i = 1\}$ satisfies $\sum W = \sum(S \setminus W)$ iff $M_{\sigma_1}^1 \cdots M_{\sigma_n}^n$ equals the identity matrix.

In the reduction we choose initial vector $v = (1/2, 1/2)^T$, final vector $f = (1/2, 1/2)^T$ and threshold $\lambda = 1/2$. In this way, the chosen chain of rotation matrices $M_{\sigma_1}^1 \cdots M_{\sigma_n}^n$ is applied to f before it is checked whether the resulting vector v satisfies $(1/2, 1/2) \cdot v \geq 1/2$. This is the case iff $v = (1/2, 1/2)^T$ and hence iff $M_{\sigma_1}^1 \cdots M_{\sigma_n}^n$ is equal to the identity matrix. For an illustration of this idea, see Figure 1.

Rotation matrices, however, may have irrational entries in general. It is shown in [7] that for any rational rotation angle φ and $\varepsilon \in \mathbb{Q}_{>0}$, a rotation matrix R^θ (rotating by angle θ) with rational entries can be computed in time polynomial in $\log(1/\varepsilon)$ such that $|\varphi - \theta| < \varepsilon$. Using such rational matrices which rotate by approximately the desired angles and by estimating the resulting precision of the matrix multiplication, we can provide a slightly smaller threshold $\lambda' < \lambda$ to complete the reduction from the partition problem to the 2-MCP with rational matrices (a detailed proof can be found in the extended version [25]).

Proposition 7. *The two-dimensional matrix-pair chain problem (2-MCP) is NP-complete.*

The proof of Proposition 7 depends crucially on the fact that negative numbers are allowed in the 2-MCP, as we would not be able to use the rotation matrices otherwise. To encode the 2-MCP into the *nonnegative* 3-MCP, one can embed the two-dimensional dynamics of a given instance of 2-MCP into the

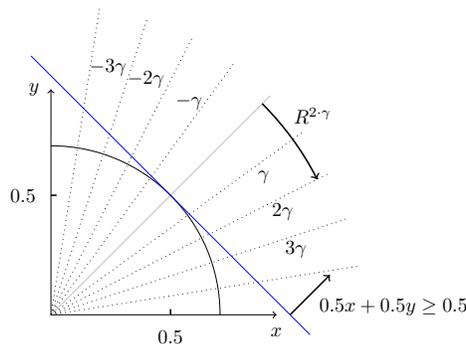


Figure 1: Sketch for Proposition 7. Matrices are rotations by angles that are multiples of γ . In order to satisfy the threshold, the final point needs to lie in the halfspace on the right of the blue line. This is only possible if the angles of the rotation matrices sum up to zero.

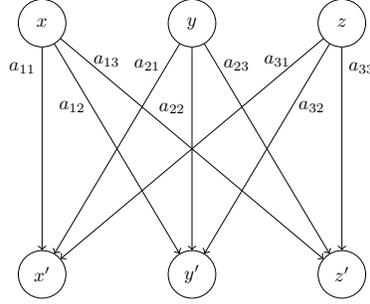


Figure 2: A gadget to encode matrix multiplication. Let M be a substochastic matrix with entries $(M)_{ij} = a_{ij} \in \mathbb{Q}_{\geq 0}$. If the probability of states (x', y', z') to reach some goal state is (v'_x, v'_y, v'_z) , then these probabilities in states (x, y, z) are $M \cdot (v'_x, v'_y, v'_z)^T$. We will abbreviate this gadget (i.e. the transitions) by a double-arrow annotated with the matrix (\xrightarrow{M}) .

two-dimensional plane with normal vector $(1, 1, 1)$ in three dimensions. To obtain nonnegative matrices, we push vectors further into the direction $(1, 1, 1)$ at each matrix multiplication step, while preserving the original 2d-dynamics when projecting onto the subspace orthogonal to $(1, 1, 1)$.

To sketch this idea in more detail, let $(M_0^1, M_1^1), \dots, (M_0^n, M_1^n)$, with $M_i^j \in \mathbb{Q}^{2 \times 2}$ for all $(i, j) \in \{0, 1\} \times \{1, \dots, n\}$, and $\iota \in \mathbb{Q}^{1 \times 2}, f \in \mathbb{Q}^{2 \times 1}$ be an instance of 2-MCP. For some $\kappa \in \mathbb{Q}$, we define

$$N_i^j = B \begin{pmatrix} M_i^j & \mathbf{0} \\ \mathbf{0} & \kappa \end{pmatrix} B^{-1}, \quad \iota' = (\iota \quad \kappa) B^{-1}, \quad f' = B \begin{pmatrix} f \\ \kappa \end{pmatrix} \quad \text{and} \quad \lambda' = \lambda + \kappa^{n+2}$$

where we use the matrix

$$B = \begin{pmatrix} 1 & 1 & 1 \\ -1 & 1 & 1 \\ 0 & -2 & 1 \end{pmatrix} \quad \text{with inverse} \quad B^{-1} = 1/6 \cdot \begin{pmatrix} 3 & -3 & 0 \\ 1 & 1 & -2 \\ 2 & 2 & 2 \end{pmatrix}$$

to change the basis. Note that the columns of B are orthogonal to each other and that the third standard basis vector is mapped to $(1, 1, 1)$ under the change of basis. For any $\sigma_1, \dots, \sigma_n \in \{0, 1\}^n$, it is now easy to compute that:

$$\iota' \cdot N_{\sigma_1}^1 \cdots N_{\sigma_n}^n \cdot f' = \iota \cdot M_{\sigma_1}^1 \cdots M_{\sigma_n}^n \cdot f + \kappa^{n+2}.$$

So, the constructed instance of the 3-MCP is a yes-instance if and only if the original instance of the 2-MCP is one. By choosing κ large enough, we furthermore can make sure that all matrices N_j^i are nonnegative. This completes the proof of the following proposition. Details can be found in the extended version [25].

Proposition 8. *The nonnegative three-dimensional matrix-pair chain problem (nonnegative 3-MCP) is NP-complete.*

4.2 Hardness of the witness problem

The aim of this section is to prove that the witness problem is NP-hard for Markov chains with bounded path-partition width. The proof goes by a polynomial reduction from the nonnegative 3-MCP. Let

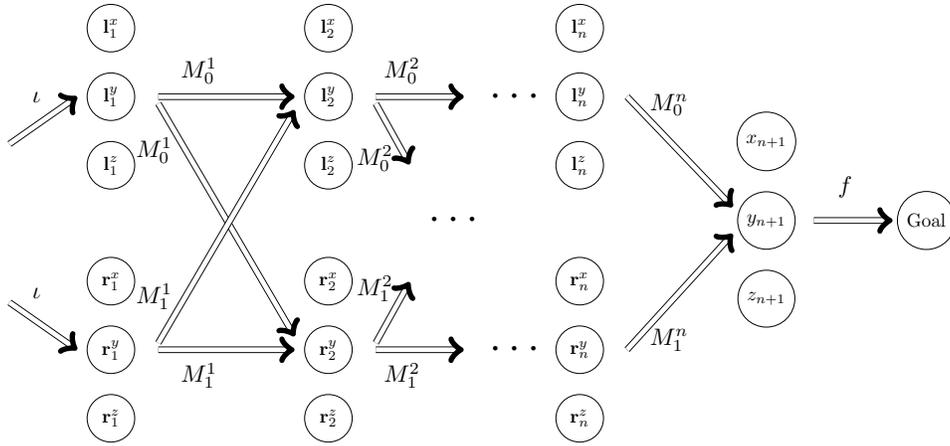


Figure 3: Main structure of the reduction. The *good subsystems* contain either the states $\{\mathbf{l}_j^x, \mathbf{l}_j^y, \mathbf{l}_j^z\}$ or $\{\mathbf{r}_j^x, \mathbf{r}_j^y, \mathbf{r}_j^z\}$ for each layer j and thereby correspond to a choice $\sigma_1 \dots \sigma_n$ in the matrix-pair chain problem.

$(M_0^1, M_1^1), \dots, (M_0^n, M_1^n)$, ι , f and λ be an instance of this problem. For technical reasons explained later, we assume that all entries of the input matrices and vectors are in the range $[1/12 - \varepsilon, 1/12]$ for some ε that satisfies:

$$0 < 12\varepsilon < 1/3 \cdot (1/12 - \varepsilon)^{n+2} \quad (1)$$

In the extended version [25], we show that the nonnegative 3-MCP problem remains NP-hard under these assumptions.

4.2.1 Structure of the reduction

As a first step, Figure 2 shows how one can encode the multiplication of a (suitable) vector with a (suitable) matrix in a Markov chain. Using this gadget, Figure 3 shows the main structure of the reduction from the nonnegative 3-MCP. The double arrows represent instances of the gadget from Figure 2. The initial distribution assigns probability $\iota(x)$ to both states \mathbf{r}_1^x and \mathbf{l}_1^x , and similarly for y and z . The final edge from state x_{n+1} to Goal has probability $f(x)$, and analogously for all other states on the final layer. The above assumption on the entries of the matrices and vectors guarantees that the sums of initial probabilities and outgoing probabilities of any state are below 1. Hence the result is indeed a Markov chain which we call \mathcal{M}_1 . Furthermore, the directed tree-partition width and directed path-partition width of \mathcal{M}_1 are both independent of the 3-MCP instance (for the proof, see the extended version [25]):

Lemma 9. $dppw(\mathcal{M}_1) = dtpw(\mathcal{M}_1) = 6$.

Let $\text{left}_i = \{\mathbf{l}_i^x, \mathbf{l}_i^y, \mathbf{l}_i^z\}$ and $\text{right}_i = \{\mathbf{r}_i^x, \mathbf{r}_i^y, \mathbf{r}_i^z\}$. A subsystem $S' \subseteq S$ is called *good* if it includes the states $\{x_{n+1}, y_{n+1}, z_{n+1}\}$ and for all $1 \leq i \leq n$:

$$\text{either } \text{left}_i \subseteq S' \text{ and } \text{right}_i \cap S' = \emptyset \quad \text{or} \quad \text{right}_i \subseteq S' \text{ and } \text{left}_i \cap S' = \emptyset$$

That is, S' chooses exactly one of the sets left_i and right_i in each layer i , with $1 \leq i \leq n$. Good subsystems have exactly $3n + 4$ states (including Goal). Subsystems that are not good are called *bad*. There is a one-to-one correspondance between good subsystems and matrix sequences in the matrix-pair chain problem. For a given sequence $\sigma = \sigma_1, \dots, \sigma_n \in \{0, 1\}^n$, let:

$$S_\sigma = \{x_{n+1}, y_{n+1}, z_{n+1}\} \cup \bigcup_{\substack{1 \leq i \leq n \\ \sigma_i = 0}} \text{left}_i \cup \bigcup_{\substack{1 \leq i \leq n \\ \sigma_i = 1}} \text{right}_i$$

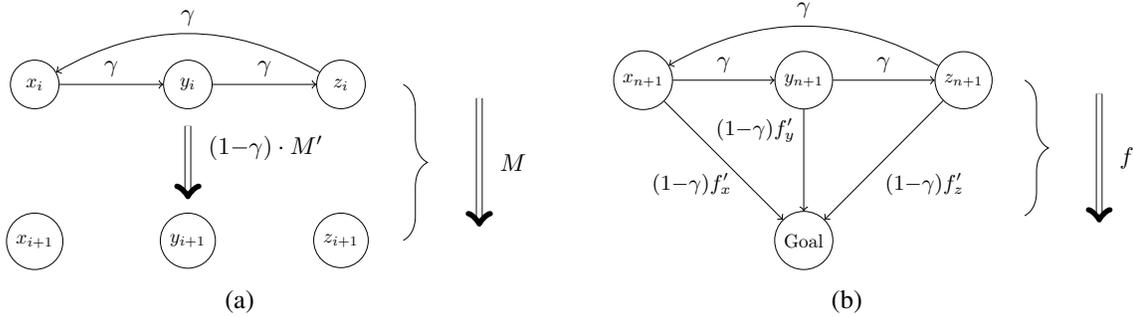


Figure 4: A γ -cycle is added to the upper states of the matrix multiplication gadget (see Figure 2) to make sure that removing any state on the cycle leads to a significant drop in probability. Figure 4a shows the construction used in all but the last layer, which is handled by the construction in Figure 4b. In Figure 4a, the matrix M' is chosen such that the probability of reaching $(x_{i+1}, y_{i+1}, z_{i+1})$ is $\theta \cdot M$, where θ is any initial distribution on states (x_i, y_i, z_i) , and similarly in Figure 4b.

In the following we denote by $\Pr_{S'}^{\mathcal{M}_1}(\diamond \text{Goal})$ the probability of reaching Goal under the subsystem induced by S' in \mathcal{M}_1 . The following lemma shows that the probability of reaching *goal* in a good subsystems coincides with the corresponding matrix product (see the extended version [25] for the proof).

Lemma 10. *For all $\sigma \in \{0, 1\}^n$ we have: $\Pr_{S_\sigma}^{\mathcal{M}_1}(\diamond \text{Goal}) = \iota \cdot M_{\sigma_1}^1 \cdots M_{\sigma_n}^n \cdot f$.*

It follows that the 3-MCP reduces to deciding whether there exists a *good* subsystem whose probability to reach goal is at least λ . However, it could still be the case that while the 3-MCP instance is a no-instance, there is some *bad* subsystem of size n that satisfies the threshold condition. We now show how \mathcal{M}_1 can be adapted such that the subsystems of size $3n + 4$ with greatest probability are the good ones.

4.2.2 Interconnecting states

The idea is to make sure that bad subsystems have decisively less probability to reach *goal*. To this end we adapt the matrix multiplication gadget from Figure 2 such that removing any state leads to a large drop in probability. This is achieved by adding a cycle which connects the upper states, as shown in Figure 4a. The states x_i, y_i, z_i represent one of the triples $\mathbf{l}_i^x, \mathbf{l}_i^y, \mathbf{l}_i^z$ or $\mathbf{r}_i^x, \mathbf{r}_i^y, \mathbf{r}_i^z$, and likewise for $x_{i+1}, y_{i+1}, z_{i+1}$. The probability of staying inside the cycle is γ in each state and the matrix that contains the pairwise probabilities of reaching the states $(x_{i+1}, y_{i+1}, z_{i+1})$ from states (x_i, y_i, z_i) is:

$$M = \underbrace{\frac{1-\gamma}{1-\gamma^3} \cdot \begin{pmatrix} 1 & \gamma & \gamma^2 \\ \gamma^2 & 1 & \gamma \\ \gamma & \gamma^2 & 1 \end{pmatrix}}_R M' \quad \text{with} \quad R^{-1} = \frac{1}{1-\gamma} \begin{pmatrix} 1 & -\gamma & 0 \\ 0 & 1 & -\gamma \\ -\gamma & 0 & 1 \end{pmatrix} \quad (2)$$

The edges between the last layer $(x_{n+1}, y_{n+1}, z_{n+1})$ and Goal are adapted in a similar way. Let us assume that we are given a matrix $(M)_{ij} = a_{ij}$ (this will be one of the input matrices of the 3-MCP) and we want to find M' such that the gadget from Figure 4a realizes the matrix multiplication M . In other words, we want the probability to reach x_{i+1} from x_i to be exactly a_{11} , and similarly for the other states. Solving the equation above for M' gives:

$$M' = R^{-1} \cdot M = \frac{1}{1-\gamma} \begin{pmatrix} a_{11} - \gamma a_{21} & a_{12} - \gamma a_{22} & a_{13} - \gamma a_{23} \\ a_{21} - \gamma a_{31} & a_{22} - \gamma a_{32} & a_{23} - \gamma a_{33} \\ a_{31} - \gamma a_{11} & a_{32} - \gamma a_{12} & a_{33} - \gamma a_{13} \end{pmatrix} \quad (3)$$

We choose γ to satisfy:

$$12\varepsilon < 1 - \gamma < 1/3 \cdot (3(1/12 - \varepsilon))^{n+2} \quad (4)$$

which is possible due to the assumption of Equation (1). This makes sure that all entries of M' are nonnegative. The argument uses that the entries a_{ij} are assumed to be in the range $[1/12 - \varepsilon, 1/12]$:

$$\frac{1}{1-\gamma}(a - \gamma a') \geq \frac{1}{1-\gamma}(1/12 - \varepsilon - \gamma/12) = 1/12 - \frac{\varepsilon}{1-\gamma} > 0 \quad \text{for all entries } a, a' \text{ of } M$$

where the last inequality follows from $12\varepsilon < 1 - \gamma$. Furthermore, we have:

$$\frac{1}{1-\gamma}(a - \gamma a') \leq \frac{1}{1-\gamma}(1/12 - \gamma(1/12 - \varepsilon)) = 1/12 + \frac{\gamma\varepsilon}{1-\gamma} < 1/6 \quad \text{for all entries } a, a' \text{ of } M$$

where the last inequality follows from $\gamma < 1$ and $12\varepsilon < 1 - \gamma$, which is equivalent to $\varepsilon/(1 - \gamma) < 1/12$. The fact that $1/6$ is an upper bound on all entries of M' implies that using the gadgets from Figure 4 in the reduction yields a DTMC (observe that all states in Figure 3 have at most 6 outgoing edges).

We call the Markov chain that is obtained by adding the γ -cycles and adapting the probabilities as discussed above \mathcal{M}_2 . The construction ensures that the good subsystems (defined as for \mathcal{M}_1) have the same probability to reach Goal in both DTMCs, and hence Lemma 11 holds as well for \mathcal{M}_2 . The main point of adding the γ -cycles was to make sure that if one state from $\{x', y', z'\}$ is excluded in a subsystem, then the probability of any state in x_i, y_i, z_i to reach the next layer drops significantly. Now this value is indeed bounded by $(1 + \gamma + \gamma^2) \cdot (1 - \gamma) < 3 \cdot (1 - \gamma)$ (as $\gamma < 1$). In a bad subsystem, both γ -cycles are interrupted on some layer. Hence, the probability to reach Goal is less than $3 \cdot (1 - \gamma)$. This value in turn is less than $(3(1/12 - \varepsilon))^{n+2}$ by Equation (4). On the other hand, the fact that all entries of matrices M_i^j (with $1 \leq i \leq n$ and $0 \leq j \leq 1$) and vectors l, f have value at least $1/12 - \varepsilon$ implies that $(3(1/12 - \varepsilon))^{n+2}$ is a lower bound on the reachability probability that is achieved by any good subsystem. A detailed discussion of these facts that lead to the following lemma can be found in the extended version [25].

Lemma 11. *Let S_1 and S_2 be a subsystems of \mathcal{M}_2 with $3n + 4$ states. If S_1 is bad and S_2 good, then*

$$\Pr_{S_1}^{\mathcal{M}_2}(\diamond \text{Goal}) \leq \Pr_{S_2}^{\mathcal{M}_2}(\diamond \text{Goal})$$

Finally, note that the directed path-partition-width and tree-partition-width of \mathcal{M}_2 is the same as of \mathcal{M}_1 , as \mathcal{M}_2 includes more edges but still allows the directed path-partition which partitions states along the layers. Hence we have $\text{dtpw}(\mathcal{M}_2) = \text{dppw}(\mathcal{M}_1) = 6$. Together with Lemma 11, Lemma 10 and the fact that the probabilities of good subsystems in \mathcal{M}_1 and \mathcal{M}_2 coincide, this proves:

Theorem 12. *The witness problem is NP-hard for Markov chains with $\text{dppw} = 6$ (and hence also for Markov chains with $\text{dtpw} = 6$).*

5 A dedicated algorithm for MDPs and a given tree partition

This section introduces an algorithm (Algorithm 2) that computes a minimal witnessing subsystem using a given *directed tree partition* of the system. The main idea is to proceed bottom-up along the induced tree order and enumerate partial subsystems for each block and to compute the values achieved by the “interface” states for each partial subsystem. Interface states are those that have incoming edges from the predecessor block in the tree partition. A domination relation between partial subsystems is used to prune away all partial subsystems that do not need to be considered further up, as a “better” one exists.

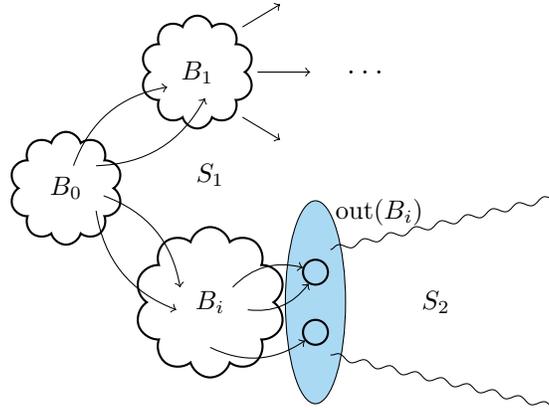


Figure 5: Visualization of the situation in Lemma 13. Some block B_i is fixed, and $\text{out}(B_i)$ are the states outside of B_i which are reachable from some state in B_i . The state set is partitioned into the sets S_1 , which includes B_i and all states that are not reachable from B_i , and S_2 , which includes all states reachable from B_i but excluding B_i . Additionally, in Lemma 13 all of these sets are intersected with a set S' in Lemma 13, which represents some subset of the entire system (this is not depicted here).

Let $\mathcal{M} = (S, \text{Act}, P, \iota)$ be a fixed MDP for the rest of this section, and $\mathcal{P} = \{B_1, \dots, B_n\}$ be a directed tree partition of \mathcal{M} . We will assume that for all $B \in \mathcal{P}$ we have $B \subseteq \text{Goal}$ or $B \cap \text{Goal} = \emptyset$. This is not a real restriction as states in Goal are trap states, which means that they can always be moved to a separate block. Furthermore, we assume that all initial states are in the root block of the tree partition. We denote by $\text{children}(B_i) \subseteq \mathcal{P}$ the children of B_i in the associated tree order, and by $\text{parent}(B_i) \in \mathcal{P}$ the unique parent of B_i . For each block B_i we denote by $\text{inc}(B_i)$ the states in B_i which have some incoming edge from a state in $\text{parent}(B_i)$ or are initial. Using this notion we define $\text{out}(B_i) = \bigcup_{B \in \text{children}(B_i)} \text{inc}(B)$. We denote by $\text{cl}(B_i)$ the union of blocks $B \in \mathcal{P}$ such that B is reachable from B_i in the tree order.

For a given partial function f from S to $[0, 1]$ and subset $S' \subseteq S$ we consider the MDP $\mathcal{M}_{S'}^f$ constructed as follows. In the subsystem $\mathcal{M}_{S'}$ induced by S' remove all outgoing edges from states $s \in \text{dom}(f)$ (the domain of f) and replace them by an action with a single transition to (some state in) Goal with probability $f(s)$ (resulting again in a sub-stochastic MDP). We define for each state $q \in S'$:

$$\min\text{-val}_{S'}^f(q) = \Pr_{\mathcal{M}_{S'}^f(q)}^{\min}(\diamond \text{Goal}) \quad \text{and} \quad \max\text{-val}_{S'}^f(q) = \Pr_{\mathcal{M}_{S'}^f(q)}^{\max}(\diamond \text{Goal})$$

We write $\min\text{-val}_{S'}$ or $\max\text{-val}_{S'}$ for the respective values in the unchanged MDP $\mathcal{M}_{S'}$. The following lemma shows that to compute the values of states in B_i under any subsystem, one can first compute the values of states in $\text{out}(B_i)$, then replace the edges of those states by an edge to Goal carrying this value, and finally compute the values for states in B_i in the adapted system.

Lemma 13. *Let $S' \subseteq S$, $S_2 = S' \cap (\text{cl}(B_i) \setminus B_i)$ (with $1 \leq i \leq n$) and $S_1 = S' \setminus S_2$. Let $v \in \{\max\text{-val}, \min\text{-val}\}$ and define f over domain $\text{out}(B_i)$ by: $f(q) = v_{S_2}(q)$ for all $q \in \text{out}(B_i)$. Finally, let $S'_1 = S_1 \cup \text{out}(B_i)$.*

Then, for all $q \in S'_1$:

$$v_{S'}(q) = v_{S'_1}^f(q)$$

5.1 The domination relation

In the following, the vector f can be thought of as an assumption on the value that is achieved in states in $\text{dom}(f)$. Different partial subsystems of the system in a subtree will correspond to different

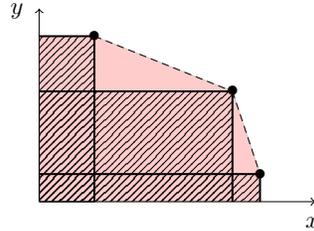


Figure 6: The black points represent three partial subsystems for $I = \{x, y\}$ via their I -points. The red area indicates I -points of partial subsystems which are dominated by these three points, while the dashed area indicates the partial subsystems *strongly* dominated by one of them. The size of the partial subsystems is not considered here, but is important in general (see Definition 15).

vectors f , where $\text{dom}(f)$ are the interface states. For two partial functions $f_1, f_2 : S \rightarrow [0, 1]$ such that $\text{dom}(f_1) = \text{dom}(f_2)$ we define $(f_1 + f_2)(q) = f_1(q) + f_2(q)$, for $q \in \text{dom}(f_1)$, and write $f_1 \leq f_2$ to mean $f_1(q) \leq f_2(q)$ for all $q \in \text{dom}(f_1)$. By $\mathbf{1}$ we denote the constant 1-function with suitable domain.

Lemma 14. *Let $T \subseteq S$, $v \in \{\text{max-val}, \text{min-val}\}$ and $f_1, f_2 : S \rightarrow [0, 1]$ be partial functions such that $\text{dom}(f_1) = \text{dom}(f_2)$ and let $I \subseteq T$ be a set of states that cannot reach Goal without seeing $\text{dom}(f_1)$ in \mathcal{M} . Then, for all $q \in I$:*

1. $f_1 \geq f_2 \implies v_T^{f_1}(q) \geq v_T^{f_2}(q)$,
2. for all $a \in \mathbb{Q}_{\geq 0}$ such that $a \cdot f \leq \mathbf{1}$: $a \cdot v_T^f(q) = v_T^{a \cdot f}(q)$,
3. if $f_1 + f_2 \leq \mathbf{1}$, then: $\text{max-val}_T^{f_1+f_2}(q) \leq \text{max-val}_T^{f_1}(q) + \text{max-val}_T^{f_2}(q)$.

Let $v \in \{\text{max-val}, \text{min-val}\}$ be fixed for the remainder of this section. For a given set $I \subseteq S$, we denote the states *reachable* from I in the underlying graph of \mathcal{M} by $\text{reach}(I)$. A *partial subsystem* for I is a set $T \subseteq \text{reach}(I)$ and the I -point corresponding to T is defined to be the vector $\text{val}_I(T) \in \mathbb{Q}^I$ with $\text{val}_I(T)(q) = v_T(q)$ if $q \in T \cap I$ (where v_T is the value vector under subsystem T) and $\text{val}_I(T)(q) = 0$ if $q \in I \setminus T$. Let π be the function which collects all possible projections of a vector $\theta \in \mathbb{Q}^I$ onto a subset of the axes as follows:

$$\pi(\theta) = \{\pi(\theta, D) \mid D \subseteq I\} \quad \text{and} \quad \pi(\theta, D)(x) = \begin{cases} \theta(x) & x \in D \\ 0 & \text{otherwise} \end{cases}$$

If a partial subsystem T for I includes no Goal state, then $\text{val}_I(T)$ will be zero in all entries, and hence we will not be interested in such T . However, T does not have to include all Goal states reachable from I .

The core of Algorithm 2 is the domination relation (see Figure 6) which is used to discard partial subsystems. Without it the algorithm would amount to an explicit enumeration of all subsystems.

Definition 15. *Let $I \subseteq S$ and $\{T\} \cup \mathcal{S}$ be a set of partial subsystems for I . We say that \mathcal{S} dominates T if there exists $\mathcal{S}' \subseteq \mathcal{S}$ such that*

1. For all $T' \in \mathcal{S}'$ we have $|T'| \leq |T|$, and
2. $\text{val}_I(T)$ is a convex combination of $\bigcup \{\pi(\text{val}_I(T')) \mid T' \in \mathcal{S}'\}$.

We say that \mathcal{S} *strongly* dominates T if there exists a singleton set $\mathcal{S}' \subseteq \mathcal{S}$ such that \mathcal{S}' dominates T .

Algorithm 1: removeDominated

```

Input: Set of partial subsystems  $\mathcal{S}$  for  $I$ , with  $I \subseteq S$ .
/* Group partial subsystems by their size. */
1  $\mathcal{S}[k] := \{S' \in \mathcal{S} \mid |S'| = k\}$ 
2  $m := \max\{|S'| \mid S' \in \mathcal{S}\}$ 
/* Initialise an empty ConvexHull object */
3  $\mathcal{H} := \text{ConvexHull}()$ 
4 for  $k = 1$  to  $m$  do
    /* Compute projections of value vectors in  $\mathcal{S}[k]$ . */
5      $\Pi := \bigcup\{\pi(\text{val}_I(S')) \mid S' \in \mathcal{S}[k]\}$ 
    /* Add  $\Pi$  to the incremental ConvexHull object. */
6      $\mathcal{H}.\text{addPoints}(\Pi)$ 
    /* Remember only subsystems in  $\mathcal{S}[k]$  that are vertices of  $\mathcal{H}$ . */
7      $R := R \cup \{S' \in \mathcal{S}[k] \mid \text{val}_I(S') \in \mathcal{H}.\text{vertices}\}$ 
8 end
9 return  $R$ 

```

Lemma 16. *Let $S' \subseteq S$, $S_2 = S' \cap (\text{cl}(B_i) \setminus B_i)$ (for some $1 \leq i \leq n$) and $S_1 = S' \setminus S_2$. Furthermore, let $I = \text{out}(B_i)$ and \mathcal{S} be a set of partial subsystems for I .*

1. *If $\mathcal{M}_{S'}$ is a witnessing subsystem for $\mathbf{Pr}^{\min}(\diamond \text{Goal}) \geq \lambda$ and \mathcal{S} strongly dominates S_2 , then there is a $T \in \mathcal{S}$ such that $\mathcal{M}_{S_1 \cup T}$ is a witnessing subsystem for $\mathbf{Pr}^{\min}(\diamond \text{Goal}) \geq \lambda$ and $|T| \leq |S_2|$.*
2. *If $\mathcal{M}_{S'}$ is a witnessing subsystem for $\mathbf{Pr}^{\max}(\diamond \text{Goal}) \geq \lambda$ and \mathcal{S} dominates S_2 , then there is a $T \in \mathcal{S}$ such that $\mathcal{M}_{S_1 \cup T}$ is a witnessing subsystem for $\mathbf{Pr}^{\max}(\diamond \text{Goal}) \geq \lambda$ and $|T| \leq |S_2|$.*

Algorithm 1 details how the domination relation can be computed using an incremental convex-hull algorithm. The ConvexHull object that is used in line 3 allows to add points incrementally, and stores the vertices of the convex hull of points added so far in the field *vertices*. The convex hull of a points in d dimensions can be computed in $O(a \cdot \log a + a^{\lfloor d/2 \rfloor})$ [12]. In our case d corresponds to the number of interface states $|I|$, and as a number of dedicated and fast algorithms exist to compute the convex hull in low dimensions [5, 9, 17] tree partitions with few interface states in each block are desirable.

Lemma 17. *Let \mathcal{S} be a set of partial subsystems for $I \subseteq S$ and $R = \text{removeDominated}(\mathcal{S})$. Then,*

- *for any $T \in \mathcal{S} \setminus R$ it holds that R dominates T .*
- *no $T \in R$ is dominated by $R \setminus \{T\}$.*

In order to avoid enumerating all subsets of a block we first apply a filter based on a Boolean condition. It requires that any state in the subset either is an interface state or has a predecessor in the subset. Likewise it should either have a successor in the subset, or an outgoing edge to another block. Consider the following Boolean formula with variables in S :

$$\phi(B_i) = \bigwedge_{s \notin \text{inc}(B_i)} \left(s \rightarrow \bigvee_{s' \in \text{pre}(s)} s' \right) \wedge \bigwedge_{s \notin \text{exit}(B_i)} \left(s \rightarrow \bigvee_{s' \in \text{post}(s)} s' \right)$$

where $\text{exit}(B_i) = \{s \in B_i \mid \text{post}(s) \setminus B_i \neq \emptyset\}$. Now any partial subsystem S' such that $S' \cap B_i$ is not a model of $\phi(B_i)$ is dominated by another subsystem, which one gets by removing unnecessary states.

Algorithm 2: A dedicated algorithm for MDPs using a given directed tree partition.

```

Input: MDP  $\mathcal{M}$ , directed tree partition  $\mathcal{P}$ , rational  $\lambda$ 
Output: Minimal witnessing subsystem for  $\Pr_{\mathcal{M}}^{\max}(\diamond \text{Goal}) \geq \lambda$ .
/* Bottom-up traversal of the tree partition. */
1 for  $B$  in  $\text{reverse}(\text{topologicalSort}(\mathcal{P}))$  do
2    $I := \text{inc}(B)$ 
3    $O := \text{out}(B)$ 
4   /* Consider only subsets of  $B$  that satisfy  $\phi(B)$  */
5   for  $S_B \subseteq B$  such that  $S_B \models \phi(B)$  do
6     /* Consider each combination of partial subsystems of the children of  $B$ . */
7     for  $(S', \text{val}_O(S'))$  in  $\text{successorPoints}(\text{psubsys}, B)$  do
8        $f := \text{val}_O(S')$ 
9       /* The new partial subsystem  $S_{\text{new}}$  for  $I$  combines  $S_B$  and  $S'$ . */
10       $S_{\text{new}} := S_B \cup S'$ 
11       $\text{val}_I(S_{\text{new}}) := (\max\text{-val}_I^f)|_I$ 
12      /* Remember the corresponding partial subsystem. */
13       $\text{psubsys}[B].\text{insert}(S_{\text{new}})$ 
14    end
15    /* Remove dominated points */
16     $\text{psubsys}[B] := \text{removeDominated}(\text{psubsys}[B])$ 
17  end
18 end
19 /* Here  $B_r$  is assumed to be the root of the tree associated with  $\mathcal{P}$ . */
20 return  $\text{argmin}\{|S'| \text{ for } S' \text{ in } \text{psubsys}[B_r] \text{ such that } t \cdot \text{val}_{\text{supp}(t)} \geq \lambda\}$ 

```

5.2 An algorithm based on the domination relation

Algorithm 2 computes a minimal witnessing subsystem of \mathcal{M} for $\Pr_{\mathcal{M}}^{\max}(\diamond \text{Goal}) \geq \lambda$, using the structure of the tree decomposition \mathcal{P} . Witnesses for $\Pr_{\mathcal{M}}^{\min}(\diamond \text{Goal}) \geq \lambda$ can be handled by replacing the call to `removeDominated` in Line 11 by a method which computes the *strong* domination relation (see Table 1 for the possible instances of the algorithm). Computing the strong domination relation requires checking whether the I -point of a new partial subsystem is pointwise smaller than that of any of the given partial subsystems. The algorithm keeps a map `psubsys` from blocks $B \in \mathcal{P}$ to partial subsystems for $\text{inc}(B)$. This map is populated in a bottom-up traversal of \mathcal{P} (Line 1). For a given block B , the models of $\phi(B)$ (which are subsets of B) are enumerated (Line 4). The method `successorPoints` in Line 5 returns all partial subsystems for $O = \text{out}(B)$ which can be obtained by combining partial subsystems in `psubsys[Bi]` for all $B_i \in \text{children}(B)$. More precisely, if $\text{children}(B_i) = \{P_1, \dots, P_k\}$, then:

$$\text{successorPoints}(\text{psubsys}, B) = \{(T, \text{val}_O(T)) \mid S_1 \in \text{psubsys}[P_1], \dots, S_k \in \text{psubsys}[P_k]\}$$

where $T = \bigcup_{1 \leq i \leq k} S_i$ and $\text{val}_O(T)$ is the vector one gets by concatenating vectors $\text{val}_{\text{inc}(P_i)}(S_i)$ (recall that $O = \text{out}(B) = \bigcup_{1 \leq i \leq k} \text{inc}(P_i)$, and the blocks are pairwise disjoint). The vectors $\text{val}_{\text{inc}(P_i)}(S_i)$ have been computed during a previous iteration of the for loop in Line 8 and are assumed to be in global memory (they are also needed in the algorithm `removeDominated`).

Proposition 18. *If Algorithm 2 returns S' on input $(\mathcal{M}, \mathcal{P}, \lambda)$, then the output S' is a minimal witness for $\Pr_{\mathcal{M}}^{\max}(\diamond \text{Goal}) \geq \lambda$. It returns within exponential time in the size of the input.*

Additional heuristics to exclude partial subsystems. In addition to the domination relation we propose two conditions on when a partial subsystem can be excluded. First, suppose we are considering block

B_i with interface $I = \text{inc}(B_i)$, and let $\text{cl}(B_i)$ be the union of blocks reachable from B_i and $R = S \setminus \text{cl}(B_i)$. If using *all states* from R together with a partial subsystem T for $\text{inc}(B_i)$ does not lead to a value above λ then T can be excluded. A sufficient condition for this which can easily be checked is if $\text{Goal} \subseteq \text{cl}(B_i)$ holds and the sum of entries of the value vector $\text{val}_I(T)$ is less than λ .

For the second condition, assume that N is an upper bound on the size of a minimal witnessing subsystem (this could have been computed by a heuristic approach) and let M be the length of a shortest path from the initial state into any state of I (these can be computed in advance and in polynomial time). Now if $M + |T| > N$, then T cannot be part of any minimal witness, and can be excluded.

5.3 Experimental evaluation

We have implemented Algorithm 2 in the tool SWITSS [24] using the convex hull library qhull². The experiments were performed on a computer with two Intel E5-2680 8 cores at 2.70 GHz running Linux, where each instance got assigned a single core, a maximum of 10GiB of memory and 900 seconds. All datasets and instructions to reproduce can be found in [23]. At the moment, the implementation only supports DTMCs as input and only returns the size of a minimal witnessing subsystem. To evaluate it, we consider the *bounded retransmission protocol* (brp) for file transfers, which is a standard benchmark included in the PRISM benchmark suite [29]. It is parametrized by N (the number of “chunks”) and K (the number of retransmissions). We fix $K = 1$ but consider increasing values for N , yielding instances of size between 151 ($N = 8$) and 1591 ($N = 80$) in terms of state numbers. We consider the probabilistic reachability constraint $\Pr(\diamond \text{Goal}) \geq \lambda$, for varying thresholds λ and fixed Goal.

The protocol maintains a counter which is only increased up to maximal value N , and using this fact one can compute a natural directed path partition for the model which essentially partitions the state space along the possible values of the counter. The directed path partitions that we get have length $N+1$ and constant width 37. After filtering out the subsets of a block B that do not satisfy $\phi(B)$ (see Section 5.1) at most 15 subsets remain. Figure 7 compares the computation time of Algorithm 2 against known *mixed-integer linear programming* (MILP) based approaches to compute minimal witnessing subsystems [16, 35]. The computation times do not include the generation of the path partition, which is straight forward in this particular case. In the figure, “min” and “max” refer to the two MILPs derived from the polytopes \mathcal{P}^{\min} and \mathcal{P}^{\max} defined in [16, Lemmas 5.1 and 6.1]. To solve the MILPs, we use the solvers GUROBI [18] (version 9.0.1) and CBC³ (version 2.9.0). More data regarding this experiment can be found in the extended version [25].

The evaluation shows that for instances which have a favourable directed path decomposition (provided it can be easily computed) it may pay off to use Algorithm 2. While a result is not returned within 900 seconds using the MILP-based approaches for the larger threshold and instances with $N \geq 30$, our implementation returns in less than 100 seconds for instances up to $N = 88$. Still, even for these instances it has an exponential increase in runtime and doesn’t scale to very large state spaces.

6 Conclusion

This paper considered the problem of computing minimal witnessing subsystems for probabilistic systems whose underlying graph has low tree width. The main result is that the corresponding decision problem remains NP-hard for systems with bounded directed tree partition width. To prove this, the *matrix-pair*

²<http://www.qhull.org/>

³<https://github.com/coin-or/Cbc>

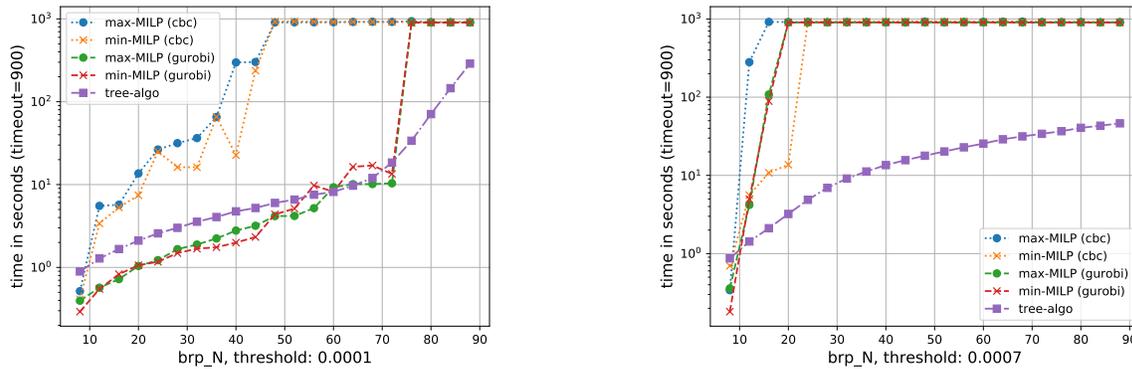


Figure 7: The computation times of the MILP approaches and Algorithm 2 for two different thresholds.

Table 1: Instances of the algorithm.

Model	value function	computing the value (Line 8)	domination relation (Line 11)
DTMC	$\Pr_{S'}(\diamond \text{Goal})$	linear equations	standard
MDP	$\Pr_{S'}^{\max}(\diamond \text{Goal})$	linear program	standard
	$\Pr_{S'}^{\min}(\diamond \text{Goal})$		strong

chain problem is introduced and shown to be NP-hard for fixed-dimension nonnegative matrices. In a second step, this problem is reduced to the witness problem. Finally, an algorithm is described which takes as input a directed tree partition of the system and computes a minimal witnessing subsystem, aiming to utilize the special structure of the system. A preliminary experimental analysis shows that it outperforms existing approaches for a standard benchmark that allows a good tree partition.

A direction for future work, which would help enabling practical usage of the algorithm described in this paper, is to study how to compute good directed tree partitions, or to characterize systems which allow for natural ones. Another direction would be to find algorithms which work on standard tree decompositions of the system, as approximation techniques exist to compute them. Furthermore, it would be interesting to consider heuristic or approximate approaches that utilize the knowledge of a given directed tree partition. For example, the algorithm described in this paper could be adapted to only store a fixed number of partial subsystems for each block.

References

- [1] Ábrahám, E., Becker, B., Dehnert, C., Jansen, N., Katoen, J.P., Wimmer, R.: Counterexample Generation for Discrete-Time Markov Models: An Introductory Survey. In: Formal Methods for Executable Software Models: 14th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, (SFM 2014), pp. 65–121. Lecture Notes in Computer Science, Springer (2014). https://doi.org/10.1007/978-3-319-07317-0_3
- [2] Andriushchenko, R., Češka, M., Junges, S., Katoen, J.P.: Inductive Synthesis for Probabilistic Programs Reaches New Horizons. In: Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2021), pp. 191–209. Lecture Notes in Computer Science, Springer (2021). https://doi.org/10.1007/978-3-030-72016-2_11

- [3] Asadi, A., Chatterjee, K., Goharshady, A.K., Mohammadi, K., Pavlogiannis, A.: Faster Algorithms for Quantitative Analysis of MCs and MDPs with Small Treewidth. In: Automated Technology for Verification and Analysis (ATVA 2020). pp. 253–270. Lecture Notes in Computer Science, Springer (2020). https://doi.org/10.1007/978-3-030-59152-6_14
- [4] Baier, C., Katoen, J.P.: Principles of Model Checking (Representation and Mind Series). The MIT Press (2008)
- [5] Barber, C.B., Dobkin, D.P., Huhdanpaa, H.: The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software* **22**(4), 469–483 (1996). <https://doi.org/10.1145/235815.235821>
- [6] Bodlaender, H.L.: Treewidth: Algorithmic techniques and results. In: Mathematical Foundations of Computer Science (MFCS 1997). pp. 19–36. Lecture Notes in Computer Science, Springer (1997). <https://doi.org/10.1007/BFb0029946>
- [7] Canny, J., Donald, B., Ressler, E.K.: A rational rotation method for robust geometric algorithms. In: Proceedings of the Eighth Annual Symposium on Computational Geometry (SCG 1992). pp. 251–260. ACM, New York, NY, USA (Jul 1992). <https://doi.org/10.1145/142675.142726>
- [8] Češka, M., Hensel, C., Junges, S., Katoen, J.P.: Counterexample-Driven Synthesis for Probabilistic Program Sketches. In: Formal Methods – The Next 30 Years. pp. 101–120. Lecture Notes in Computer Science, Springer, Cham (2019). https://doi.org/10.1007/978-3-030-30942-8_8
- [9] Chan, T.M.: Optimal output-sensitive convex hull algorithms in two and three dimensions. *Discret. Comput. Geom.* **16**(4), 361–368 (1996). <https://doi.org/10.1007/BF02712873>, <https://doi.org/10.1007/BF02712873>
- [10] Chatterjee, K., Ibsen-Jensen, R., Pavlogiannis, A.: Faster algorithms for quantitative verification in bounded treewidth graphs. *Formal Methods in System Design* (Apr 2021). <https://doi.org/10.1007/s10703-021-00373-5>
- [11] Chatterjee, K., Łački, J.: Faster Algorithms for Markov Decision Processes with Low Treewidth. In: Computer Aided Verification (CAV 2013). pp. 543–558. Lecture Notes in Computer Science, Springer (2013). https://doi.org/10.1007/978-3-642-39799-8_36
- [12] Chazelle, B.: An optimal convex hull algorithm in any fixed dimension. *Discrete & Computational Geometry* **10**(4), 377–409 (Dec 1993). <https://doi.org/10.1007/BF02573985>
- [13] Clarke, E., Veith, H.: Counterexamples Revisited: Principles, Algorithms, Applications. In: Dershowitz, N. (ed.) Verification: Theory and Practice: Essays Dedicated to Zohar Manna on the Occasion of His 64th Birthday, pp. 208–224. Lecture Notes in Computer Science, Springer (2003). https://doi.org/10.1007/978-3-540-39910-0_9
- [14] Dujmović, V., Fellows, M.R., Kitching, M., Liotta, G., McCartin, C., Nishimura, N., Ragde, P., Rosamond, F., Whitesides, S., Wood, D.R.: On the Parameterized Complexity of Layered Graph Drawing. *Algorithmica* **52**(2), 267–292 (Oct 2008). <https://doi.org/10.1007/s00453-007-9151-1>
- [15] Feige, U., Yahalom, O.: On the complexity of finding balanced oneway cuts. *Information Processing Letters* **87**(1), 1–5 (Jul 2003). [https://doi.org/10.1016/S0020-0190\(03\)00251-5](https://doi.org/10.1016/S0020-0190(03)00251-5)
- [16] Funke, F., Jantsch, S., Baier, C.: Farkas Certificates and Minimal Witnesses for Probabilistic Reachability Constraints. In: Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2020). pp. 324–345. Lecture Notes in Computer Science, Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45190-5_18
- [17] Graham, R.L.: An efficient algorithm for determining the convex hull of a finite planar set. *Information Processing Letters* **1**(4), 132–133 (1972). [https://doi.org/10.1016/0020-0190\(72\)90045-2](https://doi.org/10.1016/0020-0190(72)90045-2)
- [18] Gurobi Optimization, LLC: Gurobi Optimizer Reference Manual (2021), <https://www.gurobi.com>
- [19] Gustedt, J., Mæhle, O.A., Telle, J.A.: The Treewidth of Java Programs. In: Algorithm Engineering and Experiments. pp. 86–97. Lecture Notes in Computer Science, Springer (2002). https://doi.org/10.1007/3-540-45643-0_7

- [20] Halin, R.: Tree-partitions of infinite graphs. *Discrete Mathematics* **97**(1), 203–217 (Dec 1991). [https://doi.org/10.1016/0012-365X\(91\)90436-6](https://doi.org/10.1016/0012-365X(91)90436-6)
- [21] Han, T., Katoen, J.P.: Counterexamples in Probabilistic Model Checking. In: *Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2007)*. pp. 72–86. *Lecture Notes in Computer Science*, Springer (2007). https://doi.org/10.1007/978-3-540-71209-1_8
- [22] Hermanns, H., Wachter, B., Zhang, L.: Probabilistic CEGAR. In: *Computer Aided Verification (CAV 2008)*. pp. 162–175. *Lecture Notes in Computer Science*, Springer (2008). https://doi.org/10.1007/978-3-540-70545-1_16
- [23] Jantsch, S.: Witnessing subsystems for probabilistic systems with low treewidth - Supplementary material (2021). <https://doi.org/10.6084/m9.figshare.14915841.v1>
- [24] Jantsch, S., Harder, H., Funke, F., Baier, C.: SWITSS: Computing Small Witnessing Subsystems. In: *Formal Methods in Computer-Aided Design (FMCAD 2020)*. vol. 1, pp. 236–244. TU Wien Academic Press (2020). https://doi.org/10.34727/2020/isbn.978-3-85448-042-6_31
- [25] Jantsch, S., Piribauer, J., Baier, C.: Witnessing subsystems for probabilistic systems with low tree width (2021), arXiv preprint, arXiv:2108.08070.
- [26] Johnson, T., Robertson, N., Seymour, P.D., Thomas, R.: Directed Tree-Width. *Journal of Combinatorial Theory, Series B* **82**(1), 138–154 (May 2001). <https://doi.org/10.1006/jctb.2000.2031>
- [27] Karp, R.M.: Reducibility among combinatorial problems. In: *Complexity of Computer Computations: Proceedings of a symposium on the Complexity of Computer Computations, 1972*. pp. 85–103. Springer US, Boston, MA (1972)
- [28] Kornai, A., Tuza, Z.: Narrowness, pathwidth, and their application in natural language processing. *Discrete Applied Mathematics* **36**(1), 87–92 (Mar 1992). [https://doi.org/10.1016/0166-218X\(92\)90208-R](https://doi.org/10.1016/0166-218X(92)90208-R)
- [29] Kwiatkowska, M., Norman, G., Parker, D.: The PRISM Benchmark Suite. In: *Quantitative Evaluation of Systems (QEST 2012)*. pp. 203–204 (Sep 2012). <https://doi.org/10.1109/QEST.2012.14>
- [30] Reed, B.A.: Introducing Directed Tree Width. *Electronic Notes in Discrete Mathematics* **3**, 222–229 (May 1999). [https://doi.org/10.1016/S1571-0653\(05\)80061-7](https://doi.org/10.1016/S1571-0653(05)80061-7)
- [31] Robertson, N., Seymour, P.D.: Graph minors. I. Excluding a forest. *Journal of Combinatorial Theory, Series B* **35**(1), 39–61 (Aug 1983). [https://doi.org/10.1016/0095-8956\(83\)90079-5](https://doi.org/10.1016/0095-8956(83)90079-5)
- [32] Safari, M.A.: D-Width: A More Natural Measure for Directed Tree Width. In: Jędrzejowicz, J., Szepietowski, A. (eds.) *Mathematical Foundations of Computer Science 2005*. pp. 745–756. *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg (2005). https://doi.org/10.1007/11549345_64
- [33] Seese, D.: Tree-partite graphs and the complexity of algorithms. In: Budach, L. (ed.) *Fundamentals of Computation Theory*. pp. 412–421. *Lecture Notes in Computer Science*, Springer (1985). <https://doi.org/10.1007/BFb0028825>
- [34] Thorup, M.: All Structured Programs Have Small Tree Width and Good Register Allocation. *Information and Computation* **142**(2), 159–181 (May 1998). <https://doi.org/10.1006/inco.1997.2697>
- [35] Wimmer, R., Jansen, N., Abraham, E., Becker, B., Katoen, J.P.: Minimal Critical Subsystems for Discrete-Time Markov Models. In: *Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2012)*. pp. 299–314. *Lecture Notes in Computer Science*, Springer (2012). https://doi.org/10.1007/978-3-642-28756-5_21
- [36] Wimmer, R., Jansen, N., Abraham, E., Katoen, J.P., Becker, B.: Minimal counterexamples for linear-time probabilistic verification. *Theoretical Computer Science* **549**, 61–100 (Sep 2014). <https://doi.org/10.1016/j.tcs.2014.06.020>
- [37] Wood, D.R.: On tree-partition-width. *European Journal of Combinatorics* **30**(5), 1245–1253 (Jul 2009). <https://doi.org/10.1016/j.ejc.2008.11.010>