

A Game-based Abstraction-Refinement Framework for Markov Decision Processes

Mark Kattenbelt · Marta Kwiatkowska · Gethin Norman · David Parker

Published online: 3 August 2010

Abstract In the field of model checking, abstraction refinement has proved to be an extremely successful methodology for combating the state-space explosion problem. However, little practical progress has been made in the setting of probabilistic verification. In this paper we present a novel abstraction-refinement framework for Markov decision processes (MDPs), which are widely used for modelling and verifying systems that exhibit both probabilistic and nondeterministic behaviour. Our framework comprises an abstraction approach based on stochastic two-player games, two refinement methods and an efficient algorithm for an abstraction-refinement loop. The key idea behind the abstraction approach is to maintain a separation between nondeterminism present in the original MDP and nondeterminism introduced during the abstraction process, each type being represented by a different player in the game. Crucially, this allows lower and upper bounds to be computed for the values of reachability properties of the MDP. These give a quantitative measure of the quality of the abstraction and form the basis of the corresponding refinement methods. We describe a prototype implementation of our framework and present experimental results demonstrating automatic generation of compact, yet precise, abstractions for a large selection of real-world case studies.

1 Introduction

Numerous real-life systems from a wide range of application domains, including communication and network protocols, security protocols and distributed algorithms, exhibit both *probabilistic* and *nondeterministic* behaviour, and therefore developing techniques to verify the correctness of such systems is an important research topic. Markov decision processes

Mark Kattenbelt · Marta Kwiatkowska · David Parker
Oxford University Computing Laboratory, Wolfson Building, Parks Road, Oxford, OX1 3QD
Telephone: +44 (0) 1865 283566, Fax: +44 (0) 1865 273839
E-mail: {mark.kattenbelt, marta.kwiatkowska, david.parker}@comlab.ox.ac.uk

Gethin Norman
Department of Computing Science, 18 Lilybank Gardens, University of Glasgow, Glasgow, G12 8RZ
Telephone: +44 (0) 141 330 1630, Fax: +44 (0) 141 330 4913
E-mail: gethin@dcs.gla.ac.uk

(MDPs) are a natural and widely used model for this purpose and the automatic verification of MDPs using *probabilistic model checking* has proved successful for their analysis. Despite improvements in implementations and tool support in this area, the state-space explosion problem remains a major hurdle for the practical application of these methods.

This paper is motivated by the success of *abstraction-refinement* techniques [13], which have been established as one of the most effective ways of attacking the state-space explosion problem in non-probabilistic model checking. The basic idea is to construct a smaller *abstract model*, by removing details from the concrete system not relevant to the property of interest, which is consequently easier to analyse. This is done in such a way that when the property is verified true in the abstraction it also holds in the concrete system. On the other hand, if the property does not hold in the abstraction, information from the model checking process (typically a counterexample) is used either to show that the property is false in the concrete system or to *refine* the abstraction. This process forms the basis of a loop which refines the abstraction until the property is shown to be true or false in the concrete system.

In the probabilistic setting, it is typically necessary to consider *quantitative* properties, in which case the actual probability or expectation of some event must be determined, e.g. “the probability of reaching an error state within T time units” or “the expected power consumption before termination”. We therefore adopt a quantitative approach when defining a correspondence between properties of the concrete and abstract models. One example would be the case where quantitative results computed from the abstraction constitute conservative bounds on the actual values for the concrete model. In fact, due to the presence of nondeterminism in an MDP there is not necessarily a single value corresponding to a given quantitative measure. Instead, *best-case* and *worst-case* scenarios are analysed. More specifically, model checking of MDPs typically reduces to computation of *probabilistic reachability* and *expected reachability* properties, namely the minimum or maximum probability of reaching a set of states, and the minimum or maximum expected reward or cost incurred when reaching a set of states.

When constructing an abstraction of an MDP, the resulting model will invariably exhibit a greater degree of nondeterminism caused by the uncertainty with regards to the precise behaviour of the system that the abstraction process introduces. The key idea in our abstraction approach is to maintain a distinction between the nondeterminism from the original MDP and the nondeterminism introduced during the abstraction process. To achieve this, we model abstractions of MDPs as stochastic two-player games [40, 14], where the two players correspond to the two different forms of nondeterminism. We can then analyse these models using techniques developed for such games [15, 3, 10].

Our analysis of these abstract models results in a separate lower and upper bound for each of the minimum and maximum probabilities (or expected rewards) of reaching a set of states. This approach is particularly appealing since this information provides both a quantitative measure of the quality (or preciseness) of the abstraction and an indication of how to improve it. By comparison, if no discrimination between the two forms of nondeterminism is made, a single lower and upper bound would be obtained. Therefore, in the (common) situation where the minimum and maximum probabilities (or expected rewards) are notably different, the quantitative measure and corresponding basis for refinement would be lost. Consider, for example, the extreme case where the two-player game approach reveals that the minimum probability of reaching some set of states is in the interval $[0, \epsilon_{\min}]$ and the maximum probability is in the interval $[1 - \epsilon_{\max}, 1]$. In this case, a single pair of bounds could at best establish that both the minimum and maximum probability lie within the interval $[0, 1]$, effectively yielding no information about the utility of the abstraction or how to refine it in order to improve the bounds.

Based on the separate bounds obtained through this approach, we present two methods for automatically refining our game-based abstractions of MDPs and an efficient algorithm for an abstraction-refinement loop. In addition, using a prototype implementation of the framework, we give experimental results that demonstrate automatic generation of compact, yet precise, abstractions for a large selection of MDP case studies. For convenience, the current prototype performs model-level abstractions, first building an MDP in the probabilistic model checker PRISM [22,37] and then reducing it to a stochastic two-player game. Our framework is also designed to function with higher-level abstraction methods such as predicate abstraction [19], which has been adapted to probabilistic models [44,25]. In fact, the techniques presented in this paper have subsequently been applied to build abstractions of probabilistic models using both predicate abstraction [26] and convex polyhedra [29].

A preliminary version of this paper, introducing the game-based abstraction approach, was published in conference proceedings as [28].

Outline of the Paper. In the next section, we present background material on Markov decision processes and stochastic two-player games required in the remainder of the paper. Section 3 describes our abstraction technique and shows its correctness. Based on this, Section 4 introduces a corresponding abstraction-refinement framework. In Section 5, we describe a prototype implementation of the techniques from Sections 3 and 4, and present experimental results for a large selection of real-world case studies. Section 6 discusses related work and Section 7 concludes the paper.

2 Background

Let $\mathbb{R}_{\geq 0}$ denote the set of non-negative reals. For a finite set Q , we denote by $\text{Dist}(Q)$ the set of *discrete probability distributions* over Q , i.e. the set of functions $\mu : Q \rightarrow [0, 1]$ such that $\sum_{q \in Q} \mu(q) = 1$.

2.1 Probability Measures and Spaces

We begin by briefly introducing some relevant notions from probability and measure theory. For further details, the interested reader is referred to, for example, [8,27].

Definition 1 Let Ω be an arbitrary non-empty set and \mathcal{F} a family of subsets of Ω . We say that \mathcal{F} is a σ -field on Ω if:

1. the empty set \emptyset is in \mathcal{F} ;
2. whenever A is an element of \mathcal{F} , the complement $\Omega \setminus A$ is also in \mathcal{F} ;
3. whenever A_i is an element of \mathcal{F} for each $i \in \mathbb{N}$, the union $\bigcup_{i \in \mathbb{N}} A_i$ is also in \mathcal{F} .

The elements of \mathcal{F} are called *measurable sets*, and (Ω, \mathcal{F}) is called a *measurable space*.

Definition 2 (Probability Measure and Space) Let (Ω, \mathcal{F}) be a measurable space. A function $\text{Prob} : \mathcal{F} \rightarrow [0, 1]$ is a *probability measure* on (Ω, \mathcal{F}) and $(\Omega, \mathcal{F}, \text{Prob})$ a *probability space*, if the following conditions hold:

1. $\text{Prob}(\Omega) = 1$
2. if A_1, A_2, \dots is a disjoint sequence of elements of \mathcal{F} , then $\text{Prob}(\bigcup_i A_i) = \sum_i \text{Prob}(A_i)$.

For any family \mathcal{A} of subsets of Ω , there exists a unique smallest σ -field containing \mathcal{A} [27], which we call the σ -field *generated* by \mathcal{A} . Furthermore, given that \mathcal{A} satisfies certain properties (see e.g. [27] for details), and we define the probability for each element of \mathcal{A} with a function $Prob : \mathcal{A} \rightarrow [0, 1]$, then the function is guaranteed to extend uniquely to a probability measure on the σ -field generated by \mathcal{A} .

2.2 Discrete-Time Markov Chains

We now introduce discrete-time Markov chains (DTMCs), a model for systems with only probabilistic behaviour. A DTMC can be thought of as a state transition system, where transitions are annotated with probabilities indicating the likelihood of their occurrence.

Definition 3 A DTMC is a tuple $D = (S, s_{init}, \mathbf{P})$ where:

- S is a set of states;
- $s_{init} \in S$ is an initial state;
- $\mathbf{P} : S \times S \rightarrow [0, 1]$ is a *transition probability matrix*, with $\sum_{s' \in S} \mathbf{P}(s, s') = 1$ for all $s \in S$.

For any states $s, s' \in S$, the transition probability matrix gives the probability $\mathbf{P}(s, s')$ of making a transition from s to s' . A *path* of a DTMC represents an execution of the system that it is modelling. Formally, a path of a DTMC is a non-empty, finite or infinite sequence of states $\pi = s_0 s_1 s_2 \dots$ such that $\mathbf{P}(s_i, s_{i+1}) > 0$ for all $i \geq 0$. We denote by $\pi(i)$ the state s_i . For a finite path π_{fin} , we let $|\pi_{fin}|$ denote the length of π_{fin} (i.e. number of transitions), and let $last(\pi_{fin})$ be its final state. Finally, $\pi^{(i)}$ denotes the *prefix* of length i of π , i.e. the finite path $\pi(0) \dots \pi(i)$. The sets of all finite and infinite paths starting in state s are denoted $Path_{fin}(s)$ and $Path(s)$, respectively.

To reason about the probabilistic behaviour of the DTMC, we need to determine the probability that certain paths are taken. This is achieved by defining, for each state $s \in S$, a probability space $(Path(s), \mathcal{F}_s, Prob_s)$ over the (infinite) paths from s . Below, we give an outline of this construction. For further details, see [27].

The probability measure $Prob_s$ is induced by the transition probability matrix \mathbf{P} as follows. First, for any finite path $\pi_{fin} \in Path_{fin}(s)$ of length n , we define the probability $\mathbf{P}_s(\pi_{fin})$:

$$\mathbf{P}_s(\pi_{fin}) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } n = 0 \\ \mathbf{P}(\pi_{fin}(0), \pi_{fin}(1)) \cdots \mathbf{P}(\pi_{fin}(n-1), \pi_{fin}(n)) & \text{otherwise.} \end{cases}$$

Next, we define the *cylinder set* of π_{fin} , denoted $cyl(\pi_{fin})$, as follows:

$$cyl(\pi_{fin}) \stackrel{\text{def}}{=} \{ \pi \in Path(s) \mid \pi_{fin} \text{ is a prefix of } \pi \}$$

i.e. the cylinder set $cyl(\pi_{fin})$ is the set of all infinite paths which have π_{fin} as a prefix. We denote by $Cyl(s)$ the family of cylinder sets $cyl(\pi_{fin})$, where π_{fin} ranges over the elements of $Path_{fin}(s)$. Then, let \mathcal{F}_s be the smallest σ -field on $Path(s)$ that contains the family of sets $Cyl(s)$. Due to properties of $Cyl(s)$, we can then define the probability measure on \mathcal{F}_s as the unique measure which extends the function $Prob_s : Cyl(s) \rightarrow [0, 1]$ where $Prob_s(cyl(\pi_{fin})) = \mathbf{P}_s(\pi_{fin})$ for all $\pi_{fin} \in Path_{fin}(s)$.

2.3 Markov Decision Processes

Markov decision processes (MDPs) are a natural representation for the modelling and analysis of systems with both probabilistic and nondeterministic behaviour.

Definition 4 A *Markov decision process* is a tuple $M = (S, s_{init}, Steps, \mathbf{r})$, where

- S is a set of states;
- $s_{init} \in S$ is an initial state;
- $Steps : S \rightarrow 2^{\text{Dist}(S)}$ is a *probabilistic transition function*;
- $\mathbf{r} : S \times \text{Dist}(S) \rightarrow \mathbb{R}_{\geq 0}$ is a *reward function*.

A probabilistic transition $s \xrightarrow{\mu} s'$ is made from a state s by first nondeterministically selecting a distribution $\mu \in Steps(s)$ and then making a probabilistic choice of target state s' according to the distribution μ . As for a DTMC, a *path* of an MDP corresponds to an execution of the system that it is modelling. It represents a particular resolution of both nondeterminism and probability. Formally, a path of an MDP is a non-empty, finite or infinite sequence of probabilistic transitions:

$$\pi = s_0 \xrightarrow{\mu_0} s_1 \xrightarrow{\mu_1} s_2 \xrightarrow{\mu_2} \dots$$

such that $\mu_i(s_{i+1}) > 0$ for all $i \geq 0$. We denote by $\pi(i)$ the state s_i and by $step(\pi, i)$ the distribution μ_i . We extend the notions of length, final state and prefix of paths from DTMCs to MDPs in the standard manner.

In contrast to a path, an *adversary* (sometimes also known as a *policy*, *strategy* or *scheduler*) represents a particular resolution of nondeterminism *only*. An adversary resolves the nondeterministic choice in a state of an MDP, based on the history of its execution up to that point. More precisely, an adversary is a function mapping every finite path π_{fin} to a distribution $\mu \in Steps(last(\pi_{fin}))$. For any state $s \in S$ and adversary A , let $Path_{fin}^A(s)$ and $Path^A(s)$ denote the sets of finite and infinite paths starting in s that correspond to A .

Definition 5 An adversary A is called *memoryless* (or *simple*) if, for any finite paths π_{fin} and π'_{fin} for which $last(\pi_{fin}) = last(\pi'_{fin})$, we have $A(\pi_{fin}) = A(\pi'_{fin})$.

The behaviour of an MDP from a state s , under a given adversary A , is purely probabilistic and can be viewed as a DTMC whose states are given by the set of finite paths $Path_{fin}^A(s)$. Formally, the behaviour can be described by the infinite-state DTMC $D_s^A = (S_s^A, s, \mathbf{P}_s^A)$, where:

- $S_s^A = Path_{fin}^A(s)$;
- for any two finite paths $\pi_{fin}, \pi'_{fin} \in S_s^A$ we have

$$\mathbf{P}_s^A(\pi_{fin}, \pi'_{fin}) = \begin{cases} \mu(s') & \text{if } \pi'_{fin} \text{ is of the form } \pi_{fin} \xrightarrow{a, \mu} s' \text{ and } A(\pi_{fin}) = (a, \mu) \\ 0 & \text{otherwise.} \end{cases}$$

There is a one-to-one correspondence between the infinite paths of the DTMC D_s^A and $Path^A(s)$. It therefore follows that, using the probability measure construction for DTMC given in Section 2.2, we can construct a probability space $(Path^A(s), \mathcal{F}_s^A, Prob_s^A)$. Based on this construction, we now introduce two quantitative measures for MDPs which together form the basis for probabilistic model checking of MDPs [1, 6].

The first measure is *probabilistic reachability*, which refers to the probability of reaching a set of target states. For adversary A , the probability of reaching the target set $F \subseteq S$ from state s is given by:

$$p_s^A(F) \stackrel{\text{def}}{=} Prob_s^A \{ \pi \in Path^A(s) \mid \exists i \in \mathbb{N}. \pi(i) \in F \}.$$

The second measure we consider is *expected reachability*, which refers to the expected reward accumulated before reaching a set of target states. For an adversary A , the expected reward of reaching the target set F from state s , denoted $e_s^A(F)$, is defined as the expectation [8], with respect to the probability measure $Prob_s^A$, of the random variable $r(F, \cdot) : Path^A(s) \rightarrow \mathbb{R}$ which returns, for a given path, the total reward accumulated until a state in F is reached along the path. Formally we have:

$$e_s^A(F) \stackrel{\text{def}}{=} \int_{\pi \in Path^A(s)} r(F, \pi) dProb_s^A$$

where for any path $\pi \in Path^A(s)$:

$$r(F, \pi) \stackrel{\text{def}}{=} \begin{cases} \sum_{i=1}^{\min\{j \mid \pi(j) \in F\}} \mathbf{r}(\pi(i-1), step(\pi, i-1)) & \text{if } \exists j \in \mathbb{N}. \pi(j) \in F \\ \infty & \text{otherwise.} \end{cases}$$

For simplicity, we have defined the reward of a path that does not reach F to be ∞ , even though the total reward of the path may not be infinite. Essentially, this means that the expected reward of reaching F from s under A is finite if and only if, under the adversary A , a state in F is reached from s with probability 1.

A natural way to reason about the behaviour of an MDP is to consider the *minimum* and *maximum* values of these measures, quantifying over all adversaries.

Definition 6 For an MDP $M = (S, s_{init}, Steps, \mathbf{r})$, set of target states $F \subseteq S$ and state $s \in S$, the *minimum and maximum reachability probabilities* of reaching F from s equal:

$$p_s^{\min}(F) = \inf_A p_s^A(F) \quad \text{and} \quad p_s^{\max}(F) = \sup_A p_s^A(F)$$

and the *minimum and maximum expected rewards* of reaching F from s equal:

$$e_s^{\min}(F) = \inf_A e_s^A(F) \quad \text{and} \quad e_s^{\max}(F) = \sup_A e_s^A(F).$$

We write, for example, $\mathbf{p}^{\min}(F)$ to denote the vector of values $p_s^{\min}(F)$ over all states $s \in S$. More generally, for any vector \mathbf{x} over S , we write x_s to denote the element for state $s \in S$.

Computing values for probabilistic and expected reachability reduces to the *stochastic shortest path problem* for Markov decision processes; see for example [7, 2]. A key result in this respect is that optimality with respect to probabilistic and expected reachability can always be achieved with memoryless adversaries (see Definition 5). A consequence of this is that these quantities can be computed through a method known as *value iteration* [7, 2], which iteratively computes increasingly precise approximations to the required values. The lemma below forms the basis for value iteration.

Lemma 1 Consider an MDP $M = (S, s_{init}, Steps, \mathbf{r})$ and set of target states $F \subseteq S$.

- The sequences of vectors $\langle \mathbf{p}^{\min, (n)} \rangle_{n \in \mathbb{N}}$ and $\langle \mathbf{p}^{\max, (n)} \rangle_{n \in \mathbb{N}}$ converge to the vectors $\mathbf{p}^{\min}(F)$ and $\mathbf{p}^{\max}(F)$, where for any state $s \in S$:
 - if $s \in F$, then $p_s^{\min, (n)} = p_s^{\max, (n)} = 1$ for all $n \in \mathbb{N}$;
 - if $s \notin F$, then:

$$p_s^{\min, (n)} = \begin{cases} 0 & \text{if } n = 0 \\ \min_{\mu \in Steps(s)} \sum_{s' \in S} \mu(s') \cdot p_{s'}^{\min, (n-1)} & \text{otherwise} \end{cases}$$

$$p_s^{\max, (n)} = \begin{cases} 0 & \text{if } n = 0 \\ \max_{\mu \in Steps(s)} \sum_{s' \in S} \mu(s') \cdot p_{s'}^{\max, (n-1)} & \text{otherwise.} \end{cases}$$

- The sequences of vectors $\langle \mathbf{e}^{\min,(n)} \rangle_{n \in \mathbb{N}}$ and $\langle \mathbf{e}^{\max,(n)} \rangle_{n \in \mathbb{N}}$ converge to the vectors $\mathbf{e}^{\min}(F)$ and $\mathbf{e}^{\max}(F)$, where for any state $s \in S$:
 - if $s \in F$, then $e_s^{\min,(n)} = e_s^{\max,(n)} = 0$ for all $n \in \mathbb{N}$;
 - if $s \notin F$, then:

$$e_s^{\min,(n)} = \begin{cases} \infty & \text{if } p_s^{\max} < 1 \\ 0 & \text{if } p_s^{\max} = 1 \text{ and } n = 0 \\ \min_{\mu \in \text{Steps}(s)} \left(\mathbf{r}(s, \mu) + \sum_{s' \in S} \mu(s') \cdot e_{s'}^{\min,(n-1)} \right) & \text{otherwise} \end{cases}$$

$$e_s^{\max,(n)} = \begin{cases} \infty & \text{if } p_s^{\min} < 1 \\ 0 & \text{if } p_s^{\min} = 1 \text{ and } n = 0 \\ \max_{\mu \in \text{Steps}(s)} \left(\mathbf{r}(s, \mu) + \sum_{s' \in S} \mu(s') \cdot e_{s'}^{\max,(n-1)} \right) & \text{otherwise.} \end{cases}$$

In practice, value iteration is carried out by computing the vector of required values for increasingly large values of n , until a pre-specified convergence criterion is met. One common approach is to check that the maximum difference between the corresponding elements of successive vectors is below some fixed threshold δ . Another is to use the maximum relative difference of vector elements.

2.4 Stochastic Two-Player Games

Finally in this section, we describe (simple) stochastic games [40, 14], which are turn-based games involving two players and chance.

Definition 7 A stochastic two-player game is a tuple $G = ((V, E), v_{\text{init}}, (V_1, V_2, V_p), \delta, \bar{\mathbf{r}})$ where:

- (V, E) is a finite directed graph;
- $v_{\text{init}} \in V$ is an initial vertex;
- (V_1, V_2, V_p) is a partition of V ;
- $\delta : V_p \rightarrow \text{Dist}(V)$ is a *probabilistic transition function*;
- $\bar{\mathbf{r}} : E \rightarrow \mathbb{R}_{\geq 0}$ is a *reward function* over edges.

Vertices in V_1 , V_2 and V_p are called ‘player 1’, ‘player 2’ and ‘probabilistic’ vertices, respectively. The game proceeds in steps, moving from a vertex v to one of its neighbouring vertices v' in the game graph. The choice of v' depends on the type of the vertex v . If $v \in V_1$ then player 1 chooses v' , if $v \in V_2$ then player 2 makes the choice, and if $v \in V_p$ then v' is selected randomly according to the distribution $\delta(v)$. A Markov decision process can thus be thought of as a stochastic game in which there are no player 2 vertices and where there is a strict alternation between player 1 and probabilistic vertices.

A *play* in a game G is a sequence of vertices $\omega = v_0 v_1 v_2 \dots$ such that $(v_i, v_{i+1}) \in E$ for all $i \in \mathbb{N}$. We denote by $\omega(i)$ the vertex v_i and, for a finite play ω_{fin} , we write $\text{last}(\omega_{\text{fin}})$ for the final vertex of ω_{fin} and $|\omega_{\text{fin}}|$ for its length (the number of transitions). The prefix of length i of play ω is denoted $\omega^{(i)}$.

A *strategy* for player 1 is a function $\sigma_1 : V^* V_1 \rightarrow \text{Dist}(V)$, i.e. a function from the set of finite plays ending in a player 1 vertex to the set of distributions over vertices, such that for any $\omega_{\text{fin}} \in V^* V_1$ and $v \in V$, if $\sigma_1(\omega_{\text{fin}})(v) > 0$, then $(\text{last}(\omega_{\text{fin}}), v) \in E$. Strategies for player 2, denoted by σ_2 , are defined analogously. For a fixed pair of strategies (σ_1, σ_2) we denote by $\text{Play}_{\text{fin}}^{\sigma_1, \sigma_2}(v)$ and $\text{Play}^{\sigma_1, \sigma_2}(v)$ the set of finite and infinite plays starting in vertex v that correspond to these strategies. For strategy pair (σ_1, σ_2) , the behaviour of the

game is completely random and, for any vertex v , we can construct a probability space $(\text{Play}^{\sigma_1, \sigma_2}(v), \mathcal{F}_v^{\sigma_1, \sigma_2}, \text{Prob}_v^{\sigma_1, \sigma_2})$. This construction proceeds in similar fashion to the one described for MDPs in the previous section.

We can then, as for MDPs, define the notions of *probabilistic reachability* and *expected reachability*. For a fixed strategy pair (σ_1, σ_2) , vertex $v \in V$ and a set of target vertices $\bar{F} \subseteq V$, we define the corresponding probability of reaching \bar{F} as:

$$p_v^{\sigma_1, \sigma_2}(\bar{F}) \stackrel{\text{def}}{=} \text{Prob}_v^{\sigma_1, \sigma_2} \{ \omega \in \text{Play}^{\sigma_1, \sigma_2}(v) \mid \exists i \in \mathbb{N} \wedge \omega(i) \in \bar{F} \}$$

and the expected reward of reaching \bar{F} as:

$$e_v^{\sigma_1, \sigma_2}(\bar{F}) \stackrel{\text{def}}{=} \int_{\omega \in \text{Play}^{\sigma_1, \sigma_2}(v)} \bar{r}(\bar{F}, \omega) \, d\text{Prob}_v^{\sigma_1, \sigma_2}$$

where for any play $\omega \in \text{Play}^{\sigma_1, \sigma_2}(v)$:

$$\bar{r}(\bar{F}, \omega) \stackrel{\text{def}}{=} \begin{cases} \sum_{i=1}^{\min\{j \mid \omega(j) \in \bar{F}\}} \bar{\mathbf{r}}(\omega(i-1), \omega(i)) & \text{if } \exists j \in \mathbb{N}. \omega(j) \in \bar{F} \\ \infty & \text{otherwise.} \end{cases}$$

Definition 8 For a game $G = ((V, E), v_{\text{init}}, (V_1, V_2, V_p), \delta, \bar{\mathbf{r}})$, set of target vertices $\bar{F} \subseteq V$ and vertex $v \in V$, the optimal probabilities of the game for player 1 and player 2, with respect to \bar{F} and v , are defined as follows:

$$\sup_{\sigma_1} \inf_{\sigma_2} p_v^{\sigma_1, \sigma_2}(\bar{F}) \quad \text{and} \quad \sup_{\sigma_2} \inf_{\sigma_1} p_v^{\sigma_1, \sigma_2}(\bar{F})$$

and the optimal expected rewards are:

$$\sup_{\sigma_1} \inf_{\sigma_2} e_v^{\sigma_1, \sigma_2}(\bar{F}) \quad \text{and} \quad \sup_{\sigma_2} \inf_{\sigma_1} e_v^{\sigma_1, \sigma_2}(\bar{F}).$$

A player 1 strategy σ_1 is optimal from vertex v with respect to the probability of reaching target set \bar{F} if:

$$\inf_{\sigma_2} p_v^{\sigma_1, \sigma_2}(\bar{F}) = \sup_{\sigma_1} \inf_{\sigma_2} p_v^{\sigma_1, \sigma_2}(\bar{F}). \quad (1)$$

The optimal strategies for player 2 and for expected rewards can be defined analogously. We now summarise a number of results from [10, 14, 15].

Definition 9 A strategy σ_i is *pure* if it does not use randomisation, that is, for any finite play ω_{fin} such that $\text{last}(\omega_{\text{fin}}) \in V_i$, there exists $v' \in V$ such that $\sigma_i(\omega_{\text{fin}})(v') = 1$. A strategy σ_i is *memoryless* if its choice depends only on the current vertex, that is, $\sigma_i(\omega_{\text{fin}}) = \sigma_i(\omega'_{\text{fin}})$ for any finite plays ω_{fin} and ω'_{fin} such that $\text{last}(\omega_{\text{fin}}) = \text{last}(\omega'_{\text{fin}})$.

As is the case with MDPs, for any stochastic game, the family of pure memoryless strategies suffices for optimality with respect to probabilistic and expected reachability. Moreover, these values can be computed through an iterative processes known as value iteration [15].

Lemma 2 Consider a stochastic game $G = ((V, E), v_{\text{init}}, (V_1, V_2, V_p), \delta, \bar{\mathbf{r}})$ and set of target vertices $\bar{F} \subseteq V$.

- The sequence of vectors $\langle \mathbf{p}^{(n)} \rangle_{n \in \mathbb{N}}$ converges to the optimal probabilities for player 1 with respect to the target set \bar{F} , where for any vertex $v \in V$:
 - if $v \in \bar{F}$, then $p_v^{(n)} = 1$ for all $n \in \mathbb{N}$;

– and otherwise:

$$p_v^{(n)} = \begin{cases} 0 & \text{if } n = 0 \\ \max_{(v,v') \in E} p_{v'}^{(n-1)} & \text{if } n > 0 \text{ and } v \in V_1 \\ \min_{(v,v') \in E} p_{v'}^{(n-1)} & \text{if } n > 0 \text{ and } v \in V_2 \\ \sum_{v' \in V} \delta(v)(v') \cdot p_{v'}^{(n-1)} & \text{if } n > 0 \text{ and } v \in V_p. \end{cases}$$

- The sequence of vectors $\langle \mathbf{e}^{(n)} \rangle_{n \in \mathbb{N}}$ converges to the optimal expected rewards for player 1 with respect to the target set \bar{F} , where for any vertex $v \in V$:
 - if $v \in \bar{F}$, then $e_v^{(n)} = 0$ for all $n \in \mathbb{N}$;
 - if $\sup_{\sigma_2} \inf_{\sigma_1} p_v^{\sigma_1, \sigma_2}(\bar{F}) < 1$, then $e_v^{(n)} = \infty$ for all $n \in \mathbb{N}$;
 - and otherwise:

$$e_v^{(n)} = \begin{cases} 0 & \text{if } n = 0 \\ \max_{(v,v') \in E} (\bar{\mathbf{r}}(v, v') + e_{v'}^{(n-1)}) & \text{if } n > 0 \text{ and } v \in V_1 \\ \min_{(v,v') \in E} (\bar{\mathbf{r}}(v, v') + e_{v'}^{(n-1)}) & \text{if } n > 0 \text{ and } v \in V_2 \\ \sum_{v' \in V} (\bar{\mathbf{r}}(v, v') + \delta(v)(v') \cdot e_{v'}^{(n-1)}) & \text{if } n > 0 \text{ and } v \in V_p. \end{cases}$$

Lemma 2 forms the basis of an iterative method to compute the vector of optimal values for a game. Note that, although this concerns only the optimal probability for player 1, similar results hold for player 2. Observe the similarity between this and the value iteration method for MDP solution described in Lemma 1.

3 Abstraction for Markov Decision Processes

We now present our notion of *abstraction* for MDPs, which is based on stochastic two-player games. The key idea is to maintain a distinction between the nondeterminism in an MDP and the nondeterminism introduced by abstracting it. The abstraction of an MDP is a stochastic two-player game in which the choices made by player 1 correspond to the nondeterminism from the process of abstraction and the choices made by player 2 correspond to the nondeterminism in the original MDP.

For an MDP $M = (S, s_{init}, Steps, \mathbf{r})$, which we refer to as the *concrete model*, we construct an abstraction based on a partition $P = \{S_1, S_2, \dots, S_n\}$ of its state space S . The elements of P are referred to as *abstract states*, each comprising a set of *concrete states*. Intuitively, in the stochastic game representing the abstraction, player 1 vertices are abstract states and player 2 vertices correspond to the sets of nondeterministic choices in individual concrete states.

In order to formally define the abstraction process, we need a few preliminary definitions. For any distribution μ over S , we denote by $\bar{\mu}$ the probability distribution over P lifted from μ , i.e. $\bar{\mu}(S_i) = \sum_{s \in S_i} \mu(s)$ for all $S_i \in P$. In addition, for any state s , we denote by $\overline{Steps(s)}$ the set of reward-distribution pairs $\{(\mathbf{r}(s, \mu), \bar{\mu}) \mid \mu \in Steps(s)\}$. The inclusion of rewards in the definition of $\overline{Steps(s)}$ ensures that the reward values associated with probability distributions in the concrete model are preserved, even in the case where multiple (concrete) distributions are lifted to the same distribution over abstract states. This differs from [28] where, to simplify the presentation, we assumed that distributions that were lifted to the same abstract distribution had identical reward values.

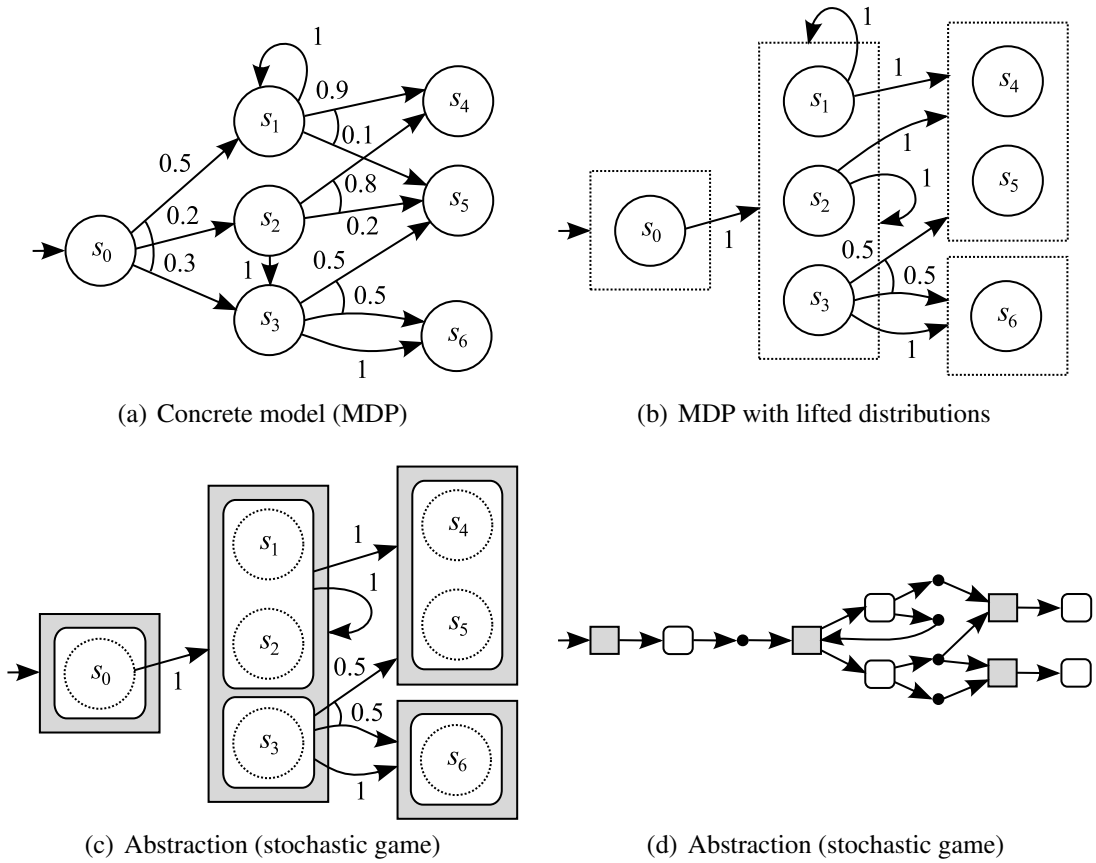


Fig. 1 Simple example of the abstraction process for Markov decision processes

Definition 10 Given an MDP $M = (S, s_{init}, Steps, \mathbf{r})$ and a partition $P = \{S_1, \dots, S_n\}$ of the state space S , we define the corresponding abstraction G_M^P as the stochastic game:

$$G_M^P = ((V, E), v_{init}, (V_1, V_2, V_p), \delta, \bar{\mathbf{r}})$$

in which:

- $V_1 = P$;
- $V_2 = \{v_2 \subseteq \mathbb{R}_{\geq 0} \times \text{Dist}(P) \mid v_2 = \overline{Steps(s)} \text{ for some } s \in S\}$;
- $V_p = \{v_p \in \mathbb{R}_{\geq 0} \times \text{Dist}(P) \mid v_p \in Steps(s) \text{ for some } s \in S\}$;
- $v_{init} = S_i$ where $s_{init} \in S_i$;
- $(v, v') \in E$ if and only if one of the following conditions holds:
 - $v \in V_1, v' \in V_2$ and $v' = \overline{Steps(s)}$ for some $s \in v$;
 - $v \in V_2, v' \in V_p$ and $v' \in v$;
 - $v \in V_p, v' \in V_1$ and $v = (r, \bar{\mu})$ such that $\bar{\mu}(v') > 0$;
- $\delta : V_p \rightarrow \text{Dist}(V)$ projects $(r, \bar{\mu}) \in V_p$ onto $\bar{\mu}$;
- $\bar{\mathbf{r}} : E \rightarrow \mathbb{R}_{\geq 0}$ is the reward function such that for any $(v, v') \in V \times V$:

$$\bar{\mathbf{r}}(v, v') = \begin{cases} r & \text{if } (v, v') \in V_2 \times V_p, \text{ and } v' = (r, \bar{\mu}) \text{ for some } \bar{\mu} \in \text{Dist}(P) \\ 0 & \text{otherwise.} \end{cases}$$

Example 1 We illustrate the abstraction process on a simple example, shown in Figure 1. The concrete model (the MDP) is given in Figure 1(a): states are drawn as circles and probability distributions as labelled arrows, grouped with arcs. For clarity, we omit rewards in

this case. We use the partition $P = \{\{s_0\}, \{s_1, s_2, s_3\}, \{s_4, s_5\}, \{s_6\}\}$. The abstraction process is illustrated in two steps. Firstly, Figure 1(b) shows the partition P , as dotted lines around groups of states, and the result of lifting each distribution μ in the MDP to $\bar{\mu}$. Secondly, Figure 1(c) shows the stochastic game representing the corresponding abstraction. Player 1 vertices are drawn as grey rectangles, each corresponding to an element of the partition P . The player 2 vertices that are successors of each player 1 vertex are drawn as white shapes, enclosed within it. The concrete states corresponding to each player 2 vertex are also depicted inside the player 2 vertices. These are states from the MDP whose outgoing choices are identical after lifting. Probabilistic vertices are not drawn explicitly, but rather shown as groups of probabilistic transitions emanating from each player 2 vertex, in the same style as the MDP. For completeness, we also show, in Figure 1(d), the same stochastic game with all vertices illustrated explicitly. Player 1, player 2 and probabilistic vertices are coloured grey, white and black, respectively.

Intuitively, the roles of the vertices and players in the abstraction can be understood as follows. A player 1 vertex corresponds to an abstract state: an element of the partition P of the states from the original MDP. Player 1 chooses a concrete state from this set and then player 2 chooses a probability distribution from those available in the concrete state (which is now a distribution over abstract states rather than concrete states).

This description perhaps misleadingly gives the impression that the stochastic game representing the abstraction is no smaller than the original concrete model. This is not the case. Firstly, note that V_2 vertices are actually sets of reward-distribution pairs that correspond to concrete states, not the concrete states themselves. Hence, concrete states with the same reward values and outgoing distributions (after lifting) are collapsed into one player 2 vertex (see, for example, s_1 and s_2 in Figure 1(c)). Furthermore, in practice there is no need to store the entire vertex set V of the stochastic game. Since we have a strict alternation between V_1 , V_2 and V_p vertices, we need only store the vertices in V_1 , the outgoing transitions comprising each reward-distribution pair from V_1 , and how these pairs are grouped (into elements of V_2). Later, in Example 3 and Section 5 we will show how, on all the case studies considered, the abstraction process brings a significant reduction in model size.

3.1 Analysis of the Abstraction

In this section we describe how, for an MDP M and partition P , the stochastic game G_M^P representing the abstraction of M yields lower and upper bounds for probabilistic reachability and expected reachability properties of M . To ease notation we let, for $x \in \{p, e\}$ and $\bar{F} \subseteq V$:

$$\begin{aligned} x_v^{lb, \min}(\bar{F}) &\stackrel{\text{def}}{=} \inf_{\sigma_1, \sigma_2} x_v^{\sigma_1, \sigma_2}(\bar{F}) \\ x_v^{ub, \min}(\bar{F}) &\stackrel{\text{def}}{=} \sup_{\sigma_1} \inf_{\sigma_2} x_v^{\sigma_1, \sigma_2}(\bar{F}) \\ x_v^{lb, \max}(\bar{F}) &\stackrel{\text{def}}{=} \sup_{\sigma_2} \inf_{\sigma_1} x_v^{\sigma_1, \sigma_2}(\bar{F}) \\ x_v^{ub, \max}(\bar{F}) &\stackrel{\text{def}}{=} \sup_{\sigma_1, \sigma_2} x_v^{\sigma_1, \sigma_2}(\bar{F}). \end{aligned}$$

The values $x_v^{ub, \min}(\bar{F})$ and $x_v^{lb, \max}(\bar{F})$ are the optimal reachability probabilities (if $x = p$) and expected rewards (if $x = e$) for players 1 and 2, respectively. As we have seen in Lemma 2, these optimal values can be generated using a simple iterative computation. This process can also be used to construct a pair of (memoryless) strategies that result in these optimal values [15]. The other values, $x_v^{lb, \min}(\bar{F})$ and $x_v^{ub, \max}(\bar{F})$, although not usually considered for

stochastic games (because the players co-operate), can be computed similarly by considering the game as an MDP [7] (see Lemma 1). The following theorem shows how these four values represent bounds on the corresponding properties of the MDP.

Theorem 1 *Let $M = (S, s_{init}, Steps, \mathbf{r})$ be an MDP, P a partition of S and $F \in P$ a set of target states. If $G_M^P = ((V, E), v_{init}, (V_1, V_2, V_p), \delta, \bar{\mathbf{r}})$ is the stochastic game constructed from M and P according to Definition 10, then for any $x \in \{e, p\}$ and $s \in S$:*

$$\begin{aligned} x_v^{lb, \min}(\bar{F}) &\leq x_s^{\min}(F) \leq x_v^{ub, \min}(\bar{F}) \\ x_v^{lb, \max}(\bar{F}) &\leq x_s^{\max}(F) \leq x_v^{ub, \max}(\bar{F}) \end{aligned}$$

where v is the unique vertex of V_1 such that $s \in v$ and $\bar{F} = \{F\}$.

We also show that using a *finer* partition results in a *more precise* abstraction, i.e. a stochastic game providing tighter lower and upper bounds.

Definition 11 If P and P' are partitions of a state space S , then P is *finer* than P' (denoted $P \preceq P'$) if, for any $v \in P$, there exists $v' \in P'$ such that $v \subseteq v'$. Equivalently, we say that P' is *coarser* than P .

Theorem 2 *Let $M = (S, s_{init}, Steps, \mathbf{r})$ be an MDP, P and P' partitions of S such that $P \preceq P'$ and $F \in P \cap P'$ a set of target states. If $G_M^P = ((V, E), v_{init}, (V_1, V_2, V_p), \delta, \bar{\mathbf{r}})$ and $G_M^{P'} = ((V', E'), v'_{init}, (V'_1, V'_2, V'_p), \delta', \bar{\mathbf{r}}')$ are the stochastic games constructed from P and P' according to Definition 10, then for any $x \in \{e, p\}$ and $v \in P$:*

$$\begin{aligned} x_{v'}^{lb, \min}(\bar{F}) &\leq x_v^{lb, \min}(\bar{F}) & \text{and} & & x_v^{ub, \min}(\bar{F}) &\leq x_{v'}^{ub, \min}(\bar{F}) \\ x_{v'}^{lb, \max}(\bar{F}) &\leq x_v^{lb, \max}(\bar{F}) & \text{and} & & x_v^{ub, \max}(\bar{F}) &\leq x_{v'}^{ub, \max}(\bar{F}) \end{aligned}$$

where v' is the unique element of P' such that $v \subseteq v'$ and $\bar{F} = \{F\}$.

3.2 Examples

In this section, we present examples illustrating the application of the results given above. Then, in the following section, we provide proofs of Theorems 1 and 2.

Example 2 Let us return to the previous simple example (see Example 1 and Figure 1). Suppose that we are interested in the minimum probability in the original MDP (Figure 1(a)) of, starting from state s_0 , reaching the target set $\{s_6\}$. This can be computed as 0.15. From the abstraction shown in Figure 1(c) and the results of Theorem 1, we can establish that the minimum probability lies within the interval $[0, 0.5]$. If, on the other hand, the abstract model had instead been constructed as an MDP, i.e. with no discrimination between the two forms of nondeterminism, we would only have been able to determine that the minimum reachability probability lay in the interval $[0, 1]$.

Example 3 To illustrate the application of our abstraction to a larger MDP, we consider a more complex example: a model of the Zeroconf protocol [12] for dynamic self-configuration of local IP addresses within a local network. Zeroconf provides a distributed, ‘plug and play’ approach to IP address configuration, managed by the individual devices of the network. The model concerns the situation where a new device joins a network of N existing hosts, in which there are a total of M IP addresses available. For full details of the MDP model of the protocol and partition used in the abstraction process, see [28].

N	M	States (transitions)		Minimum expected time			Maximum expected time		
		Concrete model	Abstraction	lb	Exact	ub	lb	Exact	ub
4	32	26,121 (50,624)	737 (1,594)	8.1088	8.1572	8.2190	8.1231	8.2465	8.3047
5		58,497 (139,104)	785 (1,678)	8.1410	8.2035	8.2839	8.1595	8.3183	8.3950
6		145,801 (432,944)	833 (1,762)	8.1758	8.2533	8.3538	8.1988	8.3956	8.4922
7		220,513 (614,976)	857 (1,806)	8.2131	8.2939	8.4293	8.2412	8.4579	8.5972
8		432,185 (1,254,480)	881 (1,850)	8.2538	8.3379	8.5110	8.2871	8.5254	8.7109
4	64	50,377 (98,080)	737 (1,594)	8.0508	8.0733	8.1022	8.0574	8.1150	8.1422
5		113,217 (270,272)	785 (1,678)	8.0645	8.0931	8.1299	8.0730	8.1457	8.1808
6		282,185 (839,824)	833 (1,762)	8.0788	8.1136	8.1586	8.0891	8.1774	8.2206
7		426,529 (1,189,792)	857 (1,806)	8.0935	8.1289	8.1883	8.1058	8.2008	8.2619
8		838,905 (2,439,600)	881 (1,850)	8.1088	8.1448	8.2190	8.1231	8.2252	8.3047

Table 1 Zeroconf protocol (Example 3): model statistics and numerical results (expected completion time)




Fig. 2 Zeroconf protocol (Example 3): maximum probability the new host has not configured successfully by time T

Table 1 shows the size of the concrete model (MDP) and its abstraction (game) for a range of values of N and M . For each model, we compute the minimum and maximum expected time for a host to complete the protocol (i.e. to configure and start using an IP address). Table 1 shows the exact results, obtained from the MDP, and the lower and upper bounds, obtained from the abstraction. In addition, Figure 2 illustrates results for the maximum probability that the new host has not configured successfully by time T .

As stated previously, an advantage of our approach is the ability to quantify the utility of the abstraction, based on the difference between the lower and upper bounds obtained. In the case of the plots in Figure 2, for a particular time bound T this difference is indicated by the vertical distance between the curves for the lower and upper bounds at the point T on the horizontal axis. Examining these differences between bounds for the presented results, it can be seen that our abstraction approach yields tight approximations for the performance characteristics of the protocol while at the same time producing a significant reduction in both the number of states and transitions. Observe also that the size of the abstract model increases linearly, rather than exponentially, in N and is independent of M .

3.3 Correctness of the Abstraction

We now prove Theorems 1 and 2, presented in Section 3.1. Before doing so, we require a number of preliminary concepts. For the remainder of this section we fix an MDP $M = (S, s_{init}, Steps, \mathbf{r})$, partitions P and P' of S such that $P \preceq P'$ and set of target states $F \in P \cap P'$. Let $G_M^P = ((V, E), v_{init}, (V_1, V_2, V_p), \delta, \bar{\mathbf{r}})$ and $G_M^{P'} = ((V', E'), v'_{init}, (V'_1, V'_2, V'_p), \delta', \bar{\mathbf{r}}')$ be the stochastic games obtained from P and P' according to Definition 10 and let $\bar{F} = \{F\}$.

To simplify notation, we will add subscripts to distributions $\bar{\mu}$ and sets of reward-distributions pairs $\overline{Steps}(s)$, indicating the partition that was used to lift them from the concrete to the abstract state space, e.g. for partition P , $\mu \in \text{Dist}(S)$, $v \in P$ and $s \in S$, we have $\bar{\mu}_P(v) = \sum_{s \in v} \mu(s)$ and $\overline{Steps}(s)_P = \{(\mathbf{r}(s, \mu), \bar{\mu}_P) \mid \mu \in Steps(s)\}$.

We next define a mapping from vertices of G_M^P to the vertices of $G_M^{P'}$.

Definition 12 Let $[\cdot] : V \rightarrow V'$ be the mapping from vertices of G_M^P to vertices of $G_M^{P'}$ such that for any $v \in V$:

- if $v \in P$, then $[v] = v'$ where v' is the unique element of P' such that $v \subseteq v'$ (such a v' exists from the fact that $P \preceq P'$);
- if $v = \overline{Steps}(s)_P$ for some $s \in S$, then $[v] = \overline{Steps}(s)_{P'}$;
- if $v = (r, \bar{\mu}_P)$ for some $r \in \mathbb{R}_{\geq 0}$ and $\mu \in \text{Dist}(S)$, then $[v] = (r, \bar{\mu}_{P'})$.

By construction, this mapping extends naturally to plays, i.e. for any play $\omega = v_0 v_1 v_2 \dots$ of G_M^P we have $[\omega] = [v_0][v_1][v_2] \dots$ is a play of $G_M^{P'}$. We denote by $[\cdot]^{-1}$ the inverse mapping from plays of $G_M^{P'}$ to sets of plays of G_M^P , i.e. for any finite play ω' of $G_M^{P'}$ we have $[\omega']^{-1} = \{\omega \in V^* \mid [\omega] = \omega'\}$.

Lemma 3 For any play ω of G_M^P , it holds that $\bar{r}(\bar{F}, \omega) = \bar{r}'(\bar{F}, [\omega])$.

Proof The proof follows directly from Definition 12 and since we assume that F is an element of both P and P' . \square

Using this mapping between plays, for any player 1 vertex v and strategy pair $\sigma = (\sigma_1, \sigma_2)$ of the game G_M^P , we now construct a (randomised) strategy pair $[\sigma]_v = ([\sigma_1]_v, [\sigma_2]_v)$ of the game $G_M^{P'}$ such that, under σ and $[\sigma]_v$, when starting at v and $[v]$ respectively, the probability of reaching F is equal. For any finite play ω' of $G_M^{P'}$ such that $\omega'(0) = [v]$ and $last(\omega') \in V'_i$, let the probability of $[\sigma_i]_v$ selecting vertex $v' \in V'$ after play ω' equal:

$$[\sigma_i]_v(\omega')(v') \stackrel{\text{def}}{=} \frac{Prob_v^\sigma([\omega'v']_v^{-1})}{Prob_v^\sigma([\omega']_v^{-1})}$$

where $[\tilde{\omega}']_v^{-1} = \{\omega \in [\tilde{\omega}']^{-1} \mid \omega(0) = v\}$ for any play $\tilde{\omega}'$ of $G_M^{P'}$ and

$$Prob_v^\sigma(\Omega) = Prob_v^\sigma\{\omega \in Play^\sigma(v) \mid \omega^{(i)} \in \Omega \text{ for some } i \in \mathbb{N}\}$$

for any set of finite plays Ω of G_M^P .

Proposition 1 For any vertex v and strategy pair $\sigma = (\sigma_1, \sigma_2)$ of G_M^P , if Ω is a measurable set of plays of $Play^{[\sigma]_v}([v])$, then $Prob_{[v]}^{[\sigma]_v}(\Omega) = Prob_v^\sigma([\Omega]_v^{-1})$.

Proof From the measure construction (see Section 2) it is sufficient to show that:

$$Prob_{[v]}^{[\sigma]_v}(\omega') = Prob_v^\sigma[\omega']_v^{-1}$$

for any finite play ω' of $G_M^{P'}$ and vertex $v \in [\omega'(0)]^{-1}$, which we now prove by induction on the length of the play ω' . If $|\omega'| = 1$, then ω' comprises a single vertex and since $[\omega']_v^{-1} = \{v\}$, by the measure construction:

$$Prob_{[v]}^{[\sigma]_v}(\omega') = Prob_v^\sigma[\omega']_v^{-1} = 1.$$

Next, suppose by induction that the lemma holds for all plays of length n . Consider any play ω' of length $n+1$ of $G_M^{P'}$. By definition ω' is of the form $\tilde{\omega}'\tilde{v}'$ for some play $\tilde{\omega}'$ of length n and vertex $\tilde{v}' \in V'$. If $last(\tilde{\omega}') \in V'_i$ (for $i = 1, 2$), then from the measure construction:

$$\begin{aligned} Prob_{[v]}^{[\sigma]_v}(\omega') &= Prob_{[v]}^{[\sigma]_v}(\tilde{\omega}') \cdot [\sigma_i]_v(\tilde{\omega}')(\tilde{v}') \\ &= Prob_v^\sigma([\tilde{\omega}']_v^{-1}) \cdot [\sigma_i]_v(\tilde{\omega}')(\tilde{v}') && \text{by induction} \\ &= Prob_v^\sigma([\tilde{\omega}']_v^{-1}) \cdot \frac{Prob_v^\sigma([\tilde{\omega}'\tilde{v}']_v^{-1})}{Prob_v^\sigma([\tilde{\omega}']_v^{-1})} && \text{by definition of } [\sigma_i]_v \\ &= Prob_v^\sigma([\tilde{\omega}'\tilde{v}']_v^{-1}) && \text{rearranging} \\ &= Prob_v^\sigma([\omega']_v^{-1}) && \text{since } \omega' = \tilde{\omega}'\tilde{v}'. \end{aligned}$$

On the other hand, if $last(\tilde{\omega}') \in V_p'$, then $last(\tilde{\omega}') = (r, \overline{\mu_{P'}})$ for some $r \in \mathbb{R}_{\geq 0}$ and $\mu \in \text{Dist}(S)$, and hence in this case, again from the measure construction, we have:

$$\begin{aligned} Prob_{[v]}^{[\sigma]_v}(\omega') &= Prob_{[v]}^{[\sigma]_v}(\tilde{\omega}') \cdot \overline{\mu_{P'}}(\tilde{v}') \\ &= Prob_v^\sigma([\tilde{\omega}']_v^{-1}) \cdot \overline{\mu_{P'}}(\tilde{v}') && \text{by induction} \\ &= Prob_v^\sigma([\tilde{\omega}']_v^{-1}) \cdot \left(\sum_{\tilde{v}' \in [\tilde{v}']^{-1}} \overline{\mu_P}(\tilde{v}') \right) && \text{by Definition 12} \\ &= \sum_{\tilde{v}' \in [\tilde{v}']^{-1}} Prob_v^\sigma([\tilde{\omega}']_v^{-1}) \cdot \overline{\mu_P}(\tilde{v}') && \text{rearranging} \\ &= \sum_{\tilde{v}' \in [\tilde{v}']^{-1}} Prob_v^\sigma([\tilde{\omega}']_v^{-1}\tilde{v}') && \text{by definition of } Prob_v^\sigma \\ &= Prob_v^\sigma([\omega']_v^{-1}) && \text{by Definition 12 and since } \omega' = \tilde{\omega}'\tilde{v}'. \end{aligned}$$

Since these are all the cases to consider, the lemma holds by induction. \square

Before we give the proofs of Theorems 1 and 2 we require the following lemma, a classical result from measure theory and two propositions.

Lemma 4 *For any vertex v and strategy pair $\sigma = (\sigma_1, \sigma_2)$ of G_M^P , the mapping $[\cdot]$ of Definition 12 is a measurable function between $(\text{Play}^\sigma(v), \mathcal{F}_v^\sigma)$ and $(\text{Play}^{[\sigma]_v}([v]), \mathcal{F}_{[v]}^{[\sigma]_v})$.*

Proof The proof follows from measure construction (see Section 2). \square

Theorem 3 ([8]) *Let (Ω, \mathcal{F}) and (Ω', \mathcal{F}') be measurable spaces, and suppose that Pr is a measure on (Ω, \mathcal{F}) and the function $T : \Omega \rightarrow \Omega'$ is measurable. If f is a real non-negative measurable function on (Ω', \mathcal{F}') , then:*

$$\int_{\omega \in \Omega} f(T\omega) dPr = \int_{\omega' \in \Omega'} f(\omega') dPrT^{-1}.$$

Proposition 2 For any vertex v and strategy pair σ of G_M^P we have $p_v^\sigma(\bar{F}) = p_{[v]^\sigma}^\sigma(\bar{F})$ and $e_v^\sigma(\bar{F}) = e_{[v]^\sigma}^\sigma(\bar{F})$.

Proof The first equation follows from Proposition 1 using standard results from measure theory. For the second equation, by definition of the expected rewards for games (see Section 2.4), we have:

$$\begin{aligned}
e_v^\sigma(\bar{F}) &= \int_{\omega \in \text{Play}^\sigma(v)} \bar{\mathbf{r}}(\bar{F}, \omega) \, d\text{Prob}_v^\sigma \\
&= \int_{\omega \in \text{Play}^\sigma(v)} \bar{\mathbf{r}}'(\bar{F}, [\omega]) \, d\text{Prob}_v^\sigma && \text{by Lemma 3} \\
&= \int_{\omega' \in \text{Play}^{[\sigma]_v}([v])} \bar{\mathbf{r}}(\bar{F}, \omega') \, d\text{Prob}_v^\sigma([\cdot]^{-1}) && \text{by Theorem 3 and Lemma 4} \\
&= \int_{\omega' \in \text{Play}^{[\sigma]_v}([v])} \bar{\mathbf{r}}(\bar{F}, \omega') \, d\text{Prob}_{[v]^\sigma}^{[\sigma]_v} && \text{by Proposition 1} \\
&= e_{[v]^\sigma}^{[\sigma]_v}(\bar{F})
\end{aligned}$$

as required. \square

Proposition 3 For any vertex v' and player 2 strategy σ_2' of $G_M^{P'}$ there exists a player 2 strategy σ_2 of G_M^P such that

$$\begin{aligned}
\sup_{\sigma_1'} x_{v'}^{\sigma_1', \sigma_2'}(\bar{F}) &\geq \sup_{\sigma_1} x_v^{\sigma_1, \sigma_2}(\bar{F}) \\
\inf_{\sigma_1'} x_{v'}^{\sigma_1', \sigma_2'}(\bar{F}) &\leq \inf_{\sigma_1} x_v^{\sigma_1, \sigma_2}(\bar{F})
\end{aligned}$$

for any vertex v of G_M^P such that $[v] = v'$ and $x \in \{e, p\}$.

Proof Let v' be a vertex and σ_2' a player 2 strategy of $G_M^{P'}$ and let v be any vertex of G_M^P such that $[v] = v'$.

In the first step of the proof we construct the player 2 strategy σ_2 . Therefore consider any finite play ω of G_M^P that ends in a player 2 vertex. By Definition 10 we have $\text{last}(\omega) = \overline{\text{Steps}(s)}_P$ for some $s \in S$ and $[\omega]$ is a finite play of $G_M^{P'}$ ending in the player 2 vertex $\overline{\text{Steps}(s)}_{P'}$. Now $\sigma_2'([\omega]) = (\mathbf{r}(s, \mu), \overline{\mu}_{P'})$ for some $\mu \in \text{Steps}(s)$ and we let $\sigma_2(\omega) = (\mathbf{r}(s, \mu), \overline{\mu}_P)$ which by Definition 10 is an element of $\overline{\text{Steps}(s)}_P$, and hence is a valid player 2 choice in G_M^P . Although the μ meeting the above requirements is not necessarily unique, for the purposes of the proof, it is sufficient to choose any such distribution. From Definition 12 it then follows that:

$$[\omega(\sigma_2(\omega))] = [\omega]\sigma_2'([\omega]) \text{ for all finite plays } \omega \text{ of } G_M^P \text{ that end in a player 2 vertex.} \quad (2)$$

In the next step of the proof we will show that, for any player 1 strategy σ_1 of G_M^P , if $([\sigma_1]_v, [\sigma_2]_v)$ is the strategy pair of $G_M^{P'}$ constructed from the strategy pair (σ_1, σ_2) of G_M^P following the method described above, then $[\sigma_2]_v(\omega) = \sigma_2'(\omega)$ for all finite plays ω' of $G_M^{P'}$ such that $\omega'(0) = [v]$ and $\text{last}(\omega')$ is a player 2 vertex. Therefore, consider any such play ω' of $G_M^{P'}$. Now, by construction of $([\sigma_1]_v, [\sigma_2]_v)$ we have that for any $v' \in V'$:

$$[\sigma_2]_v(\omega')(v') = \frac{\text{Prob}_v^\sigma([\omega'v']_v^{-1})}{\text{Prob}_v^\sigma([\omega']_v^{-1})}$$

where $[\tilde{\omega}']_v^{-1} = \{\omega \in [\tilde{\omega}']^{-1} \mid \omega(0) = v\}$ for any play $\tilde{\omega}'$ of G_M^P and

$$Prob_v^{\sigma_1, \sigma_2}(\Omega) = Prob_v^{\sigma_1, \sigma_2} \{ \omega \in Play^{\sigma_1, \sigma_2}(v) \mid \omega^{(i)} \in \Omega \text{ for some } i \in \mathbb{N} \}$$

for any set of finite plays Ω of G_M^P . Now, from (2) we have that:

$$\{ \omega \in Play^{\sigma_1, \sigma_2}(v) \mid [\omega] = \omega' v' \} = \begin{cases} \{ \omega(\sigma_2(\omega)) \in Play^{\sigma_1, \sigma_2}(v) \mid [\omega] = \omega' \} & \text{if } v' = \sigma_2'([\omega]) \\ \emptyset & \text{otherwise.} \end{cases}$$

Combining this with the definition of $[\sigma_2]_v$, it follows that:

$$[\sigma_2]_v(\omega')(v') = \begin{cases} 1 & \text{if } v' = \sigma_2'([\omega]) \\ 0 & \text{otherwise,} \end{cases}$$

and hence, since ω' was arbitrary, we have that $[\sigma_2]_v(\omega) = \sigma_2'(\omega)$ for all finite play ω' of G_M^P such that $\omega'(0) = [v]$ and $last(\omega')$ is a player 2 vertex.

In the final step of the proof, we will use these results to show that the proposition holds. Since v was arbitrary, using Proposition 2 and the above correspondence between σ_2 and $[\sigma_2]$, we have that for any player 1 strategy σ_1 of G_M^P :

$$x_{v'}^{[\sigma_1], \sigma_2'}(\bar{F}) = x_v^{\sigma_1, \sigma_2}(\bar{F})$$

for all vertices v of G_M^P such that $[v] = v'$ and $x \in \{e, p\}$. Hence, since σ_1 was arbitrary, it follows that:

$$\begin{aligned} \sup_{\sigma_1'} x_{v'}^{\sigma_1', \sigma_2'}(\bar{F}) &\geq \sup_{\sigma_1} x_v^{\sigma_1, \sigma_2}(\bar{F}) \\ \inf_{\sigma_1'} x_{v'}^{\sigma_1', \sigma_2'}(\bar{F}) &\leq \inf_{\sigma_1} x_v^{\sigma_1, \sigma_2}(\bar{F}) \end{aligned}$$

for any vertex v of G_M^P such that $[v] = v'$ and $x \in \{e, p\}$ as required. \square

We are now in a position to give the proofs of Theorems 1 and 2.

Proof (of Theorem 2) Consider any $x \in \{e, p\}$ and $v \in P$. From Proposition 2 it follows that:

$$\begin{aligned} \inf_{\sigma_1', \sigma_2'} x_{[v]}^{\sigma_1', \sigma_2'}(\bar{F}) &\leq \inf_{\sigma_1, \sigma_2} x_v^{\sigma_1, \sigma_2}(\bar{F}) \\ \sup_{\sigma_1', \sigma_2'} x_{[v]}^{\sigma_1', \sigma_2'}(\bar{F}) &\geq \sup_{\sigma_1, \sigma_2} x_v^{\sigma_1, \sigma_2}(\bar{F}) \end{aligned}$$

and hence $x_{[v]}^{lb, \min}(\bar{F}) \leq x_v^{lb, \min}(\bar{F})$ and $x_v^{ub, \max}(\bar{F}) \leq x_{[v]}^{ub, \max}(\bar{F})$.

On the other hand, by definition we have:

$$\begin{aligned} x_v^{ub, \min}(\bar{F}) &= \sup_{\sigma_1'} \inf_{\sigma_2'} x_{[v]}^{\sigma_1', \sigma_2'}(\bar{F}) \\ &= \inf_{\sigma_2'} \sup_{\sigma_1'} x_{[v]}^{\sigma_1', \sigma_2'}(\bar{F}) && \text{by [40, 14]} \\ &\geq \inf_{\sigma_2} \sup_{\sigma_1} x_v^{\sigma_1, \sigma_2}(\bar{F}) && \text{by Proposition 3} \\ &= \sup_{\sigma_1} \inf_{\sigma_2} x_v^{\sigma_1, \sigma_2}(\bar{F}) && \text{by [40, 14]} \\ &= x_{[v]}^{ub, \min}(\bar{F}) && \text{by definition.} \end{aligned}$$

Similarly, by definition and [40, 14]:

$$\begin{aligned}
 x_{[v]}^{lb, \max}(\bar{F}) &= \sup_{\sigma'_2} \inf_{\sigma'_1} x_{[v]}^{\sigma'_1, \sigma'_2}(\bar{F}) \\
 &\geq \sup_{\sigma_2} \inf_{\sigma_1} x_v^{\sigma_1, \sigma_2}(\bar{F}) && \text{by Proposition 3} \\
 &= x_v^{lb, \max}(\bar{F}) && \text{by definition and [40, 14]}
 \end{aligned}$$

which completes the proof.

Proof (of Theorem 1) The proof follows from Theorem 2 together with the following facts.

- The partition $P = \{\{s\} \mid s \in S\}$ is finer than all other partitions.
- If G_M^P is constructed from partition $P = \{\{s\} \mid s \in S\}$ according to Definition 10, then:

$$\begin{aligned}
 x_{\{s\}}^{lb, \min}(\bar{F}) &= x_s^{\min}(F) = x_{\{s\}}^{ub, \min}(\bar{F}) \\
 x_{\{s\}}^{lb, \max}(\bar{F}) &= x_s^{\max}(F) = x_{\{s\}}^{ub, \max}(\bar{F})
 \end{aligned}$$

for all $s \in S$ and $x \in \{e, p\}$. □

4 An Abstraction-Refinement Framework for Markov Decision Processes

We now present a framework for automatically generating abstractions of MDPs, which uses the game-based abstraction method of Section 3. This is inspired by the success of counterexample-guided abstraction-refinement (CEGAR) techniques, developed for non-probabilistic verification, which iteratively *refine* an abstraction based on results from model checking. Probabilistic verification, however, is typically *quantitative* in nature: probabilistic model checking of MDPs essentially involves computing minimum or maximum probabilistic or expected reachability measures. Thus, we propose a *quantitative* abstraction-refinement process.

The construction and analysis of a game-based abstraction of an MDP yields lower and upper bounds for these values (plus corresponding strategies). The difference between these bounds can be seen as a measure of the quality of the abstraction. If this difference is too high (say, above some error ϵ), then it is desirable to refine the abstraction to reduce the difference. This provides the basis for a quantitative abstraction-refinement loop, illustrated in Figure 3: starting with a simple, coarse abstraction, we refine the abstraction until the difference between the bounds is below ϵ . In fact we can consider different criteria for termination of the process: for example checking that the difference is below ϵ for a single abstract state or set of abstract states, such as the abstract states containing an initial state.

4.1 Refinement strategies

In our context, refinement corresponds to replacing the partition used to build the abstraction with a finer partition. We propose two methods, called *strategy-based* and *value-based* refinement. The first examines the difference between strategy pairs that yield the lower and upper bounds. The motivation for this is that, since the bounds are different and the actual value for the concrete model falls between the bounds, one (or both) of the strategy pairs must make choices that are not valid in the concrete model (this can be seen as analogous

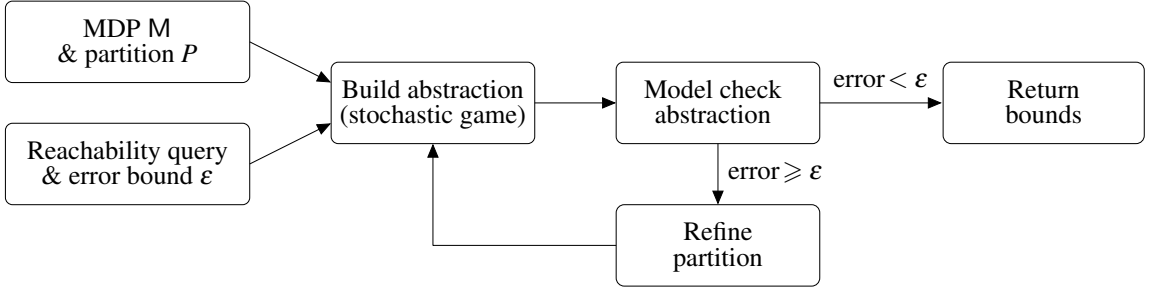


Fig. 3 A quantitative abstraction-refinement framework for MDPs

to non-probabilistic abstraction refinement, which is based on a counterexample trace). The second method differs by considering all strategy pairs yielding the bounds.

For both methods, we demonstrate that the refinement results in a strictly finer partition. For finite-state models, this guarantees that, for any user-specified error bound ϵ , the abstraction-refinement loop eventually terminates. In addition, if a partition is coarser than *probabilistic bisimulation* [31], i.e. bisimilar states are in the same element of the partition, then so is the partition after refinement. This follows from the fact that, under such a partition, bisimilar states are represented by the same player 2 vertex and, as will be seen, neither refinement method separates such states. Consequently, if the initial partition is coarser than probabilistic bisimulation, then any subsequent refinement will also be coarser. This observation guarantees that the refinement process also terminates when it is applied to infinite-state MDPs that have a finite bisimulation quotient.

To simplify presentation, we describe the refinement step when computing *minimum* reachability probabilities, i.e. we assume that we have MDP $M = (S, s_{init}, Steps, \mathbf{r})$, partition P , target set $F \in P$ and corresponding game $G_M^P = ((V, E), v_{init}, (V_1, V_2, V_p), \delta, \bar{F})$ in which $p_v^{lb, \min}(\bar{F}) < p_v^{ub, \min}(\bar{F})$ for some $v \in V_1$. The cases for maximum probabilistic and expected reachability follow in identical fashion.

Strategy-based Refinement. As mentioned in Section 2, when computing $p_v^{lb, \min}(\bar{F})$ and $p_v^{ub, \min}(\bar{F})$, a pair of strategies that obtain these values is also generated. Since a player 1 strategy's choice can be considered as a set of concrete states (the states whose choices are abstracted to the chosen player 2 vertex), the player 1 strategy from each pair provides a method for refinement. More precisely, we refine P as follows.

1. Find memoryless strategy pairs $(\sigma_1^{lb}, \sigma_2^{lb})$, $(\sigma_1^{ub}, \sigma_2^{ub})$ such that for any $v \in V$:

$$p_v^{\sigma_1^{lb}, \sigma_2^{lb}}(\bar{F}) = p_v^{lb, \min}(\bar{F}) \quad \text{and} \quad p_v^{\sigma_1^{ub}, \sigma_2^{ub}}(\bar{F}) = p_v^{ub, \min}(\bar{F}).$$

2. For each player 1 vertex $v \in V_1$ such that $\sigma_1^{lb}(v) \neq \sigma_1^{ub}(v)$ replace v in the partition P with $v_{lb}^{\sigma_1}$, $v_{ub}^{\sigma_1}$ and $v \setminus (v_{lb}^{\sigma_1} \cup v_{ub}^{\sigma_1})$ where:

$$v_{lb}^{\sigma_1} = \{s \in v \mid \sigma_1^{lb}(v) = \overline{Steps(s)}\} \quad \text{and} \quad v_{ub}^{\sigma_1} = \{s \in v \mid \sigma_1^{ub}(v) = \overline{Steps(s)}\}.$$

The following states that there always exists a vertex v such that $\sigma_1^{lb}(v) \neq \sigma_1^{ub}(v)$, and hence applying this refinement will always (strictly) refine the partition.

Lemma 5 *If there exists a vertex $v' \in V_1$ such that $p_{v'}^{lb, \min}(\bar{F}) < p_{v'}^{ub, \min}(\bar{F})$ and $(\sigma_1^{lb}, \sigma_2^{lb})$ and $(\sigma_1^{ub}, \sigma_2^{ub})$ are corresponding memoryless optimal strategy pairs, then there also exists a vertex $v \in V_1$ such that $\sigma_1^{lb}(v) \neq \sigma_1^{ub}(v)$.*

Value-based Refinement. The second refinement method is again based on the choices made by optimal strategy pairs. However, rather than looking at a single strategy pair for each of the bounds, we use all strategies that achieve one of the bounds when refining the partition. More precisely, letting $v_2^s = \overline{Steps}(s)$, we refine each element $v \in P$ as follows.

1. Construct the following subsets of v :

$$\begin{aligned} v_{lb}^{\min} &= \{s \in v \mid p_{v_2^s}^{lb, \min}(\overline{F}) = p_v^{lb, \min}(\overline{F})\} \\ v_{ub}^{\min} &= \{s \in v \mid p_{v_2^s}^{ub, \min}(\overline{F}) = p_v^{ub, \min}(\overline{F})\}. \end{aligned}$$

2. Replace v with $v_{lb}^{\min} \setminus v_{ub}^{\min}$, $v_{ub}^{\min} \setminus v_{lb}^{\min}$, $v_{lb}^{\min} \cap v_{ub}^{\min}$ and $v \setminus (v_{lb}^{\min} \cup v_{ub}^{\min})$.

The following states that this refinement method will always lead to a (strictly) finer partition. Recall that the family of pure memoryless strategies suffices for optimality with respect to probabilistic and expected reachability in stochastic games.

Lemma 6 *If there exists $v' \in V_1$ such that $p_{v'}^{lb, \min}(\overline{F}) < p_{v'}^{ub, \min}(\overline{F})$, then there exists $v \in V_1$ such that $v_{lb}^{\min} \neq v_{ub}^{\min}$.*

Proof. The proof is by contradiction. Therefore, suppose that there exists $v' \in V_1$ such that $p_{v'}^{lb, \min}(\overline{F}) < p_{v'}^{ub, \min}(\overline{F})$ and $v_{lb}^{\min} = v_{ub}^{\min}$ for all $v \in V_1$. Now, consider any memoryless strategy pairs $(\sigma_1^{lb}, \sigma_2^{lb})$ and $(\sigma_1^{ub}, \sigma_2^{ub})$ such that:

$$p_v^{\sigma_1^{lb}, \sigma_2^{lb}}(\overline{F}) = p_v^{lb, \min}(\overline{F}) \text{ and } p_v^{\sigma_1^{ub}, \sigma_2^{ub}}(\overline{F}) = p_v^{ub, \min}(\overline{F}) \text{ for all } v \in V$$

whose existence follows from [40, 14]. Now, for any player 1 vertex $v \in V_1$, by definition:

$$\begin{aligned} v_{lb}^{\min} &= \{s \in v \mid p_{v_2(s)}^{lb, \min}(\overline{F}) = p_v^{lb, \min}(\overline{F})\} \\ v_{ub}^{\min} &= \{s \in v \mid p_{v_2(s)}^{ub, \min}(\overline{F}) = p_v^{ub, \min}(\overline{F})\} \end{aligned}$$

where $v_2(s) = \overline{Steps}(s)$. Now, consider any $v \in V_1$, by construction we have $\sigma_1^{ub}(v) = v_2^{s_{ub}}$ for some $s_{ub} \in v_{ub}^{\min}$, and since by the hypothesis $v_{lb}^{\min} = v_{ub}^{\min}$, we also have that $s_{ub} \in v_{lb}^{\min}$. Therefore, since σ_2^{lb} is optimal and $v \in V_1$ was arbitrary, it follows that:

$$p_{v'}^{lb, \min}(\overline{F}) = p_{v'}^{\sigma_1^{ub}, \sigma_2^{lb}}(\overline{F}) \text{ for all } v \in V_1.$$

In particular, since v' is a player 1 vertex, we have:

$$\begin{aligned} p_{v'}^{lb, \min}(\overline{F}) &= p_{v'}^{\sigma_1^{ub}, \sigma_2^{lb}}(\overline{F}) \\ &\geq \inf_{\sigma_2} p_{v'}^{\sigma_1^{ub}, \sigma_2}(\overline{F}) && \text{since } \sigma_2^{lb} \text{ is a player 2 strategy} \\ &= \sup_{\sigma_1} \inf_{\sigma_2} p_{v'}^{\sigma_1, \sigma_2}(\overline{F}) && \text{since } \sigma_1^{ub} \text{ is an optimal strategy} \\ &= p_{v'}^{ub, \min}(\overline{F}) && \text{by definition} \end{aligned}$$

which contradicts the fact that $p_{v'}^{lb, \min}(\overline{F}) < p_{v'}^{ub, \min}(\overline{F})$. \square

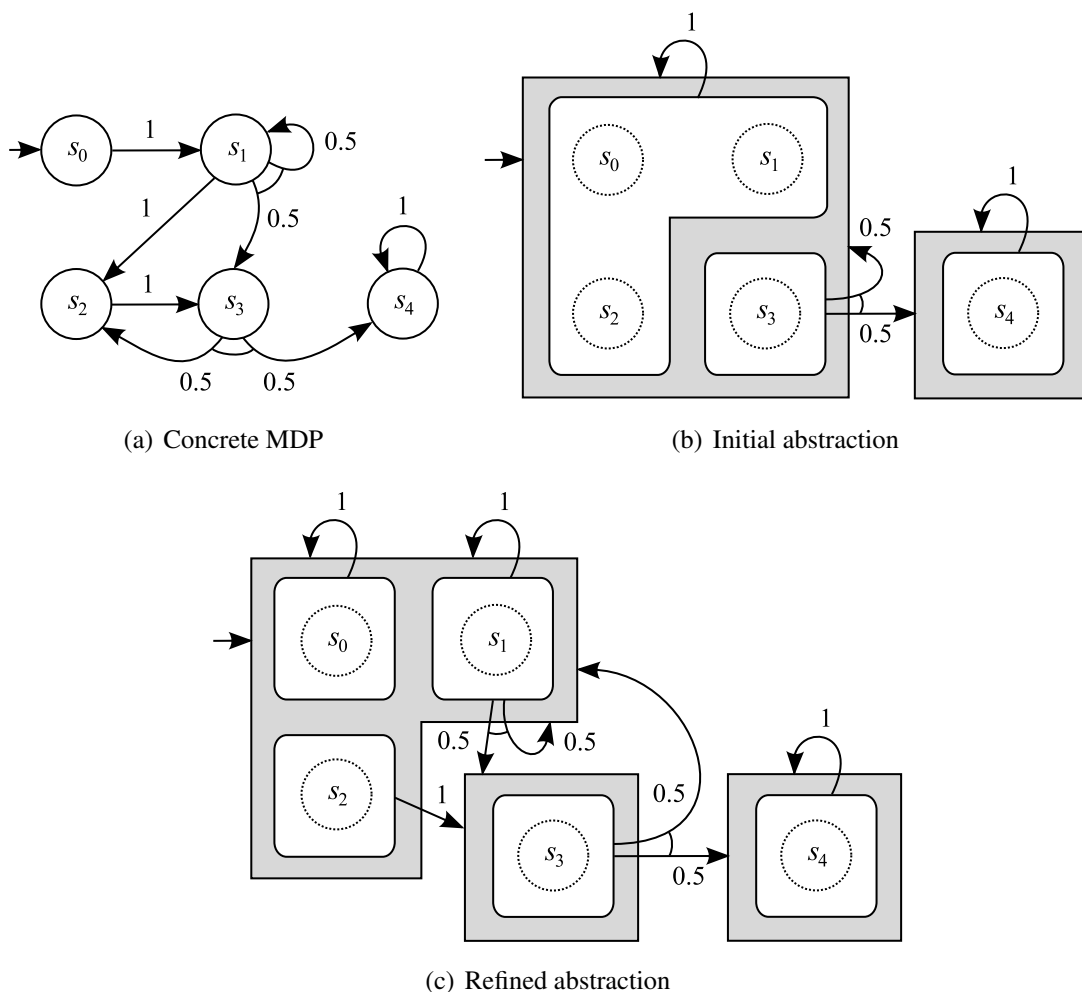


Fig. 4 Example of the abstraction-refinement process

Example 4 To illustrate the refinement strategies, we consider the simple MDP given in Figure 4(a). We assume that a reward of 1 is associated with each transition and compute the maximum expected reward of reaching the target set $\{s_4\}$. We start by partitioning the state space into the target set ($\{s_4\}$) and the remaining states, resulting in the stochastic two-player game shown in Figure 4(b). As in Example 1, player 1 vertices are drawn as grey rectangles, each one's successor player 2 vertices are drawn as white shapes enclosed within it, and probabilistic vertices are represented by sets of grouped, labelled arrows. The states of the original MDP, sets of which comprise player 1 and player 2 vertices, are also depicted.

Calculating the lower and upper bounds on the maximum expected reward to reach $\{s_4\}$, using the game in Figure 4(b), yields the interval $[2, \infty)$ for the vertex containing states s_0 – s_3 (and $[0, 0]$ for the one containing s_4). There is only one memoryless strategy resulting in each bound: the one that selects $\{s_3\}$ (for the lower bound) and the one that selects $\{s_0, s_1, s_2\}$ (for the upper bound). Consequently, both the strategy-based and value-based refinement techniques refine the partition from $\{s_0, s_1, s_2, s_3\}, \{s_4\}$ to $\{s_0, s_1, s_2\}, \{s_3\}, \{s_4\}$. The new abstraction is shown in Figure 4(c). Note that the player 1 vertex containing $\{s_0, s_1, s_2\}$ now leads to three distinct player 2 vertices, since the (lifted) distributions for MDP states s_0 , s_1 and s_2 are no longer equivalent under the new partition.

Calculating bounds with the new game now gives $[4, \infty)$ and $[3, \infty)$ for the vertices $\{s_0, s_1, s_2\}$ and $\{s_3\}$. From the vertex $\{s_0, s_1, s_2\}$, the optimal strategy for the lower bound

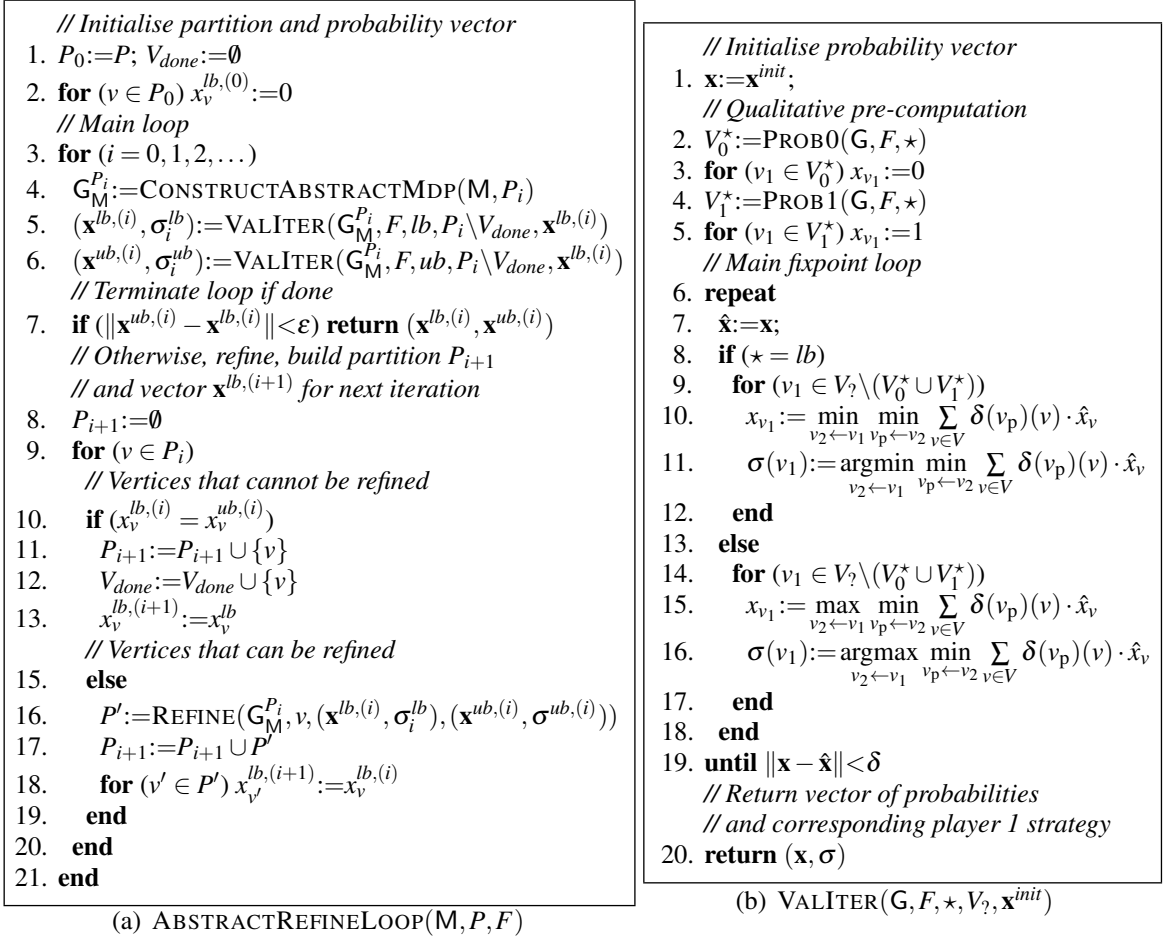


Fig. 5 Abstraction-refinement loop algorithm (minimum reachability probabilities)

chooses $\{s_2\}$ and for the upper bound can choose either $\{s_0\}$ or $\{s_1\}$. With strategy-based refinement, using either strategy for the upper bound gives the partition $\{s_0\}, \{s_1\}, \{s_2\}, \{s_3\}, \{s_4\}$. For value-based refinement, s_0 and s_1 are kept together, giving rise instead to the partition $\{s_0, s_1\}, \{s_2\}, \{s_3\}, \{s_4\}$.

4.2 The Abstraction-Refinement Algorithm

With the refinement strategies of the previous section, we can implement the abstraction-refinement loop shown in Figure 3. Starting with a coarse partition, we iteratively: construct an abstraction (according to Definition 10); compute lower and upper bounds (and corresponding strategies); and refine using either the strategy-based or value-based method of Section 4.1.

In this section, we present an optimised algorithm for the refinement loop, illustrated in Figure 5(a). At the i th refinement step, we construct the abstraction $G_M^{P_i}$ of the MDP using partition P_i . Using $G_M^{P_i}$, we then compute vectors of lower and upper bounds $\mathbf{x}^{lb,(i)}$ and $\mathbf{x}^{ub,(i)}$ (indexed over P_i , i.e. over player 1 vertices of $G_M^{P_i}$) for the minimum probability of reaching F in the MDP, using value iteration. If the bounds are sufficiently close (within ϵ), the loop terminates; otherwise, we refine the partition. The algorithm includes two optimisations, designed to improve the convergence rate of numerical solution:

- wherever possible, we re-use existing numerical results: when calculating lower bounds (line 5), we use the same vector from the previous step as the initial solution; for upper bounds (line 6), we re-use the lower bound from the current step; both are guaranteed lower bounds on the solution.
- we store (in set V_{done}) vertices for which lower and upper bounds coincide (and we thus have the exact solution for their concrete states); again, we use these values in initial solutions for value iteration.

Numerical solution of the game at each step is performed with an improved version of the standard value iteration algorithm [15], illustrated in Figure 5(b), which includes several optimisations. Firstly, we employ the qualitative algorithms of [4] to efficiently find vertices for which the solution is exactly 0 or 1 (lines 2 and 4). Secondly, we only compute new values for vertices where the answer is unknown, i.e. those not in the set V_{done} , mentioned above, and not identified by the qualitative algorithms (lines 9 and 14). Thirdly, the value iteration algorithm is given an initial vector \mathbf{x}^{init} for the fixpoint computation (normally this is a vector of zeros). This is used, as described above, to improve convergence. The correctness of the value iteration scheme is given in Section 4.3 below.

As in the previous section, the code presented in Figure 5 is for computation of minimum probabilities. The case for maximum probabilities and for expected rewards require only trivial changes to the value iteration algorithm. The REFINE function referenced in line 16 of Figure 5(a) can be either the strategy-based or value-based refinement from Section 4.1. For the latter, we require that lines 11 and 16 of the VALITER algorithm in Figure 5(b) return all (rather than just one) player 1 strategy choice which achieves the lower/upper bound.

4.3 Correctness of the Abstraction-Refinement Algorithm

We now demonstrate that the value iteration scheme of Figure 5 converges. For simplicity, we only consider computing bounds on minimum reachability probabilities and assume a fixed MDP $M = (S, s_{init}, Steps, \mathbf{r})$ and target set F . The cases for maximum probabilistic and expected reachability follow in identical fashion.

We begin with a number of preliminary concepts. For a set of vertices V , let \sqsubseteq be the partial order over $(V \rightarrow [0, 1]) \times (V \rightarrow [0, 1])$ where $\mathbf{x} \sqsubseteq \mathbf{x}'$ if and only if $x_v \leq x'_v$ for all $v \in V$. For any stochastic game $G = ((V, E), v_{init}, (V_1, V_2, V_p), \delta, \bar{\mathbf{r}})$, subset of player 1 vertices V_{done} and target set F , let $F_{V_{done}}^{lb} : (V_1 \rightarrow [0, 1]) \rightarrow (V_1 \rightarrow [0, 1])$ be the function:

$$F_{V_{done}}^{lb}(\mathbf{x})(v) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } v \in V_1^{lb} \\ 0 & \text{if } v \in V_0^{lb} \\ \inf_{\sigma_1} \inf_{\sigma_2} p_v^{\sigma_1, \sigma_2}(F) & \text{if } v \in V_{done} \setminus (V_1^{lb} \cup V_0^{lb}) \\ \min_{v_2 \leftarrow v} \min_{v_p \leftarrow v_2} \sum_{v \in V} \delta(v_p)(v) \cdot x_v & \text{otherwise} \end{cases}$$

where the notation $v' \leftarrow v$ denotes the set $\{v' \in V \mid (v, v') \in E\}$ and:

$$\begin{aligned} V_1^{lb} &= \{v \in V_1 \mid \inf_{\sigma_1} \inf_{\sigma_2} p_v^{\sigma_1, \sigma_2}(F) = 1\} \\ V_0^{lb} &= \{v \in V_1 \mid \inf_{\sigma_1} \inf_{\sigma_2} p_v^{\sigma_1, \sigma_2}(F) = 0\}. \end{aligned}$$

Furthermore, let $F_{V_{done}}^{ub} : (V_1 \rightarrow [0, 1]) \rightarrow (V_1 \rightarrow [0, 1])$ be the function:

$$F_{V_{done}}^{ub}(\mathbf{x})(v) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } v \in V_1^{ub} \\ 0 & \text{if } v \in V_0^{ub} \\ \sup_{\sigma_1} \inf_{\sigma_2} p_v^{\sigma_1, \sigma_2}(F) & \text{if } v \in V_{done} \setminus (V_1^{ub} \cup V_0^{ub}) \\ \max_{v_2 \leftarrow v} \min_{v_p \leftarrow v_2} \sum_{v \in V} \delta(v_p)(v) \cdot x_v & \text{otherwise} \end{cases}$$

where:

$$\begin{aligned} V_1^{ub} &= \{v \in V_1 \mid \sup_{\sigma_1} \inf_{\sigma_2} p_v^{\sigma_1, \sigma_2}(F) = 1\} \\ V_0^{ub} &= \{v \in V_1 \mid \sup_{\sigma_1} \inf_{\sigma_2} p_v^{\sigma_1, \sigma_2}(F) = 0\}. \end{aligned}$$

The following results are adapted from [14] to the notation used in this paper.

Lemma 7 For any stochastic game G , target set F and subset of player 1 vertices V_{done} , the function $F_{V_{done}}^{lb}$ is monotonic with respect to \sqsubseteq .

Theorem 4 For any stochastic game G , target set F and subset of player 1 vertices V_{done} , the function $F_{V_{done}}^{lb}$ has a unique fixed point \mathbf{u}^{lb} and u_v^{lb} is the optimal value for the probability of reaching F when player 1 and player 2 cooperate in order to minimise it, that is $u_v^{lb} = \inf_{\sigma_1} \inf_{\sigma_2} p_v^{\sigma_1, \sigma_2}(F)$.

Lemma 8 For any stochastic game G , target set F and subset of player 1 vertices V_{done} , the function $F_{V_{done}}^{ub}$ is monotonic with respect to \sqsubseteq .

Theorem 5 For any stochastic game G , target set F and subset of player 1 vertices V_{done} , the function $F_{V_{done}}^{ub}$ has a unique fixed point \mathbf{u}^{ub} and u_v^{ub} is the optimal value for the probability of reaching F for player 1, that is $u_v^{ub} = \sup_{\sigma_1} \inf_{\sigma_2} p_v^{\sigma_1, \sigma_2}(F)$.

Using these results and the Knaster-Tarski theorem, to prove the convergence of the value iteration scheme of Figure 5, supposing $\mathbf{x}^{lb,init}$ and $\mathbf{x}^{ub,init}$ are the initial values for computing \mathbf{u}^{lb} and \mathbf{u}^{ub} respectively, it is sufficient to show that:

$$\mathbf{x}^{\star,init} \sqsubseteq \mathbf{u}^{\star} \text{ and } \mathbf{x}^{\star,init} \sqsubseteq F_{V_{done}}^{\star}(\mathbf{x}^{\star,init}) \text{ for } \star \in \{lb, ub\}. \quad (3)$$

Therefore, suppose that P and P^{old} are the partitions from the current and previous refinement steps, G_M^P and $G_M^{P^{old}}$ are the corresponding abstractions and V_{done} is set of vertices of $G_M^{P^{old}}$ for which the lower and upper bounds are equal. Note that, by construction we have that $P \preceq P^{old}$. We now prove that (3) holds by considering the cases where $\star = lb$ and $\star = ub$ in turn.

– If $\star = lb$ then by construction, the initial vector $\mathbf{x}^{lb,init}$ is given by:

$$x_v^{lb,init} = \begin{cases} 1 & \text{if } v \in V_1^{lb} \\ 0 & \text{if } v \in V_0^{lb} \\ p_{v^{old}}^{lb, \min}(F) & \text{otherwise} \end{cases}$$

where v^{old} is the unique element of P^{old} such that $v \subseteq v^{old}$. Now, for any $v \in V_1$, by Theorem 2 we have that:

$$p_{v^{old}}^{lb, \min}(F) \leq p_v^{lb, \min}(F) (= u_v^{lb}) \quad (4)$$

and hence $\mathbf{x}^{lb,init} \sqsubseteq \mathbf{u}^{lb}$.

Therefore, it remains to show that $\mathbf{x}^{lb,init} \sqsubseteq F_{V_{done}}^{lb}(\mathbf{x}^{lb,init})$. Now consider any $v \in V_1$. First, by construction we have:

$$p_{v,old}^{lb,min}(F) \leq x_v^{lb,init}. \quad (5)$$

Next, if $v \in V_1^{lb} \cup V_0^{lb} \cup V_{done}$, then by definition of $F_{V_{done}}^{lb}$ we have $F_{V_{done}}^{lb}(\mathbf{x}^{lb,init})(v) = p_v^{lb,min}(F)$, and hence using (4) we have $F_{V_{done}}^{lb}(\mathbf{x}^{lb,init})(v) \geq x_v^{lb,init}$. On the other hand, if $v \notin V_1^{lb} \cup V_0^{lb} \cup V_{done}$, then by definition of $F_{V_{done}}^{lb}$:

$$\begin{aligned} F_{V_{done}}^{lb}(\mathbf{x}^{lb,init})(v) &= \min_{v_2 \leftarrow v} \min_{v_p \leftarrow v_2} \sum_{v_1 \in V_1} \delta(v_p)(v_1) \cdot x_{v_1}^{lb,init} \\ &\geq \min_{v_2 \leftarrow v} \min_{v_p \leftarrow v_2} \sum_{v_1 \in V_1} \delta(v_p)(v_1) \cdot p_{v_1}^{lb,min}(F) && \text{by (5)} \\ &= \min_{v_2 \leftarrow v} \min_{v_p \leftarrow v_2} \sum_{v_1^{old} \in V_1^{old}} \sum_{v_1 \subseteq v_1^{old}} \delta(v_p)(v_1) \cdot p_{v_1}^{lb,min}(F) && \text{since } P \preceq P_{old} \\ &= \min_{v_2 \leftarrow v} \min_{v_p \leftarrow v_2} \sum_{v_1^{old} \in V_1^{old}} \delta_{old}(v_p)(v) \cdot p_{v_1}^{lb,min}(F) && \text{by Definition 10} \\ &= \min_{v_2^{old} \leftarrow v} \min_{v_p^{old} \leftarrow v_2^{old}} \sum_{v_1^{old} \in V_1^{old}} \delta_{old}(v_p^{old})(v_1^{old}) \cdot p_{v_1}^{lb,min}(F) && \text{since } P \preceq P_{old} \\ &= p_{v,old}^{lb,min}(F) && \text{by Theorem 4} \\ &= x_v^{lb,init} && \text{by construction.} \end{aligned}$$

Since these are all the possible cases to consider, we have $\mathbf{x}^{lb,init} \sqsubseteq F_{V_{done}}^{lb}(\mathbf{x}^{lb,init})$ as required.

– If $\star = ub$, then by construction:

$$x_v^{ub,init} = \begin{cases} 1 & \text{if } v \in V_1^{ub} \\ 0 & \text{if } v \in V_0^{ub} \\ p_v^{ub,min}(F) & \text{if } v \in V_{done} \setminus (V_1^{ub} \cup V_0^{ub}) \\ p_v^{lb,min}(F) & \text{otherwise.} \end{cases}$$

Now for any $v \in V_1$ by definition of $p_v^{lb,min}(F)$ and $p_v^{ub,min}(F)$ and construction respectively, we have:

$$p_v^{lb,min}(F) \leq p_v^{ub,min}(F) (= u_v^{ub}) \quad (6)$$

and hence $\mathbf{x}^{ub,init} \sqsubseteq \mathbf{u}^{ub}$.

Therefore, it remains to show that $\mathbf{x}^{ub,init} \sqsubseteq F_{V_{done}}^{ub}(\mathbf{x}^{ub,init})$. Now consider any $v \in V_1$. First by construction, we have

$$p_v^{lb,min}(F) \leq x_v^{lb,init}. \quad (7)$$

Next, if $v \in V_1^{lb} \cup V_1^{ub} \cup V_{done}$ then by construction $F_{V_{done}}^{ub}(\mathbf{x}^{ub,init})(v) = p_v^{ub,min}(F)$, and hence using (6) we have $F_{V_{done}}^{ub}(\mathbf{x}^{ub,init})(v) \geq x_v^{ub,init}$. On the other hand, if $v \notin V_1^{lb} \cup V_1^{ub} \cup$

V_{done} then by definition of $F_{V_{done}}^{ub}$:

$$\begin{aligned}
F_{V_{done}}^{ub}(\mathbf{x}^{ub,init})(v) &= \max_{v_2 \leftarrow v} \min_{v_p \leftarrow v_2} \sum_{v_1 \in V_1} \delta(v_p)(v_1) \cdot x_{v_1}^{ub,init} \\
&\geq \max_{v_2 \leftarrow v} \min_{v_p \leftarrow v_2} \sum_{v_1 \in V_1} \delta(v_p)(v_1) \cdot p_{v_1}^{lb, \min}(F) && \text{by (7)} \\
&\geq \min_{v_2 \leftarrow v} \min_{v_p \leftarrow v_2} \sum_{v_1 \in V_1} \delta(v_p)(v_1) \cdot p_{v_1}^{lb, \min}(F) && \text{rearranging} \\
&= p_v^{lb, \min}(F) && \text{by Theorem 5} \\
&= x_v^{ub,init} && \text{by construction.}
\end{aligned}$$

Since these are all the cases to consider, we have $\mathbf{x}^{ub,init} \sqsubseteq F_{V_{done}}^{ub}(\mathbf{x}^{ub,init})$ as required. This completes the proof of (3), and hence that the value iteration scheme of Figure 5 converges.

5 Experimental Results

We have implemented our abstraction-refinement framework using a prototype Java implementation. This section presents experimental results for the performance of our techniques on case studies from the repository of the probabilistic model checker PRISM [37]. Below are listed the case studies and (in parentheses) the reachability properties used.¹

- the Zeroconf network configuration protocol for N configured hosts and M IP addresses (“the minimum probability that the host configures correctly”);
- the IEEE 802.11 WLAN and IEEE 802.3 CSMA/CD protocols using a backoff counter maximum of bc and a maximum packet send time of $500 \mu s$ (“the minimum probability that a station’s backoff counter reaches bc ”);
- the FireWire root contention protocol for a network transmission delay of d ns (“the maximum expected time to elect a leader”);
- the randomised consensus shared coin protocol of Aspnes and Herlihy for N processes and parameter K (“the maximum expected time until termination”).

Several of the models (FireWire, WLAN and CSMA/CD) were modelled using probabilistic timed automata and translated into MDPs using the digital clocks semantics of [30]. For all experiments, the initial partition P contained: (1) the initial state(s); (2) the target states; and (3) all remaining states. The abstraction was refined until the maximum *relative* difference between the bounds for the initial state(s) was below $\varepsilon=10^{-4}$, i.e. when

$$\max_{v \in V_{init}} \frac{x_v^{ub} - x_v^{lb}}{x_v^{ub}} < 10^{-4}$$

where V_{init} is the set of abstract states that contain at least one concrete initial state. The value iteration algorithm used the default convergence criteria of PRISM (maximum relative difference less than 10^{-6}).

Table 2 presents statistics for the performance of our implementation. For each example, we give the size of the original MDP (i.e. number of concrete states) and, for each of our two refinement techniques, the size of the final abstraction (i.e. number of abstract states/player 1 vertices), the number of refinement steps and the total time required for the entire sequence of refinements and value iteration computations.

¹ Supporting files for the models and properties used in the experimental results are available from: <http://www.prismmodelchecker.org/files/fmsd-games/>.

Case study (parameter)	States in MDP	Strategy-based refinement			Value-based refinement			
		States	Steps	Time (s)	States	Steps	Time (s)	
IPv4 Zeroconf protocol ($N M$)	4 32	26,121	1,041	206	51.66	699	112	28.27
	4 64	50,377	1,041	189	76.30	676	112	49.86
	4 128	98,889	917	179	127.3	680	112	95.80
	8 32	552,097	2,277	208	1,184	883	128	791.7
	8 64	1,065,569	2,277	188	2,412	877	128	1,584
	8 128	2,092,513	1,720	217	4,293	855	128	3,300
IEEE 802.11 WLAN (bc)	2	28,480	399	70	13.98	253	70	14.50
	3	96,302	1,141	118	84.77	704	118	83.00
	4	345,000	2,619	214	571.9	1,584	214	557.6
	5	1,295,218	5,569	406	4,512	3,605	406	4,625
IEEE 802.3 CSMA/CD (bc)	4	92,978	419	53	30.25	410	53	30.2
	5	277,493	846	78	129.8	839	77	126.9
	6	793,110	1,671	138	678.6	1,663	138	654.3
	7	2,221,189	3,297	266	3,687	3,176	267	3,580
IEEE 1394 FireWire root contention (d)	30	4,093	1,274	320	106.7	910	860	307.1
	60	8,618	2,141	244	181.9	1,495	994	588.0
	120	22,852	3,957	183	276.7	2,746	1,404	1,576
	240	84,152	7,956	203	833.3	5,572	2,379	7,681
Randomised consensus protocol ($N K$)	5 2	173,056	1,298	47	177.5	566	236	690.1
	5 4	327,936	2,529	76	801.7	1,111	346	1,940
	5 8	637,696	5,008	155	5,548	2,074	557	7,368
	5 16	1,257,216	9,987	282	39,803	4,144	1,027	34,133

Table 2 Performance statistics for the abstraction-refinement implementation

The abstractions. The first, and most important, conclusion that we draw from these results is the quality of the automatically generated abstractions. For results of a relatively high precision ($\epsilon=10^{-4}$, i.e. the bounds differ by less than 4 significant figures), the abstractions yielded state-space reductions of up to four orders of magnitude. We find it very promising that this holds for such a wide range of models.

It is also interesting to compare the abstractions we have obtained with manually developed abstractions derived for the same case studies in the literature. For the FireWire example, employing the manual abstraction process of [42] (4 refinement steps and a number of complex proofs) and the digital clocks semantics of [30] yields an abstract model with 1,212 states for $d=30$. For the Zeroconf example, the manual abstraction of [28] has 737 ($N=4, M=64$) and 881 ($N=8, M=64$) states. For both models we achieve smaller abstractions, in a fully automatic fashion, than those obtained manually. Furthermore, there is scope to combine the approaches, e.g. start with a human-derived abstraction and then refine mechanically.

The refinement techniques. Value-based refinement mostly outperforms the strategy-based variant, both in terms of the number of refinement steps and the size of the final abstraction. The faster convergence for the value-based approach is likely to be caused by the splitting of partition elements into four, rather than three, at each step. The fact that this is not at the expense of generating larger abstract models suggests that value-based refinement, as intended, avoids splitting states which should stay in the same partition element.

However, on some case studies – FireWire for example – strategy-based refinement requires much fewer steps and is thus quicker. Despite this, the final abstractions produced are slightly larger. Similar results are observed for the consensus case study but, for the largest models, the increased size of the abstraction makes the process slower overall.

Case study (parameter)		Abstraction-refinement time (s)					PRISM time (s)
		Total	Breakdown			Optimisation Saving	
			Build	Refine	Solve		
Zeroconf (<i>N M</i>)	8 64	1,584	1,574	2.91	6.41	18.83	149.5
	8 128	3,300	3,287	5.78	7.36	52.47	269.8
WLAN (<i>bc</i>)	4	557.6	527.8	0.33	29.02	153.2	19.34
	5	4,625	4,300	1.763	320.9	158.1	89.80
CSMA (<i>bc</i>)	6	654.3	616.3	3.63	33.91	61.79	41.39
	7	3,580	3,322	15.42	240.6	394.6	134.3
FireWire (<i>d</i>)	120	276.7	32.35	0.10	244.0	267.0	8.71
	240	833.3	151.4	0.23	680.9	841.0	25.37
Consensus (<i>N K</i>)	5 8	7,368	5,637	4.33	1,726	4,027	2,379
	5 16	34,133	21,837	13.28	12,278	41,878	14,956

Table 3 More detailed timing statistics for the abstraction-refinement implementation

Performance. The motivations for this work are to establish the viability of the stochastic two-player games as an abstract model for MDPs and to explore the benefits of our abstraction-refinement framework. For convenience, the implementation used in this paper performs model-level abstraction using a prototype abstraction tool which first builds the full concrete MDP (using PRISM) and then reduces it to a game based on a partition of the state space. As a result, this is currently the most time-intensive part of the process. This is illustrated by Table 3 which shows, for the largest models in each case study, a breakdown of the timing for each part of the best-performing refinement technique: abstraction construction (“Build”), refinement (“Refine”) and numerical solution using value iteration (“Solve”). Despite this, the overall performance of the framework is very encouraging.

Table 3 also gives the time-savings obtained by using the optimisations from Section 4.2 which reuse earlier numerical results to improve convergence. Comparing the time spent on numerical solution (column “Solve”) with the time savings achieved (column “Optimisation Saving”), we see that the optimisations result in substantial improvements in performance.

Lastly, Table 3 shows the time required to verify each model in PRISM (using the fastest available engine in the tool). We see that, perhaps unsurprisingly, verification with PRISM is currently faster than using the abstraction-refinement approach. We emphasise that the main goals of this work are not to improve solution time for instances where it is feasible to verify the full concrete model, but rather to develop techniques that will scale to larger models. This will be achieved by building abstractions directly from high-level languages, e.g. using predicate abstraction, and bypassing the construction of the concrete model. Even without this, in some cases (e.g. Zeroconf), the time for abstraction refinement is much faster than PRISM, when the time for explicit construction of the abstraction is disregarded.

Convergence. Finally, we look at the convergence of the refinement process, i.e., the changes in difference between the bounds at each refinement step. Previously, we argued that the difference between the bounds provided a useful quantitative measure of the quality of the abstraction. To illustrate this fact, Figure 6 presents, for two of the CSMA/CD models, the lower and upper bounds on the corresponding probabilistic reachability property after each refinement step when using the value-based refinement method and, for two of the consensus models, the lower bounds for the expected reachability property after each refinement step for both refinement methods. Although there are several sudden drops (or increases) in these bounds, the overall trends show a gradual convergence in the quantitative results. One reason for the sudden jumps is the fact that the values shown in the graphs in Figure 6 are




Fig. 6 Illustrations of convergence for the CSMA and Consensus case studies.

for the initial state of each model, whereas the refinement scheme attempts to improve lower and upper the bounds for all states. This means that not every refinement will improve the bounds for the initial state.

6 Related Work

Below, we compare our approach with the closest work in the field, namely implementations of abstraction and refinement methods for MDPs. General issues relating to abstraction in the field of probabilistic model checking are discussed in [24, 34].

D’Argenio et al. [16] introduce an approach for verifying reachability properties of MDPs based on probabilistic simulation [38] and a corresponding tool RAPTURE. Properties are analysed on abstractions obtained through successive refinements, starting from a coarse partition derived from the property under study. Since the abstractions used are themselves MDPs, this approach only produces a lower (upper) bound for the minimum

(maximum) reachability probability, and hence appears more suited to analysing Markov chains (which contain no nondeterminism and thus the minimum and maximum probabilities coincide). In [16], refinement is repeated until the probability (not the error bound) obtained falls below a given threshold.

Based on [16] and using predicate abstraction [19], Hermanns et al. [44,21] propose a CEGAR framework for MDPs described in the PRISM language and implement it in a tool called PASS [35]. The CEGAR framework is used for verifying or refuting properties of the form “the maximum probability of error is at most p ” for a given probability threshold p . Since abstractions produce only upper bounds on maximum probabilities, to refute a property, probabilistic counterexamples [20] (comprising multiple paths whose combined probability exceeds p) are generated. If these paths are spurious, then they are used to generate further predicates using interpolation.

Perhaps the most important distinguishing feature of our work is the use of two-player games as abstractions. These provide *lower and upper bounds*, avoiding enumeration of a set of paths for a counterexample (which may be large or even infinite, resulting in non-termination). In addition, the bounds enable a *quantitative* approach: we target properties without thresholds such as “what is the maximum probability of failure?”. Recent extensions to the PASS tool [43] use the game-based abstraction approach introduced in this paper and demonstrate that this is faster and yields smaller abstractions. Furthermore, the abstraction-refinement framework presented in this paper has been successfully applied in several other contexts. In [25], predicate abstraction is used to construct stochastic game abstractions of PRISM models. In [26] and [29], game-based abstraction and refinement are used to develop verification techniques for probabilistic software and probabilistic timed automata, respectively. In the former, the techniques are used to verify ANSI-C code for network utilities whose complexity is significantly beyond the scope of existing probabilistic verification tools. In the latter, methods to analyse infinite-state real-time systems are presented that outperform all existing approaches.

In [5], de Alfaro et al. propose a technique called ‘magnifying-lens abstraction’ (MLA) which computes lower and upper bounds on reachability probabilities of MDPs. The approach is based on, first, partitioning the state space into regions and then analysing (‘magnifying’) the states of each region separately by applying value iteration locally. Regions are refined adaptively until the difference between the bounds for all regions is within some specified accuracy. MLA is suited to models where there is some notion of ‘distance’ between states (states close together have similar reachability probabilities) which is used during refinement. Since probabilities are computed for each state of each region and then a minimum/maximum is stored for each region, analysing either small numbers of large regions or large number of small regions is expensive in time and memory. Furthermore, the greatest possible reduction in state-space N is $O(\sqrt{N})$.

Comparing results in [5] for Zeroconf ($N=4, M=32$), which has 26,121 concrete states, using the same accuracy of $\epsilon=10^{-3}$, the MLA approach gives 131 regions and a maximum region size of 1,005, whereas our value-based refinement yields 699 abstract states. Note that, as Figure 2 demonstrates, our abstract model is smaller for larger values of M while, for MLA, increasing M will increase either the number of regions or maximum region size (or both). Similarly, our framework is applicable to infinite-state systems, whereas [5] is not.

An alternative approach to the abstraction of MDPs is [9], which proposes a CEGAR framework for verifying a fragment of the logic PCTL using MDPs as abstract models (as in [16]). The focus of the paper is on suitable notions of counterexamples for this process, observing that existing proposals such as those used in [21] are not sufficiently expressive for general PCTL formulae. Instead, the paper proposes that counterexamples take the form

of “small” MDPs that simulate the concrete model. Algorithms to implement the proposed CEGAR framework are given, but no implementation is attempted.

Elsewhere, abstraction techniques have been considered for discrete-time Markov chains [18,39], which can be seen as MDPs without nondeterminism. These techniques abstract probabilities to intervals; however, practical applicability has yet to be demonstrated. Also relevant is the work of Chatterjee et al. [11] who propose a counterexample-based abstraction-refinement technique for planning problems on stochastic games. Again, an implementation to test this on practical examples has not yet been developed.

Another important direction for abstraction and refinement of probabilistic systems is the extension of abstract interpretation to the probabilistic setting [36,33,41], although these approaches have yet to be combined with refinement. For details on the relationship between the game-based abstraction approach and abstract interpretation, see [43].

In [17] a method for approximating continuous-state (and hence infinite-state) Markov processes by a family of finite-state Markov chains is presented. It is shown that, for simple quantitative modal logic, if the continuous Markov process satisfies a formula, then one of the approximations also satisfies the formula. Finally, McIver and Morgan have developed a framework for the refinement and abstraction of probabilistic programs using expectation transformers [32]. The proof techniques developed in this work have been implemented in the HOL theorem-proving environment [23].

7 Conclusions

We have presented a novel abstraction-refinement framework for Markov decision processes. Our abstraction technique is based on the use of stochastic two-player games, in which one player corresponds to nondeterministic choices from the MDP and the other corresponds to the nondeterminism introduced through abstraction. Using existing results and algorithms from the stochastic games literature, we have shown how the abstraction can be used to compute both lower and upper bounds on the minimum and maximum probability or expected reward of reaching a set of states in the original MDP.

In addition to providing quantitative results for properties of the MDP, the stochastic game yields a measure of the utility of the abstraction, thus giving a basis for refining it to improve its accuracy. Using this abstraction approach, we have introduced two refinement techniques and an optimised algorithm for a quantitative abstraction-refinement loop. Our experimental results illustrate how this framework provides efficient and fully automatic generation of precise yet compact abstractions for large MDPs from a wide selection of complex case studies. This demonstrates that stochastic two-player games are indeed a good choice of model for representing abstractions of MDPs.

There are many possibilities for future work on this topic. An important direction is to investigate the application of the framework at the level of MDP modelling formalisms. In fact, these techniques have subsequently been employed to develop probabilistic verification methods for imperative languages such as C [26], using predicate abstraction, and for probabilistic timed automata [29], using convex polyhedra [29]. In each case, a key challenge is developing efficient techniques and heuristics to guide refinement; this represents an important area of future work. We also plan to extend our abstraction and refinement methods to other types of models, for example, continuous-time models and hybrid automata.

Acknowledgements

This work was supported in part by EPSRC grants EP/D07956X and EP/D076625. The authors would like to thank the anonymous referees for their helpful comments.

References

1. de Alfaro, L.: Formal verification of probabilistic systems. Ph.D. thesis, Stanford University (1997)
2. de Alfaro, L.: Computing minimum and maximum reachability times in probabilistic systems. In: J. Baeten, S. Mauw (eds.) Proc. 10th Int. Conf. Concurrency Theory (CONCUR'99), *Lecture Notes in Computer Science*, vol. 1664, pp. 66–81. Springer (1999)
3. de Alfaro, L., Henzinger, T., Kupferman, O.: Concurrent reachability games. In: Proc. 39th Symp. Foundations of Computer Science (FOCS'98), pp. 564–575. IEEE CS Press (1998)
4. de Alfaro, L., Henzinger, T., Kupferman, O.: Concurrent reachability games. *Theoretical Computer Science* **386**(3), 188–217 (2007)
5. de Alfaro, L., Roy, P.: Magnifying-lens abstraction for Markov decision processes. In: W. Damm, H. Hermanns (eds.) Proc. 19th Int. Conf. Computer Aided Verification (CAV'07), *Lecture Notes in Computer Science*, vol. 4590, pp. 325–338. Springer (2007)
6. Baier, C., Kwiatkowska, M.: Model checking for a probabilistic branching time logic with fairness. *Distributed Computing* **11**(3), 125–155 (1998)
7. Bertsekas, D., Tsitsiklis, J.: An analysis of stochastic shortest path problems. *Mathematics of Operations Research* **16**(3), 580–595 (1991)
8. Billingsley, P.: Probability and Measure. John Wiley and Sons: New York (1979)
9. Chadha, R., Viswanathan, M.: A counterexample guided abstraction-refinement framework for markov decision processes. *ACM Transactions on Computational Logic* (2010). To appear
10. Chatterjee, K., de Alfaro, L., Henzinger, T.: Trading memory for randomness. In: Proc. 1st Int. Conf. Quantitative Evaluation of Systems (QEST'04), pp. 206–217. IEEE CS Press (2004)
11. Chatterjee, K., Henzinger, T., Jhala, R., Majumdar, R.: Counterexample-guided planning. In: Proc. 21st Conference in Uncertainty in Artificial Intelligence (UAI'05), pp. 104–111 (2005)
12. Cheshire, S., Adoba, B., Guttman, E.: Dynamic configuration of IPv4 link-local addresses (draft August 2002). Zeroconf Working Group of the Internet Engineering Task Force (www.zeroconf.org)
13. Clarke, E., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement. In: A. Emerson, A. Sistla (eds.) Proc. 12th Int. Conf. Computer Aided Verification (CAV'00), *Lecture Notes in Computer Science*, vol. 1855, pp. 154–169. Springer (2000)
14. Condon, A.: The complexity of stochastic games. *Information and Computation* **96**(2), 203–224 (1992)
15. Condon, A.: On algorithms for simple stochastic games. *Advances in computational complexity theory, DIMACS Series in Discrete Mathematics and Theoretical Computer Science* **13**, 51–73 (1993)
16. D'Argenio, P., Jeannot, B., Jensen, H., Larsen, K.: Reachability analysis of probabilistic systems by successive refinements. In: L. de Alfaro, S. Gilmore (eds.) Proc. 1st Joint Int. Workshop Process Algebra and Probabilistic Methods, Performance Modeling and Verification (PAPM/PROBMIV'01), *Lecture Notes in Computer Science*, vol. 2165, pp. 39–56. Springer (2001)
17. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Approximating labelled Markov processes. *Information and Computation* **184**(1), 160–200 (2003)
18. Fecher, H., Leucker, M., Wolf, V.: Don't know in probabilistic systems. In: A. Valmari (ed.) Proc. 13th Int. SPIN Workshop on Model Checking of Software (SPIN'06), *Lecture Notes in Computer Science*, vol. 3925, pp. 71–88. Springer (2006)
19. Graf, S., Saidi, H.: Construction of abstract state graphs with PVS. In: O. Grumberg (ed.) Proc. 9th Int. Conf. Computer Aided Verification (CAV'97), *Lecture Notes in Computer Science*, vol. 1254, pp. 72–83. Springer (1997)
20. Han, T., Katoen, J.P., Damman, B.: Counterexample generation in probabilistic model checking. *IEEE Transactions on Software Engineering* **35**(2), 241257 (2009)
21. Hermanns, H., Wachter, B., Zhang, L.: Probabilistic CEGAR. In: A. Gupta, S. Malik (eds.) Proc. 20th Int. Conf. Computer Aided Verification (CAV'08), *Lecture Notes in Computer Science*, vol. 5123, pp. 162–175. Springer (2008)
22. Hinton, A., Kwiatkowska, M., Norman, G., Parker, D.: PRISM: A tool for automatic verification of probabilistic systems. In: H. Hermanns, J. Palsberg (eds.) Proc. 12th Int. Conf. Tools and Algorithms for the Construction and Analysis of Systems (TACAS'06), *Lecture Notes in Computer Science*, vol. 3920, pp. 441–444. Springer (2006)

23. Hurd, J., McIver, A., Morgan, C.: Probabilistic guarded commands mechanized in HOL. *Theoretical Computer Science* **346**(1), 96–112 (2005)
24. Huth, M.: An abstraction framework for mixed nondeterministic and probabilistic systems. In: C. Baier, B. Haverkort, H. Hermanns, J.P. Katoen, M. Siegle (eds.) *Validation of Stochastic Systems, Lecture Notes in Computer Science*, vol. 2925, pp. 419–444. Springer (2004)
25. Kattenbelt, M., Kwiatkowska, M., Norman, G., Parker, D.: Game-based probabilistic predicate abstraction in PRISM. In: *Proc. 6th Workshop Quantitative Aspects of Programming Languages (QAPL'08)* (2008)
26. Kattenbelt, M., Kwiatkowska, M., Norman, G., Parker, D.: Abstraction refinement for probabilistic software. In: N. Jones, M. Müller-Olm (eds.) *Proc. 10th Int. Conf. Verification, Model Checking and Abstract Interpretation (VMCAI'09), Lecture Notes in Computer Science*, vol. 5403, pp. 182–197. Springer (2009)
27. Kemeny, J., Snell, J., Knapp, A.: *Denumerable Markov Chains*, 2nd edn. Springer-Verlag (1976)
28. Kwiatkowska, M., Norman, G., Parker, D.: Game-based abstraction for Markov decision processes. In: *Proc. 3th Int. Conf. Quantitative Evaluation of Systems (QEST'06)*, pp. 157–166. IEEE CS Press (2006)
29. Kwiatkowska, M., Norman, G., Parker, D.: Stochastic games for verification of probabilistic timed automata. In: J. Ouaknine, F. Vaandrager (eds.) *Proc. 7th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS'09), Lecture Notes in Computer Science*, vol. 5813, pp. 212–227. Springer (2009)
30. Kwiatkowska, M., Norman, G., Parker, D., Sproston, J.: Performance analysis of probabilistic timed automata using digital clocks. *Formal Methods in System Design* **29**, 33–78 (2006)
31. Larsen, K., Skou, A.: Bisimulation through probabilistic testing. *Information and Computation* **94**, 1–28 (1991)
32. McIver, A., Morgan, C.: *Abstraction, Refinement and Proof for Probabilistic Systems. Monographs in Computer Science*. Springer (2004)
33. Monniaux, D.: Abstract interpretation of programs as Markov decision processes. *Science of Computer Programming* **58**(1–2), 179–205 (2005)
34. Norman, G.: Analysing randomized distributed algorithms. In: C. Baier, B. Haverkort, H. Hermanns, J.P. Katoen, M. Siegle (eds.) *Validation of Stochastic Systems, Lecture Notes in Computer Science*, vol. 2925, pp. 384–418. Springer (2004)
35. PASS tool homepage. <http://depend.cs.uni-sb.de/PASS/>
36. Pierro, A.D., Hankin, C., Wiklicky, H.: Abstract interpretation for worst and average case analysis. In: T. Reps, M. Sagiv, J. Bauer (eds.) *Program Analysis and Compilation, Theory and Practice, Essays Dedicated to Reinhard Wilhelm on the Occasion of His 60th Birthday, Lecture Notes in Computer Science*, vol. 4444, pp. 160–174. Springer (2006)
37. PRISM web site. <http://www.prismmodelchecker.org/>
38. Segala, R.: *Modelling and verification of randomized distributed real time systems. Ph.D. thesis, Massachusetts Institute of Technology* (1995)
39. Sen, K., Viswanathan, M., Agha, G.: Model-checking Markov chains in the presence of uncertainties. In: H. Hermanns, J. Palsberg (eds.) *Proc. 12th Int. Conf. Tools and Algorithms for the Construction and Analysis of Systems (TACAS'06), Lecture Notes in Computer Science*, vol. 3920, pp. 394–410. Springer (2006)
40. Shapley, L.: Stochastic games. In: *Proc. National Academy of Science*, vol. 39, pp. 1095–1100 (1953)
41. Smith, M.: Probabilistic abstract interpretation of imperative programs using truncated normal distributions. In: A. Aldini, C. Baier (eds.) *Proc. 6th Workshop on Quantitative Aspects of Programming Languages (QAPL'08), Electronic Notes in Theoretical Computer Science*, vol. 220(3), pp. 43–59. Elsevier (2008)
42. Stoelinga, M., Vaandrager, F.: Root contention in IEEE 1394. In: J.P. Katoen (ed.) *Proc. 5th Int. AMAST Workshop Real-Time and Probabilistic Systems (ARTS'99), Lecture Notes in Computer Science*, vol. 1601, pp. 53–74. Springer (1999)
43. Wachter, B., Zhang, L.: Best probabilistic transformers. In: G. Barthe, M. Hermenegildo (eds.) *Proc. 11th Int. Conf. Verification, Model Checking and Abstract Interpretation (VMCAI'10), LNCS*, vol. 5944, pp. 362–379. Springer (2010)
44. Wachter, B., Zhang, L., Hermanns, H.: Probabilistic model checking modulo theories. In: *Proc. 4th Int. Conf. Quantitative Evaluation of Systems (QEST'07)*, pp. 129–138. IEEE CS Press (2006)