Computing Science Group

ABSTRACTION FRAMEWORK FOR MARKOV DECISION PROCESSES AND PCTL VIA GAMES

Mark Kattenbelt Michael Huth

CS-RR-09-01

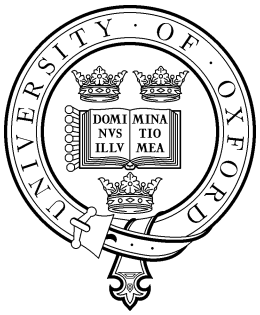

Oxford University Computing Laboratory
Wolfson Building, Parks Road, Oxford OX1 3QD

Abstract

Markov decision processes (MDPs) are natural models of computation in a wide range of applications. Probabilistic computation tree logic (PCTL) is a powerful temporal logic for reasoning about and verifying such models. Often, these models are prohibitively large or infinite-state, and so direct model checking of PCTL formulae over MDPs is infeasible. A recognised solution to this problem would be to develop finite-state abstractions of MDPs that soundly abstract the satisfaction of *arbitrary* PCTL formulae over very large or infinite-state MDPs. We state requirements for such an abstraction framework – e.g. that model checking of abstractions under-approximates generalised model checking for PCTL – and show important meta-properties that follow from these requirements. We take a notion of stochastic games from stochastic reachability analysis, adapt it, develop a simulation order for these adapted games – decidable in P – and prove that this adaptation meets all key requirements for an abstraction framework. Unlike generalised model checking, model checking our abstractions is reasonably efficient. We also show that the refinement characterised by PCTL is coarser than our simulation order.

1 Introduction

In many application areas both stochastic uncertainty and worst/best-case uncertainty coexist. Markov decision processes (MDPs) are models that capture both types of uncertainty well. This makes MDPs well-equipped to model software exhibiting both *non-deterministic* and *probabilistic* behaviour, such as randomised algorithms, or networking tools [25]. A powerful temporal logic for analysing MDPs is probabilistic computation tree logic (PCTL) [18]. For any MDP M and PCTL formula ϕ, we are therefore interested in whether M satisfies ϕ (i.e. verifying $M \models \phi$), or not (i.e. refuting $M \models \phi$). But, for software verification such direct verification is typically not feasible.

Thus, a predominant approach to software verification – used by many qualitative model checkers – is to convert M into a compact abstraction A, to verify $A \models \phi$, and obtain $M \models \phi$ from a soundness result for free. This approach is usually limited to universal fragments of temporal logics [9] and, as such fragments are not closed under logical negation, unsound for refutation.

In this setting refutation is instead realised by concretising abstract counter-examples [8]. Such an approach to refutation is less appealing for probabilistic model checking as probabilistic counter-examples typically consist of large collections of traces [17]. Thus deciding whether such a counter-example carries over to the concrete program is unlikely to scale well. It is therefore worth devising abstractions A such that for *arbitrary* formulae of the given temporal logic the verification of $A \models \phi$ soundly implies $M \models \phi$ (e.g. [10]).

This enables both abstraction-based verification and refutation for logics closed under negation, such as PCTL.

In this paper, we state requirements for such an abstraction framework for verification and refutation of MDPs and PCTL and show important meta-properties that follow from these, e.g. that model checking of abstractions under-approximates generalised model checking for PCTL. To instantiate this framework, we adapt the notion of stochastic games from stochastic reachability analysis [27], develop an efficiently decidable simulation order for these adapted games and prove that this adaptation meets all key requirements. We further demonstrate that model checking of abstractions is reasonably efficient, unlike generalised model checking. Also, we show that the refinement characterised by PCTL is coarser than our simulation order. Proofs can be found in the appendix.

Related work In [13, 19] MDP are abstracted by MDPs again, through the strong simulation preorder of [22, 30]. Reachability properties verified on abstractions are sound for the abstracted MDP. In [30] it is shown that such simulations do not soundly verify negated Until formulae.

Transition probabilities can be abstracted as sets of probabilities, e.g. intervals. Such foundations exist for (discrete and continous-time) Markov chains (e.g. [22, 20, 16, 23]). It is unclear whether, and if so how, this approach can be extended to include non-determinism (e.g. MDPs).

In [27, 24] stochastic games as abstractions of MDPs were proposed for probabilistic reachability analysis. These games separate the non-determinism stemming from MDPs from the non-determinism stemming from the abstraction process. The novelty of these abstractions is the ability to compute bounds specifically tailored to over-approximate the minimum probability and under-approximate the maximum probability of reaching a target set. This has the flavour of a three-valued abstraction for which verification and refutation are both sound [5] and has successfully been applied to probabilistic software verification [25].

For qualitative systems, sound verification and refutation of temporal logics have mostly been developed in a (sometimes implicit) three-valued setting [28, 5, 10]. Our results for sound abstraction-based verification and refutation of MDPs and PCTL are, notably, informed by work on modal/mixed transitions system [28, 10], three-valued abstraction of games [15], generalised model checking [6] and a finite model property adapted to abstractions [11].

2 Background

We write $\mathbb{N}$ for the non-negative integers and AP for a fixed set of atomic propositions. For a set X, let $\mathbb{P}(X)$ be the powerset of X. A distribution over X is a function $\lambda \in X \to [0,1]$ such that $\sum_{x \in X} \lambda(x) = 1$ and the set $\{x \in X \mid \lambda(x) > 0\}$ is countable. Let $\mathbb{D}(X)$ be the set of all distributions over X. For $x \in X$ let $\mu_x \in \mathbb{D}(X)$ be the point distribution on x, i.e. $\mu_x(x) = 1$. By abuse of notation, we write $\alpha_1 \cdot x_1 + \ldots + \alpha_n \cdot x_n$ for linear combinations of point distributions μ_{x_i}. If $X' \subseteq X$, and $\lambda \in \mathbb{D}(X')$ we will sometimes implicitly interpret λ as a distribution over X where $\lambda(x) = 0$ for all $x \in X \setminus X'$. For a set of distributions $\Lambda \in \mathbb{PD}(X)$ over X, and a distribution $\lambda_C \in \mathbb{D}(\Lambda)$ over Λ, let $(\Lambda \circ \lambda_C) \in \mathbb{D}(X)$ be defined as $(\Lambda \circ \lambda_C)(x) = \sum_{\lambda \in \Lambda} \lambda_C(\lambda) \cdot \lambda(x)$ for all $x \in X$.

For any binary relation $R \subseteq X \times Y$ let $R^{-1} \subseteq Y \times X$ be the relational inverse of R. We will sometimes use infix notation $x\,R\,y$ for $\langle x, y\rangle \in R$. For every $X' \subseteq X$ let $R.X'$ be the image of X' in R, i.e. the set $\{y \in Y \mid \exists x' \in X' \colon \langle x', y\rangle \in R\}$. We often write $R.x$ for $R.\{x\}$.

As in [22] we lift R to a relation over distributions $\mathbb{D}(R) \subseteq \mathbb{D}(X) \times \mathbb{D}(Y)$ by letting $\langle \lambda_X, \lambda_Y\rangle \in \mathbb{D}(R)$ iff there is a weight function $\delta \in X \times Y \to [0,1]$ such that:

$$\forall x \in X \colon \textstyle\sum_{y \in Y}\ \delta(x,y) = \lambda_X(x) \tag{1a}$$

$$\forall y \in Y \colon \textstyle\sum_{x \in X}\ \delta(x,y) = \lambda_Y(y) \tag{1b}$$

$$\forall \langle x, y\rangle \in X \times Y \colon (\delta(x,y) > 0 \Rightarrow \langle x, y\rangle \in R) \tag{1c}$$

Let π be an arbitrary finite or infinite (non-empty) sequence of elements $\omega_0, \omega_1, \omega_2, \ldots$. Let $|\pi|$ be the number elements of π minus one. For $i \leq |\pi|$ let $\pi(i)$ be the i+1-th element ω_i of π and, if π is finite, let $\overrightarrow{\pi}$ be the last element of π. For $i \leq |\pi|$ let π^i be the prefix of π such that $|\pi^i| = i$. We denote with $\pi \frown \pi'$ the concatenation of two sequences.

Probabilistic CTL Properties of probabilistic models are often written in probabilistic computation tree logic (PCTL) [18]. We define a minimal fragment of PCTL whose unrestricted negation makes other operators, such as a *tautology* (`tt`), *conjunction*, *eventuality* and *globality*, definable.

Definition 1 (PCTL syntax)**.** A PCTL formula is defined with the following BNF-style syntax rules where $a \in \text{AP}$, $k \in \mathbb{N} \cup \{\infty\}$, $p \in [0,1]$ and $\bowtie \in \{\leq, <, \geq, >\}$:

$$\phi ::= a \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid P_{\bowtie p}\langle\psi\rangle$$
$$\psi ::= X\phi \mid \phi_1\, U^{\leq k}\, \phi_2\ .$$

We call subformulae of the form ϕ and ψ *state* and *path* formulae, respectively. We

denote with Φ_{PCTL} and Ψ_{PCTL} the set of all PCTL state and path formulae, respectively.

Markov decision processes In quantitative software verification behaviour can be both non-deterministic and probabilistic. Markov decision processes naturally capture such semantics. Without loss of generality, our notion of MDP has propositional but no action labels.

Definition 2 (MDP)**.** A *Markov decision process* (MDP) is a tuple $\langle S, I, T, L\rangle$, where: S is a set of states; $I \subseteq S$ is a set of *initial* states; $T \in S \to \mathbb{PD}(S)$ is a *transition function* and $L \in S \to \mathbb{P}(\text{AP})$ is a *labelling function*.

Let $\mathscr{M}$ be the class of all MDPs. Let $M = \langle S, I, T, L\rangle$ denote any MDP throughout this paper. A transition originating from $s \in S$ needs to resolve both a *non-deterministic choice*, by choosing $\lambda \in T(s)$, and a probabilistic choice, by choosing s' such that $\lambda(s') > 0$. We allow $T(s)$ to be the empty set, in which case we call s a *deadlock state*. When we define probability measures and PCTL semantics we need a transformation $M \mapsto M_\perp \in \mathscr{M} \to \mathscr{M}$ which adds a sink state $\perp$ to M so that $\perp$ and every deadlock state of M deterministically transition to $\perp$ with probability 1.

A *path* of any MDP M is a sequence of transitions that strictly alternates between states and distributions as described above. Let Π_M be the set of *finite* paths, Π^∞_M the set of all *infinite* paths, and $\Pi_M(\omega)$ and $\Pi^\infty_M(\omega)$ the set of finite and infinite paths of M that start from ω (respectively).

A path resolves both non-deterministic and probabilistic choice. But, a *strategy* resolves only non-determinism. Formally, a strategy is a partial function $\sigma \in \Pi_M \to \mathbb{D}(S)$ such that $\sigma(\pi) \in \mathbb{D}(T(\overrightarrow{\pi}))$. A path π of M is consistent with σ iff for all $i \leq |\pi| - 1$ with $\pi(i) \in S$ the probability $\sigma(\pi^i)(\pi(i+1))$ is positive. Given a strategy σ and a set of paths, we will add the subscript σ to denote the set of paths consistent with σ (e.g. $\Pi_{M,\sigma}$). Finally, we denote with Σ_M the set of all strategies of M.

For any strategy $\sigma \in \Sigma_M$ and set of finite paths $\Pi \subseteq \Pi_{M,\sigma}$, let $\Pi^{\uparrow\sigma}$ be the set of infinite paths of M that are consistent with σ and have a prefix in Π. When Π is a singleton we call $\Pi^{\uparrow\sigma}$ a *cylinder set* of M.

We define probability measures over MDPs M *without* deadlock states. Using the methods from [26], every $\omega \in S \cup \mathbb{D}(S)$ and $\sigma \in \Sigma_M$ determine a unique probability measure $\mathbf{Pr}^\omega_{M,\sigma}$ over infinite paths $\Pi^\infty_{M,\sigma}(\omega)$ such that all cylinder sets constructed from finite paths in $\Pi_{M,\sigma}(\omega)$ are measurable in $\mathbf{Pr}^\omega_{M,\sigma}$ and for all zero-length paths ω' we have that $\mathbf{Pr}^\omega_{M,\sigma}(\{\omega'\}^{\uparrow\sigma}) = 1$ if $\omega' = \omega$ and 0 otherwise. Moreover, for every finite path of

non-zero length $\pi' \frown \omega' \in \Pi_{M,\sigma}(\omega)$ we have:

$$\mathbf{Pr}^{\omega}_{M,\sigma}(\{\pi' \frown \omega'\}^{\uparrow\sigma}) = \begin{cases} \mathbf{Pr}^{\omega}_{M,\sigma}(\{\pi'\}^{\uparrow\sigma}) \cdot \sigma(\pi')(\omega') & \text{if } \overrightarrow{\pi'} \in S \\ \mathbf{Pr}^{\omega}_{M,\sigma}(\{\pi'\}^{\uparrow\sigma}) \cdot \overrightarrow{\pi'}(\omega') & \text{if } \overrightarrow{\pi'} \in \mathbb{D}(S) \end{cases}$$

We will use shorthands $\mathbf{Pr}^{\omega}_{M,\sigma}(\pi)$ and $\mathbf{Pr}^{\omega}_{M,\sigma}(\Pi)$ to denote the probabilities $\mathbf{Pr}^{\omega}_{M,\sigma}(\{\pi\}^{\uparrow\sigma})$ and $\mathbf{Pr}^{\omega}_{M,\sigma}(\Pi^{\uparrow\sigma})$, respectively, and we omit the subscript M when unambiguous.

Strong probabilistic (bi)simulation We recall the definitions of *strong probabilistic simulation* (preorder $\sqsubseteq_{\mathscr{M}} \subseteq \mathscr{M} \times \mathscr{M}$) and *strong probabilistic bisimulation* (equivalence relation $\equiv_{\mathscr{M}} \subseteq \mathscr{M} \times \mathscr{M}$) over MDPs, introduced in [30].

Definition 3 (Strong probabilistic simulation)**.** Let $\hat{M} = \langle \hat{S}, \hat{I}, \hat{T}, \hat{L} \rangle$ and $M = \langle S, I, T, L \rangle$ be MDPs. We say $\hat{M}$ is a *strong probabilistic simulation* of M via relation $R \subseteq \hat{S} \times S$, denoted $\hat{M} \sqsubseteq^{R}_{\mathscr{M}} M$, if and only if $I \subseteq R.\hat{I}$ and, whenever $\langle \hat{s}, s \rangle \in R$, the following conditions hold:

(i) $\hat{L}(\hat{s}) = L(s)$

(ii) $\hat{T}(s) = \emptyset \quad \Leftarrow \quad T(s) = \emptyset$

(iii) $\forall \lambda \in T(s)\, \exists \hat{\lambda}_C \in \mathbb{D}(\hat{T}(s))\colon \langle \hat{T}(s) \circ \hat{\lambda}_C, \lambda \rangle \in \mathbb{D}(R)$

We let $\hat{M} \sqsubseteq_{\mathscr{M}} M$ iff there exists a relation $R \subseteq \hat{S} \times S$ such that $\hat{M} \sqsubseteq^{R}_{\mathscr{M}} M$.

Condition (iii) requires that for every non-deterministic choice $\lambda \in T(s)$ there is a weight distribution over non-deterministic choices in $\hat{T}(\hat{s})$ such that the resulting distributions simulate each other. The weighted abstract transition is called a *combined transition* in [30]. Condition (ii) requires that deadlock behaviour is preserved by the simulation. We define bisimulation in the style of, e.g., [29].

Definition 4 (Strong probabilistic bisimulation)**.** Let $\hat{M} = \langle \hat{S}, \hat{I}, \hat{T}, \hat{L} \rangle$ and $M = \langle S, I, T, L \rangle$ be MDPs. We say $\hat{M}$ is a *strong probabilistic bisimulation* of M via relation $R \subseteq \hat{S} \times S$, denoted $\hat{M} \equiv^{R}_{\mathscr{M}} M$, iff $\hat{M} \sqsubseteq^{R}_{\mathscr{M}} M$ and $M \sqsubseteq^{R^{-1}}_{\mathscr{M}} \hat{M}$. Let $\hat{M} \equiv_{\mathscr{M}} M$ iff there is a relation $R \subseteq \hat{S} \times S$ with $\hat{M} \equiv^{R}_{\mathscr{M}} M$.

PCTL semantics of MDPs We now formally define what it means for an MDP $M \in \mathscr{M}$ to satisfy a PCTL formula $\phi \in \Phi_{\text{PCTL}}$. We do this by first defining semantics for $M_\perp$ and then setting $M \models_{\mathscr{M}} \phi$ iff $M_\perp \models_{\mathscr{M}} \phi$.

Definition 5 (PCTL semantics)**.** Let $M = \langle S, I, T, L\rangle$ be an MDP, and let $\phi \in \Phi_{\text{PCTL}}$ be a PCTL formula. Let $\Pi \subseteq \Pi^{\infty}_{M_\perp}$ denote all infinite paths of $M_\perp$ starting with a state in S. We first define a satisfaction relation $\models \subseteq \Pi \times \Psi_{\text{PCTL}}$ for path formulae where for $k \in \mathbb{N} \cup \{\infty\}$ we have $\pi \models X\phi$ iff $\pi(2) \models \phi$ and $\pi \models (\phi_1 \, U^{\leq k} \, \phi_2)$ iff:

$$(\exists i{\leq}k : \, (\pi(2i) \models \phi_2) \,\&\, (\forall j{<}i : \, (\pi(2j) \models \phi_1))) \, .$$

Clearly, this satisfaction relation is mutually dependent on a satisfaction relation for PCTL state formulae $\models \subseteq S_\perp \times \Phi_{\text{PCTL}}$ which we define next. First, we define for every $\sigma \in \Sigma_M$, $s \in S$ and $\psi \in \Psi_{\text{PCTL}}$ the shorthand

$$\textsc{Prob}_\sigma(s, \psi) = \mathbf{Pr}^s_{M_\perp,\sigma}\left\{\pi \in \Pi^{\infty}_{M_\perp,\sigma}(s) \mid \pi \models \psi\right\}$$

denoting the probability of all paths satisfying ψ that originate from s and are consistent with σ. In the following let $s \in S$, $\triangleright \in \{>, \geq\}$, $\triangleleft \in \{<, \leq\}$ and $p \in [0, 1]$):

$$\begin{aligned}
\perp &\not\models \phi \\
s &\models a \Leftrightarrow a \in L(s) \\
s &\models \neg\phi \Leftrightarrow s \not\models \phi \\
s &\models (\phi_1 \vee \phi_2) \Leftrightarrow (s \models \phi_1 \text{ or } s \models \phi_2) \\
s &\models P_{\triangleright p}\langle\psi\rangle \Leftrightarrow \inf_{\sigma \in \Sigma_{M_\perp}} \{\textsc{Prob}_\sigma(s, \psi)\} \triangleright p \\
s &\models P_{\triangleleft p}\langle\psi\rangle \Leftrightarrow \sup_{\sigma \in \Sigma_{M_\perp}} \{\textsc{Prob}_\sigma(s, \psi)\} \triangleleft p
\end{aligned}$$

Finally, let $M_\perp \models_{\mathscr{M}} \phi$ iff for all $s \in I$ we have $s \models \phi$.

For MDPs we need all four threshold types: $P_{>p}\langle\psi\rangle$ implies the threshold p is met under *all* schedulings, whereas $\neg P_{\leq p}\langle\psi\rangle$ implies there *exists* such a scheduling.

The semantics is well-defined but non-standard, as $\perp$ satisfies no PCTL formulae and the other clauses don't apply to $\perp$. This guarantees that deadlocks, if present, don't contribute to the probability of paths sets from any state. In fact, the semantics of negation yields that our PCTL semantics are consistent and, moreover, two-valued.

3 Abstraction framework

Given an MDP $M \in \mathscr{M}$ and PCTL formula $\phi \in \Phi_{\text{PCTL}}$ we wish to decide whether $M \models_{\mathscr{M}} \phi$ or $M \not\models_{\mathscr{M}} \phi$. That is, we wish to either *verify* or *refute* the judgement $M \models_{\mathscr{M}} \phi$. In software verification, however, directly applying such a model check is intractible. Therefore, we seek a class of models, $\mathscr{A}$ say, that abstract MDPs and allow for the sound verification and refutation of PCTL formulae. Inspired by [12], we capture these

requirements for $\mathscr{A}$ abstractly. The first requirement is that MDPs are representable in $\mathscr{A}$:

R1. Domain $\mathscr{A}$ has an *embedding* function $e^{\mathscr{A}} \in \mathscr{M} \to \mathscr{A}$.

We call the elements in $e^{\mathscr{A}}(\mathscr{M})$ *implementations* of $\mathscr{A}$. Abstraction-based verification requires an abstraction relation in $\mathscr{A}$, as formalised by the following requirement:

R2. Domain $\mathscr{A}$ has a *refinement preorder* $\sqsubseteq_{\mathscr{A}} \subseteq \mathscr{A} \times \mathscr{A}$.

The meaning of $\hat{A} \sqsubseteq_{\mathscr{A}} A$ is that $\hat{A}$ abstracts A or, equivalently, that A is a refinement of $\hat{A}$. Implementations are typically maximal elements of $\mathscr{A}$, i.e. they cannot be further refined. The refinement ordering enables us to associate with each $A \in \mathscr{A}$ the set of implementations that refine A:

Definition 6 (Implementations)**.** Let $\mathcal{I} \in \mathscr{A} \to \mathbb{P}(\mathscr{M})$ be defined as $\mathcal{I}(A) = \{M \in \mathscr{M} \mid A \sqsubseteq_{\mathscr{A}} e^{\mathscr{A}}(M)\}$.

We want to understand how refinement should behave over implementations. As strong probabilistic bisimulation over MDPs preserves PCTL satisfaction [30] we will require that the refinement preorder $\sqsubseteq_{\mathscr{A}}$, when restricted to $e^{\mathscr{A}}(\mathscr{M}) \times e^{\mathscr{A}}(\mathscr{M})$, over-approximates $\equiv_{\mathscr{M}}$.

R3. For all $M, M' \in \mathscr{M}$ we have that $M \equiv_{\mathscr{M}} M'$ implies $e^{\mathscr{A}}(M) \sqsubseteq_{\mathscr{A}} e^{\mathscr{A}}(M')$.

R1 up to **R3** secure a first meta-property of abstractions:

Lemma 3.1. For any $A \in \mathscr{A}$ the set of implementations $\mathcal{I}(A)$ is a union of equivalence classes of $\equiv_{\mathscr{M}}$.

The refinement preorder $\sqsubseteq_{\mathscr{A}}$ is used in practice to deduce properties about implementations. However, we can use function $\mathcal{I}$ (defined through $\sqsubseteq_{\mathscr{A}}$) to define the largest refinement that preserves implementations [6].

Definition 7 (Thorough refinement)**.** The *thorough refinement relation* is the preorder $\sqsubseteq^{\text{th}}_{\mathscr{A}} \subseteq \mathscr{A} \times \mathscr{A}$ such that $\hat{A} \sqsubseteq^{\text{th}}_{\mathscr{A}} A$ iff $\mathcal{I}(A) \subseteq \mathcal{I}(\hat{A})$.

Therefore thorough refinements can only remove, but not add, implementations. The next meta-property states that $\sqsubseteq_{\mathscr{A}}$ soundly under-approximates $\sqsubseteq^{\text{th}}_{\mathscr{A}}$ in that very sense:

Lemma 3.2. For $\hat{A}, A \in \mathscr{A}$, $\hat{A} \sqsubseteq_{\mathscr{A}} A$ implies $\hat{A} \sqsubseteq^{\text{th}}_{\mathscr{A}} A$.

We typically expect thorough refinement to be strictly more precise than the ordinary refinement preorder, i.e. that there exist $\hat{A}, A \in \mathscr{A}$ such that $\mathcal{I}(A) \subseteq \mathcal{I}(\hat{A})$ but $\hat{A} \not\sqsubseteq_{\mathscr{A}} A$. For example, this is the case for transition systems and modal transition systems and their refinement [1].

Property verification and refutation require a PCTL semantics over $\mathscr{A}$.

R4. The domain $\mathscr{A}$ has a *satisfaction relation* $\models_{\mathscr{A}} \subseteq \mathscr{A} \times \Phi_{\text{PCTL}}$ such that $e^{\mathscr{A}}(M) \models_{\mathscr{A}} \phi$ iff $M \models_{\mathscr{M}} \phi$ for all $M \in \mathscr{M}$.

The latter part of **R4** ensures consistency of the PCTL semantics across these two representations of MDPs and implies that PCTL formulae have a two-valued semantics over embedded implementations. To decide whether $M \models_{\mathscr{M}} \phi$ or $M \not\models_{\mathscr{M}} \phi$ using abstractions, refinement $\sqsubseteq_{\mathscr{A}}$ has to mesh well with the abstract PCTL semantics $\models_{\mathscr{A}}$:

R5. For any $\hat{A}, A \in \mathscr{A}$ we have that $\hat{A} \sqsubseteq_{\mathscr{A}} A$ implies that for all $\phi \in \Phi_{\text{PCTL}}$ it holds that $\hat{A} \models_{\mathscr{A}} \phi \Rightarrow A \models_{\mathscr{A}} \phi$.

With **R5**, finally we have a method of deciding whether $M \models_{\mathscr{M}} \phi$ or $M \not\models_{\mathscr{M}} \phi$ by only considering abstractions. That is, in order to verify the judgement $M \models_{\mathscr{M}} \phi$ it is sufficient to find some $A \in \mathscr{A}$ such that $M \in \mathcal{I}(A)$ and $A \models_{\mathscr{A}} \phi$. Using **R4** and **R5**, this yields $M \models_{\mathscr{M}} \phi$. Similarly, to refute the judgement $M \models_{\mathscr{M}} \phi$ it is sufficient to find some $A \in \mathscr{A}$ such that $M \in \mathcal{I}(A)$ and $A \models_{\mathscr{A}} \neg\phi$.

As typical in abstraction-based verification, it is possible that *both* $A \not\models_{\mathscr{A}} \phi$ and $A \not\models_{\mathscr{A}} \neg\phi$ hold. In that case we can neither verify nor refute $M \models_{\mathscr{M}} \phi$ and we may have to refine A. It is also possible that *both* $A \models_{\mathscr{A}} \phi$ and $A \models_{\mathscr{A}} \neg\phi$ hold but, by **R4** and **R5**, A has then no implementations.

The abstract satisfaction relation of **R4** is the one used in practice but, it has a more precise version, analogous to the relationship between refinement and thorough refinement:

Definition 8 (Thorough satisfaction)**.** The *thorough satisfaction relation* is the relation $\models^{\text{th}}_{\mathscr{A}} \subseteq \mathscr{A} \times \Phi_{\text{PCTL}}$ such that $A \models^{\text{th}}_{\mathscr{A}} \phi$ iff $M \models_{\mathscr{M}} \phi$ for all $M \in \mathcal{I}(A)$.

Thorough satisfaction is the logical dual of generalised model checking in [6]. The next meta-property shows $\models_{\mathscr{A}}$ soundly under-approximates its thorough version $\models^{\text{th}}_{\mathscr{A}}$:

Lemma 3.3. For any $A \in \mathscr{A}$ and $\phi \in \Phi_{\text{PCTL}}$ we have $A \models_{\mathscr{A}} \phi$ implies $A \models^{\text{th}}_{\mathscr{A}} \phi$.

Thorough refinement and thorough satisfaction also constitute a method for verifying or refuting $M \models_{\mathscr{M}} \phi$:

Lemma 3.4. For any $\hat{A}, A \in \mathscr{A}$ with $\hat{A} \sqsubseteq^{\text{th}}_{\mathscr{A}} A$ we have $\hat{A} \models^{\text{th}}_{\mathscr{A}} \phi \Rightarrow A \models^{\text{th}}_{\mathscr{A}} \phi$ for all $\phi \in \Phi_{\text{PCTL}}$.

The requirements so far allow us to verify and refute $M \models_{\mathscr{M}} \phi$ through an abstraction A. For abstraction-based verification and refutation to be tractible, a minimal requirement is that A be finite. This leads to Dams and Namjoshi's [11] notion of complete abstraction frameworks:

R6. For every $M \in \mathscr{M}$ and $\phi \in \Phi_{\text{PCTL}}$ with $M \models_{\mathscr{M}} \phi$ there is a *finite* $A \in \mathscr{A}$ such that $M \in \mathcal{I}(A)$ and $\hat{A} \models_{\mathscr{A}} \phi$.

R6 makes it possible, in principle, to verify or refute $M \models_{\mathscr{M}} \phi$ through finite abstractions. Dams and Namjoshi's notion of completeness has been investigated for Markov chains and PCTL in [31]. Finally, we state another suitability requirement:

R7. Deciding $\models_{\mathscr{A}}$ and $\sqsubseteq_{\mathscr{A}}$ has relatively low computational complexity, compared to their thorough versions.

4 Game-based abstraction framework

We now develop stochastic games with two players, 1 and 2, as an instance of the abstraction framework of Section 3. We first introduce stochastic two-player games.

Stochastic two-player games Our stochastic two-player games have two distinct types of non-deterministic choice (corresponding to the two players). This is in contrast to the single notion of non-determinism in MDPs.

Definition 9 (Stochastic two-player game)**.** A *stochastic two-player game* G is a tuple $\langle S, I, T, L^!, L^? \rangle$, where: S is a set of states; $I \subseteq S$ a non-empty set of *initial* states; $T \in S \to \mathbb{PPD}(S)$ a transition function and $L^!, L^? \in S \to \mathbb{P}(\text{AP})$ labelling functions with $L^!(s) \subseteq L^?(s)$ for each $s \in S$.

$L^!(s)$ is the set of atomic propositions that *must* be true in s, whereas $L^?(s)$ are propositions that *may* be true in s. Henceforth we will refer to stochastic two-player games as '*games*', and let $\mathscr{G}$ be the class of all games. A *transition* originating from $s \in S$ requires resolving a *player 1* (non-deterministic) choice, by choosing a *set* of distributions $\Lambda \in T(s)$, a *player 2* (non-deterministic) choice, by choosing a distribution $\lambda \in \Lambda$ and a *probabilistic* choice according to λ. Like MDPs, games can deadlock.

We define a transformation $G \mapsto G_\perp \in \mathscr{G} \to \mathscr{G}$ by adding to G two sink states $\perp_1$ and $\perp_2$. We let $\perp_1$ and every $s \in S$ with $T(s) = \emptyset$ deterministically transition to $\perp_1$

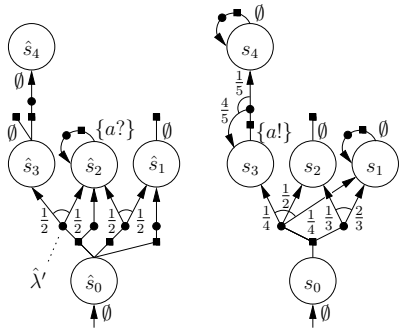

Figure 1: Game $\hat{G}$ and an implementation G with $\hat{G} \sqsubseteq^R_{\mathscr{G}} G$, for $R = \{\langle \hat{s}_0, s_0\rangle, \langle \hat{s}_1, s_2\rangle, \langle \hat{s}_3, s_2\rangle, \langle \hat{s}_2, s_3\rangle, \langle \hat{s}_2, s_1\rangle, \langle \hat{s}_2, s_4\rangle\}$.

with probability 1. For player 2 deadlocks we let $T(\bot_2) = \{\{\mu_{\bot_2}\}\}$ and we replace any potential player 2 deadlock $\emptyset \in T(s)$ with $\{\mu_{\bot_2}\} \in T(s)$.

A *play* of any game G is a sequence of transitions that strictly alternate between states, sets of distributions, and distributions as described above. We let Π_G, Π^∞_G be the set of finite and infinite plays of G, and let $\Pi_G(\omega)$ and $\Pi^\infty_G(\omega)$ be finite and infinite plays starting from ω (respectively).

In contrast to MDPs, games require *two* strategies; one for each player. Formally, a *player 1 strategy* is a partial function $\sigma_1 \in \Pi_G \to \mathbb{DPD}(S)$ such that $\sigma_1(\pi) \in \mathbb{D}(T(\overrightarrow{\pi}))$. A *player 2 strategy* is a partial function $\sigma_2 \in \Pi_G \to \mathbb{DD}(S)$ such that $\sigma_2(\pi) \in \mathbb{D}(\overrightarrow{\pi})$.

A play π is consistent with σ_1 if for every $i \leq |\pi| - 1$ with $\pi(i) \in S$ the probability $\sigma_1(\pi^i)(\pi(i+1))$ is positive. Similarly, π is consistent with σ_2 if $\sigma_2(\pi^i)(\pi(i+1))$ is positive whenever $\pi(i) \in \mathbb{PD}(S)$. We add the subscript σ_1, σ_2 to π to denote sets of consistent plays and let Σ^1_G and Σ^2_G be all player 1 and 2 strategies (respectively). For some $\sigma_1 \in \Sigma^1_G$ and $\sigma_2 \in \Sigma^2_G$ and a set of finite plays $\Pi \subseteq \Pi_{G,\sigma_1,\sigma_2}$, we denote with $\Pi^{\uparrow^{\sigma_2}_{\sigma_1}}$ the infinite plays of G that are consistent with σ_1, σ_2 and that have a prefix in Π.

We define probability measures over our games G *without* deadlocks. Analogous to that definition for MDPs, every $\omega \in S \cup \mathbb{PD}(S) \cup \mathbb{D}(S)$ and $\langle \sigma_1, \sigma_2\rangle \in \Sigma^1_G \times \Sigma^2_G$ determine a unique probability measure $\mathbf{Pr}^\omega_{G,\sigma_1,\sigma_2}$ over infinite plays in $\Pi^\infty_{G,\sigma_1,\sigma_2}(\omega)$ such that for all zero-length plays $\mathbf{Pr}^\omega_{G,\sigma_1,\sigma_2}(\omega')$ yields 1 if $\omega' = \omega$ and 0 otherwise. Moreover, for every finite play of non-zero length $\pi' \frown \omega' \in \Pi_{\sigma_1,\sigma_2}(\omega)$ we have:

$$\mathbf{Pr}^\omega_{G,\sigma_1,\sigma_2}(\{\pi' \frown \omega'\}^{\uparrow^{\sigma_2}_{\sigma_1}}) = \begin{cases} \mathbf{Pr}^\omega_{G,\sigma_1,\sigma_2}(\{\pi'\}^{\uparrow^{\sigma_2}_{\sigma_1}}) \cdot \sigma_1(\pi')(\omega') & \text{if } \overrightarrow{\pi'} \in S \\ \mathbf{Pr}^\omega_{G,\sigma_1,\sigma_2}(\{\pi'\}^{\uparrow^{\sigma_2}_{\sigma_1}}) \cdot \sigma_2(\pi')(\omega') & \text{if } \overrightarrow{\pi'} \in \mathbb{PD}(S) \\ \mathbf{Pr}^\omega_{G,\sigma_1,\sigma_2}(\{\pi'\}^{\uparrow^{\sigma_2}_{\sigma_1}}) \cdot \overrightarrow{\pi'}(\omega') & \text{if } \overrightarrow{\pi'} \in \mathbb{D}(S) \end{cases}$$

We omit the G subscript when unambiguous to do so and use shorthands $\mathbf{Pr}^{\omega}_{G,\sigma_1,\sigma_2}(\pi)$ and $\mathbf{Pr}^{\omega}_{G,\sigma_1,\sigma_2}(\Pi)$ to denote $\mathbf{Pr}^{\omega}_{G,\sigma_1,\sigma_2}(\{\pi\}^{\uparrow^{\sigma_2}_{\sigma_1}})$ and $\mathbf{Pr}^{\omega}_{G,\sigma_1,\sigma_2}(\Pi^{\uparrow^{\sigma_2}_{\sigma_1}})$, respectively.

Assumption 1. In figures we depict states of games with open circles, distributions (i.e. player 2 choices) with filled black circles and sets of distributions (i.e. player 1 choices) with filled black squares. Labels depict the probability of transitions (omitted for point distributions). We write $a!$ next to $\hat{s}$ iff $a \in \hat{L}^{!}(\hat{s})$, and write $a?$ next to $\hat{s}$ iff $a \in \hat{L}^{?}(\hat{s})$ and $a \notin \hat{L}^{!}(\hat{s})$.

Example 1. Consider the game $\hat{G}$ depicted in Figure 1 (left). Let $\hat{\lambda}' \in \mathbb{D}(\hat{S})$ be the distribution $\frac{1}{2} \cdot \hat{s}_3 + \frac{1}{2} \cdot \hat{s}_2$. As depicted $\{\hat{\lambda}', \mu_{\hat{s}_2}\} \in \hat{T}(\hat{s}_0)$. Example plays in $\hat{G}_{\perp}$ are:

$$\hat{s_1}, \{\mu_{\perp_2}\}, \mu_{\perp_2}, \perp_2, \{\mu_{\perp_2}\}, \dots$$
$$\hat{s}_0, \{\hat{\lambda}', \mu_{\hat{s}_2}\}, \hat{\lambda}', \hat{s}_3, \{\mu_{\hat{s}_4}\}, \mu_{\hat{s}_4}, \hat{s}_4, \{\mu_{\perp_1}\}, \mu_{\perp_1}, \perp_1, \dots$$

Games as abstractions We now instantiate the embedding, refinement preorder (as a simulation) and the abstract satisfaction relation of Section 3 for these games.

The intuition of our simulation is that player 1 non-determinism captures the indeterminacy introduced by abstraction and player 2 non-determinism corresponds to the non-determinism that is present in implementations. Therefore, player 1 has no power in embedded MDPs:

Definition 10 (R1). Let $e^{\mathscr{G}} \in \mathscr{M} \to \mathscr{G}$ be the function which for every MDP $M = \langle S, I, T, L \rangle$ yields a game $G = \langle S, I, \hat{T}, L, L \rangle$ such that $\hat{T}(s) = \{T(s)\}$ for every $s \in S$.

Every state of an embedded MDP G has *precisely one* corresponding player 1 choice, and hence Σ^1_G contains only a single trivial strategy. Moreover, this means it is impossible for player 1 deadlocks to occur in embedded MDPs. The *implementations* $e^{\mathscr{G}}(\mathscr{M})$ of $\mathscr{G}$ are precisely those games $G = \langle S, I, T, L^{!}, L^{?} \rangle$ where for every $s \in S$, player 1 has precisely one choice and the state labels are two-valued, that is $|T(s)| = 1$ and $L^{!}(s) = L^{?}(s)$ (as, e.g., in the game in Figure 1 (right)). All other games are *abstractions*.

We now define how abstractions and implementations are related by *strong probabilistic game-simulation*:

Definition 11 (R2). Let $\hat{G} = \langle \hat{S}, \hat{I}, \hat{T}, \hat{L}^{!}, \hat{L}^{?} \rangle$ and $G = \langle S, I, T, L^{!}, L^{?} \rangle$ be games. We call $\hat{G}$ a *strong probabilistic game-simulation* of G via relation $R \subseteq \hat{S} \times S$, denoted $\hat{G} \sqsubseteq^{R}_{\mathscr{G}} G$, iff $I \subseteq R.\hat{I}$ and, whenever $\langle \hat{s}, s \rangle \in R$, we have:

(i) $\hat{L}^{!}(\hat{s}) \subseteq L^{!}(s)$

(ii) $\hat{L}^{?}(\hat{s}) \supseteq L^{?}(s)$

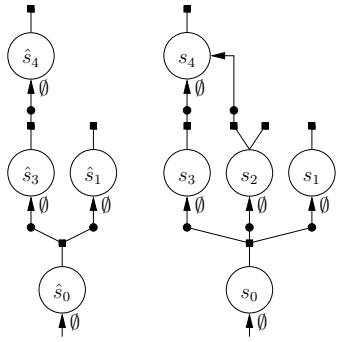

Figure 2: Games $\hat{G}$ and G such that $\hat{G} \sqsubseteq^{\text{th}}_{\mathscr{G}} G$ but $\hat{G} \not\sqsubseteq_{\mathscr{G}} G$. Example based on Fig. 5 of [1].

(iii) $\forall \Lambda \in T(s)\, \exists \hat{\Lambda} \in \hat{T}(s)\colon (\hat{\Lambda}{=}\emptyset \Leftarrow \Lambda{=}\emptyset)$ &
$(\forall \lambda \in \Lambda\, \exists \hat{\lambda}_C \in \mathbb{D}(\hat{\Lambda})\colon \langle \hat{\Lambda} \circ \hat{\lambda}_C, \lambda \rangle \in \mathbb{D}(R))$

(iv) $\forall \Lambda \in T(s)\, \exists \hat{\Lambda} \in \hat{T}(s)\colon (\hat{\Lambda}{=}\emptyset \Rightarrow \Lambda{=}\emptyset)$ &
$(\forall \hat{\lambda} \in \hat{\Lambda}\, \exists \lambda_C \in \mathbb{D}(\Lambda)\colon \langle \hat{\lambda}, \Lambda \circ \lambda_C \rangle \in \mathbb{D}(R))$

Let $\hat{G} \sqsubseteq_{\mathscr{G}} G$ iff there is a relation $R \subseteq \hat{S} \times S$ with $\hat{G} \sqsubseteq^{R}_{\mathscr{M}} G$. Clearly $\sqsubseteq_{\mathscr{G}}$ is a preorder on $\mathscr{G}$ and so **R2** is met.

Condition (iii) requires that the 'abstract game' $\hat{G}$ over-approximates the player 2 non-determinism of every player 1 choice of the 'concrete game' G. Thus, this condition corresponds with the 'must' modality. Dually, condition (iv) requires that the abstract game under-approximates the player 2 non-determism of every player 1 choice of the concrete game, corresponding to the 'may' modality.

Example 2. For games $\hat{G}$ and G in Figure 1 we have $\hat{G} \sqsubseteq^{R}_{\mathscr{G}} G$. To illustrate this, let $\lambda = \frac{2}{3} \cdot s_1 + \frac{1}{3} \cdot s_2$ of G and $\hat{\lambda} = \frac{1}{2} \cdot \hat{s}_2 + \frac{1}{2} \cdot \hat{s}_3$ of $\hat{G}$. Let $\Lambda = \{\lambda\} \in T(s_0)$ and $\hat{\Lambda} = \{\hat{\lambda}, \mu_{\hat{s}_2}\} \in \hat{T}(\hat{s}_0)$. To show condition (iii) for Λ observe that $\langle \hat{\Lambda} \circ \hat{\lambda}_C, \lambda \rangle \in \mathbb{D}(R)$ for $\hat{\lambda}_C = \frac{2}{3} \cdot \hat{\lambda} + \frac{1}{3} \cdot \mu_{\hat{s}_2} \in \mathbb{D}(\hat{\Lambda})$.

In our framework **R3** holds:

Proposition 4.1 (R3). For all $M, M' \in \mathscr{M}$ we have that $M \equiv_{\mathscr{M}} M'$ implies $e^{\mathscr{G}}(M) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}(M')$.

Instantiating Definition 6 we obtain for each game G the set of MDPs $\mathcal{I}(G)$ that refine G. As **R1** up to **R3** are met for games, by meta-property Lemma 3.1 $\mathcal{I}(G)$ is closed under strong probabilistic simulation.

Definition 7 gives rise to the thorough refinement preorder of games ($\sqsubseteq^{\text{th}}_{\mathscr{G}}$). Lemma 3.2, for games, already yields that $\sqsubseteq_{\mathscr{G}}$ must be a sound under-approximation of $\sqsubseteq^{\text{th}}_{\mathscr{G}}$. As perhaps expected, we observe that $\sqsubseteq^{\text{th}}_{\mathscr{G}}$ is strictly more precise than $\sqsubseteq_{\mathscr{G}}$ (see Figure 2):

Lemma 4.2. There are $\hat{G}, G \in \mathscr{G}$ with $\hat{G} \sqsubseteq^{\text{th}}_{\mathscr{G}} G$, $\hat{G} \not\sqsubseteq_{\mathscr{G}} G$.

Turning now to requirement **R4**, we define the abstract PCTL semantics $\models_{\mathscr{G}} \subseteq \mathscr{G} \times \Phi_{\text{PCTL}}$. Due to unrestricted negation this requires both under-approximating ($\models^!$) and over-approximating ($\models^?$) PCTL semantics.

Definition 12 (PCTL semantics)**.** Let $G = \langle S, I, T, L^!, L^? \rangle$ be a game, and let $\phi \in \Phi_{\text{PCTL}}$ be a PCTL formula. Analogous to MDPs we define PCTL semantics for $G_\perp$ and then set $G \models_{\mathscr{G}} \phi$ iff $G_\perp \models_{\mathscr{G}} \phi$.

We let $\Pi \subseteq \Pi^\infty_{G_\perp}$ denote the set of all infinite plays of $G_\perp$ starting with a state in S. We define two satisfaction relations $\models^!, \models^? \subseteq \Pi \times \Psi_{\text{PCTL}}$ for path formulae where for $\text{M} \in \{!, ?\}$ and $k \in \mathbb{N} \cup \{\infty\}$ we have $\pi \models^{\text{M}} X\phi$ iff $\pi(3) \models^{\text{M}} \phi$ and $\pi \models^{\text{M}} (\phi_1 \, U^{\leq k} \, \phi_2)$ iff:

$$(\exists i \leq k : \, (\pi(3i) \models^{\text{M}} \phi_2) \,\&\, (\forall j < i : \, (\pi(3j) \models^{\text{M}} \phi_1))) \, .$$

These satisfaction relations mutually depend on satisfaction relations for PCTL state formulae $\models^!, \models^? \subseteq S_\perp \times \Phi_{\text{PCTL}}$. First, we define for every $\text{M} \in \{!, ?\}$, $\sigma_1 \in \Sigma^1_{G_\perp}$, $\sigma_2 \in \Sigma^2_{G_\perp}$, $s \in S$ and $\psi \in \Psi_{\text{PCTL}}$ the shorthand

$$\textsc{Prob}^{\text{M}}_{\sigma_1\sigma_2}(s, \psi) = \mathbf{Pr}^s_{\sigma_1,\sigma_2} \left\{ \pi \in \Pi^\infty_{\sigma_1,\sigma_2}(s) \mid \pi \models^{\text{M}} \psi \right\} \, .$$

In the following we let $s \in S$, $\neg! = ?$, $\neg? = !$, $\triangleright \in \{>, \geq\}$, $\triangleleft \in \{<, \leq\}$ and $p \in [0, 1]$):

$$\begin{aligned}
\perp_1 &\models^! \phi & \perp_1 &\not\models^? \phi \\
\perp_2 &\not\models^! \phi & \perp_2 &\not\models^? \phi \\
s &\models^{\text{M}} a \Leftrightarrow a \in L^{\text{M}}(s) & s &\models^{\text{M}} \neg\phi \Leftrightarrow s \not\models^{\neg\text{M}} \phi \\
s &\models^{\text{M}} (\phi_1 \vee \phi_2) \Leftrightarrow (s \models^{\text{M}} \phi_1 \text{ or } s \models^{\text{M}} \phi_2)
\end{aligned}$$

and

$$\begin{aligned}
s &\models^! P_{\triangleright p}\langle\psi\rangle \Leftrightarrow \inf_{\sigma_1 \in \Sigma^1_{G_\perp}} \inf_{\sigma_2 \in \Sigma^2_{G_\perp}} \left\{ \textsc{Prob}^!_{\sigma_1\sigma_2}(s, \psi) \right\} \triangleright p \\
s &\models^? P_{\triangleright p}\langle\psi\rangle \Leftrightarrow \sup_{\sigma_1 \in \Sigma^1_{G_\perp}} \inf_{\sigma_2 \in \Sigma^2_{G_\perp}} \left\{ \textsc{Prob}^?_{\sigma_1\sigma_2}(s, \psi) \right\} \triangleright p \\
s &\models^? P_{\triangleleft p}\langle\psi\rangle \Leftrightarrow \inf_{\sigma_1 \in \Sigma^1_{G_\perp}} \sup_{\sigma_2 \in \Sigma^2_{G_\perp}} \left\{ \textsc{Prob}^!_{\sigma_1\sigma_2}(s, \psi) \right\} \triangleleft p \\
s &\models^! P_{\triangleleft p}\langle\psi\rangle \Leftrightarrow \sup_{\sigma_1 \in \Sigma^1_{G_\perp}} \sup_{\sigma_2 \in \Sigma^2_{G_\perp}} \left\{ \textsc{Prob}^?_{\sigma_1\sigma_2}(s, \psi) \right\} \triangleleft p
\end{aligned}$$

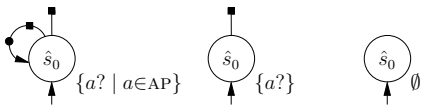

Figure 3: Three games used to illustrate some points made in the body of the paper.

Finally, let $G_\perp \models_{\mathscr{G}} \phi$ iff for all $s \in I$ we have $s \models^! \phi$.

These PCTL semantics are consistent for consistent games, i.e. those that have implementations. The semantics of $\perp_1$ cannot introduce inconsistencies since player 1 deadlocks are avoidable in consistent games. This consistency follows from **R5** and Proposition 4.3 below.

Example 3. For the game $\hat{G}$ with $\mathcal{I}(\hat{G}) = \emptyset$ in Figure 3 (right) both $\hat{G} \models_{\mathscr{G}} P_{\geq 1}\langle Xa \rangle$ and $\hat{G} \models_{\mathscr{G}} \neg P_{\geq 1}\langle Xa \rangle$ hold.

To see that the PCTL semantics over embedded MDPs coincides with the MDP semantics recall that there are no player 1 deadlocks, the proposition labels are two-valued ($L^!(s) = L^?(s)$) and there is exactly one player 1 strategy $|\Sigma^1_G| = 1$. We now validate **R5** for our framework.

Proposition 4.3 (R5). For any $\hat{G}, G \in \mathscr{G}$ with $\hat{G} \sqsubseteq_{\mathscr{G}} G$ we have that $\hat{G} \models_{\mathscr{G}} \phi \Rightarrow G \models_{\mathscr{G}} \phi$ for all $\phi \in \Phi_{\mathrm{PCTL}}$.

Proposition 4.3 shows our notion of refinement soundly preserves PCTL satisfaction. As explained in Section 3 this now gives us a method to verify or refute a PCTL property by looking only at abstractions.

Example 4. Consider the abstraction $\hat{G}$ and implementation G of Figure 1. We have that $\hat{G} \models_{\mathscr{G}} \neg P_{>0.25}\langle Xa \rangle$ and hence by Proposition 4.3 $G \models_{\mathscr{G}} \neg P_{>0.25}\langle Xa \rangle$. However, $\hat{G}$ is too abstract to verify the judgement $G \models_{\mathscr{G}} P_{\leq 0.25}\langle Xa \rangle$ as $\hat{G} \not\models_{\mathscr{G}} P_{\leq 0.25}\langle Xa \rangle$ and $\hat{G} \not\models_{\mathscr{G}} \neg P_{\leq 0.25}\langle Xa \rangle$.

We note that the reverse implication of Proposition 4.3 does not hold, as stated by the following lemma:

Lemma 4.4. There exist $\hat{G}, G \in \mathscr{G}$ such that $\hat{G} \not\sqsubseteq_{\mathscr{G}} G$ but for all $\phi \in \Phi_{\mathrm{PCTL}}$ it holds that $\hat{G} \models_{\mathscr{G}} \phi \Rightarrow G \models_{\mathscr{G}} \phi$.

Figure 4 gives an example of $\hat{G}$ and G satisfying the lemma. Lemma 4.4 shows PCTL does not characterise games up to refinement equivalence. In fact, as the games of Figure 4

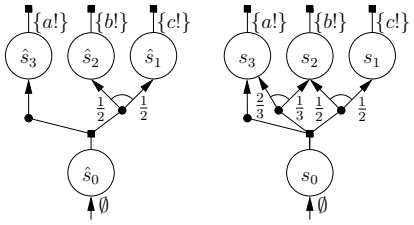

Figure 4: Games $\hat{G}$ and G with $\hat{G} \not\sqsubseteq_{\mathscr{G}} G$ such that $\hat{G} \models_{\mathscr{G}} \phi \Leftrightarrow G \models_{\mathscr{G}} \phi$ for all $\phi \in \Phi_{\text{PCTL}}$. Example due to personal correspondence with R. Segala in relation to [29].

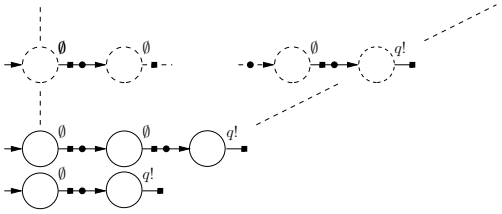

Figure 5: Game G with $G \models_{\mathscr{G}} P_{>0}\langle \mathtt{tt}\, U\, q\rangle$; no finite abstraction of G satisfies $P_{>0}\langle \mathtt{tt}\, U\, q\rangle$.

are non-bisimilar embedded MDPs, the example also shows PCTL does not characterise MDPs up to strong probabilistic bisimulation, as already suggested in [29].

Using Definition 8 we obtain thorough satisfaction for games ($\models^{\text{th}}_{\mathscr{G}}$). Meta-property Lemma 3.3, for games, ensures that $\models_{\mathscr{G}}$ soundly under-approximates $\models^{\text{th}}_{\mathscr{G}}$. The following lemma states that this is a strict under-approximation:

Lemma 4.5. There exists $G \in \mathscr{G}$ and $\phi \in \Phi_{\text{PCTL}}$ such that $G \models^{\text{th}}_{\mathscr{G}} \phi$ but $G \not\models_{\mathscr{G}} \phi$.

Lemma 4.5 follows from the fact that any *implementation* trivially satisfies the PCTL formula $a \vee \neg a$, but the abstraction depicted in Figure 3 (middle) does not.

The next lemma shows that our game-based abstractions of MDPs are incomplete for PCTL:

Lemma 4.6 (R6). There is an MDP $M \in \mathscr{M}$ and a PCTL formula $\phi \in \Phi_{\text{PCTL}}$ with $M \models_{\mathscr{M}} \phi$ such that there is no *finite* game $\hat{G} \in \mathscr{G}$ with $M \in \mathcal{I}(\hat{G})$ and $\hat{G} \models_{\mathscr{G}} \phi$.

The embedded game in Figure 5, adapted from [11, Theorem 1] to our setting, proves Lemma 4.6. This incompleteness relies on the existence of strategies that cannot reach $q!$-states. As argued in [11], fairness or more general acceptance conditions can rule out such strategies. Based on the structural and operational nature of our games we believe that the absence of such acceptence conditions is the only reason for our framework to be incomplete. We note that no complete abstraction framework for MDPs and PCTL

is known. The incompleteness of three-valued Markov chains as abstractions of Markov chains has been shown in [31].

Our final requirement **R7** concerns the computational complexity of the various refinement and satisfaction relations. We now corroborate that the computational complexity of deciding $\sqsubseteq_{\mathscr{G}}$ and $\models_{\mathscr{G}}$ is relatively low compared to deciding their thorough counterparts $\sqsubseteq^{\text{th}}_{\mathscr{G}}$ and $\models^{\text{th}}_{\mathscr{G}}$.

Proposition 4.7. Deciding $\hat{G} \sqsubseteq_{\mathscr{G}} G$ is in P.

The proof of Proposition 4.7 is based on deciding strong probabilistic simulation of probabilistic automata [32]. We turn to thorough refinement. *Qualitative* thorough refinement of games asks whether all *qualitative* implementations (consisting of only point-distributions) of one game also implement another. In the appendix we prove thorough refinement of Krikpe modal transition systems [21] – shown to be PSPACE-hard in [2] – can be reduced to this problem.

We now discuss the complexity of deciding $\models_{\mathscr{G}}$. As for MDPs and PCTL, $M \models_{\mathscr{M}} \phi$ can be computed by a bottom-up recursion on $\phi \in \Phi_{\text{PCTL}}$. For any Until subformula, e.g. $P_{>p}\langle \phi_1 \, U \, \phi_2 \rangle$ under modality ?, we first compute

$$\sup_{\sigma_1 \in \Sigma^1_G} \inf_{\sigma_2 \in \Sigma^2_G} \left\{ \textsc{Prob}^{?}_{\sigma_1\sigma_2}(s, \phi_1 \, U \, \phi_2) \right\}$$

by solving a naturally derived stochastic parity game and then convert this into a Boolean judgement by comparing the value against the threshold. Solving stochastic parity games with rational probabilities is in NP $\cap$ co-NP [7]. Also, the computational overhead for processing other types of subformulae is polynomial. Thus deciding $\models_{\mathscr{G}}$ is in NP$\cap$co-NP.

Proposition 4.8. For games G with rational probabilities, deciding $G \models_{\mathscr{G}} \phi$ is in NP $\cap$ co-NP.

We now demonstrate that thorough satisfaction of games is harder than deciding satisfiability of PCTL formulae over MDPs which is not known to be decidable. (In [4] satisfiability of a *qualitative* fragment of PCTL is shown to be EXPTIME-complete with respect to Markov chains.)

Proposition 4.9. PCTL satisfiability over MDPs can be reduced to deciding the thorough satisfaction of games.

To see this, game $\hat{G}$ depicted in Figure 3 (left) is implemented by every MDP – i.e. $\mathcal{I}(\hat{G}) = \mathscr{M}$. Hence, the model check $\hat{G} \not\models^{\text{th}}_{\mathscr{G}} \neg\phi$ decides whether ϕ is satisfiable.

To summarize, our games satisfy all requirements **R1** up to **R7**, except for the completeness property **R6**.

5 Discussion and conclusions

Our games enable sound, abstraction-based verification *and* refutation of PCTL formulae. In frameworks based on the simulation preorders of [22, 30] refutation doesn't come cheaply, as demonstrated in [19]: soundness of probabilistic counter-examples, represented as finite Markov chains, appeals to properties of the possibly infinitely many concretisations of that finite Markov chain.

In abstraction-refinement implementations it is common practice to over-approximate the transition relations of abstractions [3]. In [25], however, costly optimal abstractions are computed over a partition of an MDP. The work reported here now allows us to sythesise over-approximating games on that same partition, which can then be successively refined over this partition in the style of [14].

Let us conclude. We motivated the need for a three-valued framework for the verification and refutation of PCTL formulae on MDPs, stated requirements for such a framework, and derived some meta-properties enjoyed by any framework meeting these requirements. We then instantiated this framework by adapting a notion of stochastic games, developing a simulation order for these adapted games and proving that this adaptation meets all key requirements – including the requirement that deciding simulation and abstract satisfaction is reasonably efficient. Along the way, we showed that the simulation logically characterised by PCTL is coarser than our simulation order.

Acknowledgements This work was supported in part by the UK EPSRC grants EP/D07956X/2 and EP/E028985/1. We also thank R. Segala for his correspondence regarding [29].

References

[1] A. Antonik, M. Huth, K. Larsen, U. Nyman, and A. Wasowski. 20 Years of Modal and Mixed Specifications. *Bull. of the EATCS* 95: 94–129, July 2008.

[2] A. Antonik, M. Huth, K. G. Larsen, U. Nyman, and A. Wasowski. Complexity of decision problems for mixed and modal specifications. In *Proc. of FoSSaCS'08*, LNCS 4962, pp. 112–126. Springer, 2008.

[3] T. Ball, B. Cook, S. Das, and S. K. Rajamani. Refining approximations in software predicate abstraction. In *Proc. of TACAS'04*, LNCS 2988, pp. 388–403. Springer, 2004.

[4] T. Brázdil, V. Forejt, J. Křetínský, and A. Kučera. The satisfiability problem for probabilistic CTL. In *Proc. of LICS'08*, pp. 391–402. IEEE CS, 2008.

[5] G. Bruns and P. Godefroid. Model checking partial state spaces with 3-valued temporal logics. In *Proc. of CAV'99*, LNCS 1633, pp. 274–287. Springer, 1999.

[6] G. Bruns and P. Godefroid. Generalized model checking: Reasoning about partial state spaces. In *Proc. of CONCUR'00*, LNCS 1877, pp. 168–182. Springer, 2000.

[7] K. Chatterjee, M. Jurdzinski, and T. A. Henzinger. Quantitative stochastic parity games. In *Proc. of SODA'04*, pp. 121–130, ACM Press, 2004.

[8] E. M. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-guided abstraction refinement. In *Proc. CAV'00*, LNCS 1855, pp. 154–169. Springer, 2000.

[9] E. M. Clarke, O. Grumberg, and D. E. Long. Model checking and abstraction. *ACM Transactions on Programming Languages and Systems* 16(5): 1512–1542, 1994.

[10] D. Dams, R. Gerth, and O. Grumberg. Abstract interpretation of reactive systems. *ACM Transactions on Programming Languages and Systems* 19(2): 253 – 291, 1997.

[11] D. Dams and K. S. Namjoshi. The existence of finite abstractions for branching time model checking. In *Proc. of LICS'04*, pp. 335–344. IEEE CS, 2004.

[12] D. Dams and K. S. Namjoshi. Automata as abstractions. In *Proc. of VMCAI'05*, LNCS 3385, pp. 216–232. Springer, 2005.

[13] P. R. D'Argenio, B. Jeannet, H. E. Jensen, and K. G. Larsen. Reduction and refinement strategies for probabilistic systems. In *Proc. of PAPM-PROBMIV'02*, LNCS 2399, pp. 57–76. Springer, 2002.

[14] S. Das and D. L. Dill. Successive approximation of abstract transition relations. In *Proc. of LICS'01*, pp. 51–60. IEEE CS, 2001.

[15] L. de Alfaro, P. Godefroid, and R. Jagadeesan. Three-valued abstractions of games: Uncertainty, but with precision. In *Proc. of LICS'04*, pp. 170–179. IEEE CS, 2004.

[16] H. Fecher, M. Leucker, and V. Wolf. Don't know in probabilistic systems. In *Proc. of SPIN'06*, LNCS 3925, pp. 71–88. Springer, 2006.

[17] T. Han and J.-P. Katoen. Counterexamples in probabilistic model checking. In *Proc. of TACAS'07*, LNCS 4424, pp. 72–86. Springer, 2007.

[18] H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing* 6: 512–535, 1994.

[19] H. Hermanns, B. Wachter, and L. Zhang. Probabilistic CEGAR. In *Proc. of CAV'08*, LNCS 5123, pp. 162–175. Springer, 2008.

[20] M. Huth. On finite-state approximants for probabilistic computation tree logic. *Theoretical Computer Science* 346(1): 113–134, 2005.

[21] M. Huth, R. Jagadeesan, and D. A. Schmidt. Modal transition systems: A foundation for three-valued program analysis. In *Proc. of ESOP'01*, LNCS 2028, pp. 155–169. Springer, 2001.

[22] B. Jonsson and K. G. Larsen. Specification and refinement of probabilistic processes. In *Proc. of LICS'91*, pp. 266–277. IEEE CS, 1991.

[23] J.-P. Katoen, D. Klink, M. Leucker, and V. Wolf. Three-valued abstraction for continuous-time Markov chains. In *Proc. of CAV'07*, LNCS 4590, pp. 311–324. Springer, 2007.

[24] M. Kattenbelt, M. Kwiatkowska, G. Norman, and D. Parker. A game-based abstraction-refinement framework for Markov decision processes. *Technical Report RR-08-06*, OUCL, April 2008.

[25] M. Kattenbelt, M. Kwiatkowska, G. Norman, and D. Parker. Abstraction refinement for probabilistic programs. In *Proc. of VMCAI'09*, LNCS 5403, pp. 182–197. Springer, 2009.

[26] J. G. Kemeny, J. L. Snell, and A. W. Knapp. *Denumerable Markov Chains*. Springer-Verlag, Second edition, 1976.

[27] M. Kwiatkowska, G. Norman, and D. Parker. Game-based abstraction for Markov decision processes. In *Proc. of QEST'06*, pp. 157–166. IEEE CS, 2006.

[28] K. G. Larsen and B. Thomsen. A modal process logic. In *Proc. of LICS'88*, pp. 203–210. IEEE CS, 1988.

[29] A. Parma and R. Segala. Logical characterizations of bisimulations for discrete probabilistic systems. In *Proc. of FoSSaCS'07*, LNCS 4423, pp. 287–301. Springer, 2007.

[30] R. Segala and N. A. Lynch. Probabilistic simulations for probabilistic processes. In *Proc. of CONCUR'94*, LNCS 836, pp. 481–496. Springer, 1994.

[31] D. Wagner. *MPhil/PhD Transfer Report*. Department of Computing, Imperial College London. April 2008.

[32] L. Zhang and H. Hermanns. Deciding simulations on probabilistic automata. In *Proc. of ATVA'07*, LNCS 4762, pp. 207–222. Springer, 2007.

A Proofs for Section 3

A.1 Proof of Lemma 3.1

Lemma 3.1. For any $A \in \mathscr{A}$ the set of implementations $\mathcal{I}(A)$ is a union of equivalence classes of $\equiv_{\mathscr{M}}$.

Proof. Let $A \in \mathscr{A}$, $M \in \mathcal{I}(A)$ and let $M' \in \mathscr{M}$ be an arbitrary MDP such that $M \equiv_{\mathscr{M}} M'$. It is sufficient to prove that $M' \in \mathcal{I}(A)$. As $M \in \mathcal{I}(A)$, we have $A \sqsubseteq_{\mathscr{A}} e^{\mathscr{A}}(M)$. By **R3**, we have $e^{\mathscr{A}}(M) \sqsubseteq_{\mathscr{A}} e^{\mathscr{A}}(M')$. By transitivity of preorders we have $A \sqsubseteq_{\mathscr{A}} e^{\mathscr{A}}(M')$; hence $M' \in \mathcal{I}(A)$. □

A.2 Proof of Lemma 3.2

Lemma 3.2. For $\hat{A}, A \in \mathscr{A}$, $\hat{A} \sqsubseteq_{\mathscr{A}} A$ implies $\hat{A} \sqsubseteq^{\text{th}}_{\mathscr{A}} A$.

Proof. Let $\hat{A}, A \in \mathscr{A}$ be such that $\hat{A} \sqsubseteq_{\mathscr{A}} A$ and let $M \in \mathscr{M}$ be an MDP with $M \in \mathcal{I}(A)$. To show $\hat{A} \sqsubseteq^{\text{th}}_{\mathscr{A}} A$, by Def. 7, it is sufficient to show $M \in \mathcal{I}(\hat{A})$. By Def. 6 we have $A \sqsubseteq_{\mathscr{A}} e^{\mathscr{A}}(M)$. By transitivity of $\sqsubseteq_{\mathscr{A}}$ (as $\hat{A} \sqsubseteq_{\mathscr{A}} A$) we obtain $\hat{A} \sqsubseteq_{\mathscr{A}} e^{\mathscr{A}}(M)$; hence $M \in \mathcal{I}(\hat{A})$. □

A.3 Proof of Lemma 3.3

Lemma 3.3. For any $A \in \mathscr{A}$ and $\phi \in \Phi_{\text{PCTL}}$ we have $A \models_{\mathscr{A}} \phi$ implies $A \models^{\text{th}}_{\mathscr{A}} \phi$.

Proof. Suppose $A \models_{\mathscr{A}} \phi$, then for all $M \in \mathcal{I}(A)$, as $A \sqsubseteq_{\mathscr{A}} e^{\mathscr{A}}(M)$, by **R5** it must be that $e^{\mathscr{A}}(M) \models_{\mathscr{A}} \phi$. By **R4** this corresponds to $M \models_{\mathscr{M}} \phi$. Then, because $M \models_{\mathscr{M}} \phi$ for all $M \in \mathcal{I}(A)$, by Def. 8, it follows that $A \models^{\text{th}}_{\mathscr{A}} \phi$. □

A.4 Proof of Lemma 3.4

Lemma 3.4. For any $\hat{A}, A \in \mathscr{A}$ with $\hat{A} \sqsubseteq^{\text{th}}_{\mathscr{A}} A$ we have $\hat{A} \models^{\text{th}}_{\mathscr{A}} \phi \Rightarrow A \models^{\text{th}}_{\mathscr{A}} \phi$ for all $\phi \in \Phi_{\text{PCTL}}$.

Proof. Suppose $\hat{A} \models^{\text{th}}_{\mathscr{A}} \phi$, then by Def. 8 it must be that $M \models_{\mathscr{M}} \phi$ for all $M \in \mathcal{I}(\hat{A})$. By Def. 7 we have that $\mathcal{I}(A) \subseteq \mathcal{I}(\hat{A})$ and hence we have that $M \models_{\mathscr{M}} \phi$ for all $M \in \mathcal{I}(A)$. By Def. 8, $A \models^{\text{th}}_{\mathscr{A}} \phi$. □

B Proofs for Section 4

Given a distribution $\lambda_X \in \mathbb{D}(X)$ we let $\textsc{supp}(\lambda_X)$ be the *support* of λ_X, i.e. the countable set $\{x \in X \mid \lambda_X(x) > 0\}$. We also define the size of games:

Definition 13. Let $G = \langle S, I, T, L^!, L^? \rangle$ be a game. We denote with $|G|$ the size of G, defined as $|G| = n_G + m_G$, where:

- $n_G = |S|$, and
- $m_G = \sum_{s \in S} \sum_{\Lambda \in T(s)} (1 + |\Lambda|)$.

Note that $\sum_{s \in S} |T(s)| \leq m_G$ and $\sum_{s \in S} \sum_{\Lambda \in T(s)} |\Lambda| \leq m_G$. We assume that $\textsc{ap}$ is fixed and can therefore consider I, $L^!$ and $L^?$ as functions that we can query in polynomial time.

B.1 Proof of Proposition 4.1

Proposition 4.1 (R3). For all $M, M' \in \mathscr{M}$ we have that $M \equiv_{\mathscr{M}} M'$ implies $e^{\mathscr{G}}(M) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}(M')$.

Proof. Let $\hat{M} = \langle \hat{S}, \hat{I}, \hat{T}, \hat{L} \rangle$ and $M = \langle S, I, T, L \rangle$ be two strongly bisimilar MDPs and let $R \subseteq \hat{S} \times S$ be the relation such that $\hat{M} \equiv^R_{\mathscr{M}} M$. We will show that $e^{\mathscr{G}}(\hat{M}) \sqsubseteq^R_{\mathscr{G}} e^{\mathscr{G}}(M)$. To show this, let $e^{\mathscr{G}}(\hat{M}) = \langle \hat{S}, \hat{I}, \hat{T}', \hat{L}, \hat{L} \rangle$ and $e^{\mathscr{G}}(M) = \langle S, I, \hat{T}', L, L \rangle$ as defined by Def. 10. Using Def. 4 we immediately satisfy that $I \subseteq R.\hat{I}$. Remaining to show is that for every $\langle \hat{s}, s \rangle \in R$ we satisfy all conditions of Def. 11. Condition (i) and (ii) of Def. 11 are trivially true as due to the bisimulation we have $\hat{L}(\hat{s}) = L(s)$. To show conditions (iii) and (iv) hold, recall that $\hat{T}'(\hat{s}) = \{\hat{T}(\hat{s})\}$ and $T'(s) = \{T(s)\}$. Knowing both have precisely one player 1 choice greatly simplifies the quantifiers of the game-simulation definition and it is immediate to see condition (iii) of Def. 11 of the game-simulation $e^{\mathscr{G}}(\hat{M}) \sqsubseteq^R_{\mathscr{G}} e^{\mathscr{G}}(M)$ corresponds directly with the conditions (ii) and (iii) of Def. 3 of the simulation $\hat{M} \sqsubseteq^R_{\mathscr{M}} M$. Analogously, condition (iv) of Def. 11 of the simulation $e^{\mathscr{G}}(\hat{M}) \sqsubseteq^R_{\mathscr{G}} e^{\mathscr{G}}(M)$ corresponds with conditions (ii) and (iii) of Def. 3 of the simulation $M \sqsubseteq^{R^{-1}}_{\mathscr{M}} \hat{M}$. As $\hat{M} \equiv^R_{\mathscr{M}} M$ yields $\hat{M} \sqsubseteq^R_{\mathscr{M}} M$ and $M \sqsubseteq^{R^{-1}}_{\mathscr{M}} \hat{M}$ both (iii) and (iv) of Def. 11 must hold. □

B.2 Proof of Lemma 4.2

Lemma 4.2. There are $\hat{G}, G \in \mathscr{G}$ with $\hat{G} \sqsubseteq^{\text{th}}_{\mathscr{G}} G$, $\hat{G} \not\sqsubseteq_{\mathscr{G}} G$.

Proof. Consider the two games $\hat{G} = \langle \hat{S}, \hat{I}, \hat{T}, \hat{L}^!, \hat{L}^? \rangle$ and $G = \langle S, I, T, L^!, L^? \rangle$ depicted in Figure 2. To see that $\hat{G} \not\sqsubseteq_{\mathscr{G}} G$, observe that $s_2 \in S$ (and hence $s_0 \in S$) cannot be

game-simulated by any state in $\hat{G}$. Remaining to show is that $\hat{G} \sqsubseteq^{\text{th}}_{\mathscr{G}} G$. By Definition 7 it is sufficient to show that for any MDP $M = \langle S', I', T', L' \rangle$ such that $M \in \mathcal{I}(G)$ we have that $M \in \mathcal{I}(\hat{G})$. Clearly, the main concern is that any behaviour of M that is simulated by s_2 cannot be simulated in $\hat{G}$. Potentially this problem may cascade through the co-algebraic definition of game-simulation. We will show that due to the fact M is an implementation, this is not true.

Suppose $G \sqsubseteq^{R}_{\mathscr{G}} e^{\mathscr{G}}(M)$ with $R \subseteq S' \times S$. We will construct $\hat{R} \subseteq \hat{S} \times S'$ such that $\hat{G} \sqsubseteq^{\hat{R}}_{\mathscr{G}} e^{\mathscr{G}}(M)$. Let us start by taking $\hat{R} = \{\langle \hat{s}_i, s' \rangle \mid \langle s_i, s' \rangle \in R,\ i \neq 2\}$. One problem is that states $s' \in S'$ such that $\langle s_2, s' \rangle \in R$ may now no longer have a suitable simulation in $\hat{G}$ via $\hat{R}$.

Consider $s' \in S$ to be any state of M such that $\langle s_2, s' \rangle \in R$ and let $T'(s') = \Lambda'$. The choice Λ' is either the empty set or not. If $\Lambda' = \emptyset$ then, considering condition (iii) and (iv) of Definition 11, the presence of player 1 choice $\emptyset \in T(s_2)$ is the only requirement for us to simulate s' (that is $\{\mu_{s_4}\} \in T(s_2)$ does not contribute anything). We add $\langle \hat{s}_1, s' \rangle$ to $\hat{R}$. Note that $\langle \hat{s}_1, s' \rangle$ satisfies condition (iii) and (iv).

Suppose Λ' is not the emptyset, then we must have that for every $\lambda' \in \Lambda'$ we have that $\langle \mu_{s_4}, \lambda' \rangle \in \mathbb{D}(R)$ (the presence of $\emptyset \in T(s_2)$ does not contribute anything). We add $\langle \hat{s}_3, s' \rangle$ to $\hat{R}$. Tuple $\langle \hat{s}_3, s' \rangle$ satisfies condition (iii) and (iv).

Although now we have fixed the issue that states $s' \in S'$ such that $\langle s_2, s' \rangle \in R$ cannot be simulated with $\hat{R}$, we have not shown the absence of problems in view of the co-algebraic nature of game-simulations. This follows from the observation that s_2 can only be reached through choice $T(s_0) = \{\mu_{s_1}, \mu_{s_2}, \mu_{s_3}\}$ in G. Hence, remaining to show is that for any $s' \in S'$ such that $\langle s_0, s' \rangle \in R$ we have that $\langle \hat{s}_0, s' \rangle$ satisfies condition (iii) and (iv) of Def. 9 via $\hat{R}$.

As $T(s_0)$ 'subsumes' $\hat{T}(\hat{s}_0)$ condition (iv) is trivially satisfied. To satisfy condition (iii) it must be that any $\lambda' \in T'(s')$ is simulated by a combined transition in $\{\mu_{\hat{s}_1}, \mu_{\hat{s}_3}\}$. We know that λ' is simulated by a combined transition in $\{\mu_{s_1}, \mu_{s_2}, \mu_{s_3}\}$. Clearly, for any weight attributed to μ_{s_2} and for any $s'' \in \text{SUPP}(\lambda')$ such that $\langle s_2, s'' \rangle \in R$ we have that either $\langle \hat{s}_1, s'' \rangle \in \hat{R}$ or $\langle \hat{s}_3, s'' \rangle \in \hat{R}$. We can easily redistribute the weight of μ_{s_2} to μ_{s_1} and μ_{s_3} to obtain a suitable weight distribution. □

B.3 Proof of Proposition 4.3

In this section we will show that strong probabilistic game-refinement ($\sqsubseteq_{\mathscr{G}}$) preserves PCTL satisfaction ($\models_{\mathscr{G}}$). Our proof extends that of [24]. In this paper a game is constructed from a partition of the state space of an MDP. It is shown that with the obvious refinement order on partitions the 'abstract' game soundly approximates extremal reach-

ability probabilities of the 'concrete' game. We extend this proof to arbitrary games related by strong probabilistic game-simulation and to preservation of arbitrary PCTL formulae.

Our extensions introduce a few changes in the proof. In [24] the two games are related by means of an abstraction relation that is both right-total, left-total and left-unique. This means that any play of the 'concrete' game maps to a single play of the 'abstract' game. When considering strong probabilistic game-simulation (in which we allow arbitrary relations) we lose this correspondence. We will instead assign *weights* between 'plays' of the concrete and abstract game.

To improve presentation we will sometimes denote tuples of strategies with a single symbol and we will use the subscripts 1 and 2 to denote the player 1 and player 2 strategies, e.g. $\hat{\sigma} = \langle \hat{\sigma}_1, \hat{\sigma}_2 \rangle \in \Sigma^1_{\hat{G}_\perp} \times \Sigma^2_{\hat{G}_\perp}$.

Partial plays & Arenas Due to the presence of combined transitions, plays themselves are too fine-grained for this weight function; there is not always a clear correspondence between the distributions of the plays. We therefore introduce a new concept called *partial plays*:

Definition 14 (Partial plays)**.** Let $G = \langle S, I, T, L^!, L^? \rangle$ be an arbitrary game. A *partial play* of G is a strictly alternating sequence of states S and sets of distributions $\mathbb{PD}(S)$ such that a state $s \in S$ with $T(s) \neq \emptyset$ can be followed by $\Lambda \in \mathbb{PD}(S)$ if $\Lambda \in T(s)$, and $\Lambda \in \mathbb{PD}(S)$ can be followed by $s \in S$ iff there exists $\lambda \in \Lambda$ such that $\lambda(s) > 0$.

We denote with $\widetilde{\Pi}_G$ the set of all finite partial plays. Strategy consistency of partial plays is defined and denoted analogous to normal plays, as are partial plays starting from configurations $\omega \in S \cup \mathbb{PD}(S)$.

For $\langle \sigma_1, \sigma_2 \rangle \in \Sigma^1_G \times \Sigma^2_G$ and set of finite partial plays $\Pi \subseteq \widetilde{\Pi}_{\sigma_1,\sigma_2}$ we let $[\Pi]_{\langle \sigma_1,\sigma_2 \rangle} \subseteq \Pi^\infty_{\sigma_1,\sigma_2}$ denote the obvious mapping of finite partial plays onto *infinite* plays consistent with σ_1 and σ_2.

Definition 15 (Arena)**.** Let $\hat{G} = \langle \hat{S}, \hat{I}, \hat{T}, \hat{L}^!, \hat{L}^? \rangle$, $G = \langle S, I, T, L^!, L^? \rangle$ be games *without* deadlocks. An *arena* of $\hat{G}$ and G is a tuple $\langle \mathrm{A}_1, \mathrm{A}_2 \rangle$, where:

$$\mathrm{A}_1 \subseteq \hat{S} \times S$$
$$\mathrm{A}_2 \subseteq \mathbb{PD}(\hat{S}) \times \mathbb{PD}(S)$$

such that A_1 consists of tuples of player 1 configurations and A_2 consists of tuples of player 2 configurations.

Let $\langle \mathrm{A}_1, \mathrm{A}_2 \rangle$ be an arena, and let $\hat{\pi} \in \widetilde{\Pi}_{\hat{G}}$ and $\pi \in \widetilde{\Pi}_G$ be two partial plays. We call $\hat{\pi}$ and π $\langle \mathrm{A}_1, \mathrm{A}_2 \rangle$-*invariant* if and only if $|\pi| = |\hat{\pi}|$ and for every $i \leq |\pi| = |\hat{\pi}|$:

$$\langle \pi(i), \hat{\pi}(i) \rangle \in (\mathrm{A}_1 \cup \mathrm{A}_2)$$

Let $\langle \hat{\sigma}_1, \hat{\sigma}_2 \rangle \in \Sigma^1_{\hat{G}} \times \Sigma^2_{\hat{G}}$ and $\langle \sigma_1, \sigma_2 \rangle \in \Sigma^1_G \times \Sigma^2_G$ be strategy pairs. We call $\langle \hat{\sigma}_1, \hat{\sigma}_2 \rangle$ and $\langle \sigma_1, \sigma_2 \rangle$ $\langle \mathrm{A}_1, \mathrm{A}_2 \rangle$-*invariant* if and only if for all finite $\langle \mathrm{A}_1, \mathrm{A}_2 \rangle$-invariant partial plays $\hat{\pi} \in \widetilde{\Pi}_{\hat{\sigma}_1, \hat{\sigma}_2}$ and $\pi \in \widetilde{\Pi}_{\sigma_1, \sigma_2}$ the following conditions hold:

$$\langle \overrightarrow{\hat{\pi}}, \overrightarrow{\pi} \rangle \in \mathrm{A}_1 \Rightarrow \langle \hat{\sigma}_1(\hat{\pi}), \sigma_1(\pi) \rangle \in \mathbb{D}(\mathrm{A}_2) \tag{2a}$$

$$\langle \overrightarrow{\hat{\pi}}, \overrightarrow{\pi} \rangle \in \mathrm{A}_2 \Rightarrow \langle \overrightarrow{\hat{\pi}} \circ \hat{\sigma}_2(\hat{\pi}), \overrightarrow{\pi} \circ \sigma_2(\pi) \rangle \in \mathbb{D}(\mathrm{A}_1) \tag{2b}$$

Let $\hat{\sigma} \in \Sigma^1_{\hat{G}} \times \Sigma^2_{\hat{G}}$ and $\sigma \in \Sigma^1_G \times \Sigma^2_G$ be two $\langle \mathrm{A}_1, \mathrm{A}_2 \rangle$-invariant strategy pairs and $\pi \in \widetilde{\Pi}_\sigma$ and $\hat{\pi} \in \widetilde{\Pi}_{\hat{\sigma}}$ two finite $\langle \mathrm{A}_1, \mathrm{A}_2 \rangle$-invariant partial plays. We denote with $\delta^{\langle \hat{\sigma}, \sigma \rangle}_{\langle \hat{\pi}, \pi \rangle}$ the weight function witnessing (2a) resp. (2b).

Properties of arenas We now show some very useful properties for two deadlock-free games under the assumption that we have tuples of strategies that are invariant under an arena.

In this section we fix $\hat{G} = \langle \hat{S}, \hat{I}, \hat{T}, \hat{L}^!, \hat{L}^? \rangle$, $G = \langle S, I, T, L^!, L^? \rangle$ to be deadlock-free games and $\langle \mathrm{A}_1, \mathrm{A}_2 \rangle$ to be an arena of $\hat{G}$ and G. We first explore the relation between the partial plays of $\hat{G}$ and G.

Definition 16 (Weight)**.** Let $\sigma \in \Sigma^1_G \times \Sigma^2_G$ and $\hat{\sigma} \in \Sigma^1_{\hat{G}} \times \Sigma^2_{\hat{G}}$ be two $\langle \mathrm{A}_1, \mathrm{A}_2 \rangle$-invariant strategy pairs. We define a weight function $\mathrm{w}^{\hat{\sigma}}_{\sigma} \in \widetilde{\Pi}_{\hat{G}} \times \widetilde{\Pi}_G \to [0, 1]$ as follows:

$$\mathrm{w}^{\hat{\sigma}}_{\sigma}(\hat{\pi}, \pi) = \begin{cases} \prod_{i=0}^{|\pi|-1} \delta^{\langle \hat{\sigma}, \sigma \rangle}_{\langle \hat{\pi}^i, \pi^i \rangle}(\pi(i{+}1), \hat{\pi}(i{+}1)) & \text{if } \langle \mathrm{A}_1, \mathrm{A}_2 \rangle\text{-invariant} \\ 0 & \text{otherwise.} \end{cases}$$

Intuitively, the weight $\mathrm{w}^{\hat{\sigma}}_{\sigma}(\hat{\pi}, \pi)$ is the amount of 'probability mass' of $\hat{\pi}$ and π that the plays have attributed to simulating each other under $\hat{\sigma}$ and σ.

Lemma B.1. Let $\hat{\sigma} \in \Sigma^1_{\hat{G}} \times \Sigma^2_{\hat{G}}$ and $\sigma \in \Sigma^1_G \times \Sigma^2_G$ be two $\langle \mathrm{A}_1, \mathrm{A}_2 \rangle$-invariant strategy

pairs, then, for any $\langle \hat{\omega}, \omega \rangle \in (\mathrm{A}_1 \cup \mathrm{A}_2)$:

$$\forall \pi \in \widetilde{\Pi}_\sigma(\omega) : \sum_{\hat{\pi} \in \widetilde{\Pi}_{\hat{G}}(\hat{\omega})} \mathrm{w}_\sigma^{\hat{\sigma}}(\hat{\pi}, \pi) = \mathbf{Pr}_\sigma^{\omega}([\pi]_\sigma) \tag{3a}$$

$$\forall \hat{\pi} \in \widetilde{\Pi}_{\hat{\sigma}}(\hat{\omega}) : \sum_{\pi \in \widetilde{\Pi}_G(\omega)} \mathrm{w}_\sigma^{\hat{\sigma}}(\hat{\pi}, \pi) = \mathbf{Pr}_{\hat{\sigma}}^{\hat{\omega}}([\hat{\pi}]_{\hat{\sigma}}) \tag{3b}$$

$$\forall \pi \in \widetilde{\Pi}_G(\omega) \setminus \widetilde{\Pi}_\sigma(\omega) : \sum_{\hat{\pi} \in \widetilde{\Pi}_{\hat{G}}(\hat{\omega})} \mathrm{w}_\sigma^{\hat{\sigma}}(\hat{\pi}, \pi) = 0 \tag{3c}$$

$$\forall \hat{\pi} \in \widetilde{\Pi}_{\hat{G}}(\hat{\omega}) \setminus \widetilde{\Pi}_{\hat{\sigma}}(\hat{\omega}) : \sum_{\pi \in \widetilde{\Pi}_G(\omega)} \mathrm{w}_\sigma^{\hat{\sigma}}(\hat{\pi}, \pi) = 0 \tag{3d}$$

Proof. Let us first proof (3c). Clearly, any play $\pi \in \Pi_G(s)$ that is inconsistent with σ has to be of non-zero length and has to make a move that is inconsistent with σ; there are two types of inconsistent transitions:

- The first type of an inconsistent transition occurs when for some $i < |\pi|$, we have $\pi(i) \in S$ and $\pi(i+1) \in T(\pi(i))$ but $\sigma_1(\pi^i)(\pi(i+1)) = 0$. Due to (1b), the weight function $\delta^{\langle \hat{\sigma}, \sigma \rangle}_{\langle \hat{\pi}^i, \pi^i \rangle}(\hat{\pi}(i+1), \pi(i+1))$ will always yield 0 for any $\hat{\pi}$.
- The second type of inconsistency occurs when if some $i < |\pi|$, we have $\pi(i) \in \mathbb{PD}(S)$ and $\pi(i+1) \in S$ such that for some $\lambda \in \pi(i)$ we have $\lambda(\pi(i+1)) > 0$, but for no such λ we have that $\sigma_2(\pi^i)(\lambda) > 0$ and hence $(\pi(i) \circ \sigma_2(\pi^i))(\pi(i+1)) = 0$ and, by (1b), the weight function $\delta^{\langle \hat{\sigma}, \sigma \rangle}_{\langle \hat{\pi}^i, \pi^i \rangle}(\hat{\pi}(i+1), \pi(i+1))$ will always yield 0 for any $\hat{\pi}$.

Hence, inconsistent partial plays do not contribute anything to the the sum of (3c). The proof of (3d) follows from symmetry.

From (3c) and (3d) we learn that $\mathrm{w}_\sigma^{\hat{\sigma}}(\hat{\pi}, \pi)$ can only be non-zero when π is consistent with σ and $\hat{\pi}$ is consistent with $\hat{\sigma}$. Hence, we can reduce the sum in (3a) and (3b) to plays consistent with σ and $\hat{\sigma}$ at will, which we will use implicitly in the following proofs.

We will proof (3a) by induction on the length of partial plays under consideration. Note that the base case, partial plays of length 0, trivially satisfy (3a). Suppose (3a) holds for all partial plays of size i. We will prove (3a) also holds for plays of size $i+1$. Let $\pi' \in \widetilde{\Pi}_\sigma(s)$ be an arbitrary partial play of length $i+1$ consistent with σ. We split the proof into the following cases cases: $\pi'(i) \in S$, and $\pi'(i) \in \mathbb{PD}(S)$.

First suppose $\pi'(i) \in S$, then by definition π' is of the form $\pi \frown \Lambda$, for some $\pi \in \widetilde{\Pi}_\sigma(\omega)$

of length i such that $\overrightarrow{\pi} \in S$ and some non-empty $\Lambda \in \mathbb{PD}(S)$ with $\sigma_1(\pi)(\Lambda) > 0$.

$$
\begin{aligned}
&\mathbf{Pr}^{\omega}_{\sigma}([\pi \frown \Lambda]_{\sigma}) = \\
&\sigma_1(\pi)(\Lambda) \cdot \mathbf{Pr}^{\omega}_{\sigma}([\pi]_{\sigma}) = \\
&\quad \sum_{\hat{\pi} \in \Pi_{\hat{G}}(\hat{\omega})} \left(\mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi}, \pi) \cdot \sigma_1(\pi)(\Lambda) \right) = && \text{(Ind.)} \\
&\quad \sum_{\hat{\pi} \in \Pi_{\hat{G}}(\hat{\omega})} \left(\mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi}, \pi) \cdot \sum_{\hat{\Lambda} \in (\mathbb{PD}(\hat{S}))} \delta^{\langle \hat{\sigma}, \sigma \rangle}_{\langle \hat{\pi}, \pi \rangle}(\hat{\Lambda}, \Lambda) \right) = && \text{(Eq. 1b)} \\
&\quad \sum_{\hat{\pi} \in \Pi_{\hat{G}}(\hat{\omega})} \left(\sum_{\hat{\Lambda} \in \mathbb{PD}(\hat{S})} \left(\mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi}, \pi) \cdot \delta^{\langle \hat{\sigma}, \sigma \rangle}_{\langle \hat{\pi}, \pi \rangle}(\hat{\Lambda}, \Lambda) \right) \right) = \\
&\quad \sum_{\hat{\pi} \frown \hat{\Lambda} \in \Pi_{\hat{G}}(\hat{\omega})} \mathrm{W}^{\hat{\sigma}}_{\sigma}(\pi \frown \Lambda, \hat{\pi} \frown \hat{\Lambda}) = && \text{(Def. 16)}
\end{aligned}
$$

Now, suppose $\pi'(i) \in \mathbb{PD}(S)$, then by definition π' is of the form $\pi \frown s$, for some $\pi \in \widetilde{\Pi}_{\sigma}(\omega)$ of length i such that $\overrightarrow{\pi} \in \mathbb{PD}(S)$ and some $s \in S$ such that $(\overrightarrow{\pi} \circ \sigma_2(\pi))(s) > 0$.

$$
\begin{aligned}
&\mathbf{Pr}^{\omega}_{\sigma}([\pi \frown s]_{\sigma}) = \\
&(\overrightarrow{\pi} \circ \sigma_2(\pi))(s) \cdot \mathbf{Pr}^{\omega}_{\sigma}([\pi]_{\sigma}) = \\
&\quad \sum_{\hat{\pi} \in \Pi_{\hat{G}}(\hat{\omega})} \left(\mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi}, \pi) \cdot (\overrightarrow{\pi} \circ \sigma_2(\pi))(s) \right) = && \text{(Ind.)} \\
&\quad \sum_{\hat{\pi} \in \Pi_{\hat{G}}(\hat{\omega})} \left(\mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi}, \pi) \cdot \sum_{\hat{s} \in \hat{S}} \delta^{\langle \hat{\sigma}, \sigma \rangle}_{\langle \hat{\pi}, \pi \rangle}(\hat{s}, s) \right) = && \text{(Eq. 1b)} \\
&\quad \sum_{\hat{\pi} \in \Pi_{\hat{G}}(\hat{\omega})} \left(\sum_{\hat{s} \in \hat{S}} \left(\mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi}, \pi) \cdot \delta^{\langle \hat{\sigma}, \sigma \rangle}_{\langle \hat{\pi}, \pi \rangle}(\hat{s}, s) \right) \right) = \\
&\quad \sum_{\hat{\pi} \frown \hat{s} \in \Pi_{\hat{G}}(\hat{\omega})} \mathrm{W}^{\hat{\sigma}}_{\sigma}(\pi \frown s, \hat{\pi} \frown \hat{s}) = && \text{(Def. 16)}
\end{aligned}
$$

The proof for (3b) follows symmetrically. □

Under a pair of strategies $\hat{\sigma} \in \Sigma^1_{\hat{G}} \times \Sigma^2_{\hat{G}}$ we call a set of partial finite plays $P \subseteq \widetilde{\Pi}_{\hat{\sigma}}(\hat{s})$ *disjoint* iff for every $\pi, \pi' \in P$ we have that $[\pi]_{\hat{\sigma}} \cap [\pi']_{\hat{\sigma}} = \emptyset$.

Informally, this means we can compute the probability of $\mathbf{Pr}^{\hat{s}}_{\hat{\sigma}}([P]_{\hat{\sigma}})$ by means of the sum $\sum_{p \in P} \mathbf{Pr}^{\hat{s}}_{\hat{\sigma}}([p]_{\hat{\sigma}})$.

We now show how arenas relate probabilities of certain disjoint sets of partial plays:

Proposition B.2. Let $\hat{\sigma} \in \Sigma^1_{\hat{G}} \times \Sigma^2_{\hat{G}}$ and $\sigma \in \Sigma^1_G \times \Sigma^2_G$ be two $\langle \mathrm{A}_1, \mathrm{A}_2 \rangle$-invariant strategy pairs, let $\langle \hat{\omega}, \omega \rangle \in (\mathrm{A}_1 \cup \mathrm{A}_2)$ and let $\mathrm{CYL}^{!}_{\hat{G}}, \mathrm{CYL}^{?}_{\hat{G}} \subseteq \widetilde{\Pi}_{\hat{\sigma}}(\hat{\omega})$ and $\mathrm{CYL}^{!}_{G}, \mathrm{CYL}^{?}_{G} \subseteq \widetilde{\Pi}_{\sigma}(\omega)$ be *disjoint* sets of finite partial plays such that if for arbitrary $\langle \mathrm{A}_1, \mathrm{A}_2 \rangle$-invariant $\hat{\pi} \in \widetilde{\Pi}_{\hat{\sigma}}(\hat{\omega})$

and $\pi \in \widetilde{\Pi}_\sigma(\omega)$ we have

$$(\hat{\pi} \in \textsc{cyl}^{!}_{\hat{G}}) \Rightarrow (\exists i \leq |\pi| \colon \pi^i \in \textsc{cyl}^{!}_{G}) \tag{4a}$$

$$(\exists i \leq |\hat{\pi}| \colon \hat{\pi}^i \in \textsc{cyl}^{?}_{\hat{G}}) \Leftarrow (\pi \in \textsc{cyl}^{?}_{G})\ . \tag{4b}$$

then, the following inequalities hold

$$\mathbf{Pr}^{\hat{s}}_{\hat{\sigma}}([\textsc{cyl}^{!}_{\hat{G}}]_{\hat{\sigma}}) \leq \mathbf{Pr}^{s}_{\sigma}([\textsc{cyl}^{!}_{G}]_{\sigma}) \tag{4c}$$

$$\mathbf{Pr}^{\hat{s}}_{\hat{\sigma}}([\textsc{cyl}^{?}_{\hat{G}}]_{\hat{\sigma}}) \geq \mathbf{Pr}^{s}_{\sigma}([\textsc{cyl}^{?}_{G}]_{\sigma}) \tag{4d}$$

Proof. We extend the notion of disjointness to *sets* of tuples of $\langle \mathrm{A}_1, \mathrm{A}_2\rangle$-invariant partial plays. We call such a $P \subseteq \widetilde{\Pi}_{\hat{\sigma}}(\hat{s}) \times \widetilde{\Pi}_\sigma(s)$ *disjoint* if for every $\langle \hat{\pi}, \pi\rangle, \langle \hat{\pi}', \pi'\rangle \in P$ we have that either π and π' are disjoint *or* $\hat{\pi}$ and $\hat{\pi}'$ are disjoint.

For a disjoint set of tuples of $\langle \mathrm{A}_1, \mathrm{A}_2\rangle$-invariant partial plays P and and $\langle \mathrm{A}_1, \mathrm{A}_2\rangle$-invariant partial play $\langle \hat{\pi}', \pi'\rangle$ such that for every $\langle \hat{\pi}, \pi\rangle \in P$ we have $\langle \hat{\pi}^{|\hat{\pi}'|}, \hat{\pi}^{|\pi'|}\rangle = \langle \hat{\pi}', \pi'\rangle$ then by Definition 16:

$$\sum_{\langle \hat{\pi}, \pi\rangle \in P} \mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi}, \pi) \leq \mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi}', \pi')\ .$$

Now, for some $\hat{\pi} \in \textsc{cyl}^{!}_{\hat{G}}$ let $\textsc{img}_G(\hat{\pi})$ denote partial plays $\pi \in \widetilde{\Pi}_\sigma(\omega)$ such that $\hat{\pi}$ and π are $\langle \mathrm{A}_1, \mathrm{A}_2\rangle$-invariant. Let $P = \{\langle \hat{\pi}, \pi\rangle \mid \hat{\pi} \in \textsc{cyl}^{!}_{\hat{G}}, \pi \in \textsc{img}_G(\hat{\pi})\}$. By construction P is disjoint.

Similarly, for some $\pi' \in \textsc{cyl}^{!}_{G}$ let $\textsc{img}_{\hat{G}}(\pi')$ denote plays $\hat{\pi}' \in \widetilde{\Pi}_{\hat{\sigma}}(\hat{\omega})$ such that $\hat{\pi}'$ and π' are $\langle \mathrm{A}_1, \mathrm{A}_2\rangle$-invariant.

Observe that for each $\langle \hat{\pi}, \pi\rangle \in P$, we have $\hat{\pi} \in \textsc{cyl}^{!}_{\hat{G}}$ and $\hat{\pi}$ and π are $\langle \mathrm{A}_1, \mathrm{A}_2\rangle$-invariant. Hence, by (4a), we infer that for some $i \leq |\pi|$ we have $\pi^i \in \textsc{cyl}^{!}_{G}$. Also we know that $\hat{\pi}^i$ is $\langle \mathrm{A}_1, \mathrm{A}_2\rangle$-invariant with π^i.

We can rephrase the last paragraph as follows: for each $\langle \hat{\pi}, \pi\rangle \in P$ there exist $\pi' \in \textsc{cyl}^{!}_{G}$ and $\hat{\pi}' \in \textsc{img}_{\hat{G}}(\pi')$ such that $\langle \hat{\pi}^{|\hat{\pi}'|}, \pi^{|\pi'|}\rangle = \langle \hat{\pi}', \pi'\rangle$.

This yields the following inequality:

$$\sum_{\hat{\pi} \in \textsc{cyl}^{!}_{\hat{G}}} \sum_{\pi \in \textsc{img}_G(\hat{\pi})} \mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi}, \pi) \leq \sum_{\pi \in \textsc{cyl}^{!}_{G}} \sum_{\hat{\pi} \in \textsc{img}_{\hat{G}}(\pi)} \mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi}, \pi) \tag{4e}$$

We are now in a position to prove inequality (4c):

$$
\begin{aligned}
\mathbf{Pr}^{\hat{\omega}}_{\hat{\sigma}}([\mathrm{CYL}^{!}_{\hat{G}}]_{\hat{\sigma}}) &= \sum_{\hat{\pi}\in\mathrm{CYL}^{!}_{\hat{G}}} \mathbf{Pr}^{\hat{\omega}}_{\hat{\sigma}}([\hat{\pi}]_{\hat{\sigma}}) && \text{(Disjoint)}\\
&= \sum_{\hat{\pi}\in\mathrm{CYL}^{!}_{\hat{G}}} \sum_{\pi\in\Pi_{\sigma}(\omega)} \mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi},\pi) && \text{(Eq. 3b)}\\
&= \sum_{\hat{\pi}\in\mathrm{CYL}^{!}_{\hat{G}}} \sum_{\pi\in\mathrm{IMG}_{G}(\hat{\pi})} \mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi},\pi) && \text{(Def. } \mathrm{IMG}_{G}\text{)}\\
&\leq \sum_{\pi\in\mathrm{CYL}^{!}_{G}} \sum_{\hat{\pi}\in\mathrm{IMG}_{\hat{G}}(\pi)} \mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi},\pi) && \text{(Eq. 4e)}\\
&= \sum_{\pi\in\mathrm{CYL}^{!}_{G}} \sum_{\hat{\pi}\in\Pi_{\hat{\sigma}}(\hat{\omega})} \mathrm{W}^{\hat{\sigma}}_{\sigma}(\hat{\pi},\pi) && \text{(Def. } \mathrm{IMG}_{\hat{G}}\text{)}\\
&= \sum_{\pi\in\mathrm{CYL}^{!}_{G}} \mathbf{Pr}^{\omega}_{\sigma}([\pi]_{\sigma}) && \text{(Eq. 3a)}\\
&= \mathbf{Pr}^{\omega}_{\sigma}([\mathrm{CYL}^{!}_{G}]_{\sigma}) && \text{(Disjoint)}
\end{aligned}
$$

The proof of (4d) follows symmetrically. □

Arenas and game-simulation Let $\hat{G} = \langle \hat{S}, \hat{I}, \hat{T}, \hat{L}^{!}, \hat{L}^{?}\rangle$, $G = \langle S, I, T, L^{!}, L^{?}\rangle$ be *arbitrary* games and let $R \subseteq \hat{S} \times S$ be such that $\hat{G} \sqsubseteq^{R}_{\mathscr{G}} G$.

The results in the previous sections hold for deadlock-free games only. Fortunately, however, PCTL semantics of $\hat{G}$ and G are defined over deadlock-free $\hat{G}_{\perp} = \langle \hat{S}_{\perp}, \hat{I}, \hat{T}_{\perp}, \hat{L}^{!}_{\perp}, \hat{L}^{?}_{\perp}\rangle$ and $G_{\perp} = \langle S_{\perp}, I, T_{\perp}, L^{!}_{\perp}, L^{?}_{\perp}\rangle$. Let $R_{\perp} \subseteq \hat{S}_{\perp} \times S_{\perp}$ be such that $\langle \hat{s}, s\rangle \in R_{\perp}$ iff $\langle \hat{s}, s\rangle \in R$, $s = \perp_1$ or $\langle \hat{s}, s\rangle = \langle \perp_2, \perp_2\rangle$.

We can now show that the conditions of Definition 11 are precisely what we need to construct the following arenas:

- Let $H_{(\mathrm{iii})} \subseteq \mathbb{PD}(\hat{S}_{\perp}) \times \mathbb{PD}(S_{\perp})$ be a relation such that $\langle \hat{\Lambda}, \Lambda\rangle \in H_{(\mathrm{iii})}$ if and only if for every $\lambda \in \Lambda$ there exists a weight distribution $\hat{\Lambda}_C \in \mathbb{D}(\hat{\Lambda})$ such that $\langle \hat{\Lambda} \circ \hat{\lambda}_C, \lambda\rangle \in \mathbb{D}(R_{\perp})$.
- We let $H_{(\mathrm{iv})} \subseteq \mathbb{PD}(\hat{S}_{\perp}) \times \mathbb{PD}(S_{\perp})$ be a relation such that $\langle \hat{\Lambda}, \Lambda\rangle \in H_{(\mathrm{iv})}$ if and only if for every $\hat{\lambda} \in \hat{\Lambda}$ there exists a weight distribution $\lambda_C \in \mathbb{D}(\Lambda)$ such that $\langle \hat{\lambda}, \Lambda \circ \lambda_C\rangle \in \mathbb{D}(R_{\perp})$.

Clearly, $H_{(\mathrm{iii})}$ has a direct correspondence to condition (iii) of Def. 11. The arena $\langle R_{\perp}, H_{(\mathrm{iii})}\rangle$ will be used to show Next and Until formulae are preserved under the modality "!". In contrast, $\langle R_{\perp}, H_{(\mathrm{iv})}\rangle$ corresponds to condition (iv) of Def. 11 and will used to show Next and Until formulae are preserved under the modality "?". For this, we need

to link the notions of arenas $\langle R_\perp, H_{(\text{iii})}\rangle$ and $\langle R_\perp, H_{(\text{iv})}\rangle$ with the existence of invariant strategy pairs:

Lemma B.3. For every $\langle \sigma_1, \sigma_2\rangle \in \Sigma^1_{G_\perp} \times \Sigma^2_{G_\perp}$ there exists $\langle \hat{\sigma}_1, \hat{\sigma}_2\rangle \in \Sigma^1_{\hat{G}_\perp} \times \Sigma^2_{\hat{G}_\perp}$ such that $\langle \hat{\sigma}_1, \hat{\sigma}_2\rangle$ and $\langle \sigma_1, \sigma_2\rangle$ are $\langle R_\perp, H_{(\text{iii})}\rangle$-invariant.

Proof. We will show that for any $\langle \sigma_1, \sigma_2\rangle$ we can always define $\langle \hat{\sigma}_1, \hat{\sigma}_2\rangle$ in such a way that (2a) and (2b) are satisfied. To show this for (2a) consider that whenever $\langle \overrightarrow{\pi}, \overrightarrow{\hat{\pi}}\rangle \in R_\perp$, by the deadlock-freeness of $\hat{G}$ and G, condition (iii) yields that for any every $\Lambda \in T_\perp(\overrightarrow{\pi})$ there exists $\hat{\Lambda} \in \hat{T}_\perp(\overrightarrow{\hat{\pi}})$ such that $\langle \hat{\Lambda}, \Lambda\rangle \in H_{(\text{iii})}$. It follows that for any distribution $\sigma_1(\pi) \in \mathbb{D}(T_\perp(\overrightarrow{\pi}))$ we can construct a distribution over $\hat{\lambda} \in \hat{T}_\perp(\overrightarrow{\hat{\pi}})$ such that $\langle \hat{\lambda}, \sigma_1(\pi)\rangle \in \mathbb{D}(H_{(\text{iii})})$. If we take $\hat{\sigma}_1(\hat{\pi})$ to be $\hat{\lambda}$ we obtain $\langle \hat{\sigma}_1(\hat{\pi}), \sigma_1(\pi)\rangle \in \mathbb{D}(H_{(\text{iii})})$.

Remaining to show is that (2b) is also satisfied. Note that $\langle \overrightarrow{\hat{\pi}}, \overrightarrow{\pi}\rangle \in H_{(\text{iii})}$. As we are only considering partial plays of $\hat{G}_\perp$ and $G_\perp$ we can safely assume neither play ends in $\emptyset$. Hence, by the definition of $H_{(\text{iii})}$, for every $\lambda \in \overrightarrow{\pi}$ there exists a weight distribution $\hat{\lambda}_C \in \mathbb{D}(\overrightarrow{\hat{\pi}})$ such that $\langle \overrightarrow{\hat{\pi}} \circ \hat{\lambda}_C, \lambda\rangle =\in \mathbb{D}(R_\perp)$. From this it follows that for any *distribution* $\sigma_2(\pi) \in \mathbb{D}(\overrightarrow{\pi})$ there exists a weight distribution $\hat{\lambda}'_C \in \mathbb{D}(\overrightarrow{\hat{\pi}})$ such that $\langle \overrightarrow{\hat{\pi}} \circ \hat{\lambda}'_C, \overrightarrow{\pi} \circ \sigma_2(\pi)\rangle \in \mathbb{D}(R_\perp)$. If we take $\hat{\sigma}_2(\hat{\pi})$ to be $\hat{\lambda}'_C$ then (2b) is satisfied, meaning that $\langle \hat{\sigma}_1, \hat{\sigma}_2\rangle$ and $\langle \sigma_1, \sigma_2\rangle$ are $\langle R_\perp, H_{(\text{iii})}\rangle$-invariant. □

Lemma B.4. For every $\sigma_1 \in \Sigma^1_{G_\perp}$ there exists a $\hat{\sigma}_1 \in \Sigma^1_{\hat{G}_\perp}$, and, independently, for every $\hat{\sigma}_2 \in \Sigma^2_{\hat{G}_\perp}$ there exists a $\sigma_2 \in \Sigma^2_{G_\perp}$ such that $\langle \sigma_1, \sigma_2\rangle$ and $\langle \hat{\sigma}_1, \hat{\sigma}_2\rangle$ are $\langle R_\perp, H_{(\text{iv})}\rangle$-invariant.

Proof. Proof analogous to Lemma B.3, other than that condition (iv) of Definition 11 is used instead of condition (iii) and the ordering is swapped for the proof of (2b) (matching condition (iv)). □

Preservation of PCTL We are now finally in a position to prove Proposition 4.3, which we first recall:

Proposition 4.3. For any $\hat{G}, G \in \mathscr{G}$ with $\hat{G} \sqsubseteq_{\mathscr{G}} G$ we have that $\hat{G} \models_{\mathscr{G}} \phi \Rightarrow G \models_{\mathscr{G}} \phi$ for all $\phi \in \Phi_{\text{PCTL}}$.

Proof. Suppose $\hat{G} \sqsubseteq^R_{\mathscr{G}} G$. It is sufficient to show that if $\hat{G}_\perp \models_{\mathscr{G}} \phi$ then $G_\perp \models_{\mathscr{G}} \phi$. Let $\hat{G}_\perp = \langle \hat{S}_\perp, \hat{I}_\perp, \hat{T}_\perp, \hat{L}^!_\perp, \hat{L}^?_\perp\rangle$ and let $G_\perp = \langle S_\perp, I_\perp, T_\perp, L^!_\perp, L^?_\perp\rangle$.

Considering that $I \subseteq R.\hat{I}$ it is sufficient to show that $\langle \hat{s}, s\rangle \in R$ implies that $s \models^! \phi \Leftarrow \hat{s} \models^! \phi$ for arbitrary $\phi \in \Phi_{\text{PCTL}}$. We do this by structural induction on ϕ with the

induction hypothesis that for all $\phi \in \Phi_{\text{PCTL}}$, $\langle \hat{s}, s \rangle \in R$:

$$s \models^! \phi \;\Leftarrow\; \hat{s} \models^! \phi \tag{5a}$$

$$s \models^? \phi \;\Rightarrow\; \hat{s} \models^? \phi \tag{5b}$$

The base cases, in which $\phi \in \text{AP}$, follows immediately from (i) and (ii) of Definition 11. For negation $\neg\phi$ we have (using Definition 12 and (5a), (5b)):

$$s \models^! \neg\phi \Leftrightarrow s \not\models^? \phi \Leftarrow \hat{s} \not\models^? \phi \Leftrightarrow \hat{s} \models^! \neg\phi$$
$$s \models^? \neg\phi \Leftrightarrow s \not\models^! \phi \Rightarrow \hat{s} \not\models^! \phi \Leftrightarrow \hat{s} \models^? \neg\phi$$

For disjunction $\phi_1 \vee \phi_2$ we have (using Definition 12 and (5a), (5b)):

$$\hat{s} \models^! \phi_1 \vee \phi_2 \Leftrightarrow (\hat{s} \models^! \phi_1 \text{ or } \hat{s} \models^! \phi_1) \Rightarrow$$
$$(s \models^! \phi_1 \text{ or } s \models^! \phi_1) \Leftrightarrow s \models^! \phi_1 \vee \phi_2$$

and:

$$s \models^? \phi_1 \vee \phi_2 \Leftrightarrow (s \models^? \phi_1 \text{ or } s \models^? \phi_1) \Rightarrow$$
$$(\hat{s} \models^? \phi_1 \text{ or } \hat{s} \models^? \phi_1) \Leftrightarrow \hat{s} \models^? \phi_1 \vee \phi_2$$

Remaining to show is that the induction hypothesis holds for the probabilitic operator. From the PCTL semantics (Definition 12) it is easy to see that the implications are satisfied if and only if for every PCTL path formula $\psi \in \Psi_{\text{PCTL}}$ the following inequations hold:

$$\inf_{\hat{\sigma}_1,\hat{\sigma}_2} \left\{ \text{PROB}^!_{\hat{\sigma}_1\hat{\sigma}_2}(\hat{s},\psi) \right\} \leq \inf_{\sigma_1,\sigma_2} \left\{ \text{PROB}^!_{\sigma_1\sigma_2}(s,\psi) \right\} \tag{5c}$$

$$\sup_{\hat{\sigma}_1} \inf_{\hat{\sigma}_2} \left\{ \text{PROB}^?_{\hat{\sigma}_1\hat{\sigma}_2}(\hat{s},\psi) \right\} \geq \sup_{\sigma_1} \inf_{\sigma_2} \left\{ \text{PROB}^?_{\sigma_1\sigma_2}(s,\psi) \right\} \tag{5d}$$

$$\inf_{\hat{\sigma}_1} \sup_{\hat{\sigma}_2} \left\{ \text{PROB}^!_{\hat{\sigma}_1\hat{\sigma}_2}(\hat{s},\psi) \right\} \leq \inf_{\sigma_1} \sup_{\sigma_2} \left\{ \text{PROB}^!_{\sigma_1\sigma_2}(s,\psi) \right\} \tag{5e}$$

$$\sup_{\hat{\sigma}_1,\hat{\sigma}_2} \left\{ \text{PROB}^?_{\hat{\sigma}_1\hat{\sigma}_2}(\hat{s},\psi) \right\} \geq \sup_{\sigma_1,\sigma_2} \left\{ \text{PROB}^?_{\sigma_1\sigma_2}(s,\psi) \right\} \tag{5f}$$

In order to prove these inequalities indeed hold, we characterise the sets of plays being measured as a set of finite disjoint partial plays. For $G_\perp$ we define for every $s \in S$, strategy pair $\sigma \in \Sigma^1_{G_\perp} \times \Sigma^2_{G_\perp}$, PCTL formulae $\phi_1, \phi_2 \subseteq \Phi_{\text{PCTL}}$, modality $\text{M} \in \{!, ?\}$ and bound $k \in \mathbb{N} \cup \{\infty\}$ the following sets of finite partial plays:

– $\text{X}^{\text{M}}_\sigma(s, \phi_1) \subseteq \widetilde{\Pi}_\sigma(s)$ such that $\pi \in \text{X}^{\text{M}}_\sigma(s, \phi_1)$ iff $|\pi| = 2$ and $\pi(2) \models^{\text{M}} \phi_1$.

– $\mathrm{U}^{\mathrm{M}}_{\sigma}(s,k,\phi_1,\phi_2) \subseteq \widetilde{\Pi}_{\sigma}(s)$ such that $\pi \in \mathrm{U}^{\mathrm{M}}_{\sigma}(s,k,\phi_1,\phi_2)$ iff there exists $i \leq k$ such that $|\pi| = 2i$ and $\pi(2) \models^{\mathrm{M}} \phi_2$ and for all $j \leq i$ we have $\pi(2j) \models^{\mathrm{M}} \phi_1$ and $\pi(2j) \not\models^{\mathrm{M}} \phi_2$.

Note that these sets are disjoint by construction. Also note that it is easily verifiable that for all $\pi \in \Pi^{\infty}_{\sigma}(s)$

$$\pi \models^{\mathrm{M}} X\phi \Leftrightarrow \pi \in [\mathrm{X}^{\mathrm{M}}_{\sigma}(s,\phi)]_{\sigma}$$
$$\pi \models^{\mathrm{M}} (\phi_1 \, U^{\leq k} \, \phi_2) \Leftrightarrow \pi \in [\mathrm{U}^{\mathrm{M}}_{\sigma}(s,k,\phi_1,\phi_2)]_{\sigma}$$

Hence, it is possible to rewrite (5c), (5d), (5e) and (5f) using the disjoint sets of finite partial plays, e.g. :

$$\textsc{Prob}^{!}_{\hat{\sigma}}(\hat{s}, X\phi) = \mathbf{Pr}^{\hat{\mathrm{s}}}_{\hat{\sigma}}([\mathrm{X}^{!}_{\hat{\sigma}}(\hat{s},\phi)]_{\hat{\sigma}})$$

Now, consider an arena $\langle R_{\perp}, \mathrm{A}_2\rangle$ and $\langle R_{\perp}, \mathrm{A}_2\rangle$-invariant strategy pairs $\hat{\sigma} \in \Sigma^1_{\hat{G}} \times \Sigma^2_{\hat{G}}$ and $\sigma \in \Sigma^1_G \times \Sigma^2_G$. Then, using the induction hypothesis and $R_{\perp}$, for arbitrary $\langle R_{\perp}, \mathrm{A}_2\rangle$-invariant partial plays $\hat{\pi} \in \widetilde{\Pi}_{\hat{\sigma}}(\hat{s})$ and $\pi \in \widetilde{\Pi}_{\sigma}(s)$ we have

$$(\hat{\pi} \in \mathrm{X}^{!}_{\hat{\sigma}}(\hat{s},\phi)) \Rightarrow (\exists i{\leq}|\pi| : \pi^i \in \mathrm{X}^{!}_{\sigma}(s,\phi))$$
$$(\pi \in \mathrm{X}^{?}_{\sigma}(s,\phi)) \Rightarrow (\exists i{\leq}|\pi| : \hat{\pi}^i \in \mathrm{X}^{?}_{\hat{\sigma}}(\hat{s},\phi))$$
$$(\hat{\pi} \in \mathrm{U}^{!}_{\hat{\sigma}}(\hat{s},k,\phi_1,\phi_2)) \Rightarrow (\exists i{\leq}|\pi| : \pi^i \in \mathrm{U}^{!}_{\sigma}(s,k,\phi_1,\phi_2))$$
$$(\pi \in \mathrm{U}^{?}_{\sigma}(s,k,\phi_1,\phi_2)) \Rightarrow (\exists i{\leq}|\pi| : \hat{\pi}^i \in \mathrm{U}^{?}_{\hat{\sigma}}(\hat{s},k,\phi_1,\phi_2))$$

Hence, using Proposition B.2, if we have $\langle R_{\perp}, \mathrm{A}_2\rangle$-invariant strategy pairs, then:

$$\begin{aligned}
\textsc{Prob}^{!}_{\hat{\sigma}}(\hat{s}, X\phi) &= \mathbf{Pr}^{\hat{\mathrm{s}}}_{\hat{\sigma}}([\mathrm{X}^{!}_{\hat{\sigma}}(\hat{s},\phi)]_{\hat{\sigma}}) && \text{(Rewriting)}\\
&\leq \mathbf{Pr}^{s}_{\sigma}([\mathrm{X}^{!}_{\sigma}(s,\phi)]_{\sigma}) && \text{(4c)}\\
&= \textsc{Prob}^{!}_{\sigma}(s, X\phi) && \text{(Rewriting)}
\end{aligned}$$

Analogously we also obtain (using also (4d)):

$$\textsc{Prob}^{?}_{\hat{\sigma}}(\hat{s}, X\phi) \geq \textsc{Prob}^{?}_{\sigma}(s, X\phi)$$
$$\textsc{Prob}^{!}_{\hat{\sigma}}(\hat{s}, \phi_1 \, U^{\leq k} \, \phi_2) \leq \textsc{Prob}^{!}_{\sigma}(s, \phi_1 \, U^{\leq k} \, \phi_2)$$
$$\textsc{Prob}^{?}_{\hat{\sigma}}(\hat{s}, \phi_1 \, U^{\leq k} \, \phi_2) \geq \textsc{Prob}^{?}_{\sigma}(s, \phi_1 \, U^{\leq k} \, \phi_2)$$

Now recall Lemma B.3 shows that for arbitrary $\sigma \in \Sigma^1_{G_{\perp}} \times \Sigma^2_{G_{\perp}}$ there exists a corresponding strategy pair $\hat{\sigma} \in \Sigma^2_{\hat{G}_{\perp}} \times \Sigma^1_{\hat{G}_{\perp}}$ such that $\hat{\sigma}$ and σ are $\langle R_{\perp}, H_{(\mathrm{iii})}\rangle$-invariant, yielding that the inequality of (5c) and (5f) must be preserved for both Next and Until formulae.

Analogously, Lemma B.4 shows that for arbitrary $\sigma_1 \in \Sigma^1_{G_{\perp}}$ there exists a strategy

$\hat{\sigma}_1 \in \Sigma^1_{\hat{G}_\perp}$ such that for every strategy $\hat{\sigma}_2 \in \Sigma^2_{\hat{G}_\perp}$ there exists a strategy $\sigma_2 \in \Sigma^2_{G_\perp}$ such that $\langle \hat{\sigma}_1, \hat{\sigma}_2 \rangle$ and $\langle \sigma_1, \sigma_2 \rangle$ are $\langle R_\perp, H_{(\mathrm{iv})} \rangle$-invariant. Hence Lemma B.4 yields that the inequalities of (5d) and (5e) must be preserved for both Next and Until formulae. □

B.4 Proof of Lemma 4.4

Lemma 4.4. There exists $\hat{G}, G \in \mathscr{G}$ such that for all $\phi \in \Phi_{\mathrm{PCTL}}$ it holds that $\hat{G} \models_{\mathscr{G}} \phi \Rightarrow G \models_{\mathscr{G}} \phi$ but $\hat{G} \not\sqsubseteq_{\mathscr{G}} G$.

Proof. We will show that the two games in Figure 5 satisfying the lemma. Let $\hat{G} = \langle \hat{S}, \hat{I}, \hat{T}, \hat{L}^!, \hat{L}^? \rangle$ and $G = \langle S, I, T, L^!, L^? \rangle$. Let us define $\lambda_l = \mu_{s_3}$, $\lambda_m = \frac{2}{3} \cdot s_3 + \frac{1}{3} \cdot s_2$ and $\lambda_r = \frac{1}{2} \cdot s_1 + \frac{1}{2} \cdot s_2$, and let $\hat{\lambda}_l, \hat{\lambda}_r$ be the corresponding distributions in $\hat{G}$, e.g. $T(s_0) = \{\{\lambda_l, \lambda_m, \lambda_r\}\}$ and $\hat{T}(\hat{s}_0) = \{\{\hat{\lambda}_l, \hat{\lambda}_r\}\}$. We have that $\hat{G} \not\sqsubseteq_{\mathscr{G}} G$ as no combined transition of $\hat{T}(\hat{s}_0)$ can simulate λ_m.

It is sufficient to show that $\hat{G}_\perp$ and $G_\perp$ satisfy the same PCTL state formulae in each state (assuming the obvious mapping between states of $\hat{G}_\perp$ and $G_\perp$ and excluding $\perp_1, \perp_2$). We will prove this by structural induction over PCTL state formulae. More formally, our induction hypothesis is that for all $\phi \in \Phi_{\mathrm{PCTL}}$, $\hat{s} \in \hat{S}$ and $\mathrm{M} \in \{!, ?\}$ we have that $\hat{s}_i \models^{\mathrm{M}} \phi$ in $\hat{G}_\perp$ iff $s_i \models^{\mathrm{M}} \phi$ in $G_\perp$.

The base cases of our inductive argument are the PCTL formulae consisting of a single atomic proposition $a \in \mathrm{AP}$. Let $\hat{s}_i \in \hat{S}$ and $\mathrm{M} \in \{!, ?\}$. As $\hat{L}^{\mathrm{M}}(\hat{s}_i) = L^{\mathrm{M}}(s_i)$ we have that $\hat{s}_i \models^{\mathrm{M}} a$ iff $s_i \models^{\mathrm{M}} a$.

Let us now consider negation $\neg\phi$. Clearly, for any $\hat{s}_i \in \hat{S}$ and $\mathrm{M} \in \{!, ?\}$ we have $\hat{s}_i \models^{\mathrm{M}} \neg\phi$ if and only if $\hat{s}_i \not\models^{\neg \mathrm{M}} \phi$. Similarly, we have $\hat{s}_i \models^{\mathrm{M}} \neg\phi$ iff $s_i \not\models^{\neg \mathrm{M}} \phi$. By the induction hypothesis the satisfaction of ϕ coincides for $\hat{s}_i$ and s_i. Hence, negation preserves the induction hypothesis.

Let us now consider disjunction $\phi_1 \vee \phi_2$. Clearly, for any $\hat{s}_i \in \hat{S}$ and $\mathrm{M} \in \{!, ?\}$ we have $\hat{s}_i \models^{\mathrm{M}} \phi_1 \vee \phi_2$ if and only if ($\hat{s}_i \models^{\mathrm{M}} \phi_1$ or $\hat{s}_i \models^{\mathrm{M}} \phi_2$) and $s_i \models^{\mathrm{M}} \phi_1 \vee \phi_2$ if and only if ($s_i \models^{\mathrm{M}} \phi_1$ or $s_i \models^{\mathrm{M}} \phi_2$). By the induction hypothesis, satisfaction of both ϕ_1 and ϕ_2 coincide for $\hat{s}_i$ and s_i, therefore disjunction preserves the induction hypothesis.

Finally, consider PCTL formulae of the form $P_{\bowtie p}\langle\psi\rangle$, where $\bowtie \in \{<, \leq, >, \geq\}$, $p \in [0, 1]$ and $\psi \in \Psi_{\mathrm{PCTL}}$.

We first observe that under any pair of strategies the states in $\hat{S} \setminus \{\hat{s}_0\}$ and $S \setminus \{s_0\}$ will lead to $\perp_2$, in which no formula is satisfied. Satisfaction of $P_{\bowtie p}\langle\psi\rangle$ in these states therefore only depends on the state labelling and the modality. As both are the same for $\hat{G}$ and G, we trivially have that for all $\hat{s}_i \in \hat{S} \setminus \{\hat{s}_0\}$ and $\mathrm{M} \in \{!, ?\}$ we have $\hat{s}_i \models^{\mathrm{M}} P_{\bowtie p}\langle\psi\rangle$ iff $s_i \models^{\mathrm{M}} P_{\bowtie p}\langle\psi\rangle$. The remaining case concerns only satisfaction of the formulae of the

form $P_{\bowtie p}\langle\psi\rangle$ in $\hat{s}_0$ and s_0.

Our second observation is that, due to the absence of player 1 choice, there is only a single player 1 strategy in both games, denoted $\hat{\sigma}_1 \in \Sigma^1_{\hat{G}_\perp}$ and $\sigma_2 \in \Sigma^1_{G_\perp}$. Therefore, to show that $\hat{s}_0 \models^{\mathrm{M}} P_{\bowtie p}\langle\psi\rangle$ iff $s_0 \models^{\mathrm{M}} P_{\bowtie p}\langle\psi\rangle$ it is sufficient to show that for any $\psi \in \Psi_{\mathrm{PCTL}}$, $\mathrm{M} \in \{!, ?\}$:

$$\inf_{\hat{\sigma}_2 \in \Sigma^2_{\hat{G}_\perp}} \left\{\textsc{Prob}^{\mathrm{M}}_{\hat{\sigma}_1\hat{\sigma}_2}(\hat{s}_0, \psi)\right\} = \inf_{\sigma_2 \in \Sigma^2_{G_\perp}} \left\{\textsc{Prob}^{\mathrm{M}}_{\sigma_1\sigma_2}(s_0, \psi)\right\}$$
$$\sup_{\hat{\sigma}_2 \in \Sigma^2_{\hat{G}_\perp}} \left\{\textsc{Prob}^{\mathrm{M}}_{\hat{\sigma}_1\hat{\sigma}_2}(\hat{s}_0, \psi)\right\} = \sup_{\sigma_2 \in \Sigma^2_{G_\perp}} \left\{\textsc{Prob}^{\mathrm{M}}_{\sigma_1\sigma_2}(s_0, \psi)\right\} .$$

Clearly, every player 2 strategy of $\hat{G}_\perp$ has a corresponding strategy in $G_\perp$. By the induction hypothesis and the semantics of ψ, for any player 2 strategy $\hat{\sigma}_2 \in \Sigma^2_{\hat{G}}$ of G the corresponding strategy $\sigma_2 \in \Sigma^2_G$ yields

$$\textsc{Prob}^{\mathrm{M}}_{\hat{\sigma}_1\hat{\sigma}_2}(\hat{s}_0, \psi) = \textsc{Prob}^{\mathrm{M}}_{\sigma_1\sigma_2}(s_0, \psi) .$$

Therefore, we have

$$\inf_{\hat{\sigma}_2 \in \Sigma^2_{\hat{G}_\perp}} \left\{\textsc{Prob}^{\mathrm{M}}_{\hat{\sigma}_1\hat{\sigma}_2}(\hat{s}_0, \psi)\right\} \geq \inf_{\sigma_2 \in \Sigma^2_{G_\perp}} \left\{\textsc{Prob}^{\mathrm{M}}_{\sigma_1\sigma_2}(s_0, \psi)\right\}$$
$$\sup_{\hat{\sigma}_2 \in \Sigma^2_{\hat{G}_\perp}} \left\{\textsc{Prob}^{\mathrm{M}}_{\hat{\sigma}_1\hat{\sigma}_2}(\hat{s}_0, \psi)\right\} \leq \sup_{\sigma_2 \in \Sigma^2_{G_\perp}} \left\{\textsc{Prob}^{\mathrm{M}}_{\sigma_1\sigma_2}(s_0, \psi)\right\} .$$

Conversely, player 2 strategies of $\sigma_2 \in \Sigma^2_{G_\perp}$ of $G_\perp$ may not have a corresponding strategy in $\hat{G}_\perp$. As we are only concerned with the player 2 choices based on finite plays from s_0, the strategy σ_2 has no corresponding strategy in $\hat{G}_\perp$ precisely when the λ_m has a positive probability on the trivial path $\pi = s_0, \{\lambda_l, \lambda_m, \lambda_r\} \in \Pi^\infty_{G_\perp}(s_0)$, e.g. we have that $\sigma_2(\pi)(\lambda_m) > 0$. However, we will show this does not affect the infimum and supremum above, as there are always strategies $\sigma^!_2, \sigma^?_2 \in \Sigma^2_{G_\perp}$ such that $\sigma^!_2(\pi)(\lambda_m) = 0$ and $\sigma^?_2(\pi)(\lambda_m) = 0$ and which yield a lower respectively higher probability:

$$\textsc{Prob}^{\mathrm{M}}_{\sigma^!_1\sigma_2}(s_0, \psi) \leq \textsc{Prob}^{\mathrm{M}}_{\sigma_1\sigma_2}(s_0, \psi)$$
$$\textsc{Prob}^{\mathrm{M}}_{\sigma^?_1\sigma_2}(s_0, \psi) \geq \textsc{Prob}^{\mathrm{M}}_{\sigma_1\sigma_2}(s_0, \psi)$$

Clearly, $\sigma^!_2$ and $\sigma^!_1$ *do* have corresponding strategies in $\hat{G}_\perp$ and hence the infimum and supremum must be conserved, remaining to show is that such strategies can be constructed.

We construct $\sigma^!_2$ and $\sigma^!_1$ depending on ψ and modality M. Note that the semantics of a path formula only considers the states of a play, and not the sets of distributions and distributions in between states. As any infinite play from s_0 runs through s_1, s_2 or s_3 we

	$\sigma_2^!$	$\sigma_2^?$
$\emptyset$	λ_l, λ_r	λ_l, λ_r
s_1	λ_l	λ_l, λ_r
s_2	λ_l	λ_r
s_3	λ_r	λ_l
s_1, s_2	λ_l	λ_r
s_1, s_3	λ_r	λ_l
s_2, s_3	λ_l, λ_r	λ_l
s_1, s_2, s_3	λ_l, λ_r	λ_l, λ_r

Table 1: Redistribution of λ_m's probability according to S_ψ^{M}.

can characterise the set of infinite plays that satisfy ψ under modality M by means of a set $S_\psi^{\mathrm{M}} \subseteq \{s_1, s_2, s_3\}$.

With this knowledge we can modify σ_2's strategy on input π using table Table B.4 to obtain the strategies $\sigma_2^!$ and $\sigma_2^?$ that satisfy the requirements. More specifically, we pick a λ' according to Table B.4 (in the row of S_ψ^{M} and the column of $\sigma_2^!$). We let $\sigma_2^!(\pi)(\lambda_m) = 0$, $\sigma_2^!(\pi)(\lambda') = \sigma_2(\pi)(\lambda') + \sigma_2(\pi)(\lambda_m)$ and $\sigma_2^!(\pi)(\{\lambda_l, \lambda_r\} \setminus \{\lambda'\}) = \sigma_2(\pi)(\{\lambda_l, \lambda_r\} \setminus \{\lambda'\})$. We construct $\sigma_2^?$ similarly from Table B.4. It is easy to see that with this construction $\sigma_2^!$ and $\sigma_2^?$ yield a smaller and greater probability to satisfy ψ under M, respectively. As these strategies do have a corresponding strategy in $\hat{G}_\perp$ the infimum and supremum above are not affected by σ_2. This means that $\hat{s}_0 \models^{\mathrm{M}} P_{\bowtie p}\langle\psi\rangle$ iff $s_0 \models^{\mathrm{M}} P_{\bowtie p}\langle\psi\rangle$ and hence our induction hypothesis is preserved by the probabilistic operator. □

B.5 Proof of Lemma 4.6

Lemma 4.6 (R6). There is an MDP $M \in \mathscr{M}$ and PCTL formula $\phi \in \Phi_{\mathrm{PCTL}}$ such that $M \models_{\mathscr{M}} \phi$ and there is no *finite* game $\hat{G} \in \mathscr{G}$ such that $M \in \mathcal{I}(\hat{G})$ and $\hat{G} \models_{\mathscr{G}} \phi$.

Proof. We extend the proof of [11, Theorem 1]. Consider an MDP $M = \langle S, I, T, L\rangle$ where $S = \{\langle m, n\rangle \in \mathbb{N} \times \mathbb{N} \mid m \leq n+1\}$ and $I = \{0\} \times \mathbb{N}$. For every $n \in \mathbb{N}$ we have transitions of the form $T(\langle 0, n\rangle) = \{\mu_{\langle 1,n\rangle}\}$, $T(\langle 1, n\rangle) = \{\mu_{\langle 2,n\rangle}\}$, $\ldots$, $T(\langle n, n\rangle) = \{\mu_{\langle n+1,n\rangle}\}$. In addition, for every $n \in \mathbb{N}$ we label $L(\langle n+1, n\rangle) = \{q\}$ and for every $m \leq n$ we label $L(\langle m, n\rangle) = \emptyset$. Clearly $M \models_{\mathscr{M}} P_{>0}\langle \mathtt{tt}\, U\, q\rangle$.

We will show there does not exists a *finite* game $\hat{G} = \langle \hat{S}, \hat{I}, \hat{T}, \hat{L}^!, \hat{L}^?\rangle$ such that $\hat{G} \models_{\mathscr{G}} P_{>0}\langle \mathtt{tt}\, U\, q\rangle$ and $\hat{G} \sqsubseteq^R_{\mathscr{G}} e^{\mathscr{G}}(M)$ for some $R \subseteq \hat{S} \times S$. From Definition 12, this amounts to showing that for any such game $\hat{G}$ and $\hat{s} \in \hat{I}$ we have that

$$\inf_{\hat{\sigma}_1 \in \Sigma^1_{\hat{G}_\perp}} \inf_{\hat{\sigma}_2 \in \Sigma^2_{\hat{G}_\perp}} \left\{ \textsc{Prob}^!_{\hat{\sigma}_1\hat{\sigma}_2}(\hat{s}, \mathtt{tt}\, U\, q) \right\} = 0\,. \tag{6a}$$

Informally, we need to show that for any such $\hat{G}$ and $\hat{s} \in \hat{I}$ there exists a strategy pair under which the probability of reaching a state in which q must hold is 0. That is, such strategies must loop forever in non $q!$-states. We now construct such a strategy pair $\langle \hat{\sigma}_1, \hat{\sigma}_2 \rangle \in \Sigma^1_{\hat{G}} \times \Sigma^2_{\hat{G}}$ inductively over the lengths of plays $\hat{\pi} \in \Pi_{\hat{\sigma}_1,\hat{\sigma}_2}(\hat{s})$, such that for all $\hat{\pi}$ where $\overrightarrow{\hat{\pi}} \in \hat{S}$ we have that $\langle \overrightarrow{\hat{\pi}}, \langle l, k \rangle \rangle \in R$ for some $l \leq k$. Under such a strategy pair, as $L(\langle l, k \rangle) = \emptyset$, by Definition 11 and 12, we must have that $\overrightarrow{\hat{\pi}} \not\models^! q$ for *any* $\hat{\pi}$ consistent with $\hat{\sigma}_1$ and $\hat{\sigma}_2$, meaning that (6a) must hold.

Let k denote $|\hat{S}|$. Because $|\hat{I}|$ is finite and $|I|$ is infinite, due to the condition that $I \subseteq R.\hat{I}$, there must exist some $\hat{s} \in \hat{I}$ such that $R.\hat{s}$ contains infinitely many initial states. Hence, we must have that $\langle 0, k \rangle \in R.\hat{s}$ for some $k \geq |\hat{S}|$.

First suppose $\hat{\pi}$ has length 0 and hence $\hat{\pi} = \hat{s}$. As $\langle \hat{s}, \langle 0, k \rangle \rangle \in R$, by Definition 11 it must be that there exists a $\hat{\Lambda} \in \hat{T}(\hat{s})$ such that for some $\hat{\lambda}_C \in \mathbb{D}(\hat{\Lambda})$ we have that $\langle \hat{\Lambda} \circ \hat{\lambda}_C, \mu_{\langle 1,k \rangle} \rangle$. As the right-hand distribution is a point distribution, we have that any $\hat{\lambda} \in \text{SUPP}(\hat{\lambda}_C)$ also satisfies $\langle \hat{\lambda}, \mu_{\langle 1,k \rangle} \rangle \in \mathbb{D}(R)$. We therefore let $\hat{\sigma}_1(\hat{s}) = \mu_{\hat{\Lambda}}$ and $\hat{\sigma}_2(\hat{s}, \hat{\Lambda}) = \mu_{\hat{\lambda}}$. Clearly, with these strategies, all $\langle \hat{\sigma}_1, \hat{\sigma}_2 \rangle$-consistent plays $\hat{\pi}$ of length 3 are such that $\langle \overrightarrow{\hat{\pi}}, \langle 1, k \rangle \rangle \in R$, but it may also be the case that $\langle \overrightarrow{\hat{\pi}}, \langle 0, k \rangle \rangle \in R$, in which case we already have a loop.

Now suppose $\hat{\pi}$ is of length $3i$ for some $i \in \mathbb{N}$. We take the smallest $l \in \mathbb{N}$ such that $\langle \overrightarrow{\hat{\pi}}, \langle l, k \rangle \rangle \in R$ and can again choose $\hat{\sigma}_1$ and σ_2 such that all plays $\hat{\pi}$ of length $3(i+1)$ satisfy that $\langle \overrightarrow{\hat{\pi}}, \langle l+1, k \rangle \rangle \in R$. To show the invariant is preserved, we need to argue that we always have that $l + 1 \leq k$. First note that for plays of length $3i$, we have that the minimum l that satisfies $\langle \overrightarrow{\hat{\pi}}, \langle l, k \rangle \rangle \in R$ is at most i, as we always have that $\langle \overrightarrow{\hat{\pi}}, \langle i, k \rangle \rangle \in R$. Therefore, the only problematic case occurs with plays of length $3k$. However, we will show that in this case the minimum such l is not equal to k. In a play $\hat{\pi}$ of length $3k$ we must have $k + 1 > |\hat{S}|$ states. Clearly, $\hat{\pi}$ *must* contain a cycle, e.g. $\hat{\pi}(3k) = \hat{\pi}(3x)$ for some $x < k$. Hence, $l \leq x < k$. □

B.6 Proof of Proposition 4.7

We will construct a procedure that decides $\sqsubseteq_{\mathscr{G}}$ in polynomial time. This procedure is based on a decision procedure for strong probabilistic simulation on probabilistic automata [32].

As basic building block of deciding strong probabilistic game-refinement we use a decision procedure $\text{MATCHP2}_{X,Y}(\lambda_X, \Lambda_Y, R)$, parameterised by sets X and Y which, given a distribution $\lambda_X \in \mathbb{D}(X)$ over X, a non-empty set of distributions $\Lambda_Y \in \mathbb{PD}(Y)$ over Y and a relation $R \subseteq X \times Y$ returns **true** iff there exists a weight function $\lambda_C \in \mathbb{D}(\Lambda_Y)$ such that $\langle \lambda_X, \Lambda_Y \circ \lambda_C \rangle \in \mathbb{D}(R)$. An implementation of such a procedure is given in

[32, Section 4.4.1]. We refer to the pseudo-code below. This implementation runs in polynomial time using a reduction to a linear programming problem; other than this, no bound on the worst-case run-time complexity is provided. We will assume that in the worst case the procedure $\textsc{MatchP2}_{X,Y}(\lambda_X, \Lambda_Y, R)$ runs in $\mathcal{O}(f_{\mathrm{P1}}(|X|, |Y|, |\Lambda_Y|))$ time, where f_{P1} is polynomial in terms of $|X|$, $|Y|$ and $|\Lambda_Y|$ (we ignore the other inputs because $|\textsc{supp}(\lambda_X)|$ is bounded by $\mathcal{O}(|X|)$ and $|R|$ is bounded by $\mathcal{O}(|X| \cdot |Y|)$).

Consider the procedure $\textsc{MathP1}_{X,Y}(\Lambda_X, \Lambda_Y, R)$ specified above, which takes as input two potentially empty sets of distributions $\Lambda_X \in \mathbb{PD}(X)$ and $\Lambda_Y \in \mathbb{PD}(Y)$ and a relation $R \subseteq X \times Y$. Given the specification of $\textsc{MathP2}_{X,Y}$ the procedure $\textsc{MathP1}_{X,Y}(\Lambda_X, \Lambda_Y, R)$ returns **true** if and only if $\Lambda_X = \emptyset \Rightarrow \Lambda_Y = \emptyset$ and for all $\lambda_X \in \Lambda_X$ there exists a weight function $\lambda_C \in \mathbb{D}(\Lambda_Y)$ such that $\langle \lambda_X, \Lambda_Y \circ \lambda_C \rangle \in \mathbb{D}(R)$. Clearly, in the worst-case we call $\textsc{MathP2}_{X,Y}$ for each $\lambda_X \in \Lambda_X$, which means the run-time of $\textsc{MathP1}_{X,Y}(\Lambda_X, \Lambda_Y, R)$ is bounded by $\mathcal{O}(|\Lambda_X| \cdot f_{\mathrm{P1}}(|X|, |Y|, |\Lambda_Y|))$.

We are now in a position to provide a procedure to decide strong probabilistic game-simulation, denoted $\textsc{DecideD1}(\hat{G}, G)$, shown in the pseudo-code below. We first prove the correctness of this procedure:

Proposition B.5. For two $\hat{G}, G \in \mathscr{G}$ we have $\hat{G} \sqsubseteq_{\mathscr{G}} G$ if and only if $\textsc{DecideD1}(\hat{G}, G)$ returns **true**.

Proof. Let $\hat{G} = \langle \hat{S}, \hat{I}, \hat{T}, \hat{L}^!, \hat{L}^? \rangle$ and $G = \langle S, I, T, L^!, L^? \rangle$ be two games. We write $\hat{G} \preceq^A_{\mathscr{G}} G$ iff $I \subseteq A.\hat{I}$ implies $\hat{G} \sqsubseteq^A_{\mathscr{G}} G$, that is, if with R'' the conditions (i), (ii), (iii) and (iv) of Definition 11 hold for $\hat{G}$ and G, and hence $\hat{G}$ is a strong probabilistic game-simulation of G with A iff $I \subseteq A.\hat{I}$.

We first show that, by construction, whenever the algorithm returns **true** we have that $\hat{G} \sqsubseteq_{\mathscr{G}} G$. To show this, it is sufficient to show that whenever the algorithm reaches line 19 it must be that $\hat{G} \preceq^R_{\mathscr{G}} G$.

Procedure 1 $\textsc{MatchP1}_{X,Y}(\Lambda_X, \Lambda_Y, R)$

```
1: if (Λ_X = ∅) then
2:     return (Λ_Y = ∅)
3: else
4:     for all (λ_X ∈ Λ_X) do
5:         if (¬MatchP2_{X,Y}(λ_X, Λ_X, R)) then
6:             return false
7:         end if
8:     end for
9:     return true
10: end if
```

Conditions (i) and (ii) trivially hold at line 19 due the initialisation in line 3. For every $\langle \hat{s}, s \rangle \in R$, conditions (iii) and (iv) must also hold when line 19 is reached, because for $R = R'$ to be true, for every $\Lambda \in T(s)$ there must be some $\hat{\Lambda} \in \hat{T}(\hat{s})$ such that $\textsc{MatchP1}_{S,\hat{S}}(\Lambda, \hat{\Lambda}, R^{-1})$ (line 10) and there must be some $\hat{\Lambda} \in \hat{T}(\hat{s})$ such that $\textsc{MatchP1}_{\hat{S},S}(\hat{\Lambda}, \Lambda, R)$ (line 11). When considering the definition of MatchP1, this spells out precisely conditions (iii) and (iv).

Therefore, it must be that at line 19 we have $\hat{G} \preceq^{R}_{\mathscr{G}} G$ and, due to the expression of line 19, we have that $\hat{G} \sqsubseteq_{\mathscr{G}} G$ if the algorithm returns **true**.

Remaining to show is that the algorithm always returns **true** whenever $\hat{G} \sqsubseteq_{\mathscr{G}} G$. We let $R^{\text{MAX}} \subseteq \hat{S} \times S$ be the relation such that $\langle \hat{s}, s \rangle \in R^{\text{MAX}}$ iff there exists a relation $B \subseteq \hat{S} \times S$ such that $\hat{G} \preceq^{B}_{\mathscr{G}} G$ and $\langle \hat{s}, s \rangle \in B$. We will show that $R \supseteq R^{\text{MAX}}$ is an invariant of the main loop of the algorithm. Line 3 ensures $R \supseteq R^{\text{MAX}}$ prior to the start of the loop because for any tuple $\langle \hat{s}, s \rangle \notin R$ we also have, due to conditions (i) and (ii), that $\langle \hat{s}, s \rangle \notin R^{\text{MAX}}$.

To show the inductive invariant is preserved, we need to show that, if $\langle \hat{s}, s \rangle \in R^{\text{MAX}}$, then $\langle \hat{s}, s \rangle$ is not removed from R' in line 14. Recall that if $\langle \hat{s}, s \rangle \in R^{\text{MAX}}$ then there exists some $B \subseteq \hat{S} \times S$ such that $\hat{G} \preceq^{B}_{\mathscr{G}} G$ and $\langle \hat{s}, s \rangle \in B$. In the algorithm this corresponds to

Procedure 2 $\textsc{DecideD1}(\hat{G}, G)$

```
1: ⟨Ŝ, Î, T̂, L̂!, L̂?⟩ ← Ĝ
2: ⟨S, I, T, L!, L?⟩ ← G
3: R, R' ← {⟨ŝ, s⟩ | L̂!(ŝ) ⊆ L!(s) ∧ L̂?(ŝ) ⊇ L?(s)}
4: repeat
5:     R ← R'
6:     for all ⟨ŝ, s⟩ ∈ R do
7:         for all Λ ∈ T(s) do
8:             c3, c4 ← false
9:             for all Λ̂ ∈ T̂(ŝ) do
10:                c3 ← c3 ∨ MatchP1_{S,Ŝ}(Λ, Λ̂, R^{-1})
11:                c4 ← c4 ∨ MatchP1_{Ŝ,S}(Λ̂, Λ, R)
12:            end for
13:            if (¬c3 ∨ ¬c4) then
14:                R' ← R' \ {⟨ŝ, s⟩}
15:            end if
16:        end for
17:    end for
18: until (R' = R)
19: return (I ⊆ R.Î)
```

the idea that for every $\Lambda \in T(s)$ there exists a $\hat{\Lambda} \in \hat{T}(\hat{s})$ such that $\textsc{MatchP1}_{S,\hat{S}}(\Lambda,\hat{\Lambda},B^{-1})$ holds. Similarly, there exists a $\hat{\Lambda} \in \hat{T}(\hat{s})$ such that $\textsc{MatchP1}_{\hat{S},S}(\hat{\Lambda},\Lambda,B)$ holds.

By definition of $R^{\max}$, we have that $B \subseteq R^{\max}$. Also, by the inductive invariant, we have $R^{\max} \subseteq R$; hence $B \subseteq R$, $B^{-1} \subseteq R^{-1}$, $\mathbb{D}(B) \subseteq \mathbb{D}(R)$, $\mathbb{D}(B^{-1}) \subseteq \mathbb{D}(R^{-1})$ and

$$\textsc{MatchP1}_{S,\hat{S}}(\Lambda,\hat{\Lambda},B^{-1}) \Rightarrow \textsc{MatchP1}_{S,\hat{S}}(\Lambda,\hat{\Lambda},R^{-1})$$
$$\textsc{MatchP1}_{\hat{S},S}(\hat{\Lambda},\Lambda,B) \Rightarrow \textsc{MatchP1}_{\hat{S},S}(\hat{\Lambda},\Lambda,R)$$

Therefore we know that when considering $\langle \hat{s}, s\rangle$, we never remove $\langle \hat{s}, s\rangle$ from R' in line 14 because $c3$ and $c4$ must be true for all $\Lambda \in T(s)$. Hence the invariant is preserved.

Now, suppose there exists some $A \subseteq \hat{S} \times S$ such that $\hat{G} \sqsubseteq^A_{\mathscr{G}} G$, we will show the algorithm returns **true**. Considering that R and R' become strictly smaller with each loop iteration, the procedure must eventually reach line 19. By definition of $R^{\textsc{max}}$, as $\hat{G} \sqsubseteq^A_{\mathscr{G}} G$ (and hence $\hat{G} \preceq^A_{\mathscr{G}} G$), we must have that $A \subseteq R^{\textsc{max}}$. By the inductive invariant of the loop, we have that when we reach line 19 $R^{\textsc{max}} \subseteq R$, and hence $A \subseteq R$. As $I \subseteq A.\hat{I}$ and $A \subseteq R$, we must have that $I \subseteq R.\hat{I}$. Hence, the algorithm returns **true**. □

We are now in a position to prove Proposition 4.7, which we now recall:
Proposition 4.7. Deciding $\hat{G} \sqsubseteq_{\mathscr{G}} G$ is in P.

Proof. By Proposition B.5 it is sufficient to show $\textsc{DecideD1}(\hat{G}, G)$ runs in polynomial time with respect to the size of the games.

R and R' become strictly smaller with each loop iteration, hence the procedure must reach line 19 in $\mathcal{O}(n_{\hat{G}} \cdot n_G)$ iterations. The number of times lines 10 and 11 are executed in each loop iteration is bounded by $\mathcal{O}(m_{\hat{G}} \cdot m_G)$. Finally, the execution of lines 10 and 11 is bounded by $\mathcal{O}(m_G \cdot f_{\text{P1}}(n_G, n_{\hat{G}}, m_{\hat{G}}) + m_{\hat{G}} \cdot f_{\text{P1}}(n_{\hat{G}}, n_G, m_G))$. □

B.7 Proof of KMTS reduction

We will prove that deciding *qualitative* thorough refinement of games is at least as hard as deciding thorough refinement of Kripke modal transition systems (KMTSs) [21].

We call an MDP $M \in \mathscr{M}$ *qualitative* iff every distribution that occurs in this MDP is a point distribution. A MDP that is not qualitative is called quantitative. We let $\mathcal{I}^{\text{ql}} : \mathscr{G} \to \mathbb{P}(\mathscr{M})$ be a function yielding all qualitative implementations of a game, e.g. for every $\hat{G} \in \mathscr{G}$ we have $\mathcal{I}^{\text{ql}}(\hat{G}) = \{M \in \mathscr{M} \mid \hat{G} \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}(M),\ M \text{ is qualitative}\}$. The *qualitative* thorough refinement order $\sqsubseteq^{\text{ql}}_{\mathscr{G}} \subseteq \mathscr{G} \times \mathscr{G}$ is such that for two games $\hat{G}, G \in \mathscr{G}$ we have $\hat{G} \sqsubseteq^{\text{ql}}_{\mathscr{G}} G$ iff $\mathcal{I}^{\text{ql}}(G) \subseteq \mathcal{I}^{\text{ql}}(\hat{G})$. We will show that deciding whether two games are qualitative thorough refinements is at least as hard as deciding thorough refinement of KMTSs.

In order to show this we first introduce Kripke structures, the implementations of KMTSs, and then proceed to formally introduce KMTSs and their refinement preorder. Then we will provide an embedding from KMTSs to games and show qualitative thorough refinement of embedded KMTSs corresponds to thorough refinement.

Definition 17 (Kripke structure)**.** A *Kripke structure* is a tuple $\langle S, I, T, L\rangle$ where: S is a set of states; $I \subseteq S$ is a non-empty set of initial states; $T \subseteq S \times S$ is a transition relation, and $L \in S \to \mathbb{P}(\text{AP})$ is a labelling function.

We denote with $\mathscr{P}$ the set of all Kripke structures. The three-valued abstractions of Kripke structures are Kripke modal transition systems (KMTSs), which have two transition relations:

Definition 18 (Kripke modal transition system)**.** A *Kripke modal transition system* (KMTS) is a tuple $K = \langle S, I, T^!, T^?, L^!, L^?\rangle$: S is a set of states; $I \subseteq S$ is a non-empty set of initial states; $T^!, T^? \subseteq S \times S$ are transition relations such that $T^! \subseteq \hat{T}^?$, and $L^!, L^? \in S \to \mathbb{P}(\text{AP})$ are labelling functions such that, for every $s \in S$, we have that $L^!(s) \subseteq L^?(s)$. The size $|K|$ of the KMTS is the sum $|S| + |T^! \cup T^?|$. We denote with $\mathscr{K}$ the set of all KMTSs.

We now recall the refinement preorder of KMTSs:

Definition 19. Let $\hat{K} = \langle \hat{S}, \hat{I}, \hat{T}^!, \hat{T}^?, L^!, L^?\rangle$ and $K = \langle S, I, T^!, T^?, \hat{L}^!, \hat{L}^?\rangle$ be KMTSs. We say $\hat{K}$ is a simulation of K via relation $R \subseteq \hat{S} \times S$, denoted $\hat{K} \sqsubseteq^R_{\mathscr{K}} K$, iff $I \subseteq R.\hat{I}$ and, whenever $\langle \hat{s}, s\rangle \in R$, the following conditions hold:

(i) $\hat{L}^!(\hat{s}) \subseteq L^!(s)$

(ii) $\hat{L}^?(\hat{s}) \supseteq L^?(s)$

(iii) $\forall \langle \hat{s}, \hat{s}'\rangle \in \hat{T}^!\ \exists \langle s, s'\rangle \in T^! : \langle \hat{s}', s'\rangle \in R.$

(iv) $\forall \langle s, s'\rangle \in T^?\ \exists \langle \hat{s}, \hat{s}'\rangle \in \hat{T}^? : \langle \hat{s}', s'\rangle \in R.$

We let $\sqsubseteq_{\mathscr{K}} \subseteq \mathscr{K} \times \mathscr{K}$ be the relation such that $\hat{K} \sqsubseteq_{\mathscr{K}} K$ iff there is a relation $R \subseteq \hat{S} \times S$ with $\hat{K} \sqsubseteq^R_{\mathscr{K}} K$.

Finally, we define the embedding of Kripke structures as KMTSs:

Definition 20 (Kripke to KMTS)**.** Let $e^{\mathscr{K}} \in \mathscr{P} \to \mathscr{K}$ be the embedding function which, for every Kripke structure $\langle S, I, T, L\rangle$, yields the KMTS $\langle S, I, T, T, L, L\rangle$.

Clearly any KMTS such that $L^! = L^?$ and $T^! = T^?$ is an embedding of some Kripke structure. Akin to Definition 6, the embedding function in combination with the preorder give rise to an implementation function $\mathcal{I} \in \mathscr{K} \to \mathbb{P}(\mathscr{P})$ yielding all Kripke structures that refine a given KMTS. In addition, akin to Definition 7, we obtain a thorough refinement preorder $\sqsubseteq^{\text{th}}_{\mathscr{K}}$.

Now we have introduced KMTSs and the thorough refinement relation over KMTSs, we will now provide an embedding of KMTSs into games and show that deciding qualitative thorough refinement $\sqsubseteq^{\text{ql}}_{\mathscr{G}}$ over such embeddings corresponds to deciding $\sqsubseteq^{\text{th}}_{\mathscr{K}}$.

Definition 21 (KMTS to games)**.** Let $e^{\mathscr{G}}_{\mathscr{K}} \in \mathscr{K} \to \mathscr{G}$ be the embedding function which, for every KMTS $\langle S, I, T^!, T^?, L^!, L^? \rangle$, yields the game $\langle S, I, \hat{T}, L^!, L^? \rangle$ where for every $s \in S$:

$$\hat{T}(s) = \{\{\mu_{s'} \mid s' \in T^!.s\}\} \cup \{\{\mu_{s'} \mid s' \in T^?.s\}\} \,. \tag{7}$$

This embedding function is polynomial in both space and time. In order to show this embedding allows us to decide thorough refinement of KMTSs through games, we first proof some smaller lemmas:

Lemma B.6. The embedding $e^{\mathscr{G}}_{\mathscr{K}}$ preserves implementations. That is, for any Kripke structure $P \in \mathscr{P}$ there exists an MDP $M \in \mathscr{M}$ such that $e^{\mathscr{G}}(M) = e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P))$.

Proof. Let $P = \langle S, I, T, L \rangle$ and let $e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P)) = \langle S, I, \hat{T}, L, L \rangle$ as obtained by applying Def. 20 and Def. 21. We already have that the labelling functions are two-valued by definition. Therefore, by Def. 10 a suitable MDP exists if and only if for every $s \in S$ we have $|\hat{T}(s)| = 1$. By Def. 20 we know that in $e^{\mathscr{K}}(P)$ the may and must transition relations coincide. Therefore the transition function, as defined in Eq. 7 of Def. 21, always yields a singleton set by definition. However, note that the element of this singleton set may contain a set with multiple point distributions. □

Lemma B.7. The embedding $e^{\mathscr{G}}_{\mathscr{K}}$ preserves the refinement order. That is, for arbitrary KMTSs $\hat{K}, K \in \mathscr{K}$ we have that $\hat{K} \sqsubseteq_{\mathscr{K}} K$ iff $e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}_{\mathscr{K}}(K)$.

Proof. Let $\hat{K} = \langle \hat{S}, \hat{I}, \hat{T}^!, \hat{T}^?, \hat{L}^!, \hat{L}^? \rangle$ and $K \in \langle S, I, T^!, T^?, L^!, L^? \rangle$ be arbitrary KMTSs and let $\langle \hat{S}, \hat{I}, \hat{T}, \hat{L}^!, \hat{L}^? \rangle$ and $\langle S, I, T, L^!, L^? \rangle$ be the corresponding embedded KMTSs $e^{\mathscr{G}}_{\mathscr{K}}(\hat{K})$ and $e^{\mathscr{G}}_{\mathscr{K}}(K)$, respectively.

It is sufficient to show that for any $R \subseteq \hat{S} \times S$ we have that $\hat{K} \sqsubseteq^R_{\mathscr{K}} K$ iff $e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}) \sqsubseteq^R_{\mathscr{G}} e^{\mathscr{G}}_{\mathscr{K}}(K)$. It is immediate to see that the conditions regarding initial states and the labelling for strong probabilistic game-simulation of $e^{\mathscr{G}}_{\mathscr{K}}(\hat{K})$ and $e^{\mathscr{G}}_{\mathscr{K}}(K)$ (see Def. 11 and Def. 21) are identical to the corresponding conditions on KMTS refinement of $\hat{K}$ and K (see Def.

19). Therefore, it is sufficient to show that conditions (iii) and (iv) of Def. 19 hold iff conditions (iii) and (iv) of Def. 11 hold.

Consider any $\langle \hat{s}, s \rangle \in R$ and let $\Lambda^! = \{\mu_{s'} \mid s' \in T^!.s\}$ and $\Lambda^? = \{\mu_{s'} \mid s' \in T^?.s\}$. By definition $T(s) = \{\Lambda^!\} \cup \{\Lambda^?\}$ and $\Lambda^! \subseteq \Lambda^?$. Similarly, $\hat{\Lambda}^! = \{\mu_{\hat{s}'} \mid \hat{s}' \in \hat{T}^!.\hat{s}\}$ and $\hat{\Lambda}^? = \{\mu_{\hat{s}'} \mid \hat{s}' \in \hat{T}^?.\hat{s}\}$. We have $\hat{T}(\hat{s}) = \{\hat{\Lambda}^!\} \cup \{\hat{\Lambda}^?\}$ and $\hat{\Lambda}^! \subseteq \hat{\Lambda}^?$.

In this knowlegde, the quantifiers in condition (iii) and condition (iv) of Def. 11 can be simplified, hence, we rephrase conditions (iii) and (iv) of Def. 11 as follows:

(G.i) $\hat{\Lambda}^! = \emptyset \Leftarrow \Lambda^! = \emptyset$

(G.ii) $\hat{\Lambda}^? = \emptyset \Rightarrow \Lambda^? = \emptyset$

(G.iii) For every $\mu_{\hat{s}'} \in \hat{\Lambda}^!$ there exists a $\mu_{s'} \in \Lambda^!$ such that $\langle \mu_{\hat{s}'}, \mu_{s'} \rangle \in \mathbb{D}(R)$

(G.iv) For every $\mu_{s'} \in \Lambda^?$ there exists a $\mu_{\hat{s}'} \in \hat{\Lambda}^?$ such that $\langle \mu_{\hat{s}'}, \mu_{s'} \rangle \in \mathbb{D}(R)$

In fact, (G.i) is implied by (G.iii) and (G.ii) is implied by (G.iv). Moreover, by the definition of $\Lambda^!, \Lambda^?, \hat{\Lambda}^!$ and $\hat{\Lambda}^?$ and in the knowledge that $\langle x, y \rangle \in R$ iff $\langle \mu_x, \mu_y \rangle \in \mathbb{D}(R)$, condition (G.iii) coincides with condition (iii) of Def. 19 and condition (G.iv) coincides with condition (iv) of Def. 19. □

Note that up to now we have not needed to restrict ourselves to qualitative MDPs. However, to prove the second part of the following lemma we do need this requirement:

Lemma B.8. For every *qualitative* MDP $M \in \mathscr{M}$ such that $e^{\mathscr{G}}(M) \in \mathcal{I}^{\mathrm{ql}}(e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}))$ for some KMTS $\hat{K} \in \mathscr{K}$, there exists a Kripke structure $P \in \mathscr{P}$ such that $P \in \mathcal{I}(\hat{K})$ and $e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P)) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}(M)$.

Proof. Let $\langle S, I, T_{\mathscr{M}}, L \rangle$ be the qualitative MDP M and let $\langle S, I, T_{\mathscr{G}}, L, L \rangle$ be its embedding $e^{\mathscr{G}}(M)$. Note that for every $s \in S$ we have $T_{\mathscr{G}}(s) = \{T_{\mathscr{M}}(s)\}$. Let $\hat{K} = \langle \hat{S}, \hat{I}, \hat{T}^!, \hat{T}^?, \hat{L}^!, \hat{L}^? \rangle$ and let $\langle \hat{S}, \hat{I}, \hat{T}, \hat{L}^!, \hat{L}^? \rangle$ be its game embedding $e^{\mathscr{G}}_{\mathscr{K}}(K)$. We know that for any $\hat{s} \in \hat{S}$ we have that $\hat{T}(\hat{s}) = \{\hat{\Lambda}^!\} \cup \{\hat{\Lambda}^?\}$ where $\hat{\Lambda}^! = \{\mu_{s'} \mid \hat{s}' \in \hat{T}^!.\hat{s}\}$ and $\hat{\Lambda}^? = \{\mu_{\hat{s}'} \mid \hat{s}' \in \hat{T}^?.\hat{s}\}$.

As $e^{\mathscr{G}}(M) \in \mathcal{I}(e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}))$ there must be some $R \subseteq \hat{S} \times S$ such that $e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}) \sqsubseteq^R_{\mathscr{G}} e^{\mathscr{G}}(M)$.

From M we construct a Kripke structure $P = \langle S, I, T_{\mathscr{P}}, L \rangle$ where $\langle s, s' \rangle \in T_{\mathscr{P}}$ iff for some $\lambda \in T_{\mathscr{M}}(s)$ we have that $s' \in \textsc{supp}(\lambda)$. It is now sufficient to show that $\hat{K} \sqsubseteq^R_{\mathscr{K}} e^{\mathscr{K}}(P)$ (and hence $P \in \mathcal{I}(\hat{K})$) and $e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P)) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}(M)$. We first show $\hat{K} \sqsubseteq^R_{\mathscr{K}} e^{\mathscr{K}}(P)$ without assuming that M was qualitative.

Consider that $e^{\mathscr{K}}(P) = \langle S, I, T_{\mathscr{P}}, T_{\mathscr{P}}, L, L \rangle$. We already have that $I \subseteq R.\hat{I}$ and that for every $\langle \hat{s}, s \rangle \in R$ we have $\hat{L}^!(\hat{s}) \subseteq L(s)$ and $\hat{L}^?(\hat{s}) \supseteq L(s)$ from the fact that

$e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}) \sqsubseteq^R_{\mathscr{G}} e^{\mathscr{G}}(M)$. Remaining to show is conditions (iii) and (iv) of Def. 19. To show condition (iii) of Def. 19 we must show that whenever $\langle \hat{s}, s\rangle \in R$ and $\langle \hat{s}, \hat{s}'\rangle \in \hat{T}^!$ there exists $\langle s, s'\rangle \in T_{\mathscr{P}}$ such that $\langle \hat{s}', s'\rangle \in R$. Therefore, let $\langle \hat{s}, s\rangle \in R$ and $\langle \hat{s}, \hat{s}'\rangle \in \hat{T}^!$ be arbitrary tuples. We will show that there exists $\langle s, s'\rangle \in T_{\mathscr{P}}$ such that $\langle \hat{s}', s'\rangle \in R$.

Due to the fact that $\hat{T}(\hat{s}) = \{\hat{\Lambda}^!\} \cup \{\hat{\Lambda}^?\}$ where $\hat{\Lambda}^! = \{\mu_{s'} \mid \hat{s}' \in \hat{T}^!.\hat{s}\}$ and $\hat{\Lambda}^! \subseteq \hat{\Lambda}^?$ and due to the fact that $T_{\mathscr{G}}(s) = \{T_{\mathscr{M}}(s)\}$, the condition (iii) of Def. 11 of the simulation $e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}) \sqsubseteq^R_{\mathscr{G}} e^{\mathscr{G}}(M)$ only requires that $T_{\mathscr{M}}(s)$ is such that for every $\mu_{\hat{s}'} \in \hat{\Lambda}^!$ there exists a $\lambda \in T_{\mathscr{M}}(s)$ such that $\langle \mu_{\hat{s}'}, \lambda\rangle \in \mathbb{D}(R)$. Note that λ is not necessarily a point distribution, however, from the definition of weight functions we trivially have that $\langle \hat{s}', s'\rangle \in R$ for any $s' \in \textsc{supp}(\lambda)$. As the support of any distribution is always non-empty, and $\langle s, s'\rangle \in T_{\mathscr{P}}$ for any $s' \in \textsc{supp}(\lambda)$, we have proven there exists $\langle s, s'\rangle \in T_{\mathscr{P}}$ such that $\langle \hat{s}', s'\rangle \in R$. This means that condition (iii) of Def. 19 holds.

To show condition (iv) of Def. 19 we must show that whenever $\langle \hat{s}, s\rangle \in R$ and $\langle s, s'\rangle \in T_{\mathscr{P}}$ there exists $\langle \hat{s}, \hat{s}'\rangle \in \hat{T}^?$ such that $\langle \hat{s}', s'\rangle \in R$. Therefore, let $\langle \hat{s}, s\rangle \in R$ and $\langle s, s'\rangle \in T_{\mathscr{P}}$ be arbitrary tuples. We will show that there exists $\langle \hat{s}, \hat{s}'\rangle \in \hat{T}^?$ such that $\langle \hat{s}', s'\rangle \in R$.

Due to the fact that $\hat{T}(\hat{s}) = \{\hat{\Lambda}^!\} \cup \{\hat{\Lambda}^?\}$, where $\hat{\Lambda}^? = \{\mu_{s'} \mid \hat{s}' \in \hat{T}^?.\hat{s}\}$ and $\hat{\Lambda}^! \subseteq \hat{\Lambda}^?$, and due to the fact that $T_{\mathscr{G}}(s) = \{T_{\mathscr{M}}(s)\}$, the condition (iv) of Def. 11 of the simulation $e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}) \sqsubseteq^R_{\mathscr{G}} e^{\mathscr{G}}(M)$ requires that $\hat{\Lambda}^?$ is such that for every $\lambda \in T_{\mathscr{M}}(s)$ there exists a $\mu_{\hat{s}'} \in \hat{\Lambda}^?$ such that $\langle \mu_{\hat{s}'}, \lambda\rangle \in \mathbb{D}(R)$. Because $\langle s, s'\rangle \in T_{\mathscr{P}}$ there must exist a distribution $\lambda' \in T_{\mathscr{M}}(s)$ such that $s' \in \textsc{supp}(\lambda')$. Clearly, there exists $\mu_{\hat{s}'} \in \hat{\Lambda}^?$ such that $\langle \mu_{\hat{s}'}, \lambda'\rangle \in \mathbb{D}(R)$ (and hence $\langle \hat{s}', s'\rangle \in R$). By the definition of $\hat{\Lambda}^?$, there must exist $\langle \hat{s}, \hat{s}'\rangle \in \hat{T}^?$ such that $\langle \hat{s}', s'\rangle \in R$. This concludes that condition (iv) of Def. 19 holds and therefore that $\hat{K} \sqsubseteq^R_{\mathscr{K}} e^{\mathscr{K}}(P)$.

Remaining to show is that $e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P)) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}(M)$. Let $\langle S, I, T'_{\mathscr{G}}, L, L\rangle$ be the embedding $e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P))$. Note that $e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P))$ and $e^{\mathscr{G}}(M)$ already agree on S, I and L. Moreover, as $T_{\mathscr{M}}$ consists only of point distributions, for every $s \in S$ we have that $T'_{\mathscr{G}}(s) = \{T_{\mathscr{M}}(s)\} = T_{\mathscr{G}}(s)$. This means that $e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P)) = e^{\mathscr{G}}(M)$ and as the refinement $\sqsubseteq_{\mathscr{G}}$ is a preorder, and hence reflexive, we trivially have that $e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P)) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}(M)$. □

Proposition B.9. For every $\hat{K}, K \in \mathscr{K}$ we have that $\mathcal{I}(K) \subseteq \mathcal{I}(\hat{K})$ iff $\mathcal{I}^{\mathrm{ql}}(e^{\mathscr{G}}_{\mathscr{K}}(K)) \subseteq \mathcal{I}^{\mathrm{ql}}(e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}))$.

Proof. Let $\hat{K}, K \in \mathscr{K}$ be arbitrary KMTSs. We first prove the implication from right to left, that is: if for some Kripke structure $P \in \mathscr{P}$ we have that $P \in \mathcal{I}(K)$ and

$\mathcal{I}(e^{\mathscr{G}}_{\mathscr{K}}(K)) \subseteq \mathcal{I}(e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}))$, then $P \in \mathcal{I}(\hat{K})$.

$$
\begin{aligned}
& P \in \mathcal{I}(K) \iff && \\
& K \sqsubseteq_{\mathscr{K}} e^{\mathscr{K}}(P) \iff && \text{(Def. 6)} \\
& e^{\mathscr{G}}_{\mathscr{K}}(K) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P)) \implies && \text{(Lem. B.7)} \\
& e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P)) \iff && \text{(Lem. B.6)} \\
& \hat{K} \sqsubseteq_{\mathscr{K}} e^{\mathscr{K}}(P) \iff && \text{(Lem. B.7)} \\
& P \in \mathcal{I}(\hat{K}) && \text{(Def. 6)}
\end{aligned}
$$

We now prove the implication in the other direction. That is, if for some MDP $M \in \mathscr{M}$ we have that $M \in \mathcal{I}^{\mathrm{ql}}(e^{\mathscr{G}}_{\mathscr{K}}(K))$ and $\mathcal{I}(K) \subseteq \mathcal{I}(\hat{K})$, then $M \in \mathcal{I}^{\mathrm{ql}}(e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}))$. By Lemma B.8 and B.7, in the following, let $P \in \mathscr{P}$ be such that $e^{\mathscr{G}}_{\mathscr{K}}(K) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P)) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}(M)$:

$$
\begin{aligned}
& M \in \mathcal{I}^{\mathrm{ql}}(e^{\mathscr{G}}_{\mathscr{K}}(K)) \iff && \\
& e^{\mathscr{G}}_{\mathscr{K}}(K) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P)) \iff && \text{(Lem. B.8)} \\
& K \sqsubseteq_{\mathscr{K}} e^{\mathscr{K}}(P) \implies && \text{(Lem. B.7)} \\
& \hat{K} \sqsubseteq_{\mathscr{K}} e^{\mathscr{K}}(P) \iff && \text{(Assumption)} \\
& e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}_{\mathscr{K}}(e^{\mathscr{K}}(P)) \implies && \text{(Lem. B.7)} \\
& e^{\mathscr{G}}_{\mathscr{K}}(\hat{K}) \sqsubseteq_{\mathscr{G}} e^{\mathscr{G}}(M) \iff && \text{(Lem. B.8)} \\
& M \in \mathcal{I}^{\mathrm{ql}}(e^{\mathscr{G}}_{\mathscr{K}}(\hat{K})) && \text{(Def. 6)}
\end{aligned}
$$

□

B.8 Proof of Proposition 4.9

Proposition 4.9. PCTL satisfiability over MDPs can be reduced to deciding the thorough satisfaction of games.

Proof. Consider the game $\hat{G} = \langle \{\hat{s}_0\}, \{\hat{s}_0\}, \hat{T}, \hat{L}^!, \hat{L}^? \rangle$ depicted in Figure 3 (with $\hat{T}(\hat{s}_0) = \{\{\mu_{\hat{s}_0}\}, \emptyset\}$, $\hat{L}^!(\hat{s}_0) = \emptyset$ and $\hat{L}^?(\hat{s}_0) = \textsc{ap}$). We will show that for any PCTL formula $\phi \in \Phi_{\mathrm{PCTL}}$ we have that $\hat{G} \not\models^{\mathrm{th}}_{\mathscr{G}} \neg\phi$ iff ϕ is satisfiable.

We first prove that $\mathcal{I}(\hat{G}) = \mathscr{M}$, e.g. that $\hat{G}$ is implemented by *any* MDP. To show this, let $M \in \mathscr{M}$ be an arbitrary MDP, let $e^{\mathscr{G}}(M) = \langle S, I, T, L, L \rangle$ and let $R = \{\hat{s}_0\} \times S$. We will show that $\hat{G} \sqsubseteq^R_{\mathscr{G}} e^{\mathscr{G}}(M)$. Clearly, $I \subseteq S = R.\hat{s}_0$. Moreover, for any $\langle \hat{s}_0, s \rangle \in R$, as $\hat{L}^!(\hat{s}_0) = \emptyset$ and $\hat{L}^?(\hat{s}_0) = \textsc{ap}$ condition (i) and (ii) of Def. 11 are trivially satisfied.

For any $\langle \hat{s}_0, s \rangle \in R$, consider an arbitrary concrete player 1 choice $\Lambda \in T(s)$. We wish to show that conditions (iii) and (iv) of Def. 11 are satisfied for Λ. We have two cases: $\Lambda = \emptyset$ or $\Lambda \neq \emptyset$. If $\Lambda = \emptyset$, then taking $\emptyset \in \hat{T}(\hat{s}_0)$ satisfies both condition (iii) and (iv) tivially. If $\Lambda \neq \emptyset$, we will now show that condition (iii) and (iv) are satisfied due to the presence of $\{\mu_{\hat{s}_0}\} \in \hat{T}(\hat{s}_0)$.

For condition (iii) let $\mu_{\{\mu_{\hat{s}_0}\}} \in \mathbb{D}(\{\mu_{\hat{s}_0}\})$. Clearly, $\{\mu_{\hat{s}_0}\} \circ \mu_{\{\mu_{\hat{s}_0}\}} = \mu_{\hat{s}_0}$ and for any $\lambda \in \Lambda$ we trivially have $\langle \mu_{\hat{s}_0}, \lambda \rangle \in \mathbb{D}(R)$.

For condition (iv) let $\lambda_C \in \mathbb{D}(\Lambda)$ be an arbitrary weight function over Λ. The resulting distribution $\Lambda \circ \lambda_C \in \mathbb{D}(S)$ (or in fact any distribution over S) trivially simulates $\mu_{\hat{s}_0}$, hence $\langle \mu_{\hat{s}_0}, \Lambda \circ \lambda_C \rangle \in \mathbb{D}(R)$.

We have shown that all conditions of Def. 11 hold, and therefore $\hat{G} \sqsubseteq^R_{\mathscr{G}} e^{\mathscr{G}}(M)$ (e.g. $M \in \mathcal{I}(\hat{G})$). Moreover, as we made no assumptions on M, it must be that $\mathcal{I}(\hat{G}) = \mathscr{M}$.

Now, suppose $\phi \in \Phi_{\text{PCTL}}$ is *satisfiable* over $\mathscr{M}$, i.e. there is some $M \in \mathscr{M}$ such that $M \models_{\mathscr{M}} \phi$. Due to the two-valued semantics of MDPs, $M \not\models_{\mathscr{M}} \neg\phi$ and hence, by Definition 8, $\hat{G} \not\models^{\text{th}}_{\mathscr{G}} \neg\phi$.

To show the other direction, suppose $\hat{G} \not\models^{\text{th}}_{\mathscr{G}} \neg\phi$. By Definition 8 there exists some $M \in \mathscr{M}$ such that $M \not\models_{\mathscr{M}} \neg\phi$. Clearly, this means ϕ is satisfied by some initial state of $M_\perp$. We construct M' from M by restricting the initial states to the state satisfying ϕ. Clearly, $M' \models_{\mathscr{M}} \phi$ and hence ϕ is satisfiable. □