

Game-Based Probabilistic Predicate Abstraction in PRISM¹

Mark Kattenbelt Marta Kwiatkowska Gethin Norman
David Parker

Oxford University Computing Laboratory, Oxford, UK

Abstract

Modelling and verification of systems such as communication, network and security protocols, which exhibit both probabilistic and non-deterministic behaviour, typically use Markov Decision Processes (MDPs). For large, complex systems, abstraction techniques are essential. This paper builds on a promising approach for abstraction of MDPs based on stochastic two-player games which provides distinct lower and upper bounds for minimum and maximum probabilistic reachability properties. Existing implementations work at the model level, limiting their scalability. In this paper, we develop language-level abstraction techniques that build game-based abstractions of MDPs directly from high-level descriptions in the PRISM modelling language, using predicate abstraction and SMT solvers. For efficiency, we develop a compositional framework for abstraction. We have applied our techniques to a range of case studies, successfully verifying models larger than was possible with existing implementations. We are also able to demonstrate the benefits of adopting a compositional approach.

Keywords: Automatic verification, probabilistic model checking, predicate abstraction, Markov models.

1 Introduction

Verification of systems that exhibit both non-deterministic and probabilistic behaviour has proved to be very useful in domains such as communication and network protocols, security protocols, and randomised distributed algorithms. Markov Decision Processes (MDPs) are a natural model for such systems and several tools, such as PRISM [13] and LiQuor [4], implement efficient solution methods for these models. As in the field of non-probabilistic model checking, however, the state space explosion problem tends to limit the scalability of these approaches and techniques to counter this are an important area of research.

Of particular current interest are the development of *abstraction techniques* for the verification of MDPs [6,8,18,22]. In this paper, we use the abstraction approach of [18], which is based on stochastic two-player games. The key idea is to separate the non-determinism that is introduced by the abstraction from the

¹ Supported in part by EPSRC grants EP/D07956X and EP/D076625.

non-determinism present in the original MDP. This results in abstract models that provide distinct upper and lower bounds on minimum and maximum reachability probabilities. This is in contrast to alternative abstraction methods [6], where only an upper bound on the maximum probability and a lower bound on the minimum probability can be extracted. Besides being a more informative abstraction, these bounds also provide a measure of the quality of the abstraction. This information is potentially very useful when considering refinement.

A limitation of the existing implementation in [18] is that abstractions are performed at the model level, i.e. the full concrete model (MDP) is constructed and then reduced to the corresponding stochastic game. In this paper, we develop techniques to construct the abstraction directly from a high-level description of the MDP (in this case the modelling language of PRISM) using predicate abstraction [12,1,5], which has been very successful in the non-probabilistic setting.

Predicate abstraction for PRISM models was recently considered in [22], but using the abstraction technique of [6] which represents abstractions as MDPs. Applying predicate abstraction to the approach of [18] provides the additional benefits of the game-based approach but proves to be more involved. This is because the game-based abstraction preserves additional information which is non-trivial to extract from language-level descriptions of PRISM models.

We present a compositional variant of game-based abstraction of MDPs, explain how to apply it at the level of the PRISM modelling language, and describe an implementation of these techniques using SMT solvers and ‘on-the-fly’ abstraction. We illustrate its applicability on several examples, successfully analysing models larger than is possible with the implementation of [18] and improving performance on others. We also analyse the benefits of employing a compositional approach.

The remainder of this paper is structured as follows. Section 2 provides background material, including the PRISM modelling language and its semantics. In Section 3 we present a compositional variant of the game-based abstraction method of [18]. In Section 4, we give a game-based variant of PRISM called A-PRISM and describe a predicate abstraction procedure for PRISM models that results in A-PRISM models. Sections 5 and 6 describe our implementation and present experimental results from several case studies. We conclude with a discussion of related work and ideas for future development.

Proofs can be found in the technical report version of this paper [16].

2 Background

We assume a set of typed variables V . A *valuation* of V is a function s mapping each variable in V to a value in its domain. We let $\text{val}(V)$ denote the set of all valuations of V and, for any $s \in \text{val}(V)$ and $V' \subseteq V$, let $s|_{V'}$ denote the restriction of s to the domain V' . Furthermore, if $s_1 \in \text{val}(V_1)$, $s_2 \in \text{val}(V_2)$ and $V_1 \cap V_2 = \emptyset$, we let $s_1 \parallel s_2$ denote the valuation of $V_1 \cup V_2$ where $(s_1 \parallel s_2)|_{V_1} = s_1$ and $(s_1 \parallel s_2)|_{V_2} = s_2$. We will often refer to valuations as *states*. We also assume a finite set Act of *actions* and an additional ‘silent’ action $\tau \notin Act$.

A probability distribution over a finite set S is a function $\mu : S \rightarrow [0, 1]$ such that $\sum_{s \in S} \mu(s) = 1$. Let $\text{Dist}(S)$ denote the set of all distributions over S . For any $s \in S$,

let η_s denote the point distribution at s . If $\mu_1 \in \text{Dist}(\text{val}(V_1))$, $\mu_2 \in \text{Dist}(\text{val}(V_2))$ and $V_1 \cap V_2 = \emptyset$, let $\mu_1 \parallel \mu_2$ denote the distribution over $\text{val}(V_1 \cup V_2)$ such that $(\mu_1 \parallel \mu_2)(s) = \mu_1(s|_{V_1}) \cdot \mu_2(s|_{V_2})$ for all $s \in \text{val}(V_1 \cup V_2)$.

Definition 2.1 Let V, V' be sets of variables such that $V' \subseteq V$. A *transition* from V to V' is a tuple $\langle s, \text{step} \rangle$ where $s \in \text{val}(V)$ and $\text{step} \subseteq (\text{Act} \cup \{\tau\}) \times \text{Dist}(\text{val}(V'))$.

A transition $\langle s, \text{step} \rangle$ consists of a source state s and non-deterministic choice step between pairs comprising an action and a distribution over target states. We now define (standard CSP-style) parallel composition of transitions.

Definition 2.2 Suppose $V_1, V_2 \subseteq V$ are disjoint sets of variables, $\langle s, \text{step}_i \rangle$ is a transition from V to V_i for $i \in \{1, 2\}$ and $A \subseteq \text{Act}$. Let $\langle s, \text{step}_1 \rangle \parallel [A] \langle s, \text{step}_2 \rangle$ denote the transition $\langle s, \text{step} \rangle$ from V to $V_1 \cup V_2$ where $\langle a, \mu \rangle \in \text{step}$ if and only if one of the following conditions holds:

1. $a \notin A$ and $\mu = \mu_1 \parallel \eta_{(s|_{V_2})}$ for some $\langle a, \mu_1 \rangle \in \text{step}_1$;
2. $a \notin A$ and $\mu = \eta_{(s|_{V_1})} \parallel \mu_2$ for some $\langle a, \mu_2 \rangle \in \text{step}_2$;
3. $a \in A$ and $\mu = \mu_1 \parallel \mu_2$ for some $\langle a, \mu_1 \rangle \in \text{step}_1$ and $\langle a, \mu_2 \rangle \in \text{step}_2$.

2.1 Controlled Markov Decision Processes

The techniques introduced in this paper are for Markov Decision Processes (MDPs). However, in order to adopt a compositional approach, we use a variant called Controlled Markov Decision Processes which represent components of an MDP. These are similar to the *probabilistic modules* of [7].

Definition 2.3 A *Controlled Markov Decision Process* (CMDP) is a tuple $\mathcal{C} = \langle V, V^{\text{ctrl}}, V^{\text{ext}}, \text{Act}, s^{\text{init}}, \text{Steps} \rangle$ where:

- V is a finite set of typed variables;
- V^{ctrl} and V^{ext} partition V into *controlled* and *external* variables;
- Act is a finite set of actions;
- $s^{\text{init}} \in \text{val}(V^{\text{ctrl}})$ is the initial valuation;
- $\text{Steps} : \text{val}(V) \rightarrow \mathcal{P}((\text{Act} \cup \{\tau\}) \times \text{Dist}(\text{val}(V^{\text{ctrl}})))$ is the transition function.

A CMDP specifies the initial values of its controlled variables and how these variables are updated. These updates depend on the values of both its controlled variables and the external variables, which are assumed to be under the control of other components in the system. Given a valuation of all variables $s \in \text{val}(V)$ the set of action-distribution pairs $\text{Steps}(s)$ represents a non-deterministic choice between several behaviours. If the $\langle a, \mu \rangle$ is chosen, then the CMDP performs action a and then probabilistically selects a new valuation of its controlled variables according to μ . The transition function can equivalently be defined as the set $\{\langle s, \text{Steps}(s) \rangle \mid s \in \text{val}(V)\}$ of transitions from V to V^{ctrl} .

We now describe the parallel composition of CMDPs. Note that CMDPs can only be combined in parallel when they agree on the total set of variables and their control variables are disjoint. We call such CMDPs *composable*. Let $\mathcal{C}_i = \langle V, V_i^{\text{ctrl}}, V_i^{\text{ext}}, \text{Act}_i, s_i^{\text{init}}, \text{Steps}_i \rangle$ for $i \in \{1, 2\}$.

Definition 2.4 The parallel composition of two composable CMDPs \mathcal{C}_1 and \mathcal{C}_2 is the CMDP $\mathcal{C}_1 \parallel [A] \mathcal{C}_2 = \langle V, V^{\text{ctrl}}, V^{\text{ext}}, \text{Act}, s^{\text{init}}, \text{Steps} \rangle$ where:

- $V^{\text{ctrl}} = V_1^{\text{ctrl}} \cup V_2^{\text{ctrl}}$;
- $V^{\text{ext}} = (V_1^{\text{ext}} \cup V_2^{\text{ext}}) \setminus (V_1^{\text{ctrl}} \cup V_2^{\text{ctrl}})$;
- $\text{Act} = \text{Act}_1 \cup \text{Act}_2$;
- $s^{\text{init}} = s_1^{\text{init}} \parallel s_2^{\text{init}}$;
- if $s \in \text{val}(V)$, then $\langle s, \text{Steps}(s) \rangle = \langle s, \text{Steps}_1(s) \parallel [A] \langle s, \text{Steps}_2(s) \rangle \rangle$.

We can also define *action renaming* and *action hiding* operations for CMDPs, but for brevity we will omit these from the presentation in this paper.

Example 2.5 Consider the CMDP $\mathcal{C}_{\text{walk}}$ where $V^{\text{ctrl}} = \{\text{val}\}$, $V^{\text{ext}} = \{\text{close}\}$, val has domain $\{0, \dots, 4\}$ with initial value 2 and close has the type Boolean. For any valuation $(v, c) \in \text{val}(\{\text{val}, \text{close}\})$, $\text{Steps}_{\mathcal{C}_{\text{walk}}}(v, c) = \{\langle \tau, \frac{1}{2} \cdot \eta_{v-1} + \frac{1}{2} \cdot \eta_{v+1} \rangle, \langle \text{read}, \eta_v \rangle\}$ if $v \in \{1, 2, 3\}$ and equals $\{\langle \text{read}, \eta_v \rangle\}$ otherwise. The CMDP models a random walk that can, at any time, do a *read* action. Let $\mathcal{C}_{\text{walk}}$ be composed with the CMDP \mathcal{C}_{obs} in which $V^{\text{ctrl}} = \{\text{close}\}$, $V^{\text{ext}} = \{\text{val}\}$, close has initial value **false**, and, for $(v, c) \in \text{val}(\{\text{val}, \text{close}\})$, $\text{Steps}_{\mathcal{C}_{\text{obs}}}(v, c) = \{\langle \text{read}, \eta_{\text{close}} \rangle\}$ if $v \neq 2$ and equals $\{\langle \text{read}, \eta_{\text{close}} \rangle\}$ otherwise. This models a CMDP which performs action *read*, and updates *close* depending on whether *val* is close to the ends of the walk or not. Graphical representations of these CMDPs are given in Figure 1(a).

Note that a CMDP for which $V^{\text{ext}} = \emptyset$ (for example the parallel composition of CMDPs whose controlled variables partition V) is simply an MDP. For an MDP, we are typically interested in quantitative aspects such as probabilistic reachability. An *adversary* of an MDP is a particular resolution of the non-determinism. Given an MDP, a valuation $s \in \text{val}(V)$, a set of valuations $F \subseteq \text{val}(V)$ and an adversary \mathcal{A} we use $\mathbf{p}_s^{\mathcal{A}}(F)$ to denote the probability of reaching F from s under adversary \mathcal{A} , defined in the usual way [17]. We use $\mathbf{p}_s^-(F) = \inf_{\mathcal{A}} \mathbf{p}_s^{\mathcal{A}}(F)$ and $\mathbf{p}_s^+(F) = \sup_{\mathcal{A}} \mathbf{p}_s^{\mathcal{A}}(F)$ to denote the extremal probabilities of reaching F from s [2].

2.2 PRISM Models

We now describe the modelling language used by the PRISM [13] to describe MDPs. This language is based on guarded commands extended with probabilistic choices.

Definition 2.6 A PRISM model is a tuple $\mathbf{P} = \langle \text{var}(\mathbf{P}), \text{sys}, \{M_1, \dots, M_m\} \rangle$ consisting of a finite set of (Boolean or integer) variables $\text{var}(\mathbf{P})$, a system definition sys and a finite set modules $\{M_1, \dots, M_m\}$. The system definition sys is a process algebraic expression containing each of the m modules exactly once. Each module M consists of:

- a finite set of local variables $\text{var}(M) \subseteq \text{var}(\mathbf{P})$ such that:
 - $\text{var}(M)$ are disjoint from the local variables of all other modules;
 - each variable $\text{var} \in \text{var}(M)$ has the initial value $\text{init}(\text{var})$;
 - $\text{init}(M) \in \text{val}(\text{var}(M))$ denotes the initial values of $\text{var}(M)$;
- a finite set of commands $\text{com}(M)$ where each command $\text{cmd} \in \text{com}(M)$ includes:
 - a guard $\text{guard}(\text{cmd})$ which is a Boolean function over $\text{val}(\text{var}(\mathbf{P}))$;

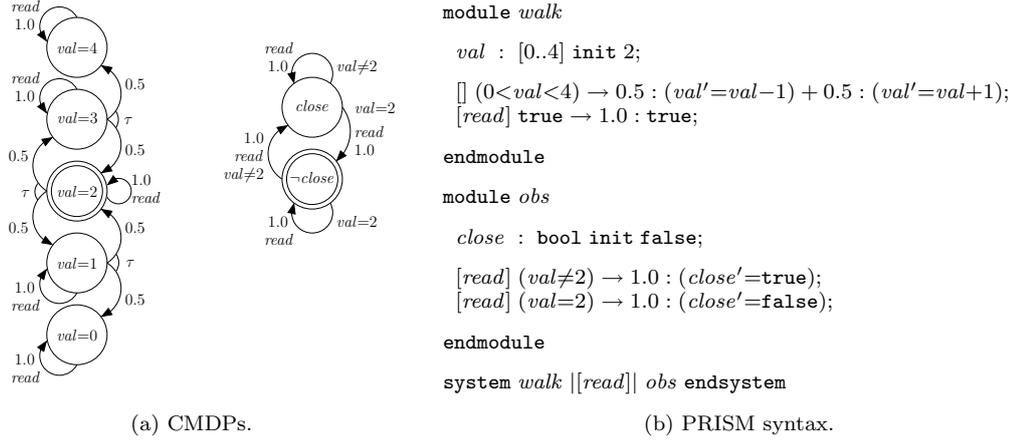


Fig. 1. Simple example: a random walk and observer process (see Examples 2.5 and 2.8).

- an action $\text{act}(cmd)$;
- a finite set of updates $\text{updates}(cmd) = \{\langle \lambda_1, u_1 \rangle, \dots, \langle \lambda_n, u_n \rangle\}$ such that $\lambda_i \in (0, 1]$, $\sum_{i=1}^n \lambda_i = 1$ and u_i is a function from $\text{val}(\text{var}(\mathbf{P}))$ to $\text{val}(\text{var}(M))$.

For each command cmd of a module M and valuation s of $\text{var}(\mathbf{P})$, supposing $\text{updates}(cmd) = \{\langle \lambda_1, u_1 \rangle, \dots, \langle \lambda_n, u_n \rangle\}$, let $\text{dist}(cmd, s)$ denote the distribution over $\text{val}(\text{var}(M))$ such that $\text{dist}(cmd, s)(s') = \sum_{u_i(s)=s'} \lambda_i$ for all $s' \in \text{val}(\text{var}(M))$. Intuitively, $\text{dist}(cmd, s)(s')$ is the probability that performing cmd in s updates the module's local variables to s' .

Definition 2.7 The semantics of a module M of a PRISM model \mathbf{P} is given by the CMDP $\llbracket M \rrbracket = \langle V, V^{\text{ctrl}}, V^{\text{ext}}, \text{Act}, s^{\text{init}}, \text{Steps} \rangle$ where:

- $V = \text{var}(\mathbf{P})$, $V^{\text{ctrl}} = \text{var}(M)$ and $V^{\text{ext}} = \text{var}(\mathbf{P}) \setminus \text{var}(M)$;
- $\text{Act} = \{\text{act}(cmd) \mid cmd \in \text{com}(M)\} \setminus \{\tau\}$;
- $s^{\text{init}} = \text{init}(M)$;
- if $s \in \text{val}(\text{var}(\mathbf{P}))$, then $\langle a, \mu \rangle \in \text{Steps}(s)$ if and only if there exists $cmd \in \text{com}(M)$ such that $\text{guard}(cmd)(s)$ holds and $\langle \text{act}(cmd), \text{dist}(cmd, s) \rangle = \langle a, \mu \rangle$.

The semantics $\llbracket \mathbf{P} \rrbracket$ of a PRISM model \mathbf{P} is defined according to its system definition sys , using the semantics $\llbracket M \rrbracket$ of each individual module M , given above, and parallel composition of CMDPs (see Definition 2.4). We assume that $\text{var}(\mathbf{P})$ is the disjoint union $\cup_{i=1}^m \text{var}(M_i)$, and hence, for the parallel composition of all modules in a PRISM model, the set V^{ext} is empty. In other words, the semantics of a PRISM model is given by an MDP.

Example 2.8 Figure 1(b) presents a PRISM model of the CMDP in Example 2.5.

3 Abstraction of CMDPs

In this section we introduce abstractions of CMDPs, using the stochastic two-player game approach of [18] and predicates. A predicate φ is *over* variables V if all valuations of V uniquely determine the truth value of φ and we write $\varphi(s)$ to denote

the value of φ for a valuation s of V . Given a set of predicates Φ , let $\text{bool}(\Phi)$ be the set of Boolean variables indexed by the predicates in Φ , i.e. the set $\{b_\varphi \mid \varphi \in \Phi\}$. Furthermore, for abstraction of a particular CMDP using Φ , we will require that every predicate is either over *only* controlled variables or *only* external variables of this component. This partitions the predicates into Φ^{ctrl} and Φ^{ext} .

3.1 Abstract Controlled Markov Decision Processes

In order to present a compositional variant of game-based MDP abstraction, we introduce Abstract Controlled Markov Decision Processes (ACMDPs) which are a variant of the class of stochastic two-player games used in [18].

Definition 3.1 An *Abstract Controlled Markov Decision Process* (ACMDP) is a tuple $\mathcal{A} = \langle \overline{V}, \overline{V}^{\text{ctrl}}, \overline{V}^{\text{ext}}, \overline{Act}, \overline{s}^{\text{init}}, \overline{Steps} \rangle$ where:

- \overline{V} is a set of typed variables;
- $\overline{V}^{\text{ctrl}}$ and $\overline{V}^{\text{ext}}$ partition \overline{V} into controlled and external variables;
- \overline{Act} is a finite set of actions;
- $\overline{s}^{\text{init}} \in \text{val}(\overline{V}^{\text{ctrl}})$ is the initial valuation;
- $\overline{Steps} : \text{val}(\overline{V}) \rightarrow \mathcal{P}(\mathcal{P}((\overline{Act} \cup \{\tau\}) \times \text{Dist}(\text{val}(\overline{V}^{\text{ctrl}}))))$ is the transition function.

The crucial difference between CMDPs and ACMDPs is that the transition function now returns *sets of sets* of action-distribution pairs. This means ACMDPs capture two levels of non-determinism: the choice of a set of action-distribution pairs, and then the choice of an element in this set. This two-level non-determinism is equivalent to that of the stochastic two-player games used in [18], where first *player 1* makes a choice, then *player 2* does, followed by a *probabilistic* choice.

We now describe the parallel composition of ACMDPs. As for CMDPs, ACMDPs can only be combined when they agree on the total set of variables and their control variables are disjoint. We call such ACMDPs *composable*. Let $\mathcal{A}_i = \langle \overline{V}_i, \overline{V}_i^{\text{ctrl}}, \overline{V}_i^{\text{ext}}, \overline{Act}_i, \overline{s}_i^{\text{init}}, \overline{Steps}_i \rangle$ for $i \in \{1, 2\}$.

Definition 3.2 The parallel composition of two composable ACMDPs \mathcal{A}_1 and \mathcal{A}_2 is the ACMDP $\mathcal{A}_1 \parallel [A] \mathcal{A}_2 = \langle \overline{V}, \overline{V}^{\text{ctrl}}, \overline{V}^{\text{ext}}, \overline{Act}, \overline{s}^{\text{init}}, \overline{Steps} \rangle$ where

- $\overline{V}^{\text{ctrl}} = \overline{V}_1^{\text{ctrl}} \cup \overline{V}_2^{\text{ctrl}}$;
- $\overline{V}^{\text{ext}} = (\overline{V}_1^{\text{ext}} \cup \overline{V}_2^{\text{ext}}) \setminus (\overline{V}_1^{\text{ctrl}} \cup \overline{V}_2^{\text{ctrl}})$;
- $\overline{Act} = \overline{Act}_1 \cup \overline{Act}_2$;
- $\overline{s}^{\text{init}} = \overline{s}_1^{\text{init}} \parallel \overline{s}_2^{\text{init}}$;
- if $\overline{s} \in \text{val}(\overline{V})$, then $\overline{\text{step}} \in \overline{Steps}(\overline{s})$ if and only if $\langle \overline{s}, \overline{\text{step}} \rangle = \langle \overline{s}, \overline{\text{step}}_1 \rangle \parallel [A] \langle \overline{s}, \overline{\text{step}}_2 \rangle$ for some $\overline{\text{step}}_1 \in \overline{Steps}_1(\overline{s})$ and $\overline{\text{step}}_2 \in \overline{Steps}_2(\overline{s})$.

Like the relation between CMDPs and MDPs, an ACMDP for which $\overline{V}^{\text{ext}} = \emptyset$ is equivalent to a stochastic two-player game from [18]. A *player 1 strategy* in such an ACMDP is a particular resolution of the first non-deterministic choice of transitions in the ACMDP, whereas a *player 2 strategy* resolves the second non-deterministic choice. Given a valuation $\overline{s} \in \text{val}(\overline{V})$, a set of valuations $\overline{F} \subseteq \text{val}(\overline{V})$ and strategy pair σ_1, σ_2 , we use $\mathbf{p}_{\overline{s}}^{\sigma_1, \sigma_2}(\overline{F})$ to denote the probability of reaching \overline{F} from \overline{s} under

the strategies σ_1, σ_2 . Like for MDPs, we define extremal values as:

$$\begin{aligned} \mathbf{p}_{\bar{s}}^{--}(\bar{F}) &= \inf_{\sigma_1} \inf_{\sigma_2} \mathbf{p}_s^{\sigma_1, \sigma_2}(\bar{F}) & \mathbf{p}_{\bar{s}}^{+-}(\bar{F}) &= \sup_{\sigma_1} \inf_{\sigma_2} \mathbf{p}_s^{\sigma_1, \sigma_2}(\bar{F}) \\ \mathbf{p}_{\bar{s}}^{-+}(\bar{F}) &= \inf_{\sigma_1} \sup_{\sigma_2} \mathbf{p}_s^{\sigma_1, \sigma_2}(\bar{F}) & \mathbf{p}_{\bar{s}}^{++}(\bar{F}) &= \sup_{\sigma_1} \sup_{\sigma_2} \mathbf{p}_s^{\sigma_1, \sigma_2}(\bar{F}) \end{aligned}$$

3.2 Predicate Abstraction for CMDPs

In this section we introduce a compositional and predicate-based extension of the abstraction procedure described in [18]. Like in non-probabilistic predicate abstraction [1], we will represent an abstract state using Boolean variables $\text{bool}(\Phi)$ indexed by a set of predicates Φ . We will denote abstractions with respect to Φ by $\alpha(\cdot, \Phi)$, which we now define for states, distributions, transitions and then CMDPs.

Definition 3.3 Given a set of variables V and predicates Φ over V , the abstractions of a valuation $s \in \text{val}(V)$ and distribution $\mu \in \text{Dist}(\text{val}(V))$ with respect to Φ are defined as follows:

- $\alpha(s, \Phi)$ is the valuation of $\text{bool}(\Phi)$ where $\alpha(s, \Phi)(b_\varphi) = \varphi(s)$ for all $\varphi \in \Phi$;
- $\alpha(\mu, \Phi)$ is the distribution over $\text{val}(\text{bool}(\Phi))$ where $\alpha(\mu, \Phi)(\bar{s}) = \sum_{\alpha(s, \Phi) = \bar{s}} \mu(s)$ for all $\bar{s} \in \text{val}(\text{bool}(\Phi))$.

Definition 3.4 Given a set of variables V , subset $V' \subseteq V$ and sets of predicates Φ and $\Phi' \subseteq \Phi$ over V and V' , the abstraction of a transition $\langle s, \text{step} \rangle$ from V to V' with respect to Φ , denoted $\alpha(\langle s, \text{step} \rangle, \Phi)$, is given by the transition $\langle \alpha(s, \Phi), \{ \langle a, \alpha(\mu, \Phi') \rangle \mid \langle a, \mu \rangle \in \text{step} \} \rangle$ from $\text{bool}(\Phi)$ to $\text{bool}(\Phi')$.

We now define an abstraction function over CMDPs. For the remainder of Section 3, we fix a CMDP $\mathcal{C} = \langle V, V^{\text{ctrl}}, V^{\text{ext}}, \text{Act}, s^{\text{init}}, \text{Steps} \rangle$ and set of predicates Φ over V .

Definition 3.5 The abstraction of CMDP \mathcal{C} with respect to the predicates Φ is the ACMDP $\alpha(\mathcal{C}, \Phi) = \langle \bar{V}, \bar{V}^{\text{ctrl}}, \bar{V}^{\text{ext}}, \text{Act}, \bar{s}^{\text{init}}, \bar{\text{Steps}} \rangle$ where:

- $\bar{V} = \text{bool}(\Phi)$;
- $\bar{V}^{\text{ctrl}} = \text{bool}(\Phi^{\text{ctrl}})$;
- $\bar{V}^{\text{ext}} = \text{bool}(\Phi^{\text{ext}})$;
- $\bar{s}^{\text{init}} = \alpha(s^{\text{init}}, \Phi^{\text{ctrl}})$;
- if $\bar{s} \in \text{val}(\bar{V})$, then $\bar{\text{step}} \in \bar{\text{Steps}}(\bar{s})$ if and only if there exists $s \in \text{val}(V)$ such that $\alpha(\langle s, \text{Steps}(s) \rangle, \Phi) = \langle \bar{s}, \bar{\text{step}} \rangle$.

Example 3.6 Consider the CMDP $\mathcal{C}_{\text{walk}}$ of Example 2.5 and set of predicates $\Phi = \{(val=0), (val=4), (close)\}$. Applying Definition 3.5, we obtain the ACMDP depicted in Figure 2(a) with $\bar{V}^{\text{ctrl}} = \{b_{val=0}, b_{val=4}\}$ and $\bar{V}^{\text{ext}} = \{b_{close}\}$. The states of the ACMDP are shown as rectangles and the initial state as a double rectangle. The first non-deterministic choice is represented by the circles within a state, and the second non-deterministic choice by outgoing distributions from a circle. Since the external variables have no influence, they are omitted.

It is straightforward to show that applying Definition 3.5 to a CMDP for which $V^{\text{ext}} = \emptyset$ yields an ACMDP (with $\bar{V}^{\text{ext}} = \emptyset$) equivalent to the stochastic two-player

game derived from the abstraction procedure described in [18]. Therefore the results of [18] carry over to this setting; in particular, by analysing the ACMDP $\alpha(\mathcal{C}, \Phi)$ we can obtain upper and lower bounds for both minimum and maximum reachability probabilities of the CMDP \mathcal{C} . More formally, for all concrete valuations $s \in \text{val}(V)$ and abstract reachability objectives $\bar{F} \subseteq \text{val}(\bar{V})$ we have:

$$\begin{aligned} \mathbf{p}_{\alpha(s, \Phi)}^{--}(\bar{F}) &\leq \mathbf{p}_s^-(\{s' \mid \alpha(s', \Phi) \in \bar{F}\}) \leq \mathbf{p}_{\alpha(s, \Phi)}^{+-}(\bar{F}) \\ \mathbf{p}_{\alpha(s, \Phi)}^{-+}(\bar{F}) &\leq \mathbf{p}_s^+(\{s' \mid \alpha(s', \Phi) \in \bar{F}\}) \leq \mathbf{p}_{\alpha(s, \Phi)}^{++}(\bar{F}) \end{aligned}$$

3.3 Compositional Abstraction of CMDPs

The abstraction of Definition 3.5 can be applied to any CMDP but, as the following example demonstrates, parallel composition and abstraction do not commute.

Example 3.7 Consider again Example 2.5 (Figure 1(a)) and set of predicates $\Phi = \{(val=0), (val=4), (close)\}$. For the abstract valuation $\bar{s} = (\neg b_{val=0}, \neg b_{val=4}, b_{close})$ from Definition 3.5 it follows that $\{\langle \tau, \eta_{(\neg b_{val=0}, \neg b_{val=4})} \rangle, \langle read, \eta_{(\neg b_{val=0}, \neg b_{val=4})} \rangle\}$ and $\{\langle read, \eta_{b_{close}} \rangle\}$ are in $\overline{Steps}_{\alpha(\mathcal{C}_{obs}, \Phi)}(\bar{s})$. Therefore, by Definition 3.2, it follows that $\{\langle \tau, \eta_{\bar{s}} \rangle, \langle read, \eta_{\bar{s}} \rangle\}$ is in $\overline{Steps}_{\alpha(\mathcal{C}_{walk}, \Phi) \parallel [read] \alpha(\mathcal{C}_{obs}, \Phi)}(\bar{s})$. However, no valuation (v, c) abstracts to \bar{s} and induces this transition in $\alpha(\mathcal{C}_{walk} \parallel [read] \mathcal{C}_{obs}, \Phi)$. More precisely, if the τ transition abstracts to a self-loop, then $v=2$, and if the $read$ transition sets $close$ to true , then $v \neq 2$.

As this example illustrates, a compositional abstraction may introduce spurious transitions, resulting in an over-approximation of the non-compositional abstraction and thus less precise lower and upper bounds for probabilistic reachability. Although such abstractions may still lead to useful results, we now introduce the notion of *abstraction preserving* CMDPs, for which compositional abstraction is precise (i.e. equivalent to the non-compositional abstraction).

Definition 3.8 The CMDP \mathcal{C} is called *abstraction preserving* with respect to the predicates Φ if for any $s, s' \in \text{val}(V)$ such that $s \upharpoonright_{V^{\text{ctrl}}} = s' \upharpoonright_{V^{\text{ctrl}}}$ and $\alpha(s, \Phi^{\text{ext}}) = \alpha(s', \Phi^{\text{ext}})$, then $\alpha(\langle s, Steps(s) \rangle, \Phi) = \alpha(\langle s', Steps(s') \rangle, \Phi)$.

Intuitively, this states that any valuations which agree on control variables and satisfy the same external predicates yield the same abstract transitions. As the following two results show, this property is both preserved under parallel composition and ensures a precise abstraction under parallel composition.

Proposition 3.9 Let \mathcal{C}_1 and \mathcal{C}_2 be composable CMDPs and A be a set of actions. If \mathcal{C}_1 and \mathcal{C}_2 are abstraction preserving with respect to the predicates Φ , then their composition $\mathcal{C}_1 \parallel [A] \mathcal{C}_2$ is also abstraction preserving with respect to Φ .

Proposition 3.10 Let \mathcal{C}_1 and \mathcal{C}_2 be composable CMDPs and A be a set of actions. If \mathcal{C}_1 and \mathcal{C}_2 are abstraction preserving with respect to the predicates Φ , then:

$$\alpha(\mathcal{C}_1, \Phi) \parallel [A] \alpha(\mathcal{C}_2, \Phi) = \alpha(\mathcal{C}_1 \parallel [A] \mathcal{C}_2, \Phi).$$

From Proposition 3.9 and Proposition 3.10, we can infer that a compositional abstraction is precise if each individual component is abstraction preserving.

Example 3.11 Consider the CMDP \mathcal{C}_{obs} from Example 2.5 and set of predicates $\Phi = \{(val=0), (val=4), (close)\}$. This CMDP is not abstraction preserving with respect to Φ . For example, the valuations $s = (2, \mathbf{true})$ and $s' = (3, \mathbf{true})$ agree on the value of $close$ and $\alpha(s, \Phi^{ext}) = \alpha(s', \Phi^{ext}) = (\neg b_{val=0}, \neg b_{val=4})$. However, $\alpha(\langle s, Steps_{\mathcal{C}_{obs}}(s) \rangle, \Phi) = \langle \alpha(s, \Phi), \{\langle read, \eta_{\neg b_{close}} \rangle\} \rangle$ while $\alpha(\langle s', Steps_{\mathcal{C}_{obs}}(s') \rangle, \Phi) = \langle \alpha(s', \Phi), \{\langle read, \eta_{b_{close}} \rangle\} \rangle$. If we extend Φ with the predicate $(val=2)$, then \mathcal{C}_{obs} is abstraction preserving.

4 Abstraction of PRISM Models

Suppose we wish to abstract a PRISM model. One possibility is to (compositionally or non-compositionally) apply the abstraction method of Section 3 to its CMDP semantics. In either case, the disadvantage of such a method is that the concrete CMDPs have to be constructed, limiting the applicability of the approach. In this section we define a language-level abstraction method to remedy the situation.

4.1 A-PRISM Models

For our language-level abstraction, we introduce the A-PRISM language, an extension of the PRISM language with an additional element of choice.

Definition 4.1 An A-PRISM model is a tuple $\mathbf{A} = \langle \overline{\mathbf{var}}(\mathbf{A}), \overline{\mathbf{sys}}, \{\overline{M}_1, \dots, \overline{M}_m\} \rangle$. The only difference between this and a PRISM model is the definition of the commands $\overline{\mathbf{com}}(\overline{M})$ for each module \overline{M} . Each command $\overline{cmd} \in \overline{\mathbf{com}}(\overline{M})$ includes:

- a guard $\overline{\mathbf{guard}}(\overline{cmd})$ which is a Boolean function over $\mathbf{val}(\overline{\mathbf{var}}(\mathbf{A}))$;
- a finite set of choices $\overline{\mathbf{choices}}(\overline{cmd})$ where each $\overline{chc} \in \overline{\mathbf{choices}}(\overline{cmd})$ consists of an action $\overline{\mathbf{act}}(\overline{chc})$ and a finite set of updates $\overline{\mathbf{updates}}(\overline{chc}) = \{\langle \lambda_1, \overline{u}_1 \rangle, \dots, \langle \lambda_n, \overline{u}_n \rangle\}$ such that $\lambda_i \in (0, 1]$, $\sum_{i=1}^n \lambda_i = 1$ and \overline{u}_i is a function from $\overline{\mathbf{var}}(\mathbf{A})$ to $\overline{\mathbf{var}}(\overline{M})$.

For a choice \overline{chc} of a command and valuation $\overline{s} \in \mathbf{val}(\overline{\mathbf{var}}(\mathbf{A}))$, supposing $\overline{\mathbf{updates}}(\overline{chc}) = \{\langle \lambda_1, \overline{u}_1 \rangle, \dots, \langle \lambda_n, \overline{u}_n \rangle\}$, let $\overline{\mathbf{dist}}(\overline{chc}, \overline{s})$ denote the distribution over $\overline{\mathbf{var}}(\overline{M})$ such that $\overline{\mathbf{dist}}(\overline{chc}, \overline{s})(\overline{s}') = \sum_{\overline{u}_i(\overline{s})=\overline{s}'} \lambda_i$ for all $\overline{s}' \in \mathbf{val}(\overline{\mathbf{var}}(\overline{M}))$.

Definition 4.2 The semantics of a module \overline{M} of A-PRISM model \mathbf{A} is given by the ACMDP $\llbracket \overline{M} \rrbracket = \langle \overline{V}, \overline{V}^{ctrl}, \overline{V}^{ext}, \overline{Act}, \overline{s}^{init}, \overline{Steps} \rangle$, where:

- $\overline{V} = \overline{\mathbf{var}}(\mathbf{A})$, $\overline{V}^{ctrl} = \overline{\mathbf{var}}(\overline{M})$ and $\overline{V}^{ext} = \overline{\mathbf{var}}(\mathbf{A}) \setminus \overline{\mathbf{var}}(\overline{M})$;
- $\overline{Act} = \{\overline{\mathbf{act}}(\overline{chc}) \mid \overline{cmd} \in \overline{\mathbf{com}}(\overline{M}), \overline{chc} \in \overline{\mathbf{choices}}(\overline{cmd})\} \setminus \{\tau\}$;
- $\overline{s}^{init} = \overline{\mathbf{init}}(\overline{M})$;
- if $\overline{s} \in \mathbf{val}(\overline{\mathbf{var}}(\mathbf{A}))$, then $\overline{\mathbf{step}} \in \overline{Steps}(\overline{s})$ if and only if there exists a command $\overline{cmd} \in \overline{\mathbf{com}}(\overline{M})$ such that $\overline{\mathbf{step}} = \{\overline{\mathbf{act}}(\overline{chc}), \overline{\mathbf{dist}}(\overline{chc}, \overline{s}) \mid \overline{chc} \in \overline{\mathbf{choices}}(\overline{cmd})\}$ and $\overline{\mathbf{guard}}(\overline{cmd})(\overline{s})$ holds.

The semantics $\llbracket \mathbf{A} \rrbracket$ of an A-PRISM model \mathbf{A} is defined according to its system definition $\overline{\mathbf{sys}}$, using the semantics $\llbracket \overline{M} \rrbracket$ of each individual module \overline{M} , given above, and parallel composition of ACMDPs (see Definition 3.2).

The ‘first’ non-deterministic choices of $\llbracket \overline{M} \rrbracket$ are caused by overlaps between the guards of commands, whereas the ‘second’ non-deterministic choices are induced by

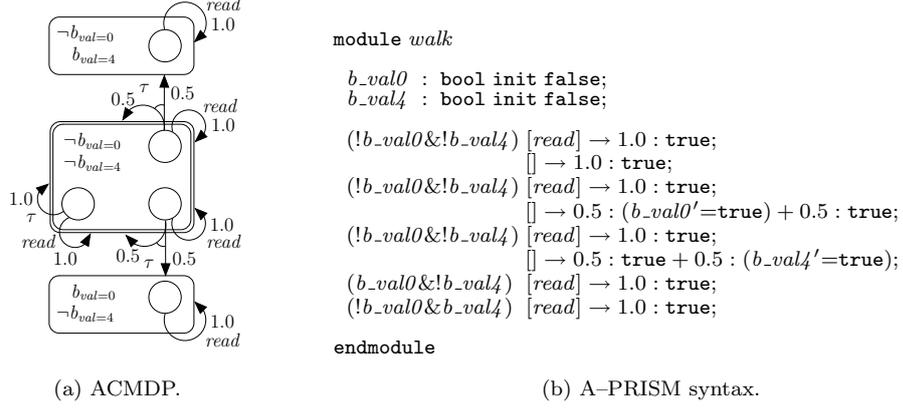


Fig. 2. Abstraction of the random walk component of Figure 1 (see Example 3.6).

the choices within commands.

Example 4.3 Figure 2 shows an A-PRISM module and its ACMDP semantics.

4.2 Language-level Abstraction of PRISM Models

In this section we introduce a language-level abstraction method for PRISM. We assume a fixed PRISM model $\mathbf{P} = \langle \text{var}(\mathbf{P}), \text{sys}, \{M_1, \dots, M_m\} \rangle$ and a set of predicates Φ which is partitioned into subsets $\Phi^{M_1}, \dots, \Phi^{M_m}$ over the local variables of the modules M_1, \dots, M_m . The abstraction of \mathbf{P} is defined as the A-PRISM model:

$$\beta(\mathbf{P}, \Phi) = \langle \text{bool}(\Phi), \beta(\text{sys}), \{\beta(M_1, \Phi), \dots, \beta(M_m, \Phi)\} \rangle$$

where the system definition $\beta(\text{sys})$ is a syntactic copy of sys and each module M is replaced by the language-level abstraction $\beta(M, \Phi)$, defined below.

Definition 4.4 The language-level abstraction of a module M of \mathbf{P} is the A-PRISM module $\beta(M, \Phi)$ where:

- the local variables $\overline{\text{var}}(\beta(M, \Phi))$ are $\text{bool}(\Phi^M)$;
- the initial value $\overline{\text{init}}(b_\varphi)$ equals $\varphi(\text{init}(M))$ for all $b_\varphi \in \text{bool}(\Phi^M)$;
- the set of commands $\overline{\text{com}}(\beta(M, \Phi))$ equals $\{\overline{\text{cmd}}_s \mid s \in \text{val}(\text{var}(\mathbf{P}))\}$ where
 - $\overline{\text{guard}}(\overline{\text{cmd}}_s) = \bigwedge_{\varphi \in \Phi} (b_\varphi = \varphi(s))$,
 - $\overline{\text{choices}}(\overline{\text{cmd}}_s) = \{\overline{\text{cmd}} \mid \text{cmd} \in \text{com}(M), \text{guard}(\text{cmd})(s)\}$ with
 - $\overline{\text{act}}(\overline{\text{cmd}}) = \text{act}(\text{cmd})$
 - if $\text{updates}(\text{cmd}) = \{\langle \lambda_1, u_1 \rangle, \dots, \langle \lambda_n, u_n \rangle\}$ and \overline{u} is the constant function that returns $\alpha(u(s), \Phi^M)$, then $\text{updates}(\overline{\text{cmd}}) = \{\langle \sum_{\overline{u}_j = \overline{u}_i} \lambda_j, \overline{u}_i \rangle \mid 1 \leq i \leq n\}$.

The results below illustrate that using the language-level method we obtain the same abstraction as with the model-level abstraction of Section 3.2.

Proposition 4.5 *If M is a module of \mathbf{P} , then $\llbracket \beta(M, \Phi) \rrbracket = \alpha(\llbracket M \rrbracket, \Phi)$.*

Combining this with the results of the previous sections we have the following.

Theorem 4.6 *If $\llbracket M \rrbracket$ is abstract preserving with respect to Φ for each module M of \mathbf{P} , then $\llbracket \beta(\mathbf{P}, \Phi) \rrbracket = \alpha(\llbracket \mathbf{P} \rrbracket, \Phi)$.*

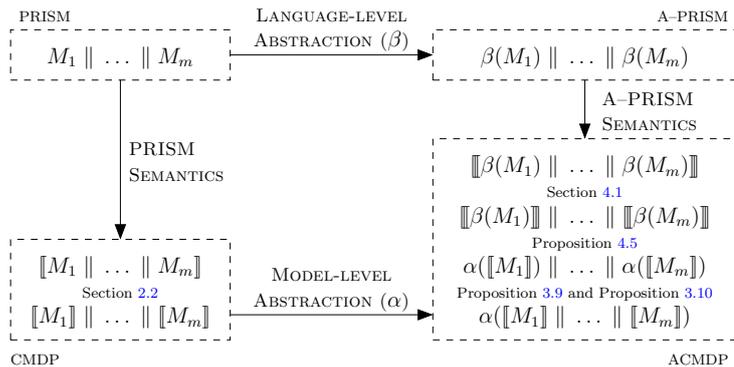


Fig. 3. Relation between model-level and language-level abstraction functions.

Figure 4.2 presents an overview of the correspondence between the model-level (α) and language-level (β) abstractions.

The remaining question is how to check abstraction preservation at the language level. We now outline a simple check which guarantees this. It is always possible to rewrite a single PRISM command into commands with disjoint guards such that the updates only contain local variables. Now, if we have a PRISM module containing only commands of this form, then to check its semantics is abstraction preserving it is sufficient to show that if $s, s' \in \text{val}(\text{var}(\mathbf{P}))$ such that $s \upharpoonright_{\text{var}(M)} = s' \upharpoonright_{\text{var}(M)}$ and $\alpha(s, \Phi) = \alpha(s', \Phi)$, then $\text{guard}(cmd)(s) = \text{guard}(cmd)(s')$ for all commands cmd .

5 Implementation

We have built prototype tools both for our language-level abstraction (translation from PRISM to A-PRISM) and model checking A-PRISM models. The model checker is a relatively simple extension of PRISM's MTBDD model checking engine for MDPs. In the remainder of this section we will discuss the abstraction tool.

5.1 Abstraction with SMT Solvers

The key step in the translation of a (concrete) PRISM model to an (abstract) A-PRISM model is the construction of abstract commands, as described in Definition 4.4. For this, the implementation uses ALL-SAT procedures over the theories of integer arithmetic and fixed-size bit-vectors through the SMT solver Yices [10]. This is based on the principles described in [5,20] for predicate abstraction of non-probabilistic systems.

In Definition 4.4 each abstract command is induced by a concrete valuation and the concrete commands enabled for this valuation. However, considering each concrete valuation individually is clearly inefficient. Our implementation therefore employs an approach which detects multiple valuations inducing identical commands. The basic idea is to instead enumerate what we call *overlaps*, which are combinations of commands that can be simultaneously enabled.

Formally, an overlap of module M from a PRISM model \mathbf{P} is a set of commands $\mathcal{O} \subseteq \text{com}(M)$ for which there exists $s \in \text{val}(\text{var}(\mathbf{P}))$ such that $cmd \in \mathcal{O}$ if and only if $\text{guard}(cmd)(s)$. Given a module M , we first find all overlaps of M with

an ALL-SAT procedure². Then, for a given overlap \mathcal{O} , we find the corresponding abstract commands, again with an ALL-SAT procedure. To optimise this approach, we remove unnecessary predicates both from the guards and updates of abstract commands. For example, we do not include any predicates in an abstract update if the corresponding concrete updates do not influence their values.

5.2 ‘On-the-Fly’ Abstraction

During prototyping, our implementation would often find a large number of overlaps, making the ALL-SAT procedures infeasible. However, further investigation revealed that the majority of these overlaps were induced by unreachable concrete valuations. Therefore, the prototype was extended with an ‘on-the-fly’ abstraction method to overcome this problem. Like in explicit-state model checking, this is achieved by keeping a stack of reachable abstract valuations of $\text{bool}(\Phi^M)$. Initially, this stack only contains the element $\alpha(s^{\text{init}}, \Phi^M)$. The method takes individual abstract valuations off the stack, constructs the abstract commands for this valuation and adds any new abstract valuations that are the target of this command to the stack. Note that, since the tool now constructs abstract commands for each abstract state \bar{s} separately, only commands that are enabled for some valuation s such that $\alpha(s, \Phi^M) = \bar{s}$ need be considered when searching for overlaps of these commands.

Although this ‘on-the-fly’ abstraction method does perform reachability over abstract states, it is important to stress that, unlike [18], it does not require the construction of the reachable concrete state space or take into account whether concrete states are reachable.

6 Experimental Results

We have tested the performance of our implementation on three case studies:³

- An extension of the sliding window protocol of [21] where channels lose messages probabilistically instead of non-deterministically and a notion of timeout is included. We fix the window size of the sender (2) and receiver (1), buffer size of the channels (2) and sequence numbers (modulo 4) while varying the number of data frames (D) in the source. We analyse ‘*the maximum probability of sending D data frames without a timeout*’ using an abstraction that removes the values of the data frames. In the compositional approach, we abstract the sender and data channel separately from the receiver and acknowledgement channel.
- IPv4 Zeroconf protocol [3], as described in [18], parameterised by the number of configured hosts (N) and with 64 IP addresses. We encode the abstraction of [18] into predicates and consider ‘*the minimum probability that the host eventually secures an IP address*’. In the compositional approach, the configuring host is abstracted separately from the channel and configured hosts.
- Israeli and Jalfon’s self-stabilisation protocol [15] for a ring with N processes. We encode the abstraction of [9] into predicates and analyse ‘*the minimum probability*

² The algorithm to check if a PRISM module is abstraction preserving can be implemented similarly.

³ Files for the case studies are available from <http://www.prismmodelchecker.org/files/qap108/>.

		Concrete Model			Abstract Models					
		Num. comm.	Num. states	Check time	Non-compositional		Compositional		Num. states	Check time
					Abstr. time	Num. Comm.	Abstr. time	Num. Comm.		
Sliding Window (D)	8	19	189,952	30.2	96.7	540	220	3,260	742	0.21
	10	19	987,136	153	126	706	336	4,870	964	0.43
	12	19	–	–	155	872	473	6,545	1,186	0.69
	14	19	–	–	200	1,038	630	8,225	1,408	1.23
	16	19	–	–	237	1,204	819	9,905	1,630	1.68
	18	19	–	–	285	1,370	962	11,585	1,852	2.47
	20	19	–	–	334	1,536	1,201	13,265	2,074	3.45
Zeroconf (N)	4	89	50,377	206	1,110	2,349	106	362	1,325	104
	5	109	113,217	355	1,480	2,523	183	396	1,421	134
	6	129	282,185	678	2,480	2,695	262	431	1,517	161
	7	149	426,529	952	3,630	2,762	434	444	1,549	175
	8	169	838,905	1,400	6,370	2,804	785	453	1,581	209
Israeli & Jalfon (N)	8	8	255	0.01	13.4	28	n/a	n/a	22	<0.01
	10	10	1,023	0.03	95.2	68	n/a	n/a	42	<0.01
	12	12	4,095	0.08	727	168	n/a	n/a	77	<0.01
	14	14	16,383	0.22	4,210	415	n/a	n/a	135	<0.01
	16	16	65,535	0.75	28,000	1,025	n/a	n/a	231	<0.01
	18	18	262,143	2.27	136,000	2,505	n/a	n/a	385	<0.01
	20	20	1,048,575	8.51	1,090,000	6,056	n/a	n/a	627	0.02

Fig. 4. Experimental results of compositional and non-compositional language-level abstraction.

that the ring eventually stabilises’. Since each predicate refers to variables from all components of the system, this model cannot be abstracted compositionally.

Figure 4 presents a summary of the performance for each case study. For the full concrete model, we give the number of states, the number of PRISM commands and the time required to perform model checking in PRISM (to ensure a fair comparison, we use PRISM’s MTBDD engine). For the abstract model, we show the number of states, the time to perform model checking with our prototype and, for each of the two abstraction approaches (non-compositional and compositional), the number of A-PRISM commands and the time required for abstraction. All times are in seconds and experiments were run on an AMD Athlon 4600+ with 2GB RAM.

Figure 5 gives quantitative results obtained from the models: lower and upper bounds from the abstract model and, where possible, exact answers from the concrete model. This, together with Figure 4, confirms that the game-based abstraction works well, in all cases providing tight lower and upper bounds from relatively small abstract models.

A key observation from the results is that we successfully managed to analyse models (for the sliding window protocol) larger than is possible with the model-level implementation of game-based abstraction from [18]. Furthermore, for the larger Zeroconf models, building and checking the abstraction is more efficient than checking the full model. In all cases, the use of ‘on-the-fly’ abstraction is essential to make the abstraction process feasible. This is because the number of potential overlaps, if reachability is not considered, is prohibitively large. The worst performance is observed for the Israeli & Jalfon’s protocol. For this model, the generation of abstract commands described in Section 5.1 needs to consider every concrete state in the model (the worst possible scenario), resulting in a large number of calls to the SMT solver and thus a very slow abstraction time.

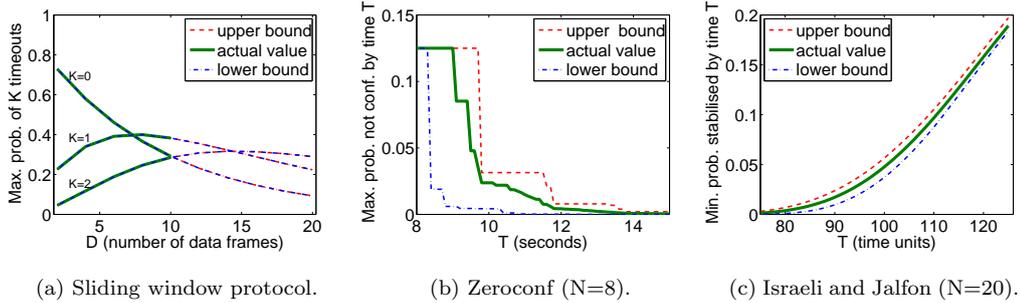


Fig. 5. Quantitative results obtained using language-level game-based abstraction.

As regards a comparison of the compositional and non-compositional approaches to abstraction, we observe varied results. For Zeroconf models, the compositional abstraction significantly outperforms the non-compositional one, both in terms of abstraction time and A-PRISM model size, but for the sliding window protocol the reverse is true. In fact, this is due to the general suitability of the models to a compositional analysis. For the sliding window protocol, each component makes no assumptions about the content (and ordering) of incoming messages and thus, when considered in isolation, its (concrete or abstract) state space is much larger. This makes the compositional approach perform poorly. What is very encouraging, however, is that for models which can be decomposed without such a blow-up (such as Zeroconf) our composition approach can exploit this and performs much better.

7 Related Work

Practical approaches for abstracting MDPs are presented in [6,18], the former using MDPs themselves as abstract models and the latter using stochastic two-player games. In [6] the tool RAPTURE is presented which performs successive abstractions and refinements for checking bounds on reachability probabilities. In [18], a prototype implementation is used to construct abstract models from the corresponding MDPs and partition of the state space and compute upper and lower bounds on reachability probabilities.

Predicate abstraction techniques [12] are prevalent in non-probabilistic verification. In the probabilistic case, the only other work we are aware of is [22] which introduces the PASS tool for language-level abstraction of PRISM models using the abstract approach of [6]. Like to our approach, PASS employs an SMT solver in the abstraction procedure. A key difference, however, is our use of stochastic two-player games rather than MDPs. While this will in general provide a more useful abstraction, it is also more difficult to apply to predicate abstraction. In [22] each command of a PRISM module can be abstracted separately. Here, as described in Section 5, we must consider overlaps between commands in order to distinguish between the two types of non-determinism. To improve efficiency, we also adopt a compositional approach to abstraction and use ‘on-the-fly’ techniques.

Also relevant is the ‘magnifying-lens abstraction’ (MLA) approach of [8], which computes lower and upper bounds for PCTL formulae on MDPs. This is done by partitioning the state space into regions and analysing each region separately. It is

still necessary to build the full MDP, however. Finally, approaches have also been proposed for abstracting discrete-time Markov chains [11,14], using interval-based extensions of Markov chains, but no implementations or results were presented.

8 Conclusions

We have introduced a method to obtain stochastic two-player game abstractions of MDPs, directly from high-level model descriptions in the PRISM language. Our approach is based on a compositional reformulation of the abstraction techniques from [18] and the use of predicate abstraction. Although a compositional abstraction is potentially an over-approximation (compared to the non-compositional version), we provide conditions which guarantee a precise abstraction. We have developed an implementation of our techniques based on the SMT solver Yices and present experimental results from a range of case studies, illustrating how our work can generate game-based abstractions for larger models than was previously possible. We also highlight the benefits of adopting a compositional approach.

In the future, we hope to improve the performance of our tool chain using symbolic decision procedures [19]. We also plan to integrate this with ongoing work to develop an abstraction-refinement loop for MDP verification. Finally, we also intend to extend the current method to imperative programming languages.

References

- [1] Ball, T., T. Millstein and S. K. Rajamani, *Automatic predicate abstraction of C programs*, SIGPLAN Notices **36** (2001), pp. 203–213.
- [2] Bianco, A. and L. de Alfaro, *Model checking of probabilistic and nondeterministic systems*, in: P. S. Thiagarajan, editor, *Proc. 15th Conf. on Foundations of Software Technology and Theoretical Computer Science (FSTTC'95)*, LNCS **1026** (1995), pp. 499–513.
- [3] Cheshire, S., B. Adoba and E. Gutterman, *Dynamic configuration of IPv4 link local addresses*, available from <http://www.ietf.org/rfc/rfc3927.txt>.
- [4] Ciesinski, F. and C. Baier, *LiQuor: A tool for qualitative and quantitative linear time analysis of reactive systems*, in: *Proc. 3rd Int. Conf. Quantitative Evaluation of Systems (QEST'06)* (2006), pp. 131–132.
- [5] Clarke, E., D. Kroening, N. Sharygina and K. Yorav, *Predicate abstraction of ANSI-C programs using SAT*, *Formal Methods in System Design* **25** (2004), pp. 105–127.
- [6] D'Argenio, P., B. Jeannet, H. Jensen and K. Larsen, *Reachability analysis of probabilistic systems by successive refinements*, in: L. de Alfaro and S. Gilmore, editors, *Proc. 1st Joint Int. Workshop on Process Algebra and Probabilistic Methods, Performance Modeling and Verification (PAPM/PROBMIV'01)*, LNCS **2165** (2001), pp. 39–56.
- [7] de Alfaro, L., T. Henzinger and R. Jhala, *Compositional methods for probabilistic systems*, in: K. Larsen and M. Nielsen, editors, *Proc. 12th Int. Conf. Concurrency Theory (CONCUR'01)*, LNCS **2154** (2001), pp. 351–365.
- [8] de Alfaro, L. and P. Roy, *Magnifying-lens abstraction for Markov decision processes*, in: W. Damm and H. Hermanns, editors, *Proc. 19th Int. Conf. Computer Aided Verification (CAV'07)*, LNCS **4590** (2007), pp. 325–338.
- [9] Dufflot, M., L. Fribourg and C. Picaconny, *Randomized finite-state distributed algorithms as Markov chains*, in: J. Welch, editor, *Proc. 15th Int. Conf. Distributed Computing (DISC'01)*, LNCS **2180** (2001), pp. 240–254.
- [10] Dutertre, B. and L. de Moura, *A fast linear-arithmetic solver for DPLL(T)*, in: T. Ball and R. Jones, editors, *Proc. 18th Int. Conf. Computer Aided Verification (CAV'06)*, LNCS **4114** (2006), pp. 81–94.

- [11] Fecher, H., M. Leucker and V. Wolf, *Don't know in probabilistic systems*, in: A. Valmari, editor, *Proc. 13th Int. SPIN Workshop (SPIN'06)*, LNCS **3925** (2006), pp. 71–88.
- [12] Graf, S. and H. Saidi, *Construction of abstract state graphs with PVS*, in: O. Grumberg, editor, *Proc. 9th Int. Conf. Computer Aided Verification (CAV'97)*, LNCS **1254** (1997), pp. 72–83.
- [13] Hinton, A., M. Kwiatkowska, G. Norman and D. Parker, *PRISM: A tool for automatic verification of probabilistic systems*, in: H. Hermanns and J. Palsberg, editors, *Proc. 12th Int. Conf. Tools and Algorithms for the Construction and Analysis of Systems (TACAS'06)*, LNCS **3920** (2006), pp. 441–444.
- [14] Huth, M., *On finite-state approximants for probabilistic computation tree logic*, *Theoretical Computer Science* **346** (2005), pp. 113–134.
- [15] Israeli, A. and M. Jalfon, *Token management schemes and random walks yield self-stabilizing mutual exclusion*, in: *Proc. 9th ACM Symp. Principles of Distributed Computing (PODC'90)* (1990), pp. 119–131.
- [16] Kattenbelt, M., M. Kwiatkowska, G. Norman and D. Parker, *Game-based probabilistic predicate abstraction in PRISM*, Technical Report RR-08-01, Oxford University Computing Laboratory (2008).
- [17] Kemeny, J. G., J. L. Snell and A. W. Knapp, “Denumerable Markov Chains,” Springer-Verlag, 1976, 2 edition.
- [18] Kwiatkowska, M., G. Norman and D. Parker, *Game-based abstraction for Markov decision processes*, in: *Proc. 3rd Int. Conf. Quantitative Evaluation of Systems (QEST'06)* (2006), pp. 157–166.
- [19] Lahiri, S., T. Ball and B. Cook, *Predicate abstraction via symbolic decision procedures*, in: K. Etessami and S. Rajamani, editors, *Proc. 17th Int. Conf. on Computer Aided Verification (CAV'05)*, LNCS **3576** (2005), pp. 24–38.
- [20] Lahiri, S. K., R. Nieuwenhuis and A. Oliveras, *SMT techniques for fast predicate abstraction*, in: T. Ball and R. B. Jones, editors, *Proc. 18th Int. Conf. Computer Aided Verification (CAV'06)*, LNCS **4144** (2006), pp. 424–437.
- [21] Stahl, K., K. Baukus, Y. Lakhnech and M. Steffen, *Divide, abstract, and model-check*, in: D. Dams, R. Gerth, S. Leue and M. Massink, editors, *Proc. 5th and 6th Int. SPIN Workshops (SPIN'99)*, LNCS **1680** (1999), pp. 57–76.
- [22] Wachter, B., L. Zhang and H. Hermanns, *Probabilistic model checking modulo theories*, in: *Proc. 4th Int. Conf. Quantitative Evaluation of Systems (QEST'07)* (2007), pp. 119–128.