

Probabilistic Alternating-Time Temporal Logic of Incomplete Information and Synchronous Perfect Recall*

Xiaowei Huang¹, Kaile Su^{2,3}, Chenyi Zhang⁴

¹ School of Computer Science and Engineering, University of New South Wales, Australia

² Institute for Integrated and Intelligent Systems, Griffith University, Brisbane, Australia

³ Key laboratory of High Confidence Software Technologies (Peking University), Ministry of Education, Beijing, China

⁴ School of Information Technology and Electrical Engineering, University of Queensland, Australia

Abstract

A probabilistic variant of ATL* logic is proposed to work with multi-player games of incomplete information and synchronous perfect recall. The semantics of the logic is settled over probabilistic interpreted system and partially observed probabilistic concurrent game structure. While unexpectedly, the model checking problem is in general undecidable even for single-group fragment, we find a fragment whose complexity is in 2-EXPTIME. The usefulness of this fragment is shown over a land search scenario.

Introduction

Alternating-time Temporal Logic (ATL) (Alur, Henzinger, and Kupferman 2002) offers the ability of reasoning about strategies over games involving multiple players. The formula $\langle\langle A \rangle\rangle\phi$ expresses the property that a coalition A (or a set A of players) has a strategy to enforce the fact ϕ .

The logic of Alur et al. assumes that players always have complete information of a game state, which is not always true in real applications. For example, it is impossible for a player to observe the local state of its opponent in some games. Notable examples include various card games. The semantics of ATL asserts that players make decision based on history. This is sometimes known as *perfect recall* (Fagin et al. 1995), as players remember all their history observations. A player assumed to have perfect recall can make maximal use of its reasoning capabilities. It serves as a sufficient condition for the adversary in designing and verifying critical systems.

Probabilistic information provides quantitative measures over the games. At each state of a game, once every player chooses a local action, a distribution exists about the next state. Moreover, it allows to specify that a player can achieve a goal with a certain minimal or maximal probability (Chen and Lu 2007).

*The authors are ordered alphabetically. This work is supported by Australian Research Council (ARC) Discovery Project DP120102489, ARC Linkage Project LP0989643, National Basic Research Program (973 Program) 2010CB328103, and ARC Future Fellowship FT0991785.

Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In this paper, we work with multi-player synchronous games in which players have perfect recall memory over observations. To the best of our knowledge, so far there are no existing works in the literature that address probabilistic variants of ATL with incomplete information and synchronous perfect recall.

The paper makes the following contributions. (1) We propose a new semantics of the PATL* logic based on a *probabilistic action interpreted system*, which generalizes the interpreted system (Fagin et al. 1995) by adding probabilistic information and explicit local actions taken by players. This probabilistic interpreted system can be obtained from a partially-observed probabilistic concurrent game structure, in which players have history-dependent strategies (as per *perfect recall* semantics). (2) We study the model checking complexity of PATL*, which is in general undecidable following the result of its nonprobabilistic variant ATL* (Bulling, Dix, and Jamroga 2010; Dima and Tiplea 2011). Although model checking ATL* is EXPTIME-complete for its single-player fragment with incomplete information (Alur, Henzinger, and Kupferman 2002), we find that after probability is introduced into the system, model checking single-player PATL* becomes undecidable. Nevertheless, we prove that a small fragment of PATL*, the $PATL_{\setminus U}^*$ logic in which the *until* operator is dropped, is decidable in 2-EXPTIME for its single-player (or equivalently, single-group) fragment.

PATL* Logic

Suppose that we are working with a probabilistic system with a finite set $Agt = \{1, \dots, n\}$ of players. Let $Prop$ be a set of propositions. To specify the properties of a probabilistic system, we present a logic PATL* that combines the temporal operators, the strategy operator and probability measures. Its syntax is given by

$$\phi ::= p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid X\phi \mid \phi_1 U \phi_2 \mid \langle\langle A \rangle\rangle^{\bowtie d} \phi$$

where $p \in Prop$, $A \subseteq Agt$, d is a rational constant in $[0, 1]$, and \bowtie is a relation symbol in the set $\{\leq, <, >, \geq\}$. Intuitively, formula $\langle\langle A \rangle\rangle^{\bowtie d} \phi$ expresses that players in A can collaboratively enforce the fact ϕ with a probability in relation \bowtie with constant d , $X\phi$ expresses that ϕ holds at the next time, $\phi_1 U \phi_2$ expresses that ϕ_1 holds until ϕ_2 becomes

true. Other operators can be obtained in the usual way, e.g., $F\phi \equiv \text{True} \cup \phi$, $G\phi \equiv \neg F\neg\phi$, etc.

Interpreted Systems Semantics

We enrich a probabilistic interpreted system (Halpern 2003; Huang, Luo, and van der Meyden 2011) by a set of players Agt and actions performed by the players, and we call the resulting system a probabilistic action interpreted system (PAIS). At all times in a PAIS, each player is assumed to be in some *local* state that records all the information that the player can access at that time. The environment e records “everything else that is relevant”. Let S be the set of environment states and let L_i be the set of local states of player $i \in \text{Agt}$. A *global* state s of a multi-player system is an $(n+1)$ -tuple (s_e, s_1, \dots, s_n) such that $s_e \in S$ and $s_i \in L_i$ for all $i \in \text{Agt}$. At a global state, each player independently takes some local action, which represents the decision it makes. The environment responds with a *deterministic* action to update its state. Let Act_e be the set of environment actions and Act_i be the set of local actions of player $i \in \text{Agt}$. A *global* action of a multi-player system in some global state is a $(n+1)$ -tuple $a = (a_e, a_1, \dots, a_n)$ such that $a_e \in \text{Act}_e$ and $a_i \in \text{Act}_i$ for all $i \in \text{Agt}$.

Time is represented discretely by using natural numbers. A *run* is a function $r : \mathbb{N} \rightarrow S \times L_1 \times \dots \times L_n \times \text{Act}_e \times \text{Act}_1 \times \dots \times \text{Act}_n$ from time to global states and actions. A pair (r, m) consisting of a run r and time m is called a *point*, which may also be written as $r(m)$. If $r(m) = (s_e, s_1, \dots, s_n, a_e, a_1, \dots, a_n)$ then we define $s_e(r, m) = s_e$, $a_e(r, m) = a_e$ and $s_i(r, m) = s_i$ and $a_i(r, m) = a_i$ for $i \in \text{Agt}$. If r is a run and m a time, we write $s_e(r, 0..m)$ for the sequence $s_e(r, 0) \dots s_e(r, m)$, and $a(r, 0..m)$ for $a(r, 0) \dots a(r, m)$. Let a system \mathcal{R} be a set of runs, and we call $\mathcal{R} \times \mathbb{N}$ the *set of points* of \mathcal{R} . Relative to a system \mathcal{R} , we define the set $\mathcal{K}_i(r, m) = \{(r', m') \in \mathcal{R} \times \mathbb{N} \mid s_i(r', m') = s_i(r, m)\}$ to be the set of points that are, for player i , indistinguishable from the point (r, m) .

We introduce some preliminary notions for probabilistic systems. A *probability space* is a triple (W, F, μ) such that W is a set, called the *carrier*, $F \subseteq \mathcal{P}(W)$ is a set of *measurable* sets in $\mathcal{P}(W)$, closed under countable union and complementation, and $\mu : F \rightarrow [0, 1]$ is a *probability measure*, such that $\mu(W) = 1$ and $\mu(U \cup V) = \mu(U) + \mu(V)$ if $U \cap V = \emptyset$. As usual, we define the conditional probability $\mu(U|V) = \mu(U \cap V)/\mu(V)$ when $\mu(V) \neq 0$.

The work of Halpern and Tuttle introduces a general notion of *adversary* to handle nondeterminism in probabilistic systems (Halpern and Tuttle 1993). It is especially the case in a game of incomplete information, where the moves of the opponents are hidden. As runs of a system are not measurable unless all the nondeterministic choices are sorted out, one needs to define adversaries to “settle” the nondeterminism. We follow the common choice by fixing the strategies of players in a game. For a system \mathcal{R} of runs, we define a *cell* $c = (\mathcal{R}_c, F_c, \mu_c)$ to be a probability space such that $\mathcal{R}_c \subseteq \mathcal{R}$ and $F_c \subseteq \mathcal{P}(\mathcal{R}_c)$. (In the game structure semantics presented in the following section, \mathcal{R}_c will be made concrete as the set of runs *compatible* with the strategies that define c .) A point (r, m) is in c if $r \in \mathcal{R}_c$. The set of indistinguishable points for player i in (r, m) assuming

c is $\mathcal{K}_i^c(r, m) = \mathcal{K}_i(r, m) \cap \{(r, m) \mid r \in \mathcal{R}_c, m \in \mathbb{N}\}$. The probability information over c is $\mathbf{P}^c = \{PR_i^c \mid i \in \text{Agt}\}$, where PR_i^c is a function mapping each point (r, m) in c to a probability space $PR_i^c(r, m) = (\mathcal{K}_i^c(r, m), F_i^c(r, m), \mu_{r,m,i}^c)$ such that $F_i^c(r, m) \subseteq \mathcal{P}(\mathcal{K}_i^c(r, m))$. Intuitively, at each point, each player has a probability space in which the carrier is the set of points $\mathcal{K}_i^c(r, m)$.

Two cells c_1 and c_2 are *strategic equivalent* for player i , denoted as $c_1 \simeq_i c_2$, if for any two points $(r, m), (r', m')$ in c_1 or c_2 , $s_i(r, m) = s_i(r', m')$ implies $a_i(r, m) = a_i(r', m')$. Note that, the relation \simeq_i is an equivalence relation, i.e., it is reflexive, symmetric, and transitive. We use $[c]_C^{\simeq_i}$ to denote the equivalence class of c in C with respect to the relation \simeq_i and $[\simeq_i]_C$ to denote the set of all equivalence classes in C with respect to the relation \simeq_i .

A *probabilistic action interpreted system* (PAIS) is a tuple $(\mathcal{R}, C, \{\mathbf{P}^c\}_{c \in C}, \{\simeq_i\}_{i \in \text{Agt}}, \pi)$, where \mathcal{R} is a system of runs, C is a set of cells in \mathcal{R} such that $\mathcal{R} = \bigcup \{\mathcal{R}_c \mid c \in C\}$, $\{\mathbf{P}^c\}_{c \in C}$ is a set of probability information for all cells in C , $\{\simeq_i\}_{i \in \text{Agt}}$ is a set of strategic equivalences over cells for all players in Agt , and $\pi : \mathcal{R} \times \mathbb{N} \rightarrow \mathcal{P}(\text{Prop})$ is an interpretation.

Let $A \subseteq \text{Agt}$ be a set of players. We define $\mathcal{K}_A(r, m) = \bigcap_{i \in A} \mathcal{K}_i(r, m)$, $\mathcal{K}_A^c(r, m) = \bigcap_{i \in A} \mathcal{K}_i^c(r, m)$. Moreover, we let $\simeq_A = \bigcap_{i \in A} \simeq_i$ and $a_A(r, m) = \{a_i(r, m) \mid i \in A\}$ be the collective action of players in A at point (r, m) . Likewise, we can define $[c]_C^{\simeq_A}$ and $[\simeq_A]_C$.

The semantics of the language in a PAIS $\mathcal{I} = (\mathcal{R}, C, \{\mathbf{P}^c\}_{c \in C}, \{\simeq_i\}_{i \in \text{Agt}}, \pi)$ is given by interpreting formulas ϕ at points (r, m) of \mathcal{I} , using a satisfaction relation $\mathcal{I}, (r, m) \models \phi$, which is defined inductively as follows.

- $\mathcal{I}, (r, m) \models p$ if $p \in \pi(r, m)$,
- $\mathcal{I}, (r, m) \models \neg\phi$ if not $\mathcal{I}, (r, m) \models \phi$
- $\mathcal{I}, (r, m) \models \phi \wedge \psi$ if $\mathcal{I}, (r, m) \models \phi$ and $\mathcal{I}, (r, m) \models \psi$
- $\mathcal{I}, (r, m) \models X\phi$ if $\mathcal{I}, (r, m+1) \models \phi$.
- $\mathcal{I}, (r, m) \models \phi U \psi$ if there exists a time $m' \geq m$ such that $\mathcal{I}, (r, m') \models \psi$ and $\mathcal{I}, (r, m'') \models \phi$ for all m'' with $m \leq m'' < m'$
- $\mathcal{I}, (r, m) \models \langle\langle A \rangle\rangle^{\bowtie d} \psi$ if there exists an equivalence class $[c]_C^{\simeq_A} \in [\simeq_A]_C$ such that for all $c' \in [c]_C^{\simeq_A}$ and for all $(r', m') \in \mathcal{K}_A^{c'}(r, m)$,

$$\mu_{r',m',A}^{c'}(\{(r'', m'') \mid (r'', m'') \in \mathcal{K}_A^{c'}(r', m') \wedge \mathcal{I}, (r'', m'') \models \psi\}) \bowtie d$$

where $\bowtie \in \{\leq, <, >, \geq\}$. Intuitively $[c]_C^{\simeq_A}$ represents a joint winning strategy of A such that for all joint opponent strategies, $[c]_C^{\simeq_A}$ enforces a win on every compatible state.

Naturally we have the following (Zhang and Pang 2010).

Proposition 1 $\mathcal{I}, (r, m) \models \langle\langle A \rangle\rangle^{\bowtie d} \psi$ iff $\mathcal{I}, (r, m) \models \langle\langle A \rangle\rangle^{\bowtie 1-d} \neg\psi$, where $\bowtie \in \{\leq, <, >, \geq\}$, and we write $\overline{\leq}$ for $>$, $\overline{<}$ for \geq , $\overline{\geq}$ for $<$, $\overline{>}$ for \leq .

The above definitions leave open two probabilistic measures: μ_c and $\mu_{r,m,i}^c$. We make μ_c concrete in the next section, by mapping from the concurrent game structure semantics. In the remaining part of this section, we present a way of obtaining $\mu_{r,m,i}^c$ from μ_c , depending on a player’s perfect recall memory.

Synchronous Perfect Recall A player i has *synchronous perfect recall*, denoted as *spr*, in system \mathcal{R} if there exists a set O (of observations) such that for each point (r, m) of \mathcal{R} , the local state $s_i(r, m)$ is a sequence of exactly $(m + 1)$ elements of O and m elements of Act_i . Formally, let $O_i : \mathcal{R} \times \mathbb{N} \rightarrow O$ be an observation function that maps each point to an observation of player i . Then the local state of i is defined by $s_i(r, 0) = O_i(r, 0)$, and $s_i(r, m + 1) = s_i(r, m) \cdot a(r, m) \cdot O_i(r, m + 1)$ for all $m \in \mathbb{N}$, where a is some action in Act_i .

Let $\mathcal{R}(U) = \{r \in \mathcal{R} \mid \exists m : (r, m) \in U\}$ be the set of runs in \mathcal{R} going through some point in the set $U \subseteq \mathcal{R} \times \mathbb{N}$. We define the measure $\mu_{r,m,i}^c$ by μ_c as

$$\mu_{r,m,i}^c(U) = \mu_c(\mathcal{R}(U) \mid \mathcal{R}(\mathcal{K}_i^c(r, m))). \quad (1)$$

Similarly, we let $\mu_{r,m,A}(U) = \mu_c(\mathcal{R}(U) \mid \mathcal{R}(\mathcal{K}_A^c(r, m)))$. The function $PR_i^c(r, m)$ is related to the prior on c by a type of temporal conditioning. The following theorem clarifies a relationship between $\mu_{r,m+1,i}$ and $\mu_{r,m,i}$. It can be generalized to have a relationship between $\mu_{r,m+1,A}$ and $\mu_{r,m,A}$.

Theorem 1 *Let $c \in C$ be a cell. Then we have*

$$\mu_{r,m+1,i}^c(U) = \mu_{r,m,i}^c(U_{r,m,i} \mid \mathcal{K}_i^c(r, m + 1)),$$

where $U_{r,m,i} = \{(r', m') \in \mathcal{K}_i^c(r, m) \mid \exists m'' : (r', m'') \in U\}$.

Proof: We first establish that $\mathcal{R}(U) \cap \mathcal{R}(\mathcal{K}_i^c(r, m + 1)) = \mathcal{R}(U_{r,m,i}) \cap \mathcal{R}(\mathcal{K}_i^c(r, m + 1))$.

(\subseteq) Let $r' \in \mathcal{R}(U) \cap \mathcal{R}(\mathcal{K}_i^c(r, m + 1))$, then there exists $m', m'' \in \mathbb{N}$ such that $(r', m') \in \mathcal{K}_i^c(r, m + 1)$ and $(r', m'') \in U$. By Lemma 1, we have $(r', m' - 1) \in \mathcal{K}_i^c(r, m)$. Then $(r', m' - 1) \in U_{r,m,i}$ and $r' \in \mathcal{R}(U_{r,m,i})$ follow from the definition. Therefore $r' \in \mathcal{R}(U_{r,m,i}) \cap \mathcal{R}(\mathcal{K}_i^c(r, m + 1))$.

(\supseteq) Since $\mathcal{R}(U_{r,m,i}) \subseteq \mathcal{R}(U)$ by definition, $\mathcal{R}(U_{r,m,i}) \cap \mathcal{R}(\mathcal{K}_i^c(r, m + 1)) \subseteq \mathcal{R}(U) \cap \mathcal{R}(\mathcal{K}_i^c(r, m + 1))$.

By (1), we have $\mu_{r,m+1,i}^c(U) = \mu_c(\mathcal{R}(U) \mid \mathcal{R}(\mathcal{K}_i^c(r, m + 1)))$, which then gives

$$\mu_{r,m+1,i}^c(U) = \frac{\mu_c(\mathcal{R}(U) \cap \mathcal{R}(\mathcal{K}_i^c(r, m + 1)))}{\mu_c(\mathcal{R}(\mathcal{K}_i^c(r, m + 1)))}$$

As $\mathcal{R}(U) \cap \mathcal{R}(\mathcal{K}_i^c(r, m + 1)) = \mathcal{R}(U_{r,m,i}) \cap \mathcal{R}(\mathcal{K}_i^c(r, m + 1))$,

$$\mu_{r,m+1,i}^c(U) = \frac{\mu_c(\mathcal{R}(U_{r,m,i}) \cap \mathcal{R}(\mathcal{K}_i^c(r, m + 1)))}{\mu_c(\mathcal{R}(\mathcal{K}_i^c(r, m + 1)))}$$

Then we get $\mu_{r,m+1,i}^c(U) = \mu_{r,m,i}^c(U_{r,m,i} \mid \mathcal{K}_i^c(r, m + 1))$. \square

Lemma 1 *Let $c \in C$ and $r \in \mathcal{R}$, then $(r, m) \in \mathcal{K}_i^c(r', m')$ implies $(r, m - 1) \in \mathcal{K}_i^c(r', m' - 1)$.*

Proof: By *spr* we have $m = m'$, $r \in \mathcal{R}^c$, and $s_i(r, m) = s_i(r', m')$. By definition of s_i in *spr* we have $s_i(r, m - 1) = s_i(r', m' - 1)$. Therefore $(r', m' - 1) \in \mathcal{K}_i^c(r, m - 1)$, and moreover, $(r', m' - 1) \in \mathcal{K}_i^c(r, m - 1)$ by $r' \in \mathcal{R}^c$. \square

Game Structure Semantics

Although PAIS provide a coherent semantic framework for PATL*, they are infinite structures which are impossible for model checking algorithms. In this section we propose a finite model called *partially observed*

probabilistic concurrent game structure (PO-PCGS). A finite PO-PCGS for a set Agt of players is a tuple $M = (S, Act_e, Act_1, \dots, Act_n, N_e, N_1, \dots, N_n, O_1, \dots, O_n, PI, PT, \pi)$, where S is a finite set of states, Act_e is the set of local actions of the environment, Act_i is the set of local actions of player $i \in Agt$, $N_i : S \rightarrow \mathcal{P}(Act_i)$ indicates the set of actions that are available to player i at a specific state, component $PI : S \rightarrow [0..1]$ is an initial distribution such that $\sum_{s \in S} PI(s) = 1$, $PT : S \times Act \times S \rightarrow [0, 1]$ is a probability transition matrix, such that $Act = Act_e \times Act_1 \times \dots \times Act_n$ and $\sum_{s' \in S} PT(s, a, s') = 1$ for all $s \in S$ and $a \in Act$, and for each player $i \in Agt$, we have an observation function $O_i : S \rightarrow O$. Finally, $\pi : S \rightarrow \mathcal{P}(Prop)$ is an interpretation of the atomic propositions $Prop$ at the states. Similar to the settings in PAIS, we suppose that the environment reacts *deterministically*. For consistency, we further require that for all states $s_1, s_2 \in S$ and $i \in Agt$, $O_i(s_1) = O_i(s_2)$ implies $N_i(s_1) = N_i(s_2)$.

We treat the set of states S as the states of the environment rather than as the set of global states, and player i 's local states are derived from the observation function O_i and the actions in Act_i that i performs. We write $k_i(s) = \{s' \in S \mid O_i(s') = O_i(s)\}$ for the set of states that are observationally indistinguishable to player i from state s .

Executions to Runs Let $s, s' \in S$ and $a \in Act$. A *path* ρ from a state s is a finite or infinite sequence of states and actions $s_0 a_0 s_1 a_1 \dots$ such that $s_0 = s$ and $PT(s_k, a_k, s_{k+1}) > 0$ for all k such that $k < |\rho| - 1$, where $|\rho|$ is the total number of states on ρ . Given a path ρ , we use $s(\rho, m)$ to denote its $(m + 1)$ -th state, $a(\rho, m)$ to denote its m -th action, in which $a_e(\rho, m)$ is its m -th environment action and $a_i(\rho, m)$ is its m -th local action of agent i . A *fullpath* from a state s is an infinite path from s . A path ρ is *initialized* if $PI(s(\rho, 0)) > 0$.

From each initialized fullpath ρ , one may define a run that corresponds to a run in a PAIS satisfying *spr* for all players. Recall that we interpret the states of the PO-PCGS as states of the environment, and the global actions of the PO-PCGS as actions of the players as well as the environment. Given an initialized fullpath ρ , we obtain a run ρ^{spr} by defining each point (ρ^{spr}, m) with $m \in \mathbb{N}$ as follows. The environment state at time m is $s_e(\rho^{\text{spr}}, m) = s(\rho, m)$. The environment action and local action are $a_e(\rho^{\text{spr}}, m) = a_e(\rho, m)$ and $a_i(\rho^{\text{spr}}, m) = a_i(\rho, m)$, respectively. The local state of player i at time m is $s_i(\rho^{\text{spr}}, m) = O_i(s(\rho, 0)) \cdot a_i(\rho, 1) \cdot \dots \cdot O_i(s(\rho, m))$, representing that the player remembers all its observations and past local actions, according to *spr*.

Complete Coalition Strategies to Cells A strategy σ_i of a player i is a function that maps each finite path $\rho = s_0 a_0 s_1 a_1 \dots s_n$ to an action in $N_i(s_n)$. A (finite or infinite) path ρ is compatible with σ_i if $a_k(i) = \sigma_i(s_0 a_0 \dots s_k)$ for all $k \leq |\rho|$ where $|\rho|$ is the number of transitions in ρ . Given a PO-PCGS M and an i strategy σ_i , write $Path(M, \sigma_i)$ for the set of infinite paths in M that are compatible with σ_i . A strategy σ_i is *uniform* if for all paths $\rho, \rho' \in Path(M, \sigma_i)$ and $m \in \mathbb{N}$, we have $s_i(\rho, m) = s_i(\rho', m)$ implies $a_i(\rho, m) = a_i(\rho', m)$, i.e., i 's reactions following σ_i respect its history observations.

Let A be a set of players. A coalition strategy σ_A fixes a

strategy σ_i for each player $i \in A$. We call σ_A a complete coalition strategy if $A = \text{Agt}$, or an incomplete coalition strategy if $A \subset \text{Agt}$. Given a complete coalition strategy $\sigma_{\text{Agt}} = \{\sigma_i \mid i \in \text{Agt}\}$, we define a cell c , and obtain a subset of runs $\mathcal{R}_c = \bigcap_{i \in \text{Agt}} \text{Path}(M, \sigma_i)$. Note the strategies of players in $\text{Agt} \setminus A$ are not required to be uniform, so that they are allowed to perform arbitrary behaviors.

We now define a probability space on \mathcal{R}_c , using a well-known construction (e.g., that of (Vardi 1985)). Given a finite initialized path ρ of $m + 1$ states and m actions, write $\mathcal{R}_c(\rho) = \{r \in \mathcal{R}_c \mid s_e(r, 0..m) = s_e(\rho, 0..m), a(r, 0..m-1) = a(\rho, 0..m-1)\}$ for the set of runs with prefix ρ . (One may view this as a *cone* of runs sharing the same prefix ρ .) Let F_c be the minimal algebra with basis the sets $\{\mathcal{R}_c(\rho) \mid \rho \text{ prefixes some } \rho' \in \mathcal{R}_c\}$, i.e., $F_{\mathcal{R}}$ is the set of all sets of infinite runs that can be constructed from the basis by using countable union and complement. We define the measure μ_c on the basis sets by $\mu_c(\mathcal{R}_c(\rho)) = \text{PI}(s(\rho, 0)) \times \prod_{i=0}^{m-1} p_i$ where $\text{PT}(s(\rho, i), a(\rho, i), s(\rho, i+1)) = p_i$ for all $0 \leq i \leq m-1$. There is a unique extension of μ_c that satisfies the constraints on probability measures (i.e., countable additivity and universality), and we also denote this extension by μ_c . The measurability of μ_c is guaranteed by the following theorem.

Proposition 2 *Each complete coalition strategy σ_{Agt} derives a subset of runs $\mathcal{R}_c \subseteq \mathcal{R}$, from which a probability space $(\mathcal{R}_c, F_c, \mu_c)$ can be uniquely defined.*

Incomplete Coalition Strategies to Equivalence Classes over Cells Let $\bar{A} = \text{Agt} \setminus A$ be the complement set of players of A . For each incomplete coalition strategy σ_A , there may exist more than one incomplete coalition strategy $\sigma_{\bar{A}}$. As a complete coalition strategy $\sigma_A \cup \sigma_{\bar{A}}$ maps the system \mathcal{R} into a cell, an incomplete coalition strategy σ_A may map \mathcal{R} into a set of cells, each of which corresponds with an incomplete coalition strategy of $\sigma_{\bar{A}}$ of players \bar{A} . The following theorem ascertains that these cells are strategic equivalent.

Proposition 3 *Let σ_A be an incomplete uniform strategy of A and $\sigma_{\bar{A}}^1$ and $\sigma_{\bar{A}}^2$ be two incomplete strategies of \bar{A} . Let c_1 and c_2 be the cells for complete strategy $\sigma_A \cup \sigma_{\bar{A}}^1$ and $\sigma_A \cup \sigma_{\bar{A}}^2$ respectively. Then we have $c_1 \simeq_A c_2$.*

Here we remark that, a single run $r \in \mathcal{R}$ may belong to different cells or even different equivalence classes. Also, there might exist more than one strategy of coalition A that are mapped to the same equivalence class over cells. Plainly, such strategies may disagree only on incompatible runs.

PO-PCGS to PAIS The system M gives us an interpretation π on its states, and we may lift this to an interpretation on the points (r, m) of \mathcal{R} by defining $\pi(r, m) = \pi(s_e(r, m))$. Using the construction above, we then obtain the probabilistic interpreted system $\mathcal{I}(M) = \mathcal{I}(\mathcal{R}, C, \{\mathbf{P}^c\}_{c \in C}, \{\simeq_i\}_{i \in \text{Agt}}, \pi)$. We will be interested in the problem of model checking formulas in this system. A formula ϕ is said to hold in M , written $M \models \phi$, if $\mathcal{I}(M), (r, 0) \models \phi$ for all $r \in \mathcal{R}$. The model checking problem is then to determine, given a PO-PCGS M and a formula ϕ , whether $M \models \phi$.




Figure 1: Example

An Example

Let $M = (S, \text{Act}_i, N_i, O_i, \text{PI}, \text{PT}, \pi)$ be a system of a single player i , where $S = \{s_0, s_1, s_2, s_3, s_4, s_5\}$ and $\text{Act}_i = \{h, t\}$. Player i can distinguish every two states except for s_2 and s_3 . s_4 is the only state in which proposition p holds. $\text{PI}(s_0) = 1$ and $N_i(s) = \{h, t\}$ for all states s . The transition matrix PT is shown in Figure 1, where states s_i and s_j are connected by an arrow labeled with (act, p_k) if $\text{PT}(s_i, act, s_j) = p_k \neq 0$.

In the following we discuss three strategies σ_1, σ_2 and σ_3 , by discarding incompatible (finite) runs and omitting irrelevant choices. The strategies are defined as follows. (1) $\sigma_1(s_0) = t$; (2) $\sigma_2(s_0) = h, \sigma_2(s_0hs_2) = \sigma_2(s_0hs_3) = t$; (3) $\sigma_3(s_0) = h, \sigma_3(s_0hs_2) = \sigma_3(s_0hs_3) = h$. Note that $\sigma(s_0hs_2) = \sigma(s_0hs_3)$ for all uniform strategies σ .

Initially, player a does not have a strategy to eventually reach state s_4 in a probability more than $1/2$.

$$M \models \neg \langle\langle i \rangle\rangle^{>1/2} (Fp) \quad (2)$$

The probabilities of satisfying Fp by taking the three strategies are 0, $1/3$ and $7/18$, respectively. However, the player has a strategy to enforce the formula Fp in a probability more than $1/3$ and by following that strategy at the next state, it has a strategy to reach state s_4 in a probability more than $1/2$, as expressed in the following expression.

$$M \models \langle\langle i \rangle\rangle^{\geq 1/3} X \langle\langle i \rangle\rangle^{\geq 1/2} Fp \quad (3)$$

One may find that the first $\langle\langle i \rangle\rangle$ can be enforced by either σ_2 or σ_3 , and the second is enforced by σ_3 . To show this, we show that $\mathcal{I}(M), (r, 1) \models \langle\langle i \rangle\rangle^{\geq 1/2} Xp$ for all infinite runs r with $s_e(r, 1) \in \{s_2, s_3\}$. By fixing strategy σ_3 that decides a unique cell c , we have $\mu_{r,1,i}^c(\{(r', m) \mid (r', m) \in \mathcal{K}_i^c(r, 1) \wedge \mathcal{I}, (r', m) \models Fp\}) = 7/12 \geq 1/2$, where $\mu_{r,1,i}^c$ assigns $1/4$ to the set of all runs r' such that $\mathcal{I}(M), (r', 1) \models Fp$ and $s_e(r', 1) = s_2$, and assigns $1/3$ to the set of all runs r'' such that $\mathcal{I}(M), (r'', 1) \models Fp$ and $s_e(r'', 1) = s_3$. Note that given incomplete information, $s_i(r, 1) = s_i(r', 1) = s_i(r'', 1)$ for all such runs r' and r'' .

Complexity results on Model Checking PATL*

In this section, we present several complexity results on model checking the PATL* logic. It is widely-believed that model checking ATL* with incomplete information and synchronous perfect recall is undecidable (Alur, Henzinger,

and Kupferman 2002; Bulling, Dix, and Jamroga 2010; Dima and Tiplea 2011). Straightforwardly, we have the following statement.

Theorem 2 *Model checking PATL* is undecidable for incomplete information and synchronous perfect recall semantics.*

The next result is somewhat surprising.

Theorem 3 *Model checking single-player PATL* is undecidable for incomplete information and synchronous perfect recall semantics.*

Single-player PATL* is the fragment of PATL* which allows only a fixed player to be quantified in a formula. Note that model checking (non-probabilistic) single-player ATL* is complete (thus decidable) in EXPTIME for incomplete information (Alur, Henzinger, and Kupferman 2002). The undecidability proof of Theorem 3 presented here is by a reduction from the emptiness problem and strict emptiness problem of probabilistic automata, both known as undecidable problems (Rabin 1963; Paz 1971; Madani, Hanks, and Condon 2003). Here we follow the notations of Gimbert et al. (Gimbert and Oualhadj 2010). A probabilistic automaton PA is a tuple $\langle Q, A, (M_a)_{a \in A}, q_0, \mathcal{F} \rangle$, where

- Q is a finite set of states and q_0 is the initial state,
- $\mathcal{F} \subseteq Q$ is a set of accepting states,
- A is the finite input alphabet, and
- $(M_a)_{a \in A}$ is the set of transition matrix.

For each $a \in A$, $M_a \in [0, 1]^{Q \times Q}$ defines transition probabilities, such that given $q, q' \in Q$, $M_a(q, q')$ is the probability that q makes a transition to q' when a is the input. For every $q \in Q$ and $a \in A$, we have $\sum_{q' \in Q} M_a(q, q') = 1$. Plainly, given a state q , an input a makes a transition to a distribution on Q , and we further extend M_a to be a transformer from distributions to distributions. Given $\Delta \in \mathcal{D}(Q)$, we write $M_a(\Delta)$ for the distribution transformed from Δ by a , such that for all $q' \in Q$, $M_a(\Delta)(q') = \sum_{q \in \text{supp}(\Delta)} \Delta(q) \cdot M_a(q, q')$. Given $w = a_1 \cdot a_2 \cdot \dots \cdot a_n \in A^*$, we write M_w for the function $M_{a_n} \circ M_{a_{n-1}} \circ \dots \circ M_{a_1}$ (we assume function application is right associative). The following problems are undecidable.

Problem 1 *Given a probabilistic automaton $PA = \langle Q, A, (M_a)_{a \in A}, q_0, \mathcal{F} \rangle$ and $\lambda \in [0, 1]$, decide whether there exists a word w such that $M_w(q_0)(\mathcal{F}) \geq \lambda$. This is known as the emptiness problem of a probabilistic automaton, where λ is called a cut-point. Replacing ' \geq ' by a strict inequality ' $>$ ' yields the strict emptiness problem.*

In the following we define a translation \mathfrak{F} mapping probabilistic automata to single player PO-PCGS. Let $PA = \langle Q, A, (M_a)_{a \in A}, q_0, \mathcal{F} \rangle$, define $M = \mathfrak{F}(PA) = (Q \cup \{s_\top, s_\perp\}, \emptyset, A, N, O, PI, PT, \pi)$ with a singleton set of proposition $Prop = \{p\}$, where

- The environment has an empty set of actions,
- The player's action set is $A \cup \{win\}$,
- $N(q) = A \cup \{win\}$ for all $q \in Q \cup \{s_\top, s_\perp\}$,
- $O(q) = \perp$ for all $q \in Q \cup \{s_\perp\}$ and $O(s_\top) = \top$,

- $PI(q_0) = 1$ and $PI(q) = 0$ for all $q \in Q \cup \{s_\top, s_\perp\} \setminus \{q_0\}$,
- PT is extended from $(M_a)_{a \in A}$ as follows
 - $PT(q, a, q') = M_a(q, q')$ for all $q, q' \in Q$ and $a \in A$,
 - If $q \in \mathcal{F}$, $PT(q, win, s_\top) = 1$ and $PT(q, a, q') = 0$ for all $q' \in Q \cup \{s_\perp\}$; otherwise, $PT(q, win, s_\perp) = 1$ and $PT(q, a, q') = 0$ for all $q' \in Q \cup \{s_\top\}$,
 - $PT(s_\top, a, s_\top) = 1$ and $PT(s_\top, a, q') = 0$ for all $a \in A \cup \{win\}$ and $q' \in Q \cup \{s_\perp\}$; and $PT(s_\perp, a, s_\perp) = 1$ and $PT(s_\perp, a, q') = 0$ for all $a \in A \cup \{win\}$ and $q' \in Q \cup \{s_\top\}$.
- $\pi(q) = \emptyset$ for all $q \in Q \cup \{s_\perp\}$ and $\pi(s_\top) = \{p\}$.

In the above translation, we have defined two additional states s_\top and s_\perp . While in a state in Q , the player is allowed to follow the transition rules as defined in the original probabilistic automaton, or to perform a special action win which has the following effect. If at the time win is performed, the system is in an accepting state, it will make a transition to s_\top with probability one; otherwise, it will make a transition to s_\perp . Both s_\top and s_\perp are absorbing states, and the player has only one chance to perform a guess in each run. The only observation of the player is to distinguish whether the current state is s_\top .

Given a probabilistic automaton PA , we study whether $\mathfrak{F}(PA)$ satisfies $\langle \mathbf{i} \rangle^{\geq \lambda} true U p$, where $\lambda \in [0, 1]$ and \mathbf{i} is the single player. Since the player cannot distinguish any states in Q , his strategy can only be a finite sequence of actions in A followed by the win action (if he wait for ever he will never get p either). Such a strategy pattern corresponds to a finite word in PA . We present such a correspondence as follows.

Lemma 2 *Given a probabilistic automaton PA and $\lambda \in [0, 1]$, there exists a word w such that $M_w(q_0)(\mathcal{F}) \geq (>)\lambda$ in PA iff $\langle \mathbf{i} \rangle^{\geq (>)\lambda} true U p$ in $\mathfrak{F}(PA)$.*

A Decidability Result without Until Operator

Lemma 2 reduces undecidable problems in probabilistic automata to model checking finite PO-PCGS with a single *until* operator. Next we show that we can achieve decidability in model checking $PATL_{\setminus U}^*$ after dropping the until operator.

Theorem 4 *Model checking single-player $PATL_{\setminus U}^*$ is in 2-EXPTIME for incomplete information and synchronous perfect recall semantics.*

Proof: (sketch) We present an algorithm for model checking $M \models \phi$, where M is a finite PO-PCGS and ϕ is a $PATL_{\setminus U}^*$ formula. Let h be the depth of nesting of X in ϕ and \mathbf{i} be the player. Our algorithm works on a set of initialized paths of length h , i.e., $R_h = \{s_e(r, 0)a(r, 0)\dots s_e(r, h) \mid r \in \mathcal{R}\}$. Note that the set R_h is of size $O(|M|^{|\phi|})$ by letting $|M|$ be the number of states.

The satisfiability of an expression $R_h, (r, k), v \models \phi$ is computed recursively as the following procedure, where (r, k) is a point in R_h and v is a sequence of observations and actions. Intuitively, this expression states that the formula ϕ holds in the point (r, k) of R_h under the observation history v .

- $\phi = X\phi'$. Then $R_h, (r, k), v \models \phi$ if $R_h, (r, k + 1), v \cdot a(r, k) \cdot O_{\mathbf{i}}(s_e(r, k + 1)) \models \phi'$
- $\phi = \neg\phi'$. Then $R_h, (r, k), v \models \phi$ if not $R_h, (r, k), v \models \phi'$

- $\phi = \langle\langle \mathbf{i} \rangle\rangle^{\leq d} \phi'$. Then $R_h, (r, k), v \models \phi$ if we can
 - existentially choose an equivalence class $[c]_{\mathcal{C}}^{\approx A} \in [\approx_A]_{\mathcal{C}}$, in which at least a run r' with $r'_1(k) = v$ exists, and then
 - universally verify the following expression for all runs r' such that $r' \in \mathcal{R}_{c'}$, $c' \in [c]_{\mathcal{C}}^{\approx A}$ and $r'_1(k) = v$

$$\mu_{r',k,\mathbf{i}}^{c'}(\{(r'', k) \mid r''_1(k) = v \wedge \mathcal{I}, (r'', k) \models \phi'\}) \geq d$$

Now, to verify $M \models \phi$ is equivalent to universally verifying $R_h, (r, 0), O_{\mathbf{i}}(s_e(r, 0)) \models \phi$ for all $r \in R_h$.

A memory of $O(|\phi| \log |M|)$ bits is allocated to store the observation history v . The computation of $\mu_{r',k,\mathbf{i}}^{c'}$ by Theorem 1 needs $O(|\phi|^2 |M|^{|\phi|})$ bits, in which there is a recursion of depth $O(|\phi|)$ levels and each level needs $O(|M|^{|\phi|} |\phi| \log |M|) = O(|\phi| |M|^{|\phi|})$ bits to store a set of points. The whole recursive procedure needs $O(|\phi|)$ alternations. Therefore, the complexity is in AEXPSPACE (that is, it can be solved by an alternating Turing machine in exponential space), which is equivalent to 2-EXPTIME. \square

The decidability result can be generalized to single-group $\text{PATL}_{\setminus U}^*$, which allows only a fixed group of players to be quantified in a formula. Instead of Theorem 1, we need the recursive relationship between $\mu_{r,m+1,A}$ and $\mu_{r,m,A}$, where A is a coalition. We have an MTBDD-based symbolic model checking algorithm for single-group $\text{PATL}_{\setminus U}^*$. It is omitted due to space limit.

Application: Land Search

We model the land search problem (Robe and Frost 2002) by a discrete-time multi-player concurrent game. The map is represented by a graph in which players may choose to move to an adjacent node in each step. They may also stay. Visibility is limited, and is modeled in the way that the probability for a player to spot other players decreases in longer distance, simulating effect of inaccurate sensors. A possible visibility setting could be that if two players i and j are in the same position (i.e., $\text{distance}(i, j) = 0$), they can see each other in 100% probability, and if the distance between them is 1, they spot each other in a probability of 50%, and so on.

Suppose we have two fixed groups of players, one as searchers and the other as intruders. Once spotted, an intruder stays at the same position for the rest of the game. A state of the game consists of the positions of players and their visibility relation vis , where $vis(i, j)$ denotes that i spots j . During each step the positions of the players may change, and the new visibility relation on previously unspotted intruders is recalculated based on the relative distance between the players following pre-defined probability values, in which sense the system transitions are probabilistic.

We propose several ways of minimizing resources usage in the land search problem by formulating a number of strategic properties of the group of searchers in $\text{PATL}_{\setminus U}^*$. In particular, we aim to determine the minimal time as well as minimal number of searchers that are required to spot all intruders in high probabilities. We introduce an atomic proposition $alsp$ to denote that all intruders have been spotted and use a constant $d \in [0, 1]$ to represent a probability threshold. Let A_P be a set of searchers.

The first specification says that there exists a coalition strategy for A_P such that after n steps, the probability that all intruders are spotted is no less than d .

$$\langle\langle A_P \rangle\rangle^{\geq d} (X^n \text{ alsp}) \quad (4)$$

Fixing A_P , if the above formula is satisfied, the following formula can be used to check if the time bound n is strict.

$$\neg \langle\langle A_P \rangle\rangle^{\geq d} (X^{n-1} \text{ alsp}) \quad (5)$$

Given a fixed time bound, one may repetitively check if the search force is more than necessary. For example, let $|A_P| = |A'_P| + 1$, the next formula tests if the size of the current search team can be reduced by 1.

$$\neg \langle\langle A'_P \rangle\rangle^{\geq d} (X^n \text{ alsp}) \quad (6)$$

Related Work

A notably rich literature has been developed since ATL (Alur, Henzinger, and Kupferman 2002) was proposed. Here we only give an overview on the works that are closely related to the topic of the paper.

To reason about incomplete information games, the paper (Hoek and Wooldridge 2002) combines strategy operator and knowledge operator (Fagin et al. 1995) into a logic called ATEL. By utilizing the interpreted system semantics, (Lomuscio and Raimondi 2006) proposes model checking algorithms for the logic. Both works assume that players can observe the environment, but are unable to remember their observation history. The semantics of ATEL is further treated in (Schobbens 2004; Jamroga and van der Hoek 2004).

The paper (Bulling and Jamroga 2009) proposes pATL to reasoning about probabilistic success over complete information games. Their work applies a special *prediction* operator to settle nondeterminism from opponents in probabilistic systems. A more recent work is (Schnoor 2010) which combines knowledge operator and strategic operator to enrich the expressiveness of ATL. Their approach is different from ours. Firstly, the semantics used in (Schnoor 2010) is observational, and we focus on perfect recall semantics. Secondly, their semantics of the strategic operator is orthogonal to the indistinguishable relation defined for the knowledge operator. That is, when interpreting strategic operators, such as in $M, \sigma_A, s \models \langle\langle A \rangle\rangle^{\leq d} \phi$, an implicit assumption is made that coalition A apply the strategy σ_A at the state s . This is uncommon in systems with incomplete information, where players may not have enough information to figure out which particular states they are in.

Several authors have worked on complexities in non-probabilistic games of incomplete information (Reif 1984; Peterson, Reif, and Azhar 2002; Chatterjee et al. 2006). Remarkably, the seminal work of Reif (Reif 1984) acquires complexities results on solving two-player games with incomplete information, which is complete in 2-EXPTIME to the configuration of the game (equivalent to EXPTIME to the state space of the game). Model checking ATL* is in general undecidable by (Alur, Henzinger, and Kupferman 2002; Bulling, Dix, and Jamroga 2010), with a decidable fragment for two-player games (but complete in EXPTIME, loosely related to the result of (Reif 1984)).

Conclusion

We have proposed a new probabilistic interpreted system semantics (PAIS) for PATL* with incomplete information and synchronous perfect recall. A probabilistic interpreted system can be generated by mapping from a probabilistic concurrent game structure (PO-PCGS). We have also studied the complexities of model checking PATL*.

References

- Alur, R.; Henzinger, T. A.; and Kupferman, O. 2002. Alternating-time temporal logic. *Journal of the ACM* 49(5):672–713.
- Bulling, N., and Jamroga, W. 2009. What agents can probably enforce. *Fundamenta Informaticae* 93:81–96.
- Bulling, N.; Dix, J.; and Jamroga, W. 2010. Model checking logics of strategic ability: Complexity. In Dastani, M.; Hindriks, K.; and Meyer, J.-J., eds., *Specification and Verification of Multi-Agent Systems*. Springer.
- Chatterjee, K.; Doyen, L.; Henzinger, T.; and Raskin, J.-F. 2006. Algorithms for omega-regular games with incomplete information. Technical Report UCB/EECS-2006-89, University of California, Berkeley.
- Chen, T., and Lu, J. 2007. Probabilistic alternating-time temporal logic and model checking algorithm. In *FSKD (2)*, 35–39.
- Dima, C., and Tiplea, F. L. 2011. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. *CoRR abs/1102.4225*.
- Fagin, R.; Halpern, J.; Moses, Y.; and Vardi, M. 1995. *Reasoning About Knowledge*. The MIT Press.
- Gimbert, H., and Oualhadj, Y. 2010. Probabilistic automata on finite words: Decidable and undecidable problems. In *ICALP (2)*, 527–538.
- Halpern, J. Y., and Tuttle, M. R. 1993. Knowledge, probability, and adversaries. *Journal of the ACM* 40:917–960.
- Halpern, J. Y. 2003. *Reasoning about Uncertainty*. MIT Press.
- Hoek, W., and Wooldridge, M. 2002. Tractable multiagent planning for epistemic goals. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'02)*, 1167–1174.
- Huang, X.; Luo, C.; and van der Meyden, R. 2011. Symbolic model checking of probabilistic knowledge. In *13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK XII)*, 177–186.
- Jamroga, W., and van der Hoek, W. 2004. Agents that know how to play. *Fundamenta Informaticae* 63(2-3):185–219.
- Lomuscio, A., and Raimondi, F. 2006. Model checking knowledge, strategies, and games in multi-agent systems. In *the proceedings of the 5th international joint conference on Autonomous agents and multiagent systems (AAMAS 2006)*, 161–168.
- Madani, O.; Hanks, S.; and Condon, A. 2003. On the undecidability of probabilistic planning and related stochastic optimization problems. *Artificial Intelligence* 147:5–34.
- Paz, A. 1971. *Introduction to probabilistic automata (Computer science and applied mathematics)*. Academic Press.
- Peterson, G. L.; Reif, J. H.; and Azhar, S. 2002. Decision algorithms for multiplayer non-cooperative games of incomplete information. *Computers and Mathematics with Applications* 43:179–206.
- Rabin, M. O. 1963. Probabilistic automata. *Information and Control* 6(3):230–245.
- Reif, J. H. 1984. The complexity of two-player games of incomplete information. *Journal of Computer and System Science* 29(2):274–301.
- Robe, R. Q., and Frost, J. R. 2002. A method for determining effective sweep widths for land searches. Technical report, The National Search and Rescue Committee, Washington, DC.
- Schnoor, H. 2010. Strategic planning for probabilistic games with incomplete information. In *the proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, 1057–1064.
- Schobbens, P.-Y. 2004. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science* 85(2):82–93.
- Vardi, M. Y. 1985. Automatic verification of probabilistic concurrent finite-state programs. In *FOCS*, 327–338.
- Zhang, C., and Pang, J. 2010. On probabilistic alternating simulations. In *Theoretical Computer Science - 6th IFIP TCS*, 71–85.