What Agents Can Achieve under Incomplete Information

Wojciech Jamroga
Institute of Computer Science
Clausthal University of Technology, Germany
wjamroga@in.tu-clausthal.de

Thomas Ågotnes
Department of Informatics
University of Bergen, Norway
agotnes@ii.uib.no

ABSTRACT

We propose a non-standard semantics for Alternating-time Temporal Logic with incomplete information, for which no commonly accepted semantics has been proposed yet. In our semantics, formulae are interpreted over sets of states rather than single states. We also propose a new epistemic operator for "constructive" knowledge, and we show that the new language is strictly more expressive than existing solutions, while retaining the same model checking complexity.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial IntelligenceMultiagent Systems; I.2.4 [**Artificial Intelligence**]: Knowledge Representation Formalisms and Methods—*Modal logic*

General Terms

Theory

Keywords

Alternating-time Temporal Logic, strategic ability, incomplete information, epistemic logic

1. INTRODUCTION

ATL [1] is probably the most important logic of strategic ability that has emerged in recent years. A combination of ATL and epistemic logic, called ATEL, was introduced in [9] to enable reasoning about agents acting under incomplete information. Still, it has been pointed out in several places that the meaning of ATEL formulae is somewhat counterintuitive. A number of ATEL updates were proposed to overcome this problem [3, 5, 8, 6, 10, 2], yet none of them seems the ultimate definitive solution. Our aim is to come up with a logic of ability under incomplete information which is both general and elegant.

In this paper, we propose a non-standard semantics for the logic of strategic ability and incomplete information. In the semantics, formulae are interpreted over *sets of states* rather than single states. This reflects the intuition that the "constructive" ability to enforce φ means that the agents in question have a single strategy that brings about φ for *all* possible initial situations – and not that a successful strategy exists for *each* initial situation (because those could be different strategies for different situations). To do it in a flexible and general way, the type of the satisfaction relation in our proposal forces one to specify the set of initial states explicitly. In consequence, we write $M, Q \models \langle\langle A\rangle\rangle\varphi$ to express the fact that A must have a strategy which is successful for all states in Q. We also propose a new epistemic operator for "practical" or "constructive" knowledge that yields the set of states for which a single evidence (i.e., a successful strategy) should be presented (instead of checking if the required property holds in each of the states separately, like standard epistemic operators do).

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.

2. WHAT AGENTS CAN ACHIEVE

In this section, we present a very brief overview of ATL and its extensions for agents with incomplete information.

Alternating-time Temporal Logic. ATL [1] can be understood as a generalization of the branching time temporal logic CTL, in which path quantifiers are replaced with so called *cooperation modalities*. Formula $\langle\langle A\rangle\rangle\varphi$, where A is a coalition of agents, expresses that A have a collective strategy to enforce φ. ATL formulae include temporal operators: "$\bigcirc$" ("in the next state"), $\Box$ ("always from now on") and $\mathcal{U}$ ("until"). The semantics can be defined using *concurrent game structures*, each including a set of agents $\mathbb{A}\mathrm{gt}$, states St, actions Act, and atomic propositions Π, plus a valuation $\pi : St \to \mathcal{P}(\Pi)$. Function $d : \mathbb{A}\mathrm{gt} \times St \to \mathcal{P}(Act)$ defines actions available to an agent in a state, and o is a transition function that assigns the outcome state $q' = o(q, \alpha_1, \ldots, \alpha_k)$ to state q and a tuple of actions $\langle\alpha_1, \ldots, \alpha_k\rangle$ that can be executed by $\mathbb{A}\mathrm{gt}$ in q. A *strategy* $s_a : St \to Act$ is a conditional plan that specifies what $a \in \mathbb{A}\mathrm{gt}$ is going to do for every possible situation. A *collective strategy* S_A is a tuple of strategies, one per agent from $A \subseteq \mathbb{A}\mathrm{gt}$. A *path* Λ in model M is an infinite sequence of states that can be effected by subsequent transitions. Function $out(q, S_A)$ returns the set of all paths that may result from agents A executing strategy S_A from state q onward. Informally speaking, $M, q \models \langle\langle A\rangle\rangle\varphi$ iff there is a collective strategy S_A such that φ holds for every $\Lambda \in out(q, S_A)$.

Alternating-time Temporal Epistemic Logic. ATEL [9] adds to ATL operators for representing agents' knowledge: $K_a\varphi$

reads as "agent a knows that φ". Additional operators $E_A\varphi$, $C_A\varphi$, and $D_A\varphi$ refer to "everybody knows", common knowledge, and distributed knowledge among the agents from A. Models for ATEL extend concurrent game structures with indistinguishability relations $\sim_1, ..., \sim_k \subseteq Q \times Q$ (one per agent) for modeling agents' uncertainty. Then: $M, q \models K_a\varphi$ iff φ holds for every q' such that $q \sim_a q'$.

Relations $\sim_A^E$, $\sim_A^C$ and $\sim_A^D$, used to model group epistemics, are derived from the individual relations of agents from A. First, $\sim_A^E$ is the union of relations $\sim_a$, $a \in A$. Next, $\sim_A^C$ is defined as the transitive closure of $\sim_A^E$. Finally, $\sim_A^D$ is the intersection of all the $\sim_a$, $a \in A$. Then, for $\mathcal{K} = C, E, D$: $M, q \models \mathcal{K}_A\varphi$ iff φ holds for every q' such that $q \sim_A^{\mathcal{K}} q'$.

Problems with ATEL. It has been pointed out in several places that the meaning of ATEL formulae is somewhat counterintuitive [3, 5, 6]. Most importantly, one would expect that an agent's ability to achieve φ should imply that the agent has enough control and knowledge to *identify* and *execute* a strategy that enforces φ. This problem is closely related to the distinction between knowledge *de re* and knowledge *de dicto*, well known in the philosophy of language [7]. Several variations on "ATL with incomplete information" have been proposed, yet none of them seems definitive. We summarize the most important proposals below.

ATL$_{ir}$. In the logic of ATL$_{ir}$ [8], cooperation modalities are presented with a subscript: $\langle\langle A\rangle\rangle_{ir}$ to indicate that they address agents with *imperfect information* and *recall*. Agents are required to use *uniform* strategies, i.e. ones that specify the same choices in indistinguishable states (if $q \sim_a q'$ then $s_a(q) = s_a(q')$). Formula $\langle\langle A\rangle\rangle_{ir}\varphi$ holds in M, q iff there is a uniform collective strategy S_A such that, for every $a \in A$, q' such that $q \sim_a q'$, and path $\Lambda \in out(q', S_A)$, we have that φ is true for Λ. In other words, there is a strategy such that *everybody in A knows* that executing this strategy will bring about φ. Note that it is not possible to express that A have common knowledge about the successful strategy, or that they can identify it if they share their knowledge etc.

Alternating-time Observational Temporal Logic. ATOL, proposed independently in [5], follows the same perspective as ATL$_{ir}$. However, it includes also epistemic modalities in the object language (like ATEL), and it offers a richer language of strategic operators to express subtle differences between various kinds of collective abilities. The reading of $\langle\langle A\rangle\rangle_{\mathcal{K}(\Gamma)}\varphi$ is: "group A has a (memoryless uniform) strategy to enforce φ, and agents Γ can identify the strategy as successful for A in the epistemic sense $\mathcal{K}$". That is, $M, q \models \langle\langle A\rangle\rangle_{\mathcal{K}(\Gamma)}\varphi$ iff there is S_A for every $a \in A$, q' such that $q \sim_\Gamma^{\mathcal{K}} q'$, and path $\Lambda \in out(q', S_A)$, we have that φ is true for Λ. We observe that model checking ATL$_{ir}$ and ATOL is NP-complete in the size of the model and the formula [8, 4, 5].

"Feasible ATEL*".* The update of ATEL from [6] extends ATEL with new modalities: $\langle\langle A\rangle\rangle^f$, $\langle\langle A\rangle\rangle^f_E$, $\langle\langle A\rangle\rangle^f_C$, $\langle\langle A\rangle\rangle^f_{K_a}$ and $\langle\langle A\rangle\rangle^f_{M_a}$, very similar to the ones of ATOL. The NP-completeness result carries over to "Feasible ATEL" (it subsumes ATL$_{ir}$ and can be seen as a subset of ATOL).

Other Approaches. Epistemic Temporal Strategic Logic [10] focuses on the concept of *undominated strategies*; in a way, $\langle\langle A\rangle\rangle\varphi$ in ETSL can be summarized as: "if A play *rationally* to achieve φ (meaning: they never play a dominated strategy), they will achieve φ". Another, very recent proposal [2] approaches the problem of strategic abilities within the framework of STIT (the logic of *seeing to it that*). We do not discuss these proposals further here due to lack of space.

In the original formulation of ATL, agents were assumed to have perfect recall of the game, in the sense that they could base their decisions on *sequences* of states rather than single states. Variants of ATL for perfect recall and incomplete information include ATL$_{iR}$ [8] and ATEL-R* [5]. However, as agents seldom have unlimited memory, and logics of strategic ability with incomplete information and perfect recall are believed to have undecidable model checking [1, 8], we do not investigate this kind of ability here.

3. NEW SEMANTICS FOR ABILITY AND KNOWLEDGE

ATOL covers more cases than ATL$_{ir}$ and "Feasible ATEL", and it is not committed to any notion of rationality (unlike ETSL). One major drawback of ATOL is that it vastly increases the number of modal operators necessary to express properties of agents. For team A, a whole family of cooperation modalities $\langle\langle A\rangle\rangle_{\mathcal{K}(\Gamma)}$ is used to specify who should identify the right strategy for A, in what way etc. It would be much more elegant to modify the semantics of "simple" cooperation modalities $\langle\langle A\rangle\rangle$ and/or epistemic operators, so that they can be composed into sufficiently expressive formulae. However, the property of a strategy being successful (under incomplete information) with respect to goal φ is *not* local to the current state; *the same* strategy must be successful in all possible "opening" states. In order to capture this feature of strategic ability, we change the type of the satisfaction relation $\models$, and define what it means for a formula φ to be satisfied in a *set of states* $Q \subseteq St$ of model M. To our best knowledge, nobody has used this kind of semantics yet. Moreover, we extend the language of ATEL with unary "constructive knowledge" operators $\mathbb{K}_a$, one for each agent a, that yield the set of states, indistinguishable from the current state from a's perspective. Constructive common, "everybody's" and distributed knowledge is formalized via operators $\mathbb{C}_A, \mathbb{E}_A$, and $\mathbb{D}_A$.

3.1 Language and Semantics

The language is defined formally as follows:

$$\varphi ::= p \mid \neg\varphi \mid {\sim}\varphi \mid \varphi \wedge \varphi \mid \langle\langle A\rangle\rangle \bigcirc \varphi \mid \langle\langle A\rangle\rangle \Box \varphi \mid \langle\langle A\rangle\rangle \varphi \,\mathcal{U}\, \varphi \mid C_A\varphi \mid E_A\varphi \mid D_A\varphi \mid \mathbb{C}_A\varphi \mid \mathbb{E}_A\varphi \mid \mathbb{D}_A\varphi.$$

The models are concurrent epistemic game structures again, and we consider only memoryless uniform strategies. Now, we define the notion of a formula φ being satisfied by a set of states Q in a model M, written $M, Q \models \varphi$. We will also write $M, q \models \varphi$ as a shorthand for $M, \{q\} \models \varphi$. Let $\mathrm{img}(q, \mathcal{R})$ be the image of state q with respect to relation $\mathcal{R}$, i.e. the set of all states q' such that $q\mathcal{R}q'$. Moreover, we use $out(Q, S_A)$ as a shorthand for $\cup_{q\in Q} out(q, S_A)$, and $\mathrm{img}(Q, \mathcal{R})$ as a shorthand for $\cup_{q\in Q}\mathrm{img}(q, \mathcal{R})$. The new semantics is given through the clauses below. Individual knowledge operators can be derived as: $K_a\varphi \equiv C_{\{a\}}\varphi$ and $\mathbb{K}_a\varphi \equiv \mathbb{C}_{\{a\}}\varphi$.

$M, Q \models p$ iff $p \in \pi(q)$ for every $q \in Q$;

$M, Q \models \neg\varphi$ iff $M, Q \not\models \varphi$;

$M, Q \models {\sim}\varphi$ iff $M, q \not\models \varphi$ for every $q \in Q$;

$M, Q \models \varphi \wedge \psi$ iff $M, Q \models \varphi$ and $M, Q \models \psi$;

$M, Q \models \langle\langle A \rangle\rangle \bigcirc \varphi$ iff there exists S_A such that, for every $\Lambda \in out(Q, S_A)$, we have that $M, \{\Lambda[1]\} \models \varphi$;[1]

$M, Q \models \langle\langle A \rangle\rangle \Box \varphi$ iff there exists S_A such that, for every $\Lambda \in out(Q, S_A)$ and $i \geq 0$, we have $M, \{\Lambda[i]\} \models \varphi$;

$M, Q \models \langle\langle A \rangle\rangle \varphi \mathcal{U} \psi$ iff there exists S_A such that, for every $\Lambda \in out(Q, S_A)$, there is $i \geq 0$ for which $M, \{\Lambda[i]\} \models \psi$ and $M, \{\Lambda[j]\} \models \varphi$ for every $0 \leq j < i$;

$M, Q \models \mathcal{K}_A \varphi$ iff $M, q \models \varphi$ for every $q \in \mathrm{img}(Q, \sim_A^{\mathcal{K}})$,

$M, Q \models \hat{\mathcal{K}}_A \varphi$ iff $M, \mathrm{img}(Q, \sim_A^{\mathcal{K}}) \models \varphi$ (where $\hat{\mathcal{K}} = \mathbb{C}, \mathbb{E}, \mathbb{D}$ and $\mathcal{K} = C, E, D$, respectively).

3.2 Expressing Agents' Strategic Abilities

$M, q \models \mathbb{K}_a \langle\langle a \rangle\rangle \varphi$ expresses the fact that a has a single strategy that enforces φ from *all* states indiscernible from q, instead of stating that φ can be achieved from *every* such state *separately*. Note that the latter property is very much in the spirit of standard epistemic logic, and indeed can be captured with the standard knowledge operator (via $K_a \langle\langle a \rangle\rangle \varphi$). More generally, the first kind of formulae refers to *having a strategy "de re"* (i.e. having a successful strategy and knowing the strategy), while the latter refers to *having a strategy "de dicto"* (i.e. only knowing that *some* successful strategy is available; cf. [5]). Note also that the property of having a winning strategy for the current state (but not necessarily even knowing *about* it) is simply expressed with $\langle\langle a \rangle\rangle \varphi$. Capturing different ability levels of coalitions is analogous, with various "epistemic modes" of collective recognizing the right strategy.

THEOREM 1. *Let φ be a formula of* ATL$_{ir}$, ATOL *or "Feasible* ATEL*", and let tr be as follows:*

$$tr(p) = p \qquad tr(\neg\varphi) = \neg tr(\varphi)$$
$$tr(\varphi \wedge \psi) = tr(\varphi) \wedge tr(\psi) \qquad tr(\bigcirc\varphi) = \bigcirc tr(\varphi)$$
$$tr(\Box\varphi) = \Box tr(\varphi) \qquad tr(\varphi \mathcal{U} \psi) = tr(\varphi) \mathcal{U} tr(\psi)$$
$$tr(\langle\langle A \rangle\rangle_{ir} \varphi) = \mathbb{E}_A \langle\langle A \rangle\rangle \varphi \qquad tr(\langle\langle A \rangle\rangle_{\mathcal{K}(\Gamma)} \varphi) = \hat{\mathcal{K}}_\Gamma \langle\langle A \rangle\rangle \varphi$$
$$tr(\langle\langle A \rangle\rangle^f \varphi) = \langle\langle A \rangle\rangle \varphi \qquad tr(\langle\langle A \rangle\rangle^f_{\mathcal{K}} \varphi) = \hat{\mathcal{K}}_A \langle\langle A \rangle\rangle \varphi$$
$$tr(\langle\langle A \rangle\rangle^f_{K_b} \varphi) = \mathbb{K}_b \langle\langle A \rangle\rangle \varphi \qquad tr(\langle\langle A \rangle\rangle^f_{M_b} \varphi) = \neg K_b \neg \langle\langle A \rangle\rangle \varphi$$
$$tr(\mathcal{K}_A \varphi) = \mathcal{K}_A tr(\varphi)$$

where $\mathcal{K} = C, E, D$ and $\hat{\mathcal{K}} = \mathbb{C}, \mathbb{E}, \mathbb{D}$, respectively. Then:

$$M, q \models \varphi \text{ iff } M, q \models tr(\varphi).$$

REMARK 2. *The new language is strictly more expressive than* ATL$_{ir}$, ATOL *etc.: for example, formula $\mathbb{E}_A \mathbb{E}_A \langle\langle A \rangle\rangle \varphi$ cannot be expressed in any of the former logics.*

3.3 Model Checking

We define *general model checking* as the problem that asks whether formula φ holds in model M and *set of states* Q. Let $mctl(\varphi, M)$ be a CTL model checker that returns the set of all states that satisfy φ in M. Below, we sketch algorithm $mcheck(\varphi, M, Q)$ that returns "yes" if $M, Q \models \varphi$ and "no" otherwise, running in nondeterministic polynomial time.

- Cases $\varphi \equiv p, \varphi \equiv \neg\psi, \varphi \equiv {\sim}\psi, \varphi \equiv \psi_1 \wedge \psi_2, \varphi \equiv \mathcal{K}_A \psi$: straightforward (proceed as usually).
- Case $\varphi \equiv \hat{\mathcal{K}}_A \psi$: return $mcheck(\psi, M, \mathrm{img}(Q, \sim_A^{\mathcal{K}}))$.
- Case $\varphi \equiv \langle\langle A \rangle\rangle \bigcirc \psi$: run $mcheck(\psi, M, q)$ for every $q \in St$, and label the states in which the answer was "yes" with a new proposition p. Then, guess the strategy of A, and "trim" model M by removing all the transitions inconsistent with the strategy (yielding a sparser model M'). Return "yes" iff $Q \subseteq mctl(\mathsf{A} \bigcirc \mathsf{p}, M)$. [For other temporal operators: analogous.]

[1] By $\Lambda[i]$, we denote the ith position on Λ (starting from 0).

Note that all the relevant strategies can be guessed *beforehand*, as a single complex (but still polynomial) witness (cf. [4]), which gives us the following result:

THEOREM 3. *General model checking for our logic is NP-complete in the size of the model and the formula.*

4. CONCLUSIONS

In this paper, we propose a non-standard semantics for the modal logic of strategic ability under incomplete information, in which formulae are interpreted over *sets of states* rather than single states. Moreover, we introduce new epistemic operators for "constructive" knowledge. It turns out that, in this new semantics, simple cooperation modalities $\langle\langle A \rangle\rangle$ can be combined with "constructive" epistemic operators into sufficiently expressive formulae. Indeed, the new logic *is strictly more expressive than most existing* ATL *versions for incomplete information*, while it retains the same model checking complexity as the least costly of them. The philosophical dimension of constructive knowledge is also natural: the constructive knowledge operators capture the notion of *knowing "de re"*, while the standard epistemic operators refer to *knowing "de dicto"*. We believe that we have finally obtained a satisfying logic of agents' strategies under uncertainty, and at the same time came up with novel, meaningful epistemic operators that capture important properties of the interaction between knowledge, action and ability.

5. REFERENCES

[1] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.

[2] A. Herzig and N. Troquard. Knowing how to play: Uniform choices in logics of agency. In *Proceedings of AAMAS-06*, 2006.

[3] W. Jamroga. Some remarks on alternating temporal epistemic logic. In *Proceedings of FAMAS 2003*, pages 133–140, 2003.

[4] W. Jamroga and J. Dix. Model checking strategic abilities of agents under incomplete information. In *Proceedings of ICTCS 2005*, pages 295–308, 2005.

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

[6] G. Jonker. Feasible strategies in Alternating-time Temporal Epistemic Logic. Master thesis,2003.

[7] W. Quine. Quantifiers and propositional attitudes. *Journal of Philosophy*, 53:177–187, 1956.

[8] P. Y. Schobbens. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science*, 85(2), 2004.

[9] W. van der Hoek and M. Wooldridge. Tractable multiagent planning for epistemic goals. In *Proceedings of AAMAS-02*, pages 1167–1174, 2002.

[10] S. van Otterloo and G. Jonker. On Epistemic Temporal Strategic Logic. In *Proceedings of LCMAS*, pages 35–45, 2004.