Strategic Planning for Probabilistic Games with Incomplete Information

Henning Schnoor
Institut für Informatik
Christian-Albrechts-Universität Kiel
24118 Kiel, Germany
schnoor@ti.informatik.uni-kiel.de

ABSTRACT

Alternating-time Temporal Logic (ATL) [1] is used to reason about strategic abilities of agents. Aiming at strategies that can realistically be implemented in software, many variants of ATL study a setting where strategies may only take available information into account [7]. Another generalization of ATL is Probabilistic ATL [4], where strategies achieve their goal with a certain probability.

We introduce a semantics of ATL that takes into account both of these aspects. We prove that our semantics allows simulation relations similar in spirit to usual bisimulations, and has a decidable model checking problem in the case of memoryless strategies (for memory-dependent strategies the problem is undecidable).

Categories and Subject Descriptors

I.2.4 [**Knowledge Representation Formalisms and Methods**]: Temporal logic

General Terms

Theory

Keywords

Alternating-time temporal logic, incomplete information, probability

1. INTRODUCTION

Alternating-time Temporal Logic (ATL) [1] is widely recognized as a suitable logic to reason about strategic abilities: The operator $\langle\langle A \rangle\rangle\, \varphi$ expresses that a coalition A has a strategy to achieve the goal specified by φ. In practice, an agent also needs to have enough information to implement the strategy. In a realistic environment, each agent will only have partial information about the current state of the system. This leads to a restriction of the available strategies to so-called *uniform* ones [7], where strategies may only take into account information that is available to the agent. Further, the existence of the strategy is not enough, there must be a way for each agent in the coalition to *determine* the correct strategy to follow. If agents are not able to freely communicate during the game, it is not guaranteed that they can "agree" on the same strategy to follow.

In addition to incomplete information, another generalization is probabilistic ATL, where strategies are required to achieve a goal with a certain minimal probability [4].

We introduce a new semantics for ATL that takes both incomplete information and probabilism into account. In order to lead to reliable strategies, our treatment of incomplete information states rather strict requirements for the admissible strategies: We demand that there is a *deterministic* way for the agents to determine their strategies for each goal they want to achieve, given only the (potentially incomplete) knowledge about the current state of the system available to them. However, we allow *prior agreement*: Coalitions may agree on a joint set of strategies *before* the start of the game.[1] This models coalitions where each agent trusts each other and can rely on information about the behavior of others. In particular, the setting applies when agents are jointly developed software programs.

During the game, we assume that agents may only communicate with each other using explicit moves. This allows to handle situations where communication is an integral part of the problem the agents want to solve, as the study of cryptographic protocols (see [11] and [12] for studies of strategic properties of cryptographic protocols in a game-theoretic setting). Similarly, we treat storing of information as an explicit move and therefore focus on *memoryless* strategies: Agents only have access to the information they can currently observe, and not to the entire history of the game. In order to be able to model storing of information, we allow infinite game structures. The three main contributions of this paper are the following:

(i) We propose a new semantics for ATL that takes into account incomplete information and probabilism at the same time. We allow agents to reach prior agreement about the strategies they will use during the game. We show that when requiring a natural "maximality" condition of the previously agreed strategies, then in the classical deterministic, complete-information setting, our semantics is equivalent to standard ATL.

(ii) We define a simulation similar to bisimulations ob-

Cite as: Strategic Planning for Probabilistic Games with Incomplete Information, Henning Schnoor, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 1057-1064

[1] As a natural situation where prior agreement is useful, consider a game in which agents are successful only if in some state while playing, both of them choose the same number. Then—without communication—the agents do not have a successful strategy; however if they anticipate this situation and agree on a number before the game is started, obviously they can be successful.

tained for ATL [2]. It allows to specify strategies on a finite "core" of a structure, but to apply them in the original infinite one (if such a finite core exists). More generally, strategies may be defined on a "simple" structure and can then be applied in a more "complicated" one, still achieving the same goals. This result paves the way for software implementations of strategies for infinite systems. Our simulation may be of independent interest, as it can be applied to standard semantics of probabilistic ATL. We are not aware of prior results on simulations or bisimulations for ATL in the probabilistic or incomplete-information setting.

(iii) We study the complexity of the model checking problem for our semantics. We prove that the problem is decidable for finite structures, which strengthens the above point of using simulations and finite structures to represent infinite systems. The model checking problem is in 3EXPTIME and is 2EXPTIME-hard, but the complexity drops to PSPACE-complete in the deterministic setting. The problem is undecidable for history-dependent strategies.

Related Work.

There is a rich literature on ATL with incomplete information, going back to the initial ATL introduction in [1]. The notion of uniform strategies that we use was first used in combination with ATL in [7] (there called incomplete information strategies), and studied in detail in [10], which also discusses methods to identify a correct strategy, and allows to separate the roles of the coalitions executing and identifying the strategy. In [16], the model checking complexity of ATL with incomplete information and both history-dependent and memoryless strategies is studied. Further, [6] discusses an extension of ATL where it is required that agents know that they have a strategy, and can identify it.

The goal of the above-mentioned works differs from ours: We do not consider planning and identifying strategies during the run of a game, but study what can be achieved by coalitions with help of an additional planning phase where coalitions may reach *prior agreement* on strategies for joint goals. This leads to a truth definition that cannot be specified in a purely local way (i.e., as a function of the game structure, the state, and the formula alone). To the best of our knowledge, prior agreement has not been addressed in combination with ATL before.

Probabilistic ATL has been studied in [3], where the success of a coalition's strategy is measured depending on a probability measure describing the (likely) actions of the remainder of the agents. In the current paper, we use the usual pessimistic worst-case assumption about the actions performed by the opponents of a coalition. In [4], a model checking algorithm for history-dependent strategies for probabilistic ATL is introduced.

To the best of our knowledge our work is the first studying ATL with incomplete information in a probabilistic setting.

The structure of the paper is as follows: In Section 2, we introduce our semantics for ATL, we discuss various aspects of it in Section 3. Section 4 introduces our notion of simulation, and explains how it allows to transfer previously agreed strategies. Section 5 contains our results on decidability and complexity of the model-checking problem. We conclude in Section 6 with some open questions. A full version of the paper containing all proofs can be found in [15].

2. PROBABILISTIC ATL* WITH INCOMPLETE INFORMATION

In this section we introduce our semantics for ATL*. We first define cuncurrent game structres, which are the objects that ATL* reasons about, and then define formulas. While both of these definitions are fairly standard, our treatment of strategies and their collections into "strategy choices" is novel and forms the heart of our semantics.

2.1 Concurrent Game Structures

The following definition of a concurrent game structure is based on the one from [1], extended to infinite structures (see also [11]), a probabilistic setting (see also [4]) and a mechanism to deal with incomplete information (see also [10]). We will give an example in Section 3.

Definition. A *concurrent game structure (CGS)* is a tuple $\mathcal{C} = (\Sigma, Q, \mathbb{P}, \pi, \Delta, \delta, \mathsf{eq})$, where

- Σ is a non-empty, finite set of *agents*,
- Q is a non-empty set of *states*,
- $\mathbb{P}$ is a finite set of *propositional variables*,
- $\pi\colon \mathbb{P} \to \mathcal{P}(Q)$ is a *propositional assignment*,
- Δ is a *move function* assigning to each state $q \in Q$ and agent $a \in \Sigma$ a nonempty set $\Delta(q, a)$ of *moves* available at state q to agent a. For $A \subseteq \Sigma$ and $q \in Q$, an (A, q)-*move* is a function c which maps each $a \in A$ to a move $c(a) \in \Delta(q, a)$. A (Σ, q)-move is a *total q-move*.
- δ is a *probabilistic transition function* which for each state q and total q-move c, returns a discrete probability distribution $\delta(q, c)$ on Q (the state obtained when in q, all agents perform their move as specified by c),
- eq is an *information function* $\mathsf{eq}\colon \{1, \ldots, n\} \times \Sigma \to \mathcal{P}(Q \times Q)$, where n is a natural number, and for each $i \in \{1, \ldots, n\}$ and $a \in \Sigma$, $\mathsf{eq}(i, a)$ is an equivalence relation on Q. We also call each $i \in \{1, \ldots, n\}$ a *degree of information*.

A subset $A \subseteq \Sigma$ is also called a *coalition of* $\mathcal{C}$. We often omit "of $\mathcal{C}$" when $\mathcal{C}$ is clear from the context. The coalition $\Sigma \setminus A$ is denoted with $\overline{A}$. We often write $\Pr(\delta(q, c) = q')$ for $(\delta(q, c))(q')$. The information function eq allows to reason about incomplete information: Often an agent a will not have complete information about the current state. Hence for each agent there is an associated "indistinguishability" relation which is an equivalence relation specifying the states between which a cannot distinguish. To be able to evaluate strategies with different degrees of information for the same agent, we specify several relations $\mathsf{eq}(1, a), \ldots, \mathsf{eq}(n, a)$ for each agent a. For the examples in this paper, we ony consider the simpler case where each agent has a fixed indistinguishability relation, i.e., there is only one degree of information. In order to simplify the presentation, we will often omit the information degree when there is only one.

We often write $q_1 \sim_{\mathsf{eq}_i(A)} q_2$ for $(q_1, q_2) \in \cap_{a \in A} \mathsf{eq}(i, a)$ (i.e., no member of the coalition A can distinguish between q_1 and q_2), and write $q_1 \sim_{\mathsf{eq}_i(a)} q_2$ for $q_1 \sim_{\mathsf{eq}_i(\{a\})} q_2$. $\mathcal{C}$ is *deterministic* if its transition function δ is (i.e., the probability distributions returned by δ assign 1 to a single state and 0 to all others). We say that an agent a in $\mathcal{C}$ has *complete* *information* if $\mathsf{eq}(i,a)$ is the equality relation on the state set for all i. A CGS has complete information if every agent has. For deterministic structures, we can simplify our definitions (for example, the notion of a "response," see below, is not needed in this case). We omit the simpler definition.

A *path* in a CGS $\mathcal{C}$ is a (possibly infinite) sequence λ of states in $\mathcal{C}$. With $\lambda[i]$ we denote the ith state in λ, and with $\lambda[i,\infty]$ the (possibly infinite) sequence $\lambda[i], \lambda[i+1], \ldots$. To shorten forthcoming examples, we will often assume that for each state q there is a propositional variable with the same name, which is true only in the state q.

2.2 ATL* formulas

We now define formulas to describe properties and strategic goals in a CGS. Our syntax is identical to ATL* ([1]), with the addition of degrees of information and probabilities.

Definition. Let $\mathcal{C}$ be a CGS with n degrees of information. Then the set of *ATL*-formulas for* $\mathcal{C}$ is defined as follows:

- A propositional variable of $\mathcal{C}$ is a state formula for $\mathcal{C}$,
- conjunctions and negations of state (path) formulas for $\mathcal{C}$ are state (path) formulas for $\mathcal{C}$,
- if A is a coalition, $1 \leq i \leq n$, $0 \leq \alpha \leq 1$, and $\blacktriangleleft$ is one of $\leq, <, \geq, >$, and ψ is a path formula for $\mathcal{C}$, then $\langle\langle A \rangle\rangle_i^{\blacktriangleleft \alpha}\,\psi$ is a state formula for $\mathcal{C}$,
- if A is a coalition, $1 \leq i \leq n$, and ψ is a state formula for $\mathcal{C}$, then $\mathcal{K}_i^A \psi$ is a state formula for $\mathcal{C}$,
- every state formula for $\mathcal{C}$ is a path formula for $\mathcal{C}$,
- If φ_1 and φ_2 are path formulas for $\mathcal{C}$, then $\mathsf{X}\varphi_1$ and $\varphi_1 \mathsf{U} \varphi_2$ are path formulas for $\mathcal{C}$.

The operator $\langle\langle . \rangle\rangle$ is called *strategy operator*. Intuitively, $\langle\langle A \rangle\rangle_i^{\blacktriangleleft \alpha}\,\varphi$ means that the agents in A have a strategy making φ true with probability $\blacktriangleleft\ \alpha$, where for the strategy, the agents may only access the knowledge available to them in information degree i, and "knowing" a strategy refers to a possible prior agreement between the agents in A. An ATL*-formula is a state formula unless specified otherwise. We define the usual abbreviations, i.e., $\varphi \vee \psi = \neg(\neg\varphi \wedge \neg\psi)$, $\Diamond\varphi = \mathsf{true}\,\mathsf{U}\varphi$, and $\Box\varphi = \neg\Diamond\neg\varphi$. We also simply write $\langle\langle a \rangle\rangle$ instead of $\langle\langle \{a\} \rangle\rangle$, etc. We say that a $\langle\langle . \rangle\rangle_i$-formula is one whose outmost operator is $\langle\langle A \rangle\rangle_i^{\blacktriangleleft \alpha}$ for some coalition A, and some $\blacktriangleleft\ \alpha$, etc. The set $ag(\varphi)$ is the set of agents mentioned in the formula φ, i.e., it contains the agent $a \in \Sigma$ if and only if $a \in A$ for some coalition A such that $\langle\langle A \rangle\rangle$ or $\mathcal{K}^A$ appears in φ. In a CGS with only one degree of information, we often omit the i subscript of the strategy operator, similarly in a deterministic CGS we usually omit the probability bound $\blacktriangleleft\ \alpha$ (and understand it to be read as ≥ 1 in deterministic structures).

2.3 Strategies and Strategy choices

In order for an agent to decide on a suitable move to achieve a certain goal, *strategies* are used. A strategy fixes, for each state, a move to be performed by the corresponding agent. When reasoning about several (possibly contradictory) goals, the question how an agent determines the correct strategy is also relevant.

Our semantics addresses the following situation: At every point in time, an agent a may decide (or be instructed to) attempt to achieve a goal G (which is specified by an ATL*-formula) as part of a coalition A'. Before the start of the game, the coaltion A' has agreed on a set of strategies to reach the goal G, where the correct strategy to choose may depend on the current state of the game. We model this agreement as a *strategy choice*. Hence there are two steps in identifying the move to perform:

1. For a given goal G to be achieved with coalition A', the agent $a \in A'$ has to *identify* the correct strategy. This is done with a function S assigning each possible state a strategy. Of course, if states are indistinguishable for a, then the same strategy should be picked in both states—otherwise agent a does not have sufficient information to determine the strategy defined by S.
2. The strategy obtained in the above is a function that, in every state, defines a move for an agent a to perform. Obviously, this move needs to be "legal" (i.e., available in the state), and the same restriction as above applies: In order for the agent to be able to follow the strategy, the same move needs to be picked in states that the agent cannot distinguish. Such strategies are called *uniform* (see also, e.g., [16]).

In the following discussion, recall that in the common case that we only specify a single indistinguishability relation for each agent, we can omit the corresponding index i and merely speak about uniform strategies, etc. Following the above ideas, the following definition of a strategy is natural.

Definition. Let $\mathcal{C}$ be a CGS with state set Q and move function Δ, and n degrees of information. For an agent a, an *a-strategy* in $\mathcal{C}$ is a function s_a assigning a move to each state such that $s_a(q) \in \Delta(q,a)$ for each $q \in Q$. For $1 \leq i \leq n$, s_a is *i-uniform*, if $q_1 \sim_{\mathsf{eq}_i(a)} q_2$ implies $s_a(q_1) = s_a(q_2)$. For a coalition A, an *A-strategy* is a family $(s_a)_{a \in A}$, where each s_a is an a-strategy.

As mentioned in the introduction, we only consider *memoryless* strategies: The action of an agent may only depend on the current state. Usually, ATL* allows *history-dependent* strategies. Since we allow CGSs to be infinite, our model canonically allows the treatment of history-dependent strategies (see Sections 3.2 and 4.2), however the price to pay is that the important model checking problem becomes undecidable (see Section 5).

Usually, a strategy for a coalition A is required to work against "all possible counter-strategies" of $\overline{A}$. Since we only consider memoryless strategies, quantifying over strategies is strictly weaker than quantifying over all possible "behaviours." Hence we quantify over "responses," where a *response* to a coalition A is a function r such that for each $t \in \mathbb{N}$ and each $q \in Q$, $r(t,q)$ is a $(\overline{A}, q)$-move: A response is an arbitrary reaction to the outcomes of a possible strategy chosen by A. Given a strategy $s_A = (s_a)_{a \in A}$ and a response r to A, the resulting "game" is a Markov process, where the transition probabilities are determined by the transition function of the CGS (the moves of the agents in A are fixed by s_A, those by $\overline{A}$ are fixed by r: When in the t-th step, the game is in the state q, then an agent $a \in A$ perform the move $s_a(q)$, and an agent $b \in \overline{A}$ performs the move $(r(t,q))(b)$). Agents in $\overline{A}$ are not bound by any strategy; they are not restricted to any uniformity conditions and also may act differently when the same state is reached twice during a run of the game. Demanding that a strategy works against all responses mirrors the usual worst-case assumptions in requiring that strategies are successful even if the opponents behave completely irrationally (i.e., do not follow any strategy at all). We define the following:

Definition. Let $\mathcal{C}$ be a CGS, let s_A be an A-strategy, let r be a response to A. For a set M of paths over $\mathcal{C}$, and a state $q \in Q$,

$$\Pr\left(q \to M \mid s_A + r\right)$$

is the probability that in the Markov process resulting from $\mathcal{C}$, s_A, and r with initial state q, the resulting path is an element of M.

Strategies allow agents to choose a move in a state. Uniform strategies ensure that an agent has sufficient information to determine the correct move. As explained above, agents also have to decide on a strategy for a given goal; we formalize this using *strategy choices*.

Definition. Let $\mathcal{C}$ be a CGS with state set Q, and let A be a coalition. A *strategy choice for A in $\mathcal{C}$* is a function S such that for each $a \in A$, $q \in Q$, each $\langle\langle . \rangle\rangle_i$-formula φ for $\mathcal{C}$ with $ag(\varphi) \subseteq A$, $\mathsf{S}(a, q, \varphi)$ is an i-uniform a-strategy in $\mathcal{C}$, and if $q_1 \sim_{\mathsf{eq}_i(a)} q_2$, then $\mathsf{S}(a, q_1, \varphi) = \mathsf{S}(a, q_2, \varphi)$.

The purpose of a strategy choice S is to model "prior agreement:" Before the game, the coalition may agree on a set of suitable strategies to achieve strategic goals specified by ATL*-formulas. These strategies are collected in S. When (during the game) in the state q, the coalition $A' \subseteq A$ decides (or is instructed to) achieve a goal by φ, each agent a in A' chooses the strategy $\mathsf{S}(a, q, \varphi)$. Hence the resulting A'-strategy is the family $(\mathsf{S}(a, q, \varphi))_{a \in A'}$, which (somewhat abusing notation) we denote with $\mathsf{S}(A', q, \varphi)$. A strategy choice is allowed to depend on the state to handle situations in which players have more information at the time of *deciding* on a strategy than they have later when *implementing* it. The uniformity conditions for strategy choices and strategies ensure that agents have "enough knowledge to *identify* and *execute*" the correct strategy (cp. [9]).

Note that the "amount" of information available to a coalition is defined in the formula ($\langle\langle A \rangle\rangle_i\, \varphi$ specifies that to reach the goal φ, A may access information of degree i, see the semantics definition below). This allows (by nesting operators) to express statements like "Coalition A has a high-knowledge strategy to reach a state where coalition B has a low-knowledge strategy to achieve φ," even when A and B are not disjoint.

2.4 Semantics Definition

We now define our semantics. The definition of truth of formulas is relative to a strategy choice: The question which strategic goals can be reached clearly depends on the agreements reached by coalitions before the state of the game.

Definition. Let $\mathcal{C} = (\Sigma, Q, \mathbb{P}, \pi, \Delta, \delta, \mathsf{eq})$ be a CGS, and let S be a strategy choice for a coalition A in $\mathcal{C}$, let φ_1, φ_2 be state formulas for $\mathcal{C}$, let ψ_1, ψ_2 be path formulas for $\mathcal{C}$, such that $ag(\varphi_1), ag(\varphi_2), ag(\psi_1), ag(\psi_2) \subseteq A$, let $q \in Q$, and let λ be a path over Q. We define

- $\mathcal{C}, \mathsf{S}, q \models p$ iff $q \in \pi(p)$ for $p \in \mathbb{P}$,
- negation and conjunction are handled as usual,
- $\lambda, \mathsf{S} \models \varphi_1$ iff $\mathcal{C}, \mathsf{S}, \lambda[0] \models \varphi_1$,
- $\lambda, \mathsf{S} \models \mathsf{X}\psi_1$ iff $\lambda[1, \infty], S \models \psi_1$,
- $\lambda, \mathsf{S} \models \psi_1 \mathsf{U} \psi_2$ iff there is some $i \geq 0$ such that $\lambda[i, \infty], \mathsf{S} \models \psi_2$ and $\lambda[j, \infty], \mathsf{S} \models \psi_1$ for all $j < i$,
- If $\varphi_1 = \langle\langle A' \rangle\rangle_i^{\blacktriangleleft \alpha}\, \psi_1$, then $\mathcal{C}, \mathsf{S}, q \models \varphi_1$ iff for every response r to A', we have $\Pr\left(q \to \{\lambda \mid \lambda, \mathsf{S} \models \psi_1\} \mid \mathsf{S}(A', q, \varphi_1) + r\right) \blacktriangleleft \alpha$,
- $\mathcal{C}, \mathsf{S}, q \models \mathcal{K}_i^A \varphi_1$ if $\mathcal{C}, \mathsf{S}, q' \models \varphi_1$ for all $q' \in Q$ with $q' \sim_{\mathsf{eq}_i(A)} q$.

The intuitive meaning of $\mathcal{C}, \mathsf{S}, q \models \langle\langle A' \rangle\rangle_i^{\blacktriangleleft \alpha}\, \psi_1$ is the following: When in the state q, the coalition A' decides (or is instructed to) try to achieve the goal ψ_1, they follow the strategies specified by the strategy choice S for ψ_1, i.e., the strategy $\mathsf{S}(A', q, \psi_1)$. The uniformity requirements for strategy choices ensure that determining and following this strategy requires the agents to access only information available to them with the specified degree i. The formula is satisfied if the thus-selected A'-strategy is successful with probability $\blacktriangleleft \alpha$, for every possible behaviour of the players $\overline{A'}$.

The *knowledge operator* $\mathcal{K}$ allows, among other possibilities, to express further requirements about the available strategies: Consider the formula $\mathcal{K}_i^{A'}\, \langle\langle A' \rangle\rangle_i^{\blacktriangleleft \alpha}\, \psi$, for which we introduce the shorthand $\langle\langle \mathcal{K} A' \rangle\rangle_i^{\blacktriangleleft \alpha}$. By the above semantics, this formula is true in a state q with respect to a strategy choice S if and only if the strategies chosen for the formula ψ by S are successful in *every* state q' such that $q' \sim_{\mathsf{eq}_i(A')} q$. The intuition behind using this operator is that in the state q, the coalition A' cannot rule out that the actual state is q', hence if their chosen strategy is unsuccessful in q', the coalition A' cannot be sure about its success in q either. A similar requirement was made in [16].

Essentially, $\langle\langle \mathcal{K} A' \rangle\rangle_i^{\blacktriangleleft \alpha}$ expresses that the coalition A' has a strategy to ensure that ψ is true with probability $\blacktriangleleft \alpha$, and with information degree i, the coalition can identify on the correct strategy, each agent can execute the strategy, and the coalition has sufficient (distributed) knowledge to "know" that the strategy is successful—note that this does *not* imply that every single agent knows this fact (the latter can be expressed by $\bigwedge_{a \in A} \mathcal{K}_i^a\, \langle\langle A \rangle\rangle_i^{\blacktriangleleft \alpha}$).

3. EXAMPLE AND DISCUSSION

We provide a short example, for simplicity we only give a deterministic one. We also show how history-dependent strategies can be handled in our framework, and discuss a subtle issue of our semantics.

3.1 An Example

Consider the CGS $\mathcal{C}$ shown in Figure 1, which revisits the classical "blind and lame agent" example: There is a blind agent a who can turn a switch for a light bulb, but does not know whether the light is on or off. A second agent b can see, but cannot influence the switch. Formally, the moves of agent a are 0 (do nothing), which does not change state, and 1 (turn the switch), which alternates between the states "On" to "Off." The moves for agent b are irrelevant (i.e., disregarded by the transition function). Agent b has complete information, while for agent a, both states are indistinguishable (there only is a single degree of information).

We now evaluate strategies for the formula XOn, i.e., to turn on the light. Obviously, agent b alone does not have a strategy to achieve XOn, as b cannot perform a relevant action. Since both states are indistinguishable for agent a, any uniform a-strategy has to perform the same move $\beta \in \{0,1\}$ in both the On and the Off state. We first consider the strategy s_a^1, which always performs the move 1 (i.e., toggles the switch). Let S_1 be the strategy choice that for the formula $\mathsf{X}\mathrm{On}$ returns this strategy.

Obviously, the strategy will be successful in the state Off, and will fail in the state On. Since On $\sim_{\mathsf{eq}(a)}$ Off, this implies that $\mathcal{C}, \mathsf{S}_1, \mathrm{Off} \not\models \langle\langle \mathcal{K}a \rangle\rangle\, \mathsf{X}\mathrm{On}$, even though the selected strategy would be successful in the state—a does not have sufficient information to know that he will be successful.

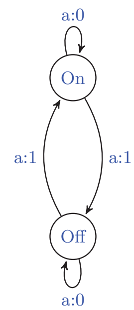

Figure 1: "Blind and Lame Agent"

Now consider the coalition $A = \{a, b\}$. Since the moves of b are irrelevant, we use the same strategy choice as above (formally, add a dummy move for b), which is still unsuccessful in the state On. However, the only state q' with $q' \sim_{\mathsf{eq}_1(A)}$ Off is the state Off itself—hence (since the move is successful in Off), it follows that $\mathcal{C}, \mathsf{S}_1, \mathrm{Off} \models \langle\langle \mathcal{K}A \rangle\rangle\, \mathsf{X}\mathrm{On}$. This expresses that *together*, the coalition A has enough information to know that their strategy is successful: When asked by an external environment whether they have a strategy to turn on the light, in the state Off, they would reply as follows (when S_1 is the previously agreed set of strategies):

Agent a: "If and only if agent b says so."

Agent b: "Yes."

When asked the same question in the state On, agent a has to give the same answer (since he does not know whether the state is On or Off), and agent b would reply with "No." Hence the coalition possesses sufficient knowledge to determine whether the strategy is successful in both states.

Also note that when considering the strategy choice S_0, that always chooses the move 0 for a instead, the formula $\langle\langle a \rangle\rangle\, \mathsf{X}\mathrm{On}$ is satisfied in On, but not in Off. In particular this establishes the claim made in the introduction: Our semantics cannot be defined "locally," without fixing the set of previously agreed strategies in a strategy choice (also note that of these two strategy choices, none is strictly "better" than the other, both are equally valid).

Finally note that strategy choices are free to (somewhat counter-intuitively) define different strategies for φ and for $\varphi \wedge \varphi$. In some situations this can be used to model an external "environment" passing information to the agents, due to space reasons we do not discuss this issue further.

3.2 History Dependence

ATL* usually allows so-called *history-dependent* strategies, where the action of an agent in a state may depend on the entire previous history of the game. In our framework, history-dependence can be expressed by defining a "history-dependent version" $\mathcal{C}^{hst}$ of a given CGS $\mathcal{C}$, which simply encodes history into the states themselves in the straightforward way: States of $\mathcal{C}^{hst}$ are sequences of states of $\mathcal{C}$, two histories are distinguishable if they have different length or they are distinguishable in at least one position, the definition of the other components follows canonically. In the remainder of the paper, $\mathcal{C}^{hst}$ will be used as a running example and as a means to formally state the undecidability result in Section 5.

3.3 Maximal Strategy Choices

Our semantics for ATL* admits strategy choices satisfying formulas which intuitively should be "unsatisfiable:"

Consider the CGS $\mathcal{C}$ shown in Figure 2. $\mathcal{C}$ is a deterministic, complete-information CGS with a single agent a. In q_0, there is a single move leading to q_1, in q_1, there are 2 moves leading to q_2 or q_3 respectively. Define φ as $\langle\langle a \rangle\rangle\, \mathsf{X}\neg \langle\langle a \rangle\rangle\, \mathsf{X} q_3$. Intuitively, φ expresses that there is a move by a such that in the resulting state, a cannot reach q_3. Intuitively (and in standard ATL*), φ is not satisfied in q_0, since the only available move for a leads to the state q_1, from which the state q_3 can be reached by the move 1. However, for the strategy choice S always returning the strategy that chooses the move 0 in every state, it follows that $\mathcal{C}, \mathsf{S}, q_0 \models \varphi$.

Obviously, S is not interesting, as it fails to achieve the goal $\mathsf{X} q_3$ (and hence succeeds in achieving φ) *deliberately*, by choosing an unsuccessful strategy although a successful one is available. In particular, the satisfaction of φ does not allow us to conclude that there is a move for a in q_0 such that in the next state, a cannot reach q_3 anymore *even when trying*[2]

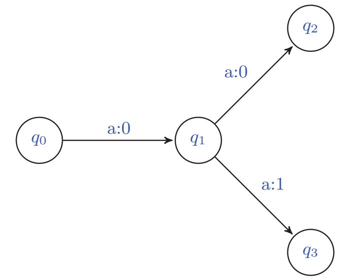

Figure 2: Example

Usually, one assumes strategy choices to not deliberately choose "bad" strategies in "innermost" formulas in order to make "outermost" operators true; it should prioritize "innermost" formulas. To formalize this intuition, let $\mathsf{sd}(\varphi)$ (the *strategic depth* of a formula φ) be the maximal nesting degree of strategic operators in φ. For a strategy choice S for a coalition A in a CGS $\mathcal{C}$ and a formula φ with $ag(\varphi) \subseteq A$, with $\mathsf{sat}_\varphi(\mathsf{S}, j)$ we denote the set of pairs (q, ψ) such that ψ is a $\langle\langle . \rangle\rangle$-subformula of φ, $\mathsf{sd}(\psi) = j$, and $\mathcal{C}, S, q \models \psi$. Using this notation, we can define an order on strategy choices, where $\mathsf{S}_1 \leq \mathsf{S}_2$ should mean that S_2 does a a better job of prioritizing formulas with small depth than S_1 does.

Definition. Let S_1 and S_2 be strategy choices for a coalition A in a CGS $\mathcal{C}$, let φ be an ATL*-formula for $\mathcal{C}$ with $ag(\varphi) \subseteq A$. Then $\mathsf{S}_1 \leq_\varphi \mathsf{S}_2$ if 1. $\mathsf{sat}_\varphi(\mathsf{S}_1, i) = \mathsf{sat}_\varphi(\mathsf{S}_2, i)$ for all $i \leq \mathsf{sd}(\varphi)$, or 2. for the minimal i such that $\mathsf{sat}_\varphi(\mathsf{S}_1, i) \neq \mathsf{sat}_\varphi(\mathsf{S}_2, i)$, we have $\mathsf{sat}_\varphi(\mathsf{S}_1, i) \subsetneq \mathsf{sat}_\varphi(\mathsf{S}_2, i)$

The definition captures the above-mentioned intuitive requirement: If $\mathsf{S}_1 \leq_\varphi \mathsf{S}_2$, then for the first degree of strategic depth where S_1 and S_2 actually differ, S_2 satisfies strictly "more" (either "more formulas," or the same formulas at

[2]Note that the example of a coalition A trying to achieve a situation where it is unable to reach a certain goal is not at all contrived: A might want to *commit* to, let's say, a secret value. Then a selection of strategies that could—but does not—violate the commitment is clearly unsatisfying: It is required that the commitment *cannot* be violated anymore, without the assumption of the good-will of A.

"more states," or both). Note that clearly, there are strategy choices S_1 and S_2 such that neither $\mathsf{S}_1 \leq_\varphi \mathsf{S}_2$ nor $\mathsf{S}_2 \leq_\varphi \mathsf{S}_1$ is true.

As an example, in the above-described $\mathcal{C}$ from Figure 2, consider the strategy choice S' which always returns the strategy that in q_0 chooses the move 0 and in q_1, chooses the move 1. It can easily be verified that $\mathsf{S} \leq_\varphi \mathsf{S}'$.

As usual, a strategy choice S is $\leq_\varphi$-maximal if $\mathsf{S} \leq_\varphi \mathsf{S}'$ implies $\mathsf{S}' \leq_\varphi \mathsf{S}$. In many situations, a restriction to maximal strategy choices is natural. In particular, for the deterministic, complete-information setting, our semantics restricted to maximal strategy choices is equivalent to the usual semantics of ATL* [1]. In the following with $\mathcal{C}, q \models_{\text{ATL-ml}} \varphi$, we mean that the formula φ (with all indices i of an $\langle\langle A \rangle\rangle_i$-operator removed) is satisfied at the state q of the CGS $\mathcal{C}$ in the standard ATL* semantics restricted to memoryless strategies. The following is very easy to show:

Proposition 3.1. *Let $\mathcal{C}$ be a deterministic CGS with complete information, let φ be a formula for $\mathcal{C}$, and let S be a $\leq_\varphi$-maximal strategy choice for $ag(\varphi)$ in $\mathcal{C}$. Then for all subformulas ψ of φ, the following are equivalent: 1. $\mathcal{C}, \mathsf{S}, q \models \psi$, 2. $\mathcal{C}, q \models_{ATL\text{-}ml} \psi$.*

An analogous result also holds when considering probabilistic CGSs and the semantics as defined in [4], restricted to pure memoryless strategies.

For a natural class of game structures, maximal strategy choices always exist: We say that a CGS has *finite index*, if for every equivalence relation $\mathsf{eq}(i, a)$, every equivalence class has finitely many elements. This criterion is clearly satisfied by finite CGSs, further if $\mathcal{C}$ has finite index, then $\mathcal{C}^{hst}$ has finite index as well. Trivially a CGS with complete information has finite index. In such CGSs, every strategy choice can be "enhanced" to obtain a maximal one:

Theorem 3.2. *Let $\mathcal{C}$ be a CGS with a countable state set and finite index, let φ be a formula for $\mathcal{C}$ with $ag(\varphi) \subseteq A$, and let S be a strategy choice for A in $\mathcal{C}$. Then there is a $\leq_\varphi$-maximal strategy choice S_{max} for A in $\mathcal{C}$ such that $\mathsf{S} \leq_\varphi \mathsf{S}_{max}$.*

Theorem 3.2 is false without requiring finite index.

4. SIMULATION RELATIONS

Simulations and Bisimulations are often used to relate structures to one another in a way preserving "interesting" features: A bisimulation between structures S_1 and S_2 with state sets Q_1 and Q_2 is usually a relation $Z \subseteq Q_1 \times Q_2$ such that when $(q_1, q_2) \in Z$, then q_1 and q_2 satisfy the same properties (e.g., the same logical formulas). In our case, a simulation allows to "translate" a strategy choice from one (potentially "easy") CGS to another (potentially "complicated") one. This allows agents to construct their joint strategy choice on an "easy" structure and apply it in the "complicated" one (when the description of the simulation relation itself is of manageable size). This feature is particularly attractive since model checking for finite structures is decidable, see Section 5.

Bisimulations for ATL* were originally defined in [2] (see also [13] for a definition which is closer to ours). The additional requirements that we make of our simulations are needed to deal with incomplete information, probabilism, and explicit strategies:

1. We require certain *uniformity* conditions similar to the ones required for strategies and strategy choices,
2. we demand that moves between related states can be transferred in a deterministic (and uniform) way,
3. we handle probabilities in the natural way,
4. for to transferring strategy choices, our simulations only need to work for a particular coalition (and then allows to transfer strategy choices for that coalition).

We only state definitions for (unidirectional) simulations; a bisimulation analogously can be defined as a relation Z that is a simulation in both directions *simultanously*.

4.1 Definition and Properties

We now give the definition of our simulation relation. We note that for the deterministic case, the definition can be relaxed (in particular, in that case it is not required that Z^{-1} is a function), we omit the details and only treat the general case. In the following, for a binary relation Z and a state q, we write $Z(q)$ to denote the set $\{q' \mid (q, q') \in Z\}$.

Definition. Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be CGSs with state sets Q_1 and Q_2, the same set of agents, the same set of propositional variables, and n degrees of information. Then a relation $Z \subseteq Q_1 \times Q_2$ is a *probabilistic uniform strong alternating simulation for a coalition A from $\mathcal{C}_1$ to $\mathcal{C}_2$* if there are functions $\Delta^{1\to 2}_{(i,a,q_1,q_2)}$ such that for all $(q_1, q_2) \in Z$, all $i \in \{1, \ldots, n\}$, all agents $a \in A$, and all $A' \subseteq A$ we have

- *Propositional equivalence:* q_1 and q_2 satisfy the same propositional variables,
- for all (A', q_1)-moves c_1, the (A', q_2)-move c_2 with $c_2(a) = \Delta^{1\to 2}_{(i,a,q_1,q_2)}(c_1(a))$ has the *Forward Move Property* ($i = 1$, $\bar{i} = 2$) and the *Backward Move Property* ($i = 2$, $\bar{i} = 1$): For each $(\overline{A'}, q_i)$-move $c_i^{\overline{A'}}$, there is a $(\overline{A'}, q_{\bar{i}})$-move $c_{\bar{i}}^{\overline{A'}}$ such that for all $q_i' \in Q_i$, we have

$$\Pr\left(\delta(q_1, c_1 \cup c_1^{\overline{A'}}) = q_1'\right) = \Pr\left(\delta(q_2, c_2 \cup c_2^{\overline{A'}}) \in Z(q_1')\right).$$

- *Move Uniformity:* If $(q_1', q_2') \in Z$ with $q_1 \sim_{\mathsf{eq}_1^i(a)} q_1'$ and $q_2 \sim_{\mathsf{eq}_1^i(a)} q_2'$, then $\Delta^{1\to 2}_{(i,a,q_1,q_2)} = \Delta^{1\to 2}_{(i,a,q_1',q_2')}$,
- *Uniformity:* for all $a \in A$, and all $(q_1', q_2') \in Z$, if $q_2 \sim_{\mathsf{eq}_2^i(a)} q_2'$, then $q_1 \sim_{\mathsf{eq}_1^i(a)} q_1'$.
- *Knowledge Transfer:* if $q_1' \sim_{\mathsf{eq}_i^1(A')} q_1$, then there is some $q_2' \in Q_2$ such that $q_2' \sim_{\mathsf{eq}_i^2(A')} q_2$ and $(q_1', q_2') \in Z$.
- *Uniqueness:* For all $q_2 \in Q_2$, there is exactly one $q_1 \in Q_1$ with $(q_1, q_2) \in Z$ (i.e., $Z^{-1}\colon Q_2 \to Q_1$ is a function).

As mentioned above, the requirements in the definition are naturally needed for dealing with the probabilistic, incomplete-information setting with prior agreement:

- Propositional equivalence is obviously necessary if Z-related states should have the same properties.
- The requirement of the existence of $\Delta^{1\to 2}_{\ldots}$ ensures that in states related via Z, agents have a method to transfer their moves from $\mathcal{C}_1$ to $\mathcal{C}_2$ that does not depend on how other agents act in the same state, and (due to the move uniformity requirement) only depends on the equivalence class of the current state; this ensures

that an agent has enough information to determine the move suggested by applying the simulation. The existence of this function, together with the forward and backward move properties forms the "core" of the simulation: These requirements ensure that every move in one of the structures can be "mirrored" in the other such that for a potential follow-up state $q_1' \in Q_1$, the probability of reaching q_1' in $\mathcal{C}_1$ is the same as the one for reaching a state Z-related to q_1' in $\mathcal{C}_2$.

- Uniformity is a basic "compatibility" requirement between the involved equivalence relations and the relation Z. It implies that the function Z^{-1} can be "computed" by the agents in question given their available information: Given a state q_2 of $\mathcal{C}_2$, an agent a can determine the equivalence class of $Z^{-1}(q_2)$ with respect to its own indistinguishability relation.
- Knowledge transfer ensures that if $(q_1, q_2) \in Z$, and a group of principals cannot distinguish between q_1 and q_1' in the "simulation" $\mathcal{C}_1$, then in the "simulated world" $\mathcal{C}_2$ there also is a pair of states (related in the same way by Z) that the principals cannot distinguish. This ensures that in the simulation, there is as much "uncertainty" as in the "simulated world," i.e., knowledge is transferred from $\mathcal{C}_2$ to $\mathcal{C}_1$. When removing the knowledge operator from the language, requiring knowledge transfer is unnecessary.
- Z^{-1} needs to be a function to allow precise statements about the involved probabilities; this is used in the proof of Theorem 4.1 in an essential way. This requirement can be omitted in the deterministic case.

Simulations allow to transfer strategies in the canonical way. We show the following theorem:

THEOREM 4.1. *Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be CGSs, let A be a coalition, and Z a probabilistic uniform strong alternating simulation for A from $\mathcal{C}_1$ to $\mathcal{C}_2$. Then for all strategy choices S_1 for A in $\mathcal{C}_1$, there is a strategy choice S_2 for A in $\mathcal{C}_2$ such that for all formulas φ for $\mathcal{C}_1/\mathcal{C}_2$ with $ag(\varphi) \subseteq A$, and for all pairs $(q_1, q_2) \in Z$, it holds that $\mathcal{C}_1, \mathsf{S}_1, q_1 \models \varphi$ iff $\mathcal{C}_2, \mathsf{S}_2, q_2 \models \varphi$.*

The above theorem should *not* be read as stating that $\mathcal{C}_1$ and $\mathcal{C}_2$ are "strategically equivalent:" this is only the case when there are simulations in *both* directions. The construction used in the proof is the canonical one arising from the definitions; the complexity of the description of S_2 is the sum of the complexities of the descriptions of S_1, the simulation Z, and the associated move transfer function $\Delta^{1\to 2}_{\dots}$.

4.2 Discussion of Simulation Properties

Due to space reasons, we do not give a detailed example of a simulation, but only a generic one: We state the following trivial result, which on first sight may be surprising:

PROPOSITION 4.2. *For every CGS $\mathcal{C}$ and every coalition A, there is a probabilistic uniform strong alternating simulation for A from $\mathcal{C}$ to $\mathcal{C}^{hst}$.*

A (false) way of reading the above is that $\mathcal{C}$ and $\mathcal{C}^{hst}$ are strategically equivalent. However, this is completely incorrect: A probabilistic uniform strong alternating simulation allows to *transfer* a strategy choice (see Theorem 4.1), but since the translation is only in one direction, no equivalence is obtained. Hence Proposition 4.2 merely states that if a group of agents has agreed on a set of joint strategies to achieve their respective goals in the basic CGS $\mathcal{C}$, then they are free to apply the same strategies even if they are given the additional ability to remember the history of the game, thereby ignoring this capability. Stated in this way, Proposition 4.2 is entirely unsurprising. In particular, it does *not* state that with the additional capabilities, the agents could not achieve *more* in $\mathcal{C}^{hst}$ than in the original CGS $\mathcal{C}$.

As an example, consider the CGS $\mathcal{C}$ with three states q_0, q_1, q_2, where in q_0, the agent can freely choose whether the successor state should be q_1 or q_2, and from the latter two states, every move leads back to q_0. Assume that the only agent a in the game has complete information. Now consider the formula $\varphi = \langle\langle a\rangle\rangle \Box(\Diamond q_1 \wedge \Diamond q_2)$. There is no strategy (and hence no strategy choice) satisfying φ in $\mathcal{C}$: When following a (memoryless) strategy, a has to make the same move every time the game is in q_0, and hence only one of the states q_1, q_2 is visited infinitely often. In the history-dependent version $\mathcal{C}^{hst}$, the agent remembers the choice made last time, and can act accordingly. This not only shows that (as is well-known) history-dependent strategies are strictly stronger than memoryless ones, but also implies that in general, there is no probabilistic uniform strong alternating simulation from $\mathcal{C}^{hst}$ to $\mathcal{C}$, i.e., the converse of Proposition 4.2 does not hold.

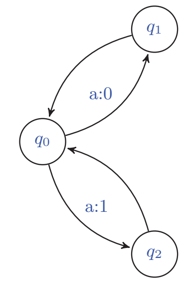

Figure 3: Example

We now show that simulations between CGSs canonically transfer to their history-dependent versions. This implies that it is sufficient to specify a simulation between history-dependent CGSs ${\mathcal{C}_1}^{hst}$ and ${\mathcal{C}_1}^{hst}$ on their memoryless cores $\mathcal{C}_1$ and $\mathcal{C}_2$ (if a simulation exists between these).

PROPOSITION 4.3. *If there is a probabilistic uniform strong alternating simulation for a coalition A from $\mathcal{C}_1$ to $\mathcal{C}_2$, then there also is one from ${\mathcal{C}_1}^{hst}$ to ${\mathcal{C}_2}^{hst}$.*

Again, the converse does not hold: For every CGS $\mathcal{C}$, the structures $\mathcal{C}^{hst}$ and ${\mathcal{C}^{hst}}^{hst}$ are essentially identical, in particular there is a probabilistic uniform strong alternating simulation from ${\mathcal{C}^{hst}}^{hst}$ to $\mathcal{C}^{hst}$, but due to the above there is not necessarily a simulation from $\mathcal{C}^{hst}$ to $\mathcal{C}$.

Finally, we mention a successful application of our notion of simulation: In parallel work, it has been shown that the game structures naturally arising when studying certain families of cryptographic protocols are infinite, but can be simulated by a finite structure. This result was then used to show that model checking for certain security properties of protocols is decidable, and that the involved strategies always have a finite representation.

5. COMPLEXITY AND DECIDABILITY

Strategy choices represent agreement of a coalition prior to a game: The coalition has to decide on a suitable strategy for every relevant goal, these strategies are then pooled in the strategy choice. Hence the "planning" of suitable strategies consists of determining a strategy choice achieving these goals for a given a CGS and a set of goals. In this section, we study the computational complexity of this problem. This situation is an example for the approach known as *planning as model checking*, see also [8]. Formally, we consider the following decision problems—depending on whether we allow all strategy choices or are only interested in maximal ones (see Section 3.3). We note that our decision algorithms are constructive.

Problem:	$\exists$Choice ($\exists$maxChoice)
Input:	A CGS $\mathcal{C}$, a state q of $\mathcal{C}$, a state-formula φ
Question:	Is there a ($\leq_\varphi$-maximal) strategy choice S for $ag(\varphi)$ in $\mathcal{C}$ such that $\mathcal{C}, \mathsf{S}, q \models \varphi$?

For studying the complexity of these problems, we assume that the transition function is specified as a complete table. For finite structures, the model checking problem is decidable, where the complexity in the deterministic case is considerably lower than in the probabilistic setting:

THEOREM 5.1. $\exists$maxChoice *and* $\exists$Choice *are*

1. PSPACE*-complete for deterministic structures,*
2. *solvable in* 3EXPTIME *and* 2EXPTIME*-hard for probabilistic structures.*

The above and the result from Schobbens [16] that model checking for memoryless ATL* is PSPACE-complete, shows that the model-checking complexity of our semantics comes at no additional cost compared to that of standard ATL* with memoryless strategies in the deterministic setting (recall that due to Proposition 3.1, our semantics are a generalization of memoryless ATL*). As expected due to results from Courcoubetis and Yannakakis [5], model checking for the probabilistic case is significantly more complex.

The situation is different for history-dependent strategies: Let $\exists\text{Choice}^{hst}$ and $\exists\text{Choice}^{hst}$ be the variations obtained from $\exists$Choice and $\exists$maxChoice by asking whether a corresponding strategy choice exists in $\mathcal{C}^{hst}$. These problems are undecidable, although the analogous problem is 2EXPTIME-complete for standard ATL*:

THEOREM 5.2. $\exists\text{Choice}^{hst}$ *and* $\exists\text{Choice}^{hst}$ *are undecidable for the deterministic and the probabilistic case.*

6. CONCLUSION

We have considered the situation in which a coalition A agrees on a set of strategies prior to the game, which are collected in a *strategy choice*, and can be identified and followed using only the information that is available to the agent. In the evaluation of the success probability of such a set of strategies, we adopted the usual pessimistic convention that the remaining agents follow their best-possible strategy, which is allowed to use information not available to the coalition A (including being allowed to be history-dependent). It would be interesting to relax this worst-case assumption and assume that the counter-coalition also has only bounded resources available.

Further, mixed strategies are an interesting issue. Note that a basic form of mixed strategies is possible in our semantics: One can introduce intermediate states in which the next move by an agent is chosen at random, where the probability distribution may be different in each state. Hence the agents may be given a certain amount of control over the distributions. However, a general treatment of mixed strategies remains open.

7. REFERENCES

[1] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.

[2] R. Alur, T. A. Henzinger, O. Kupferman, and M. Y. Vardi. Alternating refinement relations. In D. Sangiorgi and R. de Simone, editors, *CONCUR*, volume 1466 of *Lecture Notes in Computer Science*, pages 163–178. Springer, 1998.

[3] N. Bulling and W. Jamroga. What agents can probably enforce. *Fundamenta Informaticae*, 93(1-3):81–96, 2009.

[4] T. Chen and J. Lu. Probabilistic alternating-time temporal logic and model checking algorithm. In J. Lei, editor, *FSKD (2)*, pages 35–39. IEEE Computer Society, 2007.

[5] C. Courcoubetis and M. Yannakakis. Verifying temporal properties of finite-state probabilistic programs. In *FOCS*, pages 338–345. IEEE, 1988.

[6] A. Herzig and N. Troquard. Knowing how to play: uniform choices in logics of agency. In Nakashima et al. [14], pages 209–216.

[7] W. Jamroga. Some remarks on alternating temporal epistemic logic. In *Proceedings of Formal Approaches to Multi-Agent Systems (FAMAS 2003*, pages 133–140, 2004.

[8] W. Jamroga. Strategic planning through model checking of ATL formulae. In L. Rutkowski, J. H. Siekmann, R. Tadeusiewicz, and L. A. Zadeh, editors, *ICAISC*, volume 3070 of *Lecture Notes in Computer Science*, pages 879–884. Springer, 2004.

[9] W. Jamroga and T. Ågotnes. What agents can achieve under incomplete information. In Nakashima et al. [14], pages 232–234.

[10] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2-3):185–219, 2004.

[11] D. Kähler, R. Küsters, and T. Truderung. Infinite state AMC-model checking for cryptographic protocols. In *LICS*, pages 181—192. IEEE Computer Society, 2007.

[12] S. Kremer and J.-F. Raskin. A game-based verification of non-repudiation and fair exchange protocols. *Journal of Computer Security*, 11(3):399–430, 2003.

[13] F. Laroussinie, N. Markey, and G. Oreiby. On the expressiveness and complexity of ATL. *Logical Methods in Computer Science*, 4(2:7):1–25, 2008.

[14] H. Nakashima, M. P. Wellman, G. Weiss, and P. Stone, editors. *5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006), Hakodate, Japan, May 8-12, 2006*. ACM, 2006.

[15] H. Schnoor. Probabilistic ATL with incomplete information. Technical Report 0918, Institut für Informatik, Christian-Albrechts-Universität zu Kiel, 2009.

[16] P.-Y. Schobbens. Alternating-time logic with imperfect recall. *Electronis Notes in Theoretical Computer Science*, 85(2):82–93, 2004.