

Symblicit algorithms for optimal strategy synthesis in monotonic Markov decision processes*

Aaron Bohy

Université de Mons
Belgium

Véronique Bruyère

Jean-François Raskin

Université Libre de Bruxelles
Belgium

When treating Markov decision processes (MDPs) with large state spaces, using explicit representations quickly becomes unfeasible. Lately, Wimmer et al. have proposed a so-called symblicit algorithm for the synthesis of optimal strategies in MDPs, in the quantitative setting of expected mean-payoff. This algorithm, based on the strategy iteration algorithm of Howard and Veinott, efficiently combines symbolic and explicit data structures, and uses binary decision diagrams as symbolic representation. The aim of this paper is to show that the new data structure of pseudo-antichains (an extension of antichains) provides another interesting alternative, especially for the class of monotonic MDPs. We design efficient pseudo-antichain based symblicit algorithms (with open source implementations) for two quantitative settings: the expected mean-payoff and the stochastic shortest path. For two practical applications coming from automated planning and LTL synthesis, we report promising experimental results w.r.t. both the run time and the memory consumption.

1 Introduction

Markov decision processes [34, 2] (MDPs) are rich models that exhibit both nondeterministic choices and stochastic transitions. Model-checking and synthesis algorithms for MDPs exist for logical properties expressible in the logic PCTL [23], a stochastic extension of CTL [15], and are implemented in tools like PRISM [29], MODEST [24], MRMC [27]. . . There also exist algorithms for *quantitative properties* such as the long-run average reward (mean-payoff) or the stochastic shortest path, that have been implemented in tools like QUASY [14] and PRISM [19].

There are two main families of algorithms for MDPs. First, *value iteration* algorithms assign values to states of the MDPs and refine locally those values by successive approximations. If a fixpoint is reached, the value at a state s represents a probability or an expectation that can be achieved by an optimal strategy that resolves the choices present in the MDP starting from s . This value can be, for example, the maximal probability to reach a set of goal states. Second, *strategy iteration* algorithms start from an arbitrary strategy and iteratively improve the current strategy by local changes up to the convergence to an optimal strategy. Both methods have their advantages and disadvantages. Value iteration algorithms usually lead to easy and efficient implementations, but in general the fixpoint is not guaranteed to be reached in a finite number of iterations, and so only approximations are computed. On the other hand, strategy iteration algorithms have better theoretical properties as convergence towards an optimal strategy in a finite number of steps is usually ensured, but they often require to solve systems of linear equations, and so they are more difficult to implement efficiently.

When considering large MDPs obtained from high level descriptions or as the product of several components, explicit methods often exhaust available memory and are thus impractical. This is the manifestation of the well-known *state explosion problem*. In non-probabilistic systems, symbolic data struc-

*This work has been partly supported by ERC Starting Grant (279499: inVEST), ARC project (number AUWB-2010-10/15-UMONS-3) and European project Cassting (FP7-ICT-601148). An extended version of this paper is available at [10].

tures like binary decision diagrams (BDDs) have been investigated [13] to mitigate this phenomenon. For probabilistic systems, multi-terminal BDDs (MTBDDs) are useful but they are usually limited to systems with around 10^{10} or 10^{11} states only [33]. Also, as mentioned above, some algorithms for MDPs rely on solving linear systems, and there is no easy use of BDD like structures for implementing such algorithms.

Recently, Wimmer et al. [38] have proposed a method that *mixes* symbolic and explicit representations to efficiently implement the Howard and Veinott strategy iteration algorithm [25, 36] to synthesize optimal strategies for mean-payoff objectives in MDPs. Their solution is as follows. First, the MDP is represented and handled symbolically using MTBDDs. Second, a strategy is fixed symbolically and the MDP is transformed into a Markov chain (MC). To analyze this MC, a linear system needs to be constructed from its state space. As this state space is potentially huge, the MC is first reduced by *lumping* [28, 12] (bisimulation reduction), and then a (hopefully) compact linear system can be constructed and solved. Solutions to this linear system allow to show that the current strategy is optimal, or to obtain sufficient information to improve it. A new iteration is then started. The main difference between this method and the other methods proposed in the literature is its *hybrid nature*: it is symbolic for handling the MDP and for computing the lumping, and it is explicit for the analysis of the reduced MC. This is why the authors of [38] have coined their approach *symblicit*.

Contributions. In this paper, we build on the symblicit approach described above. Our contributions are threefold. First, we show that the symblicit approach and strategy iteration can also be efficiently applied to the *stochastic shortest path* problem. We start from an algorithm proposed by Bertsekas and Tsitsiklis [4] with a preliminary step of de Alfaro [1], and we show how to cast it in the symblicit approach. Second, we show that alternative data structures can be more efficient than BDDs or MTBDDs for implementing a symblicit approach, both for mean-payoff and stochastic shortest path objectives. In particular, we consider a natural class of MDPs with *monotonic properties* on which our alternative data structure is more efficient. For such MDPs, as for subset constructions in automata theory [39, 17], antichain based data structures usually behave better than BDDs. The application of antichains to monotonic MDPs requires nontrivial extensions: for instance, to handle the lumping step, we need to generalize existing antichain based data structures in order to be closed under negation. To this end, we introduce a new data structure called *pseudo-antichain*. Third, we have implemented our algorithms and we show that they are more efficient than existing solutions on natural examples of monotonic MDPs. We show that monotonic MDPs naturally arise in probabilistic planning [6] and when optimizing controllers synthesized from LTL specifications with mean-payoff objectives [9].

Structure of the paper. In Section 2, we recall useful definitions and we introduce the notion of monotonic MDP. In Section 3, we recall strategy iteration algorithms for mean-payoff and shortest path objectives, and we present the symblicit version of those algorithms. We introduce the notion of pseudo-antichains in Section 4, and we describe our pseudo-antichain based symblicit algorithms in Section 5. In Section 6, we propose two applications of the symblicit algorithms and give experimental results. Finally in Section 7, we summarize our results.

2 Background and studied problems

In this section, we recall useful definitions and we introduce the notion of monotonic Markov decision process. We also state the problems that we study.

Stochastic models. We denote by $\text{Dom}(f)$ the domain of definition of a function f , and by $\mathcal{F}_{\text{tot}}(A, B)$ the set of total functions from A to B . A *probability distribution* over a finite set A is a total function $\pi : A \rightarrow [0, 1]$ such that $\sum_{a \in A} \pi(a) = 1$. We denote by $\mathcal{D}(A)$ the set of probability distributions over A .

A *discrete-time Markov chain (MC)* is a tuple (S, \mathbf{P}) where S is a finite set of states and $\mathbf{P} : S \rightarrow \mathcal{D}(S)$ is a total stochastic transition function. In the sequel, \mathbf{P} is sometimes seen as a matrix, and for all $s, s' \in S$, we write $\mathbf{P}(s, s')$ for $\mathbf{P}(s)(s')$. A *path* is an infinite sequence of states $\rho = s_0 s_1 s_2 \dots$ such that $\mathbf{P}(s_i, s_{i+1}) > 0$ for all $i \geq 0$. Finite paths are defined similarly, and \mathbf{P} is naturally extended to finite paths.

A *Markov decision process (MDP)* is a tuple (S, Σ, \mathbf{P}) where S is a finite set of states, Σ is a finite set of actions, and $\mathbf{P} : S \times \Sigma \rightarrow \mathcal{D}(S)$ is a partial stochastic transition function. We often write $\mathbf{P}(s, \sigma, s')$ for $\mathbf{P}(s, \sigma)(s')$. For each $s \in S$, we denote by $\Sigma_s \subseteq \Sigma$ the set of enabled actions in s , where an action $\sigma \in \Sigma$ is *enabled* in s if $(s, \sigma) \in \text{Dom}(\mathbf{P})$. We require $\forall s \in S, \Sigma_s \neq \emptyset$, that is, the MDP is Σ -*non-blocking*. For each $\sigma \in \Sigma$, S_σ denotes the set of states in which σ is enabled.

Let $M = (S, \Sigma, \mathbf{P})$ be an MDP. A *memoryless strategy* is a total function $\lambda : S \rightarrow \Sigma$ mapping each state s to an enabled action $\sigma \in \Sigma_s$. We denote by Λ the set of all memoryless strategies. A memoryless strategy λ induces an MC $M_\lambda = (S, \mathbf{P}_\lambda)$ such that for all $s, s' \in S$, $\mathbf{P}_\lambda(s, s') = \mathbf{P}(s, \lambda(s), s')$.

Costs and value functions. In addition to an MDP $M = (S, \Sigma, \mathbf{P})$, we consider a partial *cost function* $\mathbf{C} : S \times \Sigma \rightarrow \mathbb{R}$ with $\text{Dom}(\mathbf{C}) = \text{Dom}(\mathbf{P})$ that associates a cost with $s \in S$ and $\sigma \in \Sigma_s$. A memoryless strategy λ assigns a total cost function $\mathbf{C}_\lambda : S \rightarrow \mathbb{R}$ to the induced MC M_λ , such that $\mathbf{C}_\lambda(s) = \mathbf{C}(s, \lambda(s))$. Given a path $\rho = s_0 s_1 s_2 \dots$ in M_λ , the *mean-payoff* of ρ is $\text{MP}(\rho) = \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} \mathbf{C}_\lambda(s_i)$. Given a subset $G \subseteq S$ of *goal* states and a finite path ρ reaching a state of G , the *truncated sum up to G* of ρ is $\text{TS}_G(\rho) = \sum_{i=0}^{n-1} \mathbf{C}_\lambda(s_i)$ where n is the first index such that $s_n \in G$.

Given an MDP with a cost function \mathbf{C} , and a memoryless strategy λ , we consider two classical value functions of λ . For all states $s \in S$, the *expected mean-payoff* of λ is $\mathbb{E}_\lambda^{\text{MP}}(s) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} \mathbf{P}_\lambda^i \mathbf{C}_\lambda(s)$. Given a subset $G \subseteq S$, and assuming that λ reaches G from state s with probability 1, the *expected truncated sum up to G* of λ is $\mathbb{E}_\lambda^{\text{TS}_G}(s) = \sum_\rho \mathbf{P}_\lambda(\rho) \text{TS}_G(\rho)$ where the sum is over all finite paths $\rho = s_0 s_1 \dots s_n$ such that $s_0 = s$, $s_n \in G$, and $s_0, \dots, s_{n-1} \notin G$. Let λ^* be a memoryless strategy. Given a value function $\mathbb{E}_\lambda \in \{\mathbb{E}_\lambda^{\text{MP}}, \mathbb{E}_\lambda^{\text{TS}_G}\}$, we say that λ^* is *optimal* if $\mathbb{E}_{\lambda^*}(s) = \inf_{\lambda \in \Lambda} \mathbb{E}_\lambda(s)$ for all $s \in S$, and \mathbb{E}_{λ^*} is called the *optimal* value function.¹ Note that we might have considered other classes of strategies, but for these value functions, there always exists a memoryless strategy minimizing the expected value of all states [2, 34].

Studied problems. In this paper, we study algorithms for solving MDPs for two quantitative settings: the expected mean-payoff and the stochastic shortest path. Let M be an MDP with a cost function \mathbf{C} . (i) The *expected mean-payoff (EMP) problem* is to synthesize an optimal strategy for the expected mean-payoff value function. (ii) When \mathbf{C} is restricted to *strictly positive* values in $\mathbb{R}_{>0}$, and a subset $G \subseteq S$ of goal states is given, the *stochastic shortest path (SSP) problem* is to synthesize an optimal strategy for the expected truncated sum value function, among the set of strategies that reach G with probability 1, provided such strategies exist. For all $s \in S$, we denote by Λ_s^P the set of *proper strategies* that lead from s to G with probability 1. Solving the SSP problem consists of two steps. The first step is to determine the set $S^P = \{s \in S \mid \Lambda_s^P \neq \emptyset\}$ of *proper states* that have at least one proper strategy. The second step consists in synthesizing an optimal strategy λ^* such that $\mathbb{E}_{\lambda^*}^{\text{TS}_G}(s) = \inf_{\lambda \in \Lambda_s^P} \mathbb{E}_\lambda^{\text{TS}_G}(s)$ for all $s \in S^P$. It is known that both problems can be solved in polynomial time via linear programming, with memoryless optimal strategies [34, 20, 4, 5].

¹An alternative objective is to maximize the value function, in which case λ^* is optimal if $\mathbb{E}_{\lambda^*}(s) = \sup_{\lambda \in \Lambda} \mathbb{E}_\lambda(s)$, $\forall s \in S$.

Monotonic MDPs. In this paper, we need a slightly different, but equivalent, definition of MDPs based on the next idea. Instead of a transition function $\mathbf{P} : S \times \Sigma \rightarrow \mathcal{D}(S)$, we rather use two functions in a way to separate the probabilities from the successors as indicated in Figure 1. In this new definition, an MDP is a tuple $M = (S, \Sigma, T, \mathbf{E}, \mathbf{D})$ where S is a finite set of states, Σ and T are two disjoint finite sets of actions, $\mathbf{E} : S \times \Sigma \rightarrow \mathcal{F}_{\text{tot}}(T, S)$ is a partial successor function, and $\mathbf{D} : S \times \Sigma \rightarrow \mathcal{D}(T)$ is a partial stochastic function such that $\text{Dom}(\mathbf{E}) = \text{Dom}(\mathbf{D})$. In this context, the notion of MC M_λ induced by a strategy λ is defined as $(S, T, \mathbf{E}_\lambda, \mathbf{D}_\lambda)$ with functions $\mathbf{E}_\lambda, \mathbf{D}_\lambda$ naturally defined. In the sequel, depending on the context, we will use both definitions.




Figure 1: Illustration of the new definition of MDPs for a state $s_0 \in S$ and an action $\sigma \in \Sigma_{s_0}$.

Let S be a finite set equipped with a partial order \preceq such that (S, \preceq) is a *semilattice*, i.e. for all $s, s' \in S$, their greatest lower bound $s \sqcap s'$ always exists. A set $L \subseteq S$ is *closed* for \preceq if for all $s \in L$ and all $s' \preceq s$, we have $s' \in L$.

A *monotonic MDP* is an MDP $M_\preceq = (S, \Sigma, T, \mathbf{E}, \mathbf{D})$ such that:

1. S is equipped with a partial order \preceq such that (S, \preceq) is a semilattice, and
2. \preceq is *compatible* with \mathbf{E} , i.e. for all $s, s' \in S$, if $s \preceq s'$, then for all $\sigma \in \Sigma$, $\tau \in T$, for all $t' \in S$ such that $\mathbf{E}(s', \sigma)(\tau) = t'$, there exists $t \in S$ such that $\mathbf{E}(s, \sigma)(\tau) = t$ and $t \preceq t'$.

Note that since (S, \preceq) is a semilattice, S is trivially closed for \preceq . With this definition, we have the next important properties: (1) for all $s, s' \in S$, if $s \preceq s'$ then $\Sigma_{s'} \subseteq \Sigma_s$, and (2) for all $\sigma \in \Sigma$, S_σ is closed.

Remark 1. In this definition, by monotonic MDPs, we mean MDPs that are built on state spaces *already equipped with a natural partial order*. For instance, this is the case for the two classes of MDPs studied in Section 6. The same kind of approach has already been proposed in [22].

Note that all MDPs can be seen monotonic. Indeed, let $(S, \Sigma, T, \mathbf{E}, \mathbf{D})$ be a given MDP and let \preceq be a partial order such that all states in S are pairwise incomparable with respect to \preceq . By adding a new state t such that $t \preceq s$, for all $s \in S$, and such that t is an isolated state with a self-loop, we have that $(S \cup \{t\}, \preceq)$ is a semilattice and \preceq is compatible with \mathbf{E} . However, such a partial order would not lead to efficient algorithms in the sense studied in this paper.

3 Strategy iteration algorithms and symblicit approach

In this section, we first present strategy iteration algorithms for synthesizing optimal strategies for both the SSP and EMP problems. We then present a symblicit version of those algorithms that mixes symbolic and explicit data representations. Our presentation is inspired from the one given in [38] for the EMP problem.

Algorithm for the SSP problem. Let $M = (S, \Sigma, \mathbf{P})$ be an MDP, $\mathbf{C} : S \times \Sigma \rightarrow \mathbb{R}_{>0}$ be a strictly positive cost function, and $G \subseteq S$ be a set of goal states. As explained in Section 2, the set S^P of proper states is

Algorithm 1 SSP_SI(MDP M , Strictly positive cost function \mathbf{C} , Goal states G)

```

1:  $n := 0, \lambda_n := \text{INITIALPROPERSTRAT}(M, G)$ 
2: repeat
3: Obtain  $v_n$  by solving
      
$$\mathbf{C}_{\lambda_n} + (\mathbf{P}_{\lambda_n} - \mathbf{I})v_n = 0$$

4:  $\widehat{\Sigma}_s := \arg \min_{\sigma \in \Sigma_s} (\mathbf{C}(s, \sigma) + \sum_{s' \in S} \mathbf{P}(s, \sigma, s') \cdot v_n(s')), \forall s \in S$ 
5: Choose  $\lambda_{n+1}$  s.t.  $\lambda_{n+1}(s) \in \widehat{\Sigma}_s, \forall s \in S$ ,
      setting  $\lambda_{n+1}(s) := \lambda_n(s)$  if possible.
6:  $n := n + 1$ 
7: until  $\lambda_n = \lambda_{n-1}$ 
8: return  $(\lambda_{n-1}, v_{n-1})$ 

```

Algorithm 2 SYMBLICIT(MDP \mathcal{M} , [Strictly positive] cost function \mathcal{C} [, Goal states \mathcal{G}])

```

1:  $n := 0, \lambda_n := \text{INITIALSTRAT}(\mathcal{M}[, \mathcal{G}])$ 
2: repeat
3:  $(\mathcal{M}_{\lambda_n}, \mathcal{C}_{\lambda_n}) := \text{INDUCEDMC\&COST}(\mathcal{M}, \mathcal{C}, \lambda_n)$ 
4:  $(\mathcal{M}'_{\lambda_n}, \mathcal{C}'_{\lambda_n}) := \text{LUMP}(\mathcal{M}_{\lambda_n}, \mathcal{C}_{\lambda_n})$ 
5:  $(M'_{\lambda_n}, C'_{\lambda_n}) := \text{EXPLICIT}(\mathcal{M}'_{\lambda_n}, \mathcal{C}'_{\lambda_n})$ 
6:  $x_n := \text{SOLVELINEARSYSTEM}(M'_{\lambda_n}, C'_{\lambda_n})$ 
7:  $\mathcal{X}_n := \text{SYMBOLIC}(x_n)$ 
8:  $\lambda_{n+1} := \text{IMPROVESTRAT}(\mathcal{M}, \lambda_n, \mathcal{X}_n)$ 
9:  $n := n + 1$ 
10: until  $\lambda_n = \lambda_{n-1}$ 
11: return  $(\lambda_{n-1}, \mathcal{X}_{n-1})$ 

```

computed in a preliminary step (see [1] for a quadratic algorithm). The strategy iteration algorithm [25, 4] (see Algorithm 1) is then applied under the typical assumption that all cycles in the underlying graph of M have strictly positive cost [4]. This assumption holds in our case by definition of the cost function \mathbf{C} . Algorithm 1 starts with an arbitrary proper strategy λ_0 , that can be easily computed with the algorithm of [1], and improves it until an optimal strategy is found. The expected truncated sum v_n up to G of the current strategy λ_n is computed by solving the system of linear equations in line 3, and used to improve the strategy (if possible) at each state. Note that the strategy λ_n is improved at a state s to an action $\sigma \in \Sigma_s$ only if the new expected truncated sum is strictly smaller than the expected truncated sum of the action $\lambda_n(s)$, i.e. only if $\lambda_n(s) \notin \arg \min_{\sigma \in \Sigma_s} (\mathbf{C}(s, \sigma) + \sum_{s' \in S} \mathbf{P}(s, \sigma, s') \cdot v_n(s'))$. If no improvement is possible for any state, an optimal strategy is found and the algorithm terminates in line 7. Otherwise, it restarts by solving the new equation system, tries to improve the strategy using the new values computed, and so on.

Algorithm for the EMP problem. The strategy iteration algorithm for the EMP problem works similarly (see [36, 34] for more details). The algorithm starts with an arbitrary strategy λ_0 . Solving the related linear system leads to two values: the gain value g_n and bias value b_n of strategy λ_n . The gain corresponds to the expected mean-payoff, while the bias can be interpreted as the expected total difference between the cost and the expected mean-payoff. The computed gain value is used to locally improve the strategy. If such an improvement is not possible for any state, the bias value is used to locally improve the strategy such that only actions that also optimize the gain are considered.

Bisimulation lumping. When treating MDPs and induced MCs with large state spaces, using explicit representations quickly becomes unfeasible for the algorithms presented above. Given an MC (S, \mathbf{P}) and a cost function $\mathbf{C} : S \rightarrow \mathbb{R}$, the *bisimulation lumping* technique [28, 30, 12] consists in gathering certain states of S which behave equivalently according to the class of properties under consideration. Let \sim be an equivalence relation on S and S_{\sim} be the induced *partition*. We call *block* of S_{\sim} any equivalence class of \sim . We say that \sim is a *bisimulation* if for all $s, t \in S$ such that $s \sim t$, we have $\mathbf{C}(s) = \mathbf{C}(t)$ and $\mathbf{P}(s, C) = \mathbf{P}(t, C)$ for all block $C \in S_{\sim}$, where $\mathbf{P}(s, C) = \sum_{s' \in C} \mathbf{P}(s, s')$. The *bisimulation quotient* is the MC $(S_{\sim}, \mathbf{P}_{\sim})$ such that $\mathbf{P}_{\sim}(C, C') = \mathbf{P}(s, C')$, where $s \in C$ and $C, C' \in S_{\sim}$. The cost function $\mathbf{C}_{\sim} : S_{\sim} \rightarrow \mathbb{R}$ is transferred to the quotient such that $\mathbf{C}_{\sim}(C) = \mathbf{C}(s)$, where $s \in C$ and $C \in S_{\sim}$. This quotient is a minimized model equivalent to the original, since it keeps the expected truncated sum and expected mean-payoff as in the original model [3]. Usually, we are interested in computing the unique *largest* bisimulation, denoted \sim_L , which leads to the smallest bisimulation quotient $(S_{\sim_L}, \mathbf{P}_{\sim_L})$ (see [16] for an algorithm).

Symblicit algorithm. The symblicit algorithm is described in Algorithm SYMBLICIT for both the SSP and EMP (see Algorithm 2). It combines symbolic² and explicit representations of the manipulated data as follows. The MDP \mathcal{M} , the cost function \mathcal{C} , the strategies λ_n , the induced MCs \mathcal{M}_{λ_n} with cost functions \mathcal{C}_{λ_n} , and the set \mathcal{G} of goal states for the SSP, are symbolically represented. The lumping algorithm is applied on symbolic MCs and produces a symbolic representation of the bisimulation quotient \mathcal{M}'_{λ_n} with cost function \mathcal{C}'_{λ_n} (line 4). The computed quotient is converted to a sparse matrix representation (line 5). As it is in general much smaller than the original model, there is no memory issue by storing it explicitly, and the linear system can be solved. The computed value functions x_n (corresponding to v_n for the SPP, and g_n and b_n for the EMP) are then converted into symbolic representations \mathcal{X}_n , and transferred back to the original MDP (line 7). Finally, the update of the strategy is performed symbolically.

4 Pseudo-antichains

We want to develop a symblicit algorithm for solving the SSP and EMP problems for monotonic MDPs. To this end, we here introduce a new data structure extended from antichains, called pseudo-antichains.

Closed sets and antichains. Let (S, \preceq) be a semilattice. The *closure* $\downarrow L$ of a set $L \subseteq S$ is the set $\downarrow L = \{s' \in S \mid \exists s \in L \cdot s' \preceq s\}$. A set α is an *antichain* if all its elements are pairwise incomparable with respect to \preceq . For $L \subseteq S$, we denote by $\lceil L \rceil$ the antichain of its maximal elements. If L is closed, then $\downarrow \lceil L \rceil = L$, and $\lceil L \rceil$ is called the *canonical representation* of L . The interest of antichains is that they are *compact representations of closed sets*. Antichain based algorithms have proved worthy for solving classical problems in game theory, but also in logic and automata theory (e.g. [39, 17, 11, 18]). We have the following classical properties on antichains (see for instance [21]):

Proposition 1. *Let $\alpha_1, \alpha_2 \subseteq S$ be two antichains and $s \in S$. Then:*

- $s \in \downarrow \alpha_1$ iff $\exists a \in \alpha_1 \cdot s \preceq a$
- $\downarrow \alpha_1 \cup \downarrow \alpha_2 = \downarrow [\alpha_1 \cup \alpha_2]$
- $\downarrow \alpha_1 \cap \downarrow \alpha_2 = \downarrow [\alpha_1 \sqcap \alpha_2]$, where $\alpha_1 \sqcap \alpha_2 \stackrel{\text{def}}{=} \{a_1 \sqcap a_2 \mid a_1 \in \alpha_1, a_2 \in \alpha_2\}$
- $\downarrow \alpha_1 \subseteq \downarrow \alpha_2$ iff $\forall a_1 \in \alpha_1 \cdot \exists a_2 \in \alpha_2 \cdot a_1 \preceq a_2$

For convenience, when α_1 and α_2 are antichains, we use notation $\alpha_1 \dot{\cup} \alpha_2$ (resp. $\alpha_1 \dot{\cap} \alpha_2$) for the antichain $\lceil \downarrow \alpha_1 \cup \downarrow \alpha_2 \rceil$ (resp. $\lceil \downarrow \alpha_1 \cap \downarrow \alpha_2 \rceil$). Let $L_1, L_2 \subseteq S$ be two closed sets. Unlike the union or intersection, the difference $L_1 \setminus L_2$ is not necessarily a closed set. There is thus a need for a new structure that “represents” $L_1 \setminus L_2$ in a compact way.

Pseudo-elements and pseudo-closures. A *pseudo-element* is a pair (x, α) where $x \in S$ and $\alpha \subseteq S$ is an antichain such that $x \notin \downarrow \alpha$. The *pseudo-closure* of a pseudo-element (x, α) , denoted by $\uparrow(x, \alpha)$, is the set $\uparrow(x, \alpha) = \{s \in S \mid s \preceq x \text{ and } s \notin \downarrow \alpha\} = \downarrow \{x\} \setminus \downarrow \alpha$. Notice that $\uparrow(x, \alpha)$ is non empty since $x \notin \downarrow \alpha$ by definition of a pseudo-element. The following example illustrates these notions.

Example 1. Let $\mathbb{N}_{\leq 3}^2$ be the set of pairs of natural numbers in $[0, 3]$ and let \preceq be a partial order on $\mathbb{N}_{\leq 3}^2$ such that $(n_1, n'_1) \preceq (n_2, n'_2)$ iff $n_1 \leq n_2$ and $n'_1 \leq n'_2$. Then, $(\mathbb{N}_{\leq 3}^2, \preceq)$ is a complete lattice with least upper bound \sqcup such that $(n_1, n'_1) \sqcup (n_2, n'_2) = (\max(n_1, n_2), \max(n'_1, n'_2))$, and greatest lower bound \sqcap such that $(n_1, n'_1) \sqcap (n_2, n'_2) = (\min(n_1, n_2), \min(n'_1, n'_2))$. With $x = (3, 2)$ and $\alpha = \{(2, 1), (0, 2)\}$, the pseudo-closure of the pseudo-element (x, α) is the set $\uparrow(x, \alpha) = \{(3, 2), (3, 1), (3, 0), (2, 2), (1, 2)\} = \downarrow \{x\} \setminus \downarrow \alpha$ (see Figure 2).




Figure 2: Pseudo-closure of a pseudo-element over $(\mathbb{N}_{\leq 3}^2, \preceq)$.




Figure 3: Inclusion of pseudo-closures of pseudo-elements.

There may exist two pseudo-elements (x, α) and (y, β) such that $\uparrow(x, \alpha) = \uparrow(y, \beta)$. A pseudo-element (x, α) is said to be in *canonical form* if $\forall a \in \alpha \cdot a \preceq x$ (as in Example 1). Notice that for all pseudo-elements (x, α) , there exists a pseudo-element in canonical form (y, β) such that $\uparrow(x, \alpha) = \uparrow(y, \beta)$: it is equal to $(x, \{x\} \dot{\cap} \alpha)$ and called the *canonical representation* of $\uparrow(x, \alpha)$. The next corollary of Proposition 2 shows that the canonical form is unique.

Proposition 2. *Let (x, α) and (y, β) be two pseudo-elements. Then $\uparrow(x, \alpha) \subseteq \uparrow(y, \beta)$ iff $x \preceq y$ and $\forall b \in \beta \cdot b \sqcap x \in \downarrow \alpha$.*

Corollary 1. *Let (x, α) and (y, β) be two pseudo-elements in canonical form. Then $\uparrow(x, \alpha) = \uparrow(y, \beta)$ iff $x = y$ and $\alpha = \beta$.*

The following example illustrates Proposition 2.

Example 2. Let (S, \preceq) be a semilattice and let $(x, \{a\})$ and $(y, \{b_1, b_2\})$, with $x, y, a, b_1, b_2 \in S$, be two pseudo-elements as depicted in Figure 3. The pseudo-closure of $(x, \{a\})$ is depicted in dark gray, whereas the pseudo-closure of $(y, \{b_1, b_2\})$ is depicted in (light and dark) gray. We have $x \preceq y$, $b_1 \sqcap x = b_1 \in \downarrow \{a\}$ and $b_2 \sqcap x = b_2 \in \downarrow \{a\}$. Therefore $\uparrow(x, \{a\}) \subseteq \uparrow(y, \{b_1, b_2\})$.

Pseudo-antichains. A *pseudo-antichain* A is a finite set of pseudo-elements, that is $A = \{(x_i, \alpha_i) \mid i \in I\}$ with I finite. The *pseudo-closure* $\uparrow A$ of A is defined as the set $\uparrow A = \bigcup_{i \in I} \uparrow(x_i, \alpha_i)$. Given $(x_i, \alpha_i), (x_j, \alpha_j) \in A$, we observe that: (1) if $x_i = x_j$, then (x_i, α_i) and (x_j, α_j) can be replaced in A by the pseudo-element $(x_i, \alpha_i \dot{\cap} \alpha_j)$, and (2) if $\uparrow(x_i, \alpha_i) \subseteq \uparrow(x_j, \alpha_j)$, then (x_i, α_i) can be removed from A . Therefore, we say that a pseudo-antichain $A = \{(x_i, \alpha_i) \mid i \in I\}$ is *simplified* if $\forall i \cdot (x_i, \alpha_i)$ is in canonical form, and $\forall i \neq j \cdot x_i \neq x_j$ and $\uparrow(x_i, \alpha_i) \not\subseteq \uparrow(x_j, \alpha_j)$. Notice that two distinct pseudo-antichains A and B can have the same pseudo-closure $\uparrow A = \uparrow B$ even if they are simplified. We thus say that A is a *PA-representation*³ of $\uparrow A$ (without saying that it is a canonical representation), and that $\uparrow A$ is *PA-represented* by A .

Any antichain α can be seen as the pseudo-antichain $A = \{(x, \emptyset) \mid x \in \alpha\}$. Furthermore, notice that *any set* X is PA-represented by $A = \{(x, \alpha_x) \mid x \in X\}$, with $\alpha_x = \lceil \{s \in S \mid s \preceq x \text{ and } s \neq x\} \rceil$. Indeed $\uparrow(x, \alpha_x) = \{x\}$ for all x , and thus $X = \uparrow A$.

The interest of pseudo-antichains is that they behave well with respect to *all Boolean operations*, as shown by Proposition 3 on pseudo-elements (and easily extended to pseudo-antichains). From the algorithmic point of view, it is important to note that the computations only manipulate (pseudo-)antichains instead of their (pseudo-)closure. Note also that the pseudo-antichains computed in this proposition are

²We use calligraphic style for symbols denoting a symbolic representation.

³“PA-representation” means pseudo-antichain based representation.




Figure 4: Intersection (left) and difference (right) of two pseudo-closures of pseudo-elements.

not necessarily simplified. However, our algorithms implementing those operations always simplify the computed pseudo-antichains for the sake of efficiency.

Proposition 3. *Let $(x, \alpha), (y, \beta)$ be two pseudo-elements. Then:*

- $\Downarrow(x, \alpha) \cup \Downarrow(y, \beta) = \Downarrow\{(x, \alpha), (y, \beta)\}$
- $\Downarrow(x, \alpha) \cap \Downarrow(y, \beta) = \Downarrow\{(x \sqcap y, \alpha \dot{\cup} \beta)\}$
- $\Downarrow(x, \alpha) \setminus \Downarrow(y, \beta) = \Downarrow(\{(x, \{y\} \dot{\cup} \alpha)\} \cup \{(x \sqcap b, \alpha) \mid b \in \beta\})$

The following example illustrates the second and third statements of the previous proposition.

Example 3. Let (S, \preceq) be a lower semilattice and let $(x, \{a\})$ and $(y, \{b\})$, with $x, y, a, b \in S$, be two pseudo-elements as depicted in Figure 4. We have $\Downarrow(x, \{a\}) \cap \Downarrow(y, \{b\}) = \Downarrow(x \sqcap y, \{a, b\})$. We also have $\Downarrow(x, \{a\}) \setminus \Downarrow(y, \{b\}) = \Downarrow(\{(x, \{y\} \dot{\cup} \{a\}), (x \sqcap b, \{a\})\}) = \Downarrow\{(x, \{y\}), (b, \{a\})\}$. Note that $(x, \{y\})$ and $(b, \{a\})$ are not in canonical form. The canonical representation of $\Downarrow(x, \{y\})$ (resp. $\Downarrow(b, \{a\})$) is given by $(x, \{x \sqcap y\})$ (resp. $(b, \{b \sqcap a\})$).

5 Pseudo-antichain based algorithms

In this section, we propose a pseudo-antichain based version of the symblicit algorithm described in Section 3 for solving the SSP and EMP problems for monotonic MDPs.

5.1 Operator $\text{Pre}_{\sigma, \tau}$

We begin by presenting a new operator $\text{Pre}_{\sigma, \tau}$ that is useful for our algorithms. Let $M_{\preceq} = (S, \Sigma, T, \mathbf{E}, \mathbf{D})$ be a monotonic MDP. Given $L \subseteq S$, $\sigma \in \Sigma$ and $\tau \in T$, we denote by $\text{Pre}_{\sigma, \tau}(L)$ the set of states that reach L by σ, τ in M_{\preceq} , that is

$$\text{Pre}_{\sigma, \tau}(L) = \{s \in S \mid \mathbf{E}(s, \sigma)(\tau) \in L\}.$$

The elements of $\text{Pre}_{\sigma, \tau}(L)$ are called *predecessors* of L for σ, τ in M_{\preceq} . This operator has the nice following properties: (1) if L is closed, then $\text{Pre}_{\sigma, \tau}(L)$ is closed, and (2) for all sets L_1, L_2 and for all $\cdot \in \{\cup, \cap, \setminus\}$, $\text{Pre}_{\sigma, \tau}(L_1 \cdot L_2) = \text{Pre}_{\sigma, \tau}(L_1) \cdot \text{Pre}_{\sigma, \tau}(L_2)$. Moreover, we have:

Proposition 4. *Let (x, α) be a pseudo-element with $x \in S$ and $\alpha \subseteq S$. Let $A = \{(x_i, \alpha_i) \mid i \in I\}$ be a pseudo-antichain with $x_i \in S$ and $\alpha_i \subseteq S$ for all $i \in I$. Then, for all $\sigma \in \Sigma$ and $\tau \in T$,*

- $\text{Pre}_{\sigma, \tau}(\Downarrow(x, \alpha)) = \bigcup_{x' \in \lceil \text{Pre}_{\sigma, \tau}(\downarrow \{x\}) \rceil} \Downarrow(x', \lceil \text{Pre}_{\sigma, \tau}(\downarrow \alpha) \rceil)$
- $\text{Pre}_{\sigma, \tau}(\Downarrow A) = \bigcup_{i \in I} \text{Pre}_{\sigma, \tau}(\Downarrow(x_i, \alpha_i))$

From Proposition 4, we can efficiently compute pseudo-antichains w.r.t. the $\text{Pre}_{\sigma,\tau}$ operator if we have an efficient algorithm to compute antichains w.r.t. $\text{Pre}_{\sigma,\tau}$ (see the first statement). We make the following assumption that we can compute the predecessors of a closed set by only considering the antichain of its maximal elements. Together with Proposition 4, it implies that the computation of $\text{Pre}_{\sigma,\tau}(\uparrow A)$, for all pseudo-antichains A , does not need to treat the whole pseudo-closure $\uparrow A$.

Assumption 1. *There exists an algorithm taking any state $x \in S$ as input and returning $[\text{Pre}_{\sigma,\tau}(\downarrow\{x\})]$ as output.*

Remark 2. Assumption 1 is a realistic and natural assumption when considering partially ordered state spaces. For instance, it holds for the two classes of MDPs considered in Section 6 for which the given algorithm is straightforward. Assumptions in the same flavor are made in [22] (see Definition 3.2).

5.2 Symbolic representations

We detail in this section the symbolic representations based on pseudo-antichains that we are going to use in our algorithms. Recall from Section 4 that PA-representations are not unique. For efficiency reasons, it will be necessary to work with PA-representations that are as *compact* as possible, as suggested in the sequel.

Representation of the stochastic models. Let $M_{\preceq} = (S, \Sigma, T, \mathbf{E}, \mathbf{D})$ be a monotonic MDP and $M_{\preceq,\lambda} = (S, T, \mathbf{E}_\lambda, \mathbf{D}_\lambda)$ be the MC induced by a strategy λ . For algorithmic purposes, in addition to Assumption 1, we make the following assumption⁴ on M_{\preceq} .

Assumption 2. *There exists an algorithm taking as input any state $s \in S$ and actions $\sigma \in \Sigma_s, \tau \in T$, and returning as output $\mathbf{E}(s, \sigma)(\tau)$ and $\mathbf{D}(s, \sigma)(\tau)$.*

By definition of M_{\preceq} , S is a closed set, and can thus be represented by the pseudo-antichain $\{(x, \emptyset) \mid x \in [S]\}$. By Assumption 2, we have a PA-representation of M_{\preceq} , in the sense that S is PA-represented and we can compute $\mathbf{E}(s, \sigma)(\tau)$ and $\mathbf{D}(s, \sigma)(\tau)$ on demand.

Given λ , we denote by \sim_λ the equivalence relation on S such that $s \sim_\lambda s'$ iff $\lambda(s) = \lambda(s')$. We denote by S_{\sim_λ} the induced partition of S . Given a block $B \in S_{\sim_\lambda}$, we denote by $\lambda(B)$ the unique action $\lambda(s)$, for all $s \in B$. As any set can be represented by a pseudo-antichain, each block of S_{\sim_λ} is PA-represented. Therefore by Assumption 2, we have a PA-representation of $M_{\preceq,\lambda}$.

Representation of a subset of goal states. Recall that a subset $G \subseteq S$ of goal states is required for the SSP problem. Our algorithm will manipulate G when computing the set of proper states. A natural assumption is to require that G is *closed* (like S), as it is the case for the two classes of monotonic MDPs studied in Section 6. Under this assumption, we have a compact representation of G as the one proposed above for S . Otherwise, we take for G any PA-representation.

Representation for \mathbf{D} and \mathbf{C} . For the needs of our algorithm, we introduce symbolic representations for \mathbf{D}_λ and \mathbf{C}_λ . Similarly to \sim_λ , let $\sim_{\mathbf{D},\lambda}$ be the equivalence relation on S such that $s \sim_{\mathbf{D},\lambda} s'$ iff $\mathbf{D}_\lambda(s) = \mathbf{D}_\lambda(s')$. We denote by $S_{\sim_{\mathbf{D},\lambda}}$ the induced partition, and for each block $B \in S_{\sim_{\mathbf{D},\lambda}}$, by $\mathbf{D}_\lambda(B)$ the unique probability distribution $\mathbf{D}_\lambda(s)$, with $s \in B$. We use similar notations for the relation $\sim_{\mathbf{C},\lambda}$ on S such that $s \sim_{\mathbf{C},\lambda} s'$ iff $\mathbf{C}_\lambda(s) = \mathbf{C}_\lambda(s')$. Each block of $S_{\sim_{\mathbf{D},\lambda}}$ and $S_{\sim_{\mathbf{C},\lambda}}$ is PA-represented.

⁴Remark 2 also holds for Assumption 2.

For each $\sigma \in \Sigma$, we also need the next equivalence relations $\sim_{\mathbf{D},\sigma}$ and $\sim_{\mathbf{C},\sigma}$ on S , such that $s \sim_{\mathbf{D},\sigma} s'$ iff $\mathbf{D}(s, \sigma) = \mathbf{D}(s', \sigma)$, and that $s \sim_{\mathbf{C},\sigma} s'$ iff $\mathbf{C}(s, \sigma) = \mathbf{C}(s', \sigma)$. Recall that \mathbf{D} and \mathbf{C} are partial functions, there may thus exist one block in their corresponding relation gathering all states s such that $\sigma \notin \Sigma_s$. Each block of the induced partitions $S_{\sim_{\mathbf{D},\sigma}}$ and $S_{\sim_{\mathbf{C},\sigma}}$ is PA-represented.

For the two classes of MDPs studied in Section 6, both functions \mathbf{D} and \mathbf{C} are independent of S . It follows that the previously described equivalence relations have only one or two blocks, leading to compact symbolic representations of these relations.

Now that the operator $\text{Pre}_{\sigma,\tau}$ and the used symbolic representations have been introduced, we come back to the steps⁵ of the symblicit approach of Section 3 (see Algorithm 2) and show how to derive a pseudo-antichain based algorithm.

5.3 Bisimulation lumping

Algorithm LUMP. Let $M_{\leq} = (S, \Sigma, T, \mathbf{E}, \mathbf{D})$ be a monotonic MDP and $M_{\leq,\lambda} = (S, T, \mathbf{E}_{\lambda}, \mathbf{D}_{\lambda})$ be the MC induced by a strategy λ .⁶ In Algorithm 2, Algorithm LUMP is called to compute the largest bisimulation \sim_L of the MC $M_{\leq,\lambda}$ (line 4 with $\hat{\lambda} = \lambda_n$). This algorithm (see [16]) first computes the initial partition $P = S_{\sim_{\mathbf{C},\lambda}}$ such that two states of the MC are in the same block iff they have the same cost. It then repeatedly splits blocks B of P according to their probability of reaching a given block C , for all C . It stops as soon as for all $B, B' \in P$ and $s, s' \in B$, $\mathbf{P}_{\lambda}(s, B') = \mathbf{P}_{\lambda}(s', B')$. The operation of splitting blocks is performed with Algorithm SPLIT (see Algorithm 3). Before describing it, we need a new operator Pre_{λ} . Given $L \subseteq S$ and $\tau \in T$, we define

$$\text{Pre}_{\lambda}(L, \tau) = \{s \in S \mid \mathbf{E}_{\lambda}(s)(\tau) \in L\}$$

as the set of states from which L is reached by τ in M_{\leq} under the selection made by λ .

Algorithm SPLIT. Given two blocks $B, C \subseteq S$, Algorithm SPLIT splits B into a partition P composed of sub-blocks B_1, \dots, B_k according to the probability of reaching C , i.e. for all $s, s' \in B$, $s, s' \in B_l$ for some l iff $\mathbf{P}_{\lambda}(s, C) = \mathbf{P}_{\lambda}(s', C)$. Given $T = \{\tau_1, \dots, \tau_m\}$, it computes intermediate partitions P of B such that at step i , B is split according to the probability of reaching C when T is restricted to $\{\tau_1, \dots, \tau_i\}$.

Initially, T is restricted to \emptyset , and the partition P is composed of one block B (see line 1). At step i with $i \geq 1$, each block B_l of P computed at step $i-1$ is split into several sub-blocks according to its intersection with $\text{Pre}_{\lambda}(C, \tau_i)$ and each $D \in S_{\sim_{\mathbf{D},\lambda}}$. We take into account intersections with $D \in S_{\sim_{\mathbf{D},\lambda}}$ in a way to know which stochastic function $\mathbf{D}_{\lambda}(D)$ is associated with the states we are considering. Suppose that at step $i-1$ the probability for any state of block B_l of reaching C is p . Then at step i , it is equal to $p + \mathbf{D}_{\lambda}(D)(\tau_i)$ if this state belongs to $D \cap \text{Pre}_{\lambda}(C, \tau_i)$, with $D \in S_{\sim_{\mathbf{D},\lambda}}$, and to p if it does not belong to $\text{Pre}_{\lambda}(C, \tau_i)$ (lines 5-7). See Figure 5 for intuition. Notice that some newly created sub-blocks could have the same probability, they are therefore merged.

The intermediate partitions P (or P_{new}) are represented by hash tables for efficiency reasons: each entry (p, block) is stored as $P[p] = \text{block}$ such that block is the set of states that reach C with probability p . Algorithm INITTABLE is called




Figure 5: Step i of Algorithm 3 on a block B_l .

⁵Due to lack of space, some steps are not detailed. See [10] for a complete description.

⁶Equivalently, with the usual definition of MCs, $M_{\leq,\lambda} = (S, \mathbf{P}_{\lambda})$ with \mathbf{P}_{λ} derived from \mathbf{E}_{λ} and \mathbf{D}_{λ} (see Section 2).

Algorithm 3 SPLIT(B, C, λ)

```

1:  $P[0] := B$ 
2: for  $i$  in  $[1, m]$  do
3:    $P_{\text{new}} := \text{INITTABLE}(P, \tau_i)$ 
4:   for all  $(p, \text{block})$  in  $P$  do
5:      $P_{\text{new}}[p] := P_{\text{new}}[p] \cup (\text{block} \setminus \text{Pre}_\lambda(C, \tau_i))$ 
6:     for all  $D \in S_{\sim_{\mathbf{D}, \lambda}}$  do
7:        $P_{\text{new}}[p + \mathbf{D}_\lambda(D)(\tau_i)] := P_{\text{new}}[p + \mathbf{D}_\lambda(D)(\tau_i)] \cup (\text{block} \cap D \cap \text{Pre}_\lambda(C, \tau_i))$ 
8:    $P := \text{REMOVEEMPTYBLOCKS}(P_{\text{new}})$ 
9: return  $P$ 

```

Algorithm 4 IMPROVESTRAT($\mathcal{L}, S_{\sim_\lambda}$)

```

1: for  $C \in \mathcal{L}$  do
2:    $S_{\sim_{\lambda'}} := \emptyset$ 
3:   for  $B \in S_{\sim_\lambda}$  do
4:     if  $B \cap C \neq \emptyset$  then
5:        $S_{\sim_{\lambda'}} := S_{\sim_{\lambda'}} \cup \{B \cap C, B \setminus C\}$ 
6:     else
7:        $S_{\sim_{\lambda'}} := S_{\sim_{\lambda'}} \cup B$ 
8:    $S_{\sim_\lambda} := S_{\sim_{\lambda'}}$ 
9: return  $S_{\sim_{\lambda'}}$ 

```

to initialize a new partition P_{new} from a previous partition P and symbol τ_i : $P_{\text{new}}[p] := \emptyset$ and $P_{\text{new}}[p + \mathbf{D}_\lambda(D)(\tau_i)] := \emptyset$, for all $D \in S_{\sim_{\mathbf{D}, \lambda}}$ and all (p, block) in P . Algorithm REMOVEEMPTYBLOCKS(P) removes from P each pair (p, block) with $\text{block} = \emptyset$.

Pseudo-antichain approach. Notice that we have a pseudo-antichain version of Algorithm LUMP as soon as the given blocks B and C are PA-represented. Indeed, this algorithm uses boolean operations and the Pre_λ operator which can be computed as $\text{Pre}_\lambda(C, \tau) = \bigcup \{ \text{Pre}_{\sigma, \tau}(C) \cap B \mid \sigma \in \Sigma, B \in S_{\sim_\lambda}, \lambda(B) = \sigma \}$. By Propositions 1, 3-4 and Assumptions 1-2, all these operations can be performed efficiently on pseudo-closures of pseudo-antichains, by limiting the computations to the related pseudo-antichains.

5.4 Improving strategies

Given an MDP M_{\leq} with cost function \mathbf{C} and the MC $M_{\leq, \lambda}$ induced by a strategy λ , we now present a pseudo-antichain based algorithm to improve strategy λ for the SSP (see line 8 of Algorithm 2). Recall that for all $s \in S$, we have to compute the set $\widehat{\Sigma}_s$ of actions $\sigma \in \Sigma_s$ that minimize the expression $l_\sigma(s) = \mathbf{C}(s, \sigma) + \sum_{s' \in S} \mathbf{P}(s, \sigma, s') \cdot v(s')$, and then we improve the strategy based on the computed $\widehat{\Sigma}_s$ (see Algorithm 1 with $v = v_n$ and $\lambda = \lambda_n$). We proceed in two steps: (1) we compute for all $\sigma \in \Sigma$, an equivalence relation \sim_{l_σ} such that the value $l_\sigma(s)$ is constant on each block of the relation, (2) we use the relations \sim_{l_σ} , with $\sigma \in \Sigma$, to improve the strategy.

Computing value l_σ . Let $\sigma \in \Sigma$ be a fixed action. We are looking for an equivalence relation \sim_{l_σ} on the set S_σ of states where action σ is enabled, such that

$$\forall s, s' \in S_\sigma : s \sim_{l_\sigma} s' \Rightarrow l_\sigma(s) = l_\sigma(s').$$

Given \sim_L the largest bisimulation for $M_{\leq, \lambda}$ and the induced partition S_{\sim_L} , we have $l_\sigma(s) = \mathbf{C}(s, \sigma) + \sum_{C \in S_{\sim_L}} \mathbf{P}(s, \sigma, C) \cdot v(C)$ for each $s \in S_\sigma$, since the value v is constant on each block C . Therefore to get relation \sim_{l_σ} , it is enough to have $s \sim_{l_\sigma} s' \Rightarrow \mathbf{C}(s, \sigma) = \mathbf{C}(s', \sigma)$ and $\mathbf{P}(s, \sigma, C) = \mathbf{P}(s', \sigma, C), \forall C \in S_{\sim_L}$.

We proceed by defining the following equivalence relations on S_σ . For the cost part, we use relation $\sim_{\mathbf{C}, \sigma}$ defined in Section 5.2. For the probabilities part, for each block C of S_{\sim_L} , we define relation $\sim_{\mathbf{P}, \sigma, C}$ such that $s \sim_{\mathbf{P}, \sigma, C} s'$ iff $\mathbf{P}(s, \sigma, C) = \mathbf{P}(s', \sigma, C)$. The required relation \sim_{l_σ} on S_σ is then defined as the relation

$$\sim_{l_\sigma} = \sim_{\mathbf{C}, \sigma} \cap \bigcap_{C \in S_{\sim_L}} \sim_{\mathbf{P}, \sigma, C} = \sim_{\mathbf{C}, \sigma} \cap \sim_{\mathbf{P}, \sigma}$$

Relation \sim_{l_σ} induces a partition of S_σ that we denote $(S_\sigma)_{\sim_{l_\sigma}}$. For each block $D \in (S_\sigma)_{\sim_{l_\sigma}}$, we denote by $l_\sigma(D)$ the unique value $l_\sigma(s)$, for $s \in D$.

Let us now explain how to compute \sim_{l_σ} with a pseudo-antichain based approach. (1) The set S_σ is obtained as $S_\sigma = \text{Pre}_{\sigma,\tau}(S)$ with τ an arbitrary action of T . (2) Each relation $\sim_{\mathbf{P},\sigma,C}$ is the output returned by $\text{SPLIT}(S_\sigma, C, \lambda)$ where λ is defined on S_σ by $\lambda(s) = \sigma$ for all $s \in S_\sigma$ ⁷ (see Algorithm 3). (3) Let us detail a way to compute $\sim_{\mathbf{P},\sigma}$ from $\sim_{\mathbf{P},\sigma,C}$, for all $C \in S_{\sim_L}$. Let $S_{\sim_{\mathbf{P},\sigma,C}} = \{B_{C,1}, B_{C,2}, \dots, B_{C,k_C}\}$ be the partition of S_σ induced by $\sim_{\mathbf{P},\sigma,C}$. For each $B_{C,i} \in S_{\sim_{\mathbf{P},\sigma,C}}$, we denote by $\mathbf{P}(B_{C,i}, \sigma, C)$ the unique value $\mathbf{P}(s, \sigma, C)$, for all $s \in B_{C,i}$. Then, computing a block D of $\sim_{\mathbf{P},\sigma}$ consists in picking, for all $C \in S_{\sim_L}$, one block D_C among $B_{C,1}, B_{C,2}, \dots, B_{C,k_C}$, such that the intersection $D = \bigcap_{C \in S_{\sim_L}} D_C$ is non empty. As $\sum_{s' \in S} \mathbf{P}(s, \sigma, s') = 1$, if $D \neq \emptyset$, then $\sum_{C \in S_{\sim_L}} \mathbf{P}(D_C, \sigma, C) = 1$. (4) Finally \sim_{l_σ} is the intersection of $\sim_{C,\sigma}$ and $\sim_{\mathbf{P},\sigma}$.

Improving the strategy. We now propose a pseudo-antichain based algorithm for improving strategy λ by using relations \sim_L , \sim_λ , and \sim_{l_σ} , $\forall \sigma \in \Sigma$ (see Algorithm 4). We first compute for all $\sigma \in \Sigma$, the equivalence relation $\sim_{l_\sigma \wedge L} = \sim_{l_\sigma} \cap \sim_L$ on S_σ . Given $B \in (S_\sigma)_{\sim_{l_\sigma \wedge L}}$, we denote by $l_\sigma(B)$ the unique value $l_\sigma(s)$ and by $v(B)$ the unique value $v(s)$, for all $s \in B$. Let $\sigma \in \Sigma$, we denote by $(S_\sigma)_{\sim_{l_\sigma \wedge L}}^< \subseteq (S_\sigma)_{\sim_{l_\sigma \wedge L}}$ the set of blocks C for which the value $v(C)$ is improved by setting $\lambda(C) = \sigma$, that is

$$(S_\sigma)_{\sim_{l_\sigma \wedge L}}^< = \{C \in (S_\sigma)_{\sim_{l_\sigma \wedge L}} \mid l_\sigma(C) < v(C)\}.$$

We then compute an ordered global list \mathcal{L} made of the blocks of all sets $(S_\sigma)_{\sim_{l_\sigma \wedge L}}^<$, for all $\sigma \in \Sigma$. It is ordered according to the decreasing value $l_\sigma(C)$. In this way, when traversing \mathcal{L} , we have more and more promising blocks to decrease v .

From input \mathcal{L} and \sim_λ , Algorithm 4 outputs an equivalence relation $\sim_{\lambda'}$ for a new strategy λ' improving λ . Given $C \in \mathcal{L}$, suppose that C comes from the relation $\sim_{l_\sigma \wedge L}$ (σ is considered). For each $B \in S_{\sim_\lambda}$ with $B \cap C \neq \emptyset$ (line 4), we improve the strategy by setting $\lambda'(B \cap C) = \sigma$, while λ' is kept unchanged for $B \setminus C$. Algorithm 4 outputs a partition $S_{\sim_{\lambda'}}$ such that $s \sim_{\lambda'} s' \Rightarrow \lambda'(s) = \lambda'(s')$ for the improved strategy λ' . If necessary, for efficiency reasons, we can compute a coarser relation for the new strategy λ' by gathering blocks B_1, B_2 of $S_{\sim_{\lambda'}}$, for all B_1, B_2 such that $\lambda'(B_1) = \lambda'(B_2)$.

6 Experiments

We present two applications of the symblicit algorithm of Section 5, one for the SSP problem in the context of automated planning, and the other for the EMP problem in the context of LTL synthesis. In both cases, we have a reduction, described in [10], to monotonic MDPs equipped with a natural partial order and that satisfy Assumptions 1 and 2. Our experiments have been done on a Linux platform with a 3.2GHz CPU (Intel Core i7), with a timeout of 10 hours and a memory usage limited to 4GB.

6.1 Stochastic shortest path on STRIPSs

Monotonic STRIPS. A *monotonic STRIPS (MS)* is a tuple (P, I, M, O) where P is a finite set of *conditions* (i.e. propositional variables), $I \subseteq P$ is a subset of conditions that are initially true (all others are assumed to be false), $M \subseteq P$ specifies which conditions must be true in a goal state, and O is a finite set of *operators*. An operator $o \in O$ is a pair $(\gamma, (\alpha, \delta))$ where $\gamma \subseteq P$ is the *guard* of o , that is, the set of conditions that must be true for o to be executable, and (α, δ) is the *effect* of o , that is, $\alpha \subseteq P$ (resp. $\delta \subseteq P$) is the set of conditions that are made true (resp. false) by the execution of o . *Monotonic stochastic*

⁷As Algorithm SPLIT only works on S_σ , it is not a problem if λ is not defined on $S \setminus S_\sigma$.

STRIPS (MSS) are MSs extended with stochastic aspects as follows [6]. Each operator $o = (\gamma, \pi) \in O$ consists of a guard γ as before, and an effect given as a probability distribution $\pi : 2^P \times 2^P \rightarrow [0, 1]$ on the set of pairs (α, δ) . If additionally, we have a *cost function* $C : O \rightarrow \mathbb{R}_{>0}$, then the problem of planning from MSSs is to minimize the expected truncated sum up to the set of goal states from the initial state. One can check that MSSs naturally define monotonic MDPs on which our algorithm of Section 5 can be applied for solving the mentioned SSP problem.

Experiments. Our implementation is available at <http://lit2.ulb.ac.be/STRIPSSolver/> together with the two benchmarks presented in this section. It is compared with the purely explicit⁸ strategy iteration algorithm implemented in the development release 4.1.dev.r7712 of the tool PRISM [29]. This explicit implementation exists primarily to prototype new techniques and is thus not fully optimized [32]. While value iteration algorithms are usually efficient, they only compute approximations. As a consequence, for the sake of a fair comparison, we consider here only the performances of strategy iteration algorithms.

In the benchmark Monkey inspired from [35], a monkey has several items (boxes, stones, pieces of sticks...) at its disposal to reach a bunch of bananas, with the condition that it has to assemble some pieces to get a stick. The monkey has multiple ways to build the stick with varying building times. The operators of getting some items and the bananas are stochastic. For instance, the probability of getting the bananas varies according to the owned items. The benchmark is parameterized in the pair (p, s) where p is the number of pieces to build a stick, and $p \cdot s$ is the total number of pieces.

The benchmark Moats and castles is an adaptation of a benchmark of [31] as proposed in [6]⁹. The goal is to build a sand castle on the beach; a moat can be dug in a way to protect it. We consider up to 7 discrete depths of moat. The operator of building the castle is stochastic: there is a strictly positive probability for the castle to be demolished by the waves. However, the deeper the moat is, the higher the probability of success is. To increase the difficulty of the problem, we consider building several castles, each one having its own moat. The benchmark is parameterized in the pair (d, c) where d is the number of depths of moat that can be dug, and c is the numbers of castles to be built.

Results are given in Table 1. On those two benchmarks, the explicit implementation quickly runs out of memory when the state space of the MDP grows. Indeed, with this method, we were not able to solve MDPs with more than 65536 states. On the other hand, the symblicit algorithm behaves well on large models: the memory consumption never exceeds 150Mo and this even for MDPs with hundreds of millions of states¹⁰.

6.2 Expected mean-payoff with LTL_{MP} synthesis

LTL_{MP} synthesis problem. Let ϕ be an LTL formula over a set $P = I \uplus O$. Let $\Sigma_P = 2^P$, $\Sigma_O = 2^O$ and $\Sigma_I = 2^I$, and $w : \Sigma_O \mapsto \mathbb{Z}$ be a weight function over Σ_O . Consider the next infinite game between Player O and Player I . At each turn k , Player O gives $o_k \in \Sigma_O$ and Player I responds by giving $i_k \in \Sigma_I$. The outcome of the game is the word $u = (o_0 \cup i_0)(o_1 \cup i_1) \cdots \in \Sigma_P^\omega$. A *value* $\text{Val}(u)$ is associated with u such that

⁸To the best of our knowledge, there is no tool implementing an MTBDD based symblicit algorithm for the SSP problem. However, a comparison with an MTBDD based symblicit algorithm is done in the second application for the EMP problem.

⁹In [6], the authors study the problem of maximizing the probability of reaching the goal within a given number of steps.

¹⁰On our benchmarks, the value iteration algorithm of PRISM performs better than the strategy iteration one w.r.t. the run time and memory consumption. But it still consumes more memory than the pseudo-antichain based algorithm, and runs out of memory on several examples.

Table 1: Stochastic shortest path on two benchmarks of MSSs. The column *example* gives the parameters (s, p) (resp. (c, d)) for the Monkey (resp. Moats and castles) benchmark. The column $\mathbb{E}_\lambda^{\text{TS}_G}$ gives the expected truncated sum of the computed strategy, and $|M_S|$ the number of states of the MDP. For the pseudo-antichain based implementation (PA), #it is the number of iterations of the strategy iteration algorithm, $|S_{\sim_L}|$ the maximum size of computed bisimulation quotients, and *lump*, *syst* and *impr* the total times (in seconds) spent respectively for lumping, solving the linear systems and improving the strategies. For both implementations, *total* is the total execution time (in seconds) and *mem* the total memory consumption (in megabytes).

example	$\mathbb{E}_\lambda^{\text{TS}_G}$	$ M_S $	PA							Explicit		
			#it	$ S_{\sim_L} $	lump	syst	impr	total	mem	total	mem	
Monkey	(3,2)	35.75	4096	4	23	0.09	0.00	0.07	0.16	16.0	60.59	1626
	(3,3)	35.75	65536	5	43	1.14	0.00	0.43	1.57	17.3		>4000
	(3,4)	35.75	1048576	6	57	12.89	0.01	4.92	17.83	21.7		>4000
	(3,5)	36.00	16777216	7	88	208.33	0.05	63.73	272.13	37.5		>4000
	(5,2)	35.75	65536	4	31	0.36	0.00	0.18	0.54	16.6	20316.17	2343
	(5,3)	35.75	4194304	5	56	5.71	0.02	2.47	8.20	19.5		>4000
	(5,4)	35.75	268435456	6	97	95.49	0.04	101.27	196.83	31.3		>4000
	(5,5)	36.00	17179869184	7	152	1813.78	0.08	5284.31	7098.40	81.3		>4000
Moats and castles	(2,5)	32.2222	4096	3	49	1.36	0.00	0.45	1.82	17.3	133.66	1202
	(2,6)	32.2222	16384	3	66	9.71	0.01	1.95	11.68	19.3		2966.80
	(3,3)	59.0000	4096	3	84	12.58	0.03	2.73	15.35	20.2	149.64	1205
	(3,4)	52.0000	32768	3	219	129.17	0.05	21.56	150.83	30.7		14660.69
	(3,5)	48.3333	262144	3	357	658.86	0.13	81.08	740.17	49.1	>4000	>4000
	(3,6)	48.3333	2097152	3	595	10730.09	0.42	865.48	11597.71	145.8		>4000
	(4,2)	96.8889	4096	3	132	31.61	0.03	12.06	43.72	26.5	173.62	1211
	(4,3)	78.6667	65536	3	464	1376.94	0.21	217.06	1594.48	82.2		>4000

$$\text{Val}(u) = \begin{cases} \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} w(o_k) & \text{if } u \models \phi \\ -\infty & \text{otherwise} \end{cases}$$

i.e. $\text{Val}(u)$ is the mean-payoff value of u if u satisfies ϕ , otherwise, it is $-\infty$. Given a threshold value $v \in \mathbb{Z}$, the LTL_{MP} realizability problem asks whether Player O has a strategy against any strategy of Player I such that $\text{Val}(u) \geq v$ for the produced outcome u . The LTL_{MP} synthesis problem is to produce such a strategy for Player O .

In [8, 9], we propose an antichain based algorithm for solving the LTL_{MP} realizability and synthesis problems, that is incremental in some parameters (K, C) , and uses a reduction to a two-player turn-based safety game G (see [8] for details). This game restricted to the winning positions of Player O is a representation of a subset W (depending on (K, C)) of all winning strategies for Player O for LTL_{MP} realizability problem. From this set W , we want to compute a strategy that behaves *the best against a stochastic opponent*. Let $\pi_I : \Sigma_I \rightarrow]0, 1]$ be a probability distribution such that $\pi_I(i) > 0$ for all $i \in \Sigma_I$ (to make sense with the worst-case). By replacing Player I by π_I in G , we can derive a monotonic MDP M_G equipped with a natural partial order, and computing the best strategy among strategies in W reduces to solving the EMP problem for the MDP M_G ¹¹.

Experiments. We have integrated the symblicit algorithm presented in Section 5 for the EMP problem into Acacia+ (v2.2) [7], a tool for solving the LTL_{MP} realizability and synthesis problems. The latest version of Acacia+ can be downloaded at <http://lit2.ulb.ac.be/acaciaplus/>, together

¹¹More precisely, it reduces to the EMP problem where the objective is to maximize the expected mean-payoff (see footnote 1).

Table 2: Expected mean-payoff on the Stochastic shared resource arbiter benchmark with 2 clients. The column v gives the threshold, $|M_G^R|$ the number of reachable states in the MDP, and all other columns have the same meaning as in Table 1. The expected mean-payoff $\mathbb{E}_\lambda^{\text{MP}}$ of the optimal strategy λ for all the examples is -0.130435 .

v	Acacia+								PRISM		
	$ M_G^R $	#it	$ S_{\sim\gamma} $	lump	LS	impr	total	mem	$ M_G^R $	total	mem
-1.1	5259	2	22	0.12	0.01	0.02	0.15	17.4	691	0.50	168.1
-1.04	35750	2	52	1.63	0.02	0.13	1.79	18.1	3325	2.06	264.1
-1.02	530299	2	102	16.62	0.11	0.64	17.39	20.2	11641	7.33	343.4
-1.01	4120599	2	202	237.78	0.50	3.94	242.30	26.2	43891	31.52	642.5
-1.004	63251499	2	502	7357.72	5.68	52.81	7416.77	60.5	264391	278.01	2544.0
-1.003	450012211	2	670	23455.44	12.72	120.25	23589.49	93.6			> 4000

with the examples considered in this section. We compared our implementation with an MTBDD based symblicit algorithm implemented in PRISM [37] (in the sequel, our implementation is simply called Acacia+ whereas the other one is called PRISM). In the used benchmark [9], a server has to grant exclusive access to a resource to two clients. We set a probability distribution such that requests of client 1 (probability $\frac{3}{5}$) are more likely to happen than requests of client 2 (probability $\frac{1}{5}$), and the benchmark is parameterized in the threshold value v .

Results are given in Table 2. Note that the number of states in the MDPs depends on the implementation. Indeed, for PRISM, it is the number of reachable states of the MDP, denoted $|M_G^R|$, that is, the states that are actually taken into account by the algorithm, while for Acacia+, it is the total number of states since unlike PRISM, our implementation does not prune unreachable states. For this application scenario, the ratio (number of reachable states)/(total number of states) is in general quite small¹². On this benchmark, PRISM is faster than Acacia+ on large models, but Acacia+ is more efficient regarding the memory consumption and this in spite of considering the whole state space. Note that the surprisingly large amount of memory consumption of both implementations on small instances is due to Python libraries loaded in memory for Acacia+, and to the JVM and the CUDD package for PRISM [26].

Finally, in the majority of experiments we performed for both the EMP and the SSP problems, we observe that most of the execution time of the pseudo-antichain based symblicit algorithms is spent for lumping. It is also the case for the MTBDD based symblicit algorithm [38].

7 Conclusion

We have presented the interesting class of monotonic MDPs, and the new data structure of pseudo-antichains. We have shown how monotonic MDPs can be exploited by symblicit algorithms using pseudo-antichains (instead of MTBDDs) for two quantitative settings: the expected mean-payoff and the stochastic shortest path. Those algorithms have been implemented, and we have reported promising experimental results for two applications coming from automated planning and LTL_{MP} synthesis. We are convinced that pseudo-antichains can be used in the design of efficient algorithms in other contexts like for instance model-checking or synthesis of non-stochastic models, as soon as a natural partial order can be exploited.

Acknowledgments. We would like to thank Mickael Randour for his fruitful discussions, Marta Kwiatkowska, David Parker and Christian Von Essen for their help regarding the tool PRISM, and Holger Hermanns and Ernst Moritz Hahn for sharing with us their prototypical implementation.

¹²For all the MDPs considered in Table 1, this ratio is 1.

References

- [1] Luca de Alfaro (1999): *Computing Minimum and Maximum Reachability Times in Probabilistic Systems*. In Jos C. M. Baeten & Sjouke Mauw, editors: *CONCUR, Lecture Notes in Computer Science 1664*, Springer, pp. 66–81, doi:10.1007/3-540-48320-9_7.
- [2] Christel Baier & Joost-Pieter Katoen (2008): *Principles of model checking*. MIT Press.
- [3] Christel Baier, Joost-Pieter Katoen, Holger Hermanns & Verena Wolf (2005): *Comparative branching-time semantics for Markov chains*. *Inf. Comput.* 200(2), pp. 149–214, doi:10.1016/j.ic.2005.03.001.
- [4] D. P. Bertsekas & J.N. Tsitsiklis (1996): *Neuro-Dynamic Programming*. Anthropological Field Studies, Athena Scientific.
- [5] Dimitri P Bertsekas & John N Tsitsiklis (1991): *An analysis of stochastic shortest path problems*. *Mathematics of Operations Research* 16(3), pp. 580–595, doi:10.1287/moor.16.3.580.
- [6] Avrim L Blum & John C Langford (2000): *Probabilistic planning in the graphplan framework*. In: *Recent Advances in AI Planning*, Springer, pp. 319–332, doi:10.1007/10720246_25.
- [7] Aaron Bohy, Véronique Bruyère, Emmanuel Filiot, Naiyong Jin & Jean-François Raskin (2012): *Acacia+, a Tool for LTL Synthesis*. In P. Madhusudan & Sanjit A. Seshia, editors: *CAV, Lecture Notes in Computer Science 7358*, Springer, pp. 652–657, doi:10.1007/978-3-642-31424-7_45.
- [8] Aaron Bohy, Véronique Bruyère, Emmanuel Filiot & Jean-François Raskin (2012): *Synthesis from LTL Specifications with Mean-Payoff Objectives*. *CoRR* abs/1210.3539. Available at <http://arxiv.org/abs/1210.3539>.
- [9] Aaron Bohy, Véronique Bruyère, Emmanuel Filiot & Jean-François Raskin (2013): *Synthesis from LTL Specifications with Mean-Payoff Objectives*. In Nir Piterman & Scott A. Smolka, editors: *TACAS, Lecture Notes in Computer Science 7795*, Springer, pp. 169–184, doi:10.1007/978-3-642-36742-7_12.
- [10] Aaron Bohy, Véronique Bruyère & Jean-François Raskin (2014): *Symblicit algorithms for optimal strategy synthesis in monotonic Markov decision processes (extended version)*. *CoRR* abs/1402.1076. Available at <http://arxiv.org/abs/1402.1076>.
- [11] Ahmed Bouajjani, Peter Habermehl, Lukás Holík, Tayssir Touili & Tomáš Vojnar (2008): *Antichain-Based Universality and Inclusion Testing over Nondeterministic Finite Tree Automata*. In Oscar H. Ibarra & Bala Ravikumar, editors: *CIAA, Lecture Notes in Computer Science 5148*, Springer, pp. 57–67, doi:10.1007/978-3-540-70844-5_7.
- [12] Peter Buchholz (1994): *Exact and ordinary lumpability in finite Markov chains*. *Journal of applied probability*, pp. 59–75, doi:10.2307/3215235.
- [13] Jerry R. Burch, Edmund M. Clarke, Kenneth L. McMillan, David L. Dill & L. J. Hwang (1992): *Symbolic Model Checking: 10²⁰ States and Beyond*. *Inf. Comput.* 98(2), pp. 142–170, doi:10.1016/0890-5401(92)90017-A.
- [14] Krishnendu Chatterjee, Thomas A. Henzinger, Barbara Jobstmann & Rohit Singh (2011): *QUASY: Quantitative Synthesis Tool*. In Parosh Aziz Abdulla & K. Rustan M. Leino, editors: *TACAS, Lecture Notes in Computer Science 6605*, Springer, pp. 267–271, doi:10.1007/978-3-642-19835-9_24.
- [15] Edmund M. Clarke & E. Allen Emerson (1981): *Design and Synthesis of Synchronization Skeletons Using Branching-Time Temporal Logic*. In Dexter Kozen, editor: *Logic of Programs, Lecture Notes in Computer Science 131*, Springer, pp. 52–71, doi:10.1007/BFb0025774.
- [16] Salem Derisavi, Holger Hermanns & William H. Sanders (2003): *Optimal state-space lumping in Markov chains*. *Inf. Process. Lett.* 87(6), pp. 309–315, doi:10.1016/S0020-0190(03)00343-0.
- [17] Laurent Doyen & Jean-François Raskin (2007): *Improved Algorithms for the Automata-Based Approach to Model-Checking*. In Orna Grumberg & Michael Huth, editors: *TACAS, Lecture Notes in Computer Science 4424*, Springer, pp. 451–465, doi:10.1007/978-3-540-71209-1_34.

- [18] Laurent Doyen & Jean-François Raskin (2010): *Antichain Algorithms for Finite Automata*. In Javier Esparza & Rupak Majumdar, editors: *TACAS, Lecture Notes in Computer Science* 6015, Springer, pp. 2–22, doi:10.1007/978-3-642-12002-2_2.
- [19] Christian von Essen & Barbara Jobstmann (2012): *Synthesizing Efficient Controllers*. In Viktor Kuncak & Andrey Rybalchenko, editors: *VMCAI, Lecture Notes in Computer Science* 7148, Springer, pp. 428–444, doi:10.1007/978-3-642-27940-9_28.
- [20] Jerzy Filar & Koos Vrieze (1997): *Competitive Markov decision processes*. Springer.
- [21] Emmanuel Filiot, Naiyong Jin & Jean-François Raskin (2011): *Antichains and compositional algorithms for LTL synthesis*. *Formal Methods in System Design* 39(3), pp. 261–296, doi:10.1007/s10703-011-0115-3.
- [22] Alain Finkel & Ph. Schnoebelen (2001): *Well-structured transition systems everywhere!* *Theor. Comput. Sci.* 256(1-2), pp. 63–92, doi:10.1016/S0304-3975(00)00102-X.
- [23] Hans Hansson & Bengt Jonsson (1994): *A Logic for Reasoning about Time and Reliability*. *Formal Asp. Comput.* 6(5), pp. 512–535, doi:10.1007/BF01211866.
- [24] Arnd Hartmanns (2012): *MODEST - A unified language for quantitative models*. In: *FDL, IEEE*, pp. 44–51.
- [25] Ronald A Howard (1960): *Dynamic Programming and Markov Processes*. John Wiley and Sons.
- [26] David N. Jansen, Joost-Pieter Katoen, Marcel Oldenkamp, Mariëlle Stoelinga & Ivan S. Zapreev (2007): *How Fast and Fat Is Your Probabilistic Model Checker? An Experimental Performance Comparison*. In Karen Yorav, editor: *Haifa Verification Conference, Lecture Notes in Computer Science* 4899, Springer, pp. 69–85, doi:10.1007/978-3-540-77966-7_9.
- [27] Joost-Pieter Katoen, Ivan S. Zapreev, Ernst Moritz Hahn, Holger Hermanns & David N. Jansen (2011): *The ins and outs of the probabilistic model checker MRMC*. *Perform. Eval.* 68(2), pp. 90–104, doi:10.1016/j.peva.2010.04.001.
- [28] John G Kemeny & J. L. Snell (1960): *Finite Markov Chains*. Van Nostrand Company, Inc.
- [29] M. Kwiatkowska, G. Norman & D. Parker (2011): *PRISM 4.0: Verification of Probabilistic Real-time Systems*. In G. Gopalakrishnan & S. Qadeer, editors: *Proc. 23rd International Conference on Computer Aided Verification (CAV’11), LNCS* 6806, Springer, pp. 585–591, doi:10.1007/978-3-642-22110-1_47.
- [30] Kim Guldstrand Larsen & Arne Skou (1991): *Bisimulation through Probabilistic Testing*. *Inf. Comput.* 94(1), pp. 1–28, doi:10.1016/0890-5401(91)90030-6.
- [31] Stephen M. Majercik & Michael L. Littman (1998): *MAXPLAN: A New Approach to Probabilistic Planning*. In Reid G. Simmons, Manuela M. Veloso & Stephen F. Smith, editors: *AIPS, AAAI*, pp. 86–93.
- [32] David Parker (2013-11-20): personal communication.
- [33] PRISM Model Checker - Frequently asked questions: <http://www.prismmodelchecker.org/manual/FrequentlyAskedQuestions/PRISMModelling>. Accessed: 2014-06-20.
- [34] Martin L. Puterman (1994): *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley and Sons, doi:10.1002/9780470316887.
- [35] Stuart J. Russell & Peter Norvig (1995): *Artificial intelligence: a modern approach*. Prentice hall Englewood Cliffs.
- [36] Arthur F Veinott (1966): *On finding optimal policies in discrete dynamic programming with no discounting*. *The Annals of Mathematical Statistics* 37(5), pp. 1284–1294, doi:10.1214/aoms/1177699272.
- [37] Christian Von Essen (2013-11-20): personal communication.
- [38] Ralf Wimmer, Bettina Brautling, Bernd Becker, Ernst Moritz Hahn, Pepijn Crouzen, Holger Hermanns, Abhishek Dhama & Oliver E. Theel (2010): *Symblicit Calculation of Long-Run Averages for Concurrent Probabilistic Systems*. In: *QEST, IEEE Computer Society*, pp. 27–36, doi:10.1109/QEST.2010.12.
- [39] Martin De Wulf, Laurent Doyen, Thomas A. Henzinger & Jean-François Raskin (2006): *Antichains: A New Algorithm for Checking Universality of Finite Automata*. In Thomas Ball & Robert B. Jones, editors: *CAV, Lecture Notes in Computer Science* 4144, Springer, pp. 17–30, doi:10.1007/11817963_5.