

Agents With Truly Perfect Recall in Alternating-Time Temporal Logic

(Extended Abstract)

Nils Bulling
Department of Informatics
Clausthal University of
Technology
bulling@in.tu-
clausthal.de

Wojciech Jamroga
Computer Science and
Communication
and Interdisciplinary Centre on
Security and Trust
University of Luxembourg
wojtek.jamroga@uni.lu

Matei Popovici
Computer Science
Department
POLITEHNICA University of
Bucharest
matei.popovici@cs.pub.ro

Categories and Subject Descriptors

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

General Terms

Theory, Verification, Languages

Keywords

alternating-time temporal logic, multi-player games, reasoning about knowledge and time

1. INTRODUCTION

Mathematical logic and game theory are important for theoretical foundations of multi-agent systems (MAS). The *alternating-time temporal logic* ATL* [2] is a formalism where the two traditions meet. The logic combines the versatile specification and modeling framework of temporal logic with basic game-theoretic notions to reason about agents' choices and their potential outcomes. As a consequence, ATL* provides a useful framework for the specification, verification and reasoning about properties of MAS. However, in order to verify/reason about MAS, we need to capture the desired properties in the right way.

Many semantic variants of ATL* have been proposed which reflect different assumptions about agents' capabilities. For example, the agents can have perfect or imperfect information about the state of the game, and perfect or imperfect recall of past observations [6, 5, 4]. Also, strategies can come with or without long-term commitment [1, 3], etc. In this paper we look closer at how agents remember the past in the semantics of ATL*. Typically, one distinguishes between agents with no memory beyond what is encapsulated in the current state of the system, and ones with perfect recall. Agents of the former kind base their decisions only on what they see of the current situation; agents of the latter kind—on the whole history of observations. We show that the semantics of perfect recall in ATL* is problematic: agents are assumed to forget their observations when they proceed to re-

alize a sub-goal in the game. As an example, consider the formula $\langle\langle b, c \rangle\rangle \diamond \langle\langle a, b \rangle\rangle \bigcirc \text{married}_{ab}$ which expresses that Bob (b) and Charles (c) have a joint strategy to ensure that, at some point in the future, Alice (a) and Bob will be able to get married. Agents' abilities rely on their knowledge; in case of perfect recall, one would assume that each agent can use all its past observations to determine its subsequent actions. However, the semantics of ATL* interprets the subformula $\langle\langle a, b \rangle\rangle \bigcirc \text{married}_{ab}$ in the original model. This amounts to assuming that Bob, when looking for his best strategy to make $\langle\langle a, b \rangle\rangle \bigcirc \text{married}_{ab}$ true, must ignore (or forget) all the observations that he has made while executing his strategy for $\langle\langle b, c \rangle\rangle \diamond \langle\langle a, b \rangle\rangle \bigcirc \text{married}_{ab}$. The same remark also applies to Alice. As a remedy, we propose a new semantics for ATL*, which avoids the “forgetting” phenomenon.

2. ATL*: WHAT AGENTS CAN ACHIEVE

ATL* [2] can be seen as a generalization of the branching time logic CTL* where path quantifiers E, A are replaced by *cooperation modalities* $\langle\langle A \rangle\rangle$. Formula $\langle\langle A \rangle\rangle \gamma$ expresses that group A has a *collective strategy* to enforce the (temporal) property γ where γ can include the temporal operators \bigcirc (“next”), and \mathcal{U} (“until”), as well as cooperation modalities. We interpret ATL* formulae over *imperfect information concurrent game structures* (iCGS). A *strategy* of agent a is a conditional plan that specifies what a is going to do in each situation. A *perfect information strategy* (I -strategy for short) is a function that maps possible finite histories of the game (i.e., finite sequences of states of the system) to available actions. An *imperfect information strategy* (i -strategy) must be additionally *uniform*, in the sense that it specifies the same action for indistinguishable histories. We refer to [2, 6] for formal details, and instead present the intuitions on the following example.

EXAMPLE 1 (SHELL GAME). Consider the model M_1 from Figure 1 which depicts a simple version of the shell game. There are two players: the shuffler (s) and the guesser (g). Initially, the shuffler places a ball in one of two shells, and possibly moves it from shell to shell (q_0 and q'_0). The shells are open, and the guesser can see the location of the ball. When the guesser says “stop!”, the shuffler turns the shells over, so that the ball becomes hidden; the shuffler can also decide to stop shuffling on his own. The guesser wins if he picks up the shell containing the ball. Clearly, $M_1, q_1 \models_I \langle\langle g \rangle\rangle \diamond \text{win}$: under perfect information (indicated by \models_I), the guesser can win by choosing the left shell in q_1 . On the other hand, $M_1, q_1 \models_i \neg \langle\langle g \rangle\rangle \diamond \text{win}$: under imperfect information (indicated by \models_i), the guesser cannot distinguish q_1 and q'_1 and has

Appears in: Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, May 5-9, 2014, Paris, France.

Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

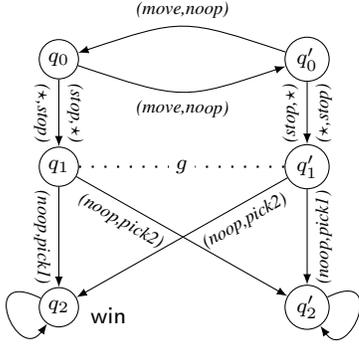


Figure 1: The iCGS M_1 describing the shell game.

no uniform strategy that succeeds from both states. Finally, if the game begins in q_0 , the guesser can win ($M_1, q_0 \models_i \langle\langle g \rangle\rangle \diamond \text{win}$) by using the uniform perfect recall strategy s_g : “play stop after histories ending with q_0 or q_0' , then play pick1 after histories ending with q_0q_1 , and play pick2 after ones ending in $q_0'q_1$.”

3. TRULY PERFECT RECALL

The standard semantics of ATL^* displays a peculiar “forgetting” phenomenon, even for agents with perfect recall. In formulae containing nested cooperation modalities, such as $\langle\langle a \rangle\rangle \diamond \langle\langle b \rangle\rangle \square p$, it requires that b starts collecting observations from scratch when executing his strategy for the subgoal $\square p$. This leads to counterintuitive effects, as the following example shows.

EXAMPLE 2 (FORGETTING IN PERFECT RECALL). *Recall that $M_1, q_0 \models_i \langle\langle g \rangle\rangle \diamond \text{win}$, that is, the guesser has a uniform strategy to win the shell game starting in q_0 . On the other hand, $M_1, q_1 \models_i \neg \langle\langle g \rangle\rangle \diamond \text{win}$. Since the shuffler in q_0 can easily enforce the next state to be q_1 , we have $M_1, q_0 \models_i \langle\langle s \rangle\rangle \circ \neg \langle\langle g \rangle\rangle \diamond \text{win}$. Thus, in M_1, q_0 , the guesser has the ability to win no matter what the shuffler does, and at the same time the shuffler has a strategy to deprive the guesser of the ability no matter what the guesser does!*

As a remedy, we propose a variant of ATL^* with “no forgetting” or “truly perfect recall”. To this end, we make the history explicit in the semantics, and update it as the temporal operators in the formula are evaluated.

3.1 ATL^* with Truly Perfect Recall

We introduce the no forgetting semantics \models_x^{nf} , $x \in \{i, I\}$ for the language of ATL^* . Formulae are interpreted over triples consisting of a model, a path and an index $k \in \mathbb{N}_0$ which indicates the current state of the path. Intuitively, the subhistory of the path up to k encodes the past, and the subpath starting after k , the future. The crucial part of this semantics is that the agents always remember the sequence of past events—they can learn from these past events.

We assume the function $\text{plays}_M^x(h, s_A)$ returns the set of relevant paths¹ (i.e. infinite sequences of states resulting from subsequent transitions) for strategy s_A executed from h on. For perfect information, $\text{plays}_M^I(h, s_A)$ returns the set of all paths that can occur when s_A is executed after the initial history h has taken place. For imperfect information, $\text{plays}_M^i(h, s_A)$ also includes the paths that A think might occur, i.e., ones starting from histories indistinguishable from h for A .

Let M be an iCGS, λ a path in M , $k \in \mathbb{N}_0$, and $x \in \{i, I\}$. The semantics of “ ATL^* with truly perfect recall” is defined as

¹We use λ to denote paths and write $\lambda[i]$ to refer to the i th state on λ and $\lambda[0, i]$ to the finite prefix of λ up to position i .

follows (the clauses for propositions, negation, and conjunction are standard and are omitted due to space limitations):

$M, \lambda, k \models_x^{nf} \langle\langle A \rangle\rangle \varphi$ iff there exists an x -strategy s_A such that, for all paths $\lambda' \in \text{plays}_x^x(\lambda[0, k], s_A)$, $M, \lambda', k \models_x^{nf} \varphi$;

$M, \lambda, k \models_x^{nf} \circ \varphi$ iff $M, \lambda, k + 1 \models_x^{nf} \varphi$;

$M, \lambda, k \models_x^{nf} \varphi_1 \mathcal{U} \varphi_2$ iff there exists $i \geq k$ such that $M, \lambda, i \models_x^{nf} \varphi_2$ and $M, \lambda, j \models_x^{nf} \varphi_1$ for all $k \leq j < i$.

Our new semantics differs from the standard semantics of ATL^* only in that it keeps track of the history by incorporating it into λ and plays_x^x . This affects the set of paths that are relevant when evaluating a strategy: instead of starting with the current state of the game (as in the standard semantics) we look at paths λ that describe the play from the very beginning.

We illustrate the new semantics by the following example.

EXAMPLE 3 (SHELL GAME CTD.). *Consider the pointed iCGS (M_1, q_0) again. Whatever the shuffler does in the first step, g can adapt his choice to win the game. In particular, the history-based uniform strategy s_g from Example 1 can be used to demonstrate that for all $\lambda \in \text{plays}_{M_1}^i(q_0, s_g)$ —for every strategy of s —we have $M_1, \lambda, 0 \models_i^{nf} \circ \langle\langle g \rangle\rangle \diamond \text{win}$. As a consequence, $M_1, q_0 \models_i^{nf} \neg \langle\langle s \rangle\rangle \circ \neg \langle\langle g \rangle\rangle \diamond \text{win}$.*

4. CONCLUSION

In this paper we propose a semantics for ATL^* that allows to model agents with perfect recall that *truly* recall all past events. Although it sounds paradoxical, the standard perfect recall semantics of ATL^* does not guarantee this property: agents forget past events when strategic operators are nested.

Due to space constraints, we cannot include any technical results in this extended abstract. We only mention in passing that the “forgetting” and “no-forgetting” semantics yield distinctly different logics of ability; in particular, it can be shown that both semantics induce logics of incomparable expressive and distinguishing power under incomplete information. Moreover, validities for “truly perfect recall” refine those of “standard perfect recall” (also cf. [4]) Finally, it can be shown that the results on expressivity and validities carry over to semantics with persistent strategies [1, 3].

Acknowledgements. The authors thank Thomas Ågotnes and Valentin Goranko for discussions and comments. Wojciech Jamroga acknowledges the support of the FNR (National Research Fund) Luxembourg under project GALOT – INTER/DFG/12/06.

5. REFERENCES

- [1] T. Ågotnes, V. Goranko, and W. Jamroga. Alternating-time temporal logics with irrevocable strategies. In D. Samet, editor, *Proceedings of TARK XI*, pages 15–24, 2007.
- [2] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.
- [3] T. Brihaye, A. Da Costa Lopes, F. Laroussinie, and N. Markey. ATL with strategy contexts and bounded memory. In *Proceedings of LFCS*, volume 5407 of *Lecture Notes in Computer Science*, pages 92–106. Springer, 2009.
- [4] N. Bulling and W. Jamroga. Comparing variants of strategic ability. *Journal of Autonomous Agents and Multi-Agent Systems*, 28(3):474–518, 2014.
- [5] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2–3):185–219, 2004.
- [6] P. Y. Schobbens. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science*, 85(2):82–93, 2004.