Some Remarks on Alternating Temporal Epistemic Logic

Corrected version: July 2003

Wojciech Jamroga

Parlevink Group, University of Twente, Netherlands
Institute of Mathematics, University of Gdansk, Poland
`jamroga@cs.utwente.nl`
`http://www.cs.utwente.nl/~jamroga`

Abstract. Alternating-Time Epistemic Logic (ATEL) has been proposed recently for specification and verification of multi-agent systems properties in situations of incomplete information. Some minor problems with ATEL semantics are characterized in this paper: an agent can 'access' the current state of the whole system when making up his strategy (even when he should be uncertain about the state); moreover, no explicit representation of actions in ATEL models makes some natural situations harder to model. A few small changes are suggested in consequence, mostly to make ATEL models consistent with the incomplete information assumption.

`Keywords:` multiagent systems, temporal logic, transition systems, games with incomplete information.

Beware: the solutions proposed in this paper turned out not to be final and completely satisfying. Several new problems and research questions concerning ATEL semantics have been identified since the publication of the paper. In this version, some technical mistakes are corrected; however, the more fundamental issues are still left untouched. A new paper that tackles both old and newly identified problems is under preparation now.

1 Introduction: Alternating-Time Logics

Two formalisms based on temporal logic have been proposed recently to tackle the verification of multi-agent systems properties. The first one was Alternating-Time Logic (ATL and ATL* [1–3]), which offered a possibility of expressing the capabilities of autonomous agents in a way similar to CTL and CTL*. The second, Alternating-Time Epistemic Logic (ATEL and ATEL* [4]), enriched the picture with an epistemic component.

There are some minor problems with ATEL semantics: an agent can 'access' the current state of the whole system when making up his strategy (even when he should be uncertain about the state), and no explicit representation of actions in ATEL models makes things a little bit difficult. The problems are characterized in more detail in section 2. In section 3 some solutions are proposed.

1.1 ATL and ATEL

ATL is an extension of CTL in which a class of *cooperation modalities* $\langle\langle A\rangle\rangle$ ($A \subseteq \Sigma$, where Σ is the set of all agents or players) replace the simple path quantifiers $\exists$ and $\forall$. The common-sense reading of $\langle\langle A\rangle\rangle\Phi$ is: "the group of agents A have a collective strategy to bring about Φ regardless of what all the other agents do". The original CTL operators $\exists$ and $\forall$ can be expressed in ATL with $\langle\langle\Sigma\rangle\rangle$ and $\langle\langle\emptyset\rangle\rangle$, respectively, but between both extremes we can express much more about the abilities of particular agents and groups of agents. ATL* extends CTL* in the same way ATL extends CTL. The dual operator $[\![A]\!]$ can be defined in the usual way as $[\![A]\!]\Phi \equiv \neg\langle\langle A\rangle\rangle\neg\Phi$, meaning that A can't avoid Φ on their own.

ATEL (ATEL*) adds to ATL (ATL*, respectively) operators for representing knowledge in the world of incomplete information. $K_a\varphi$ reads as "agent a knows that φ".

Computational complexity of model checking is the main virtue of ATL and ATEL. For both ATL and ATEL, the problem of checking whether φ is valid in M can be solved in polynomial time. Thus the semantic structures in which the formulae are verified are even more important than usually.

1.2 Concurrent Game Structures and AETS

A model for ATL is defined as a *concurrent game structure* [3]:

$$S = \langle k, Q, \Pi, \pi, d, \delta\rangle$$

where k is a natural number defining the set of players (the players are identified with numbers $1, ..., k$), Q is the set of states of the system, Π is the set of atomic propositions (observables), and $\pi : Q \to 2^{\Pi}$ is the observation function, specifying which propositions are true in which states. The decisions available to player a at state q are labeled with natural numbers $1, ..., d_a(q)$; finally, a complete tuple of decisions $\langle j_1, ..., j_k\rangle$ in state q implies a deterministic transition according to the transition function $\delta(q, j_1, ..., j_k)$.[1]

Models for ATEL – *alternating epistemic transition systems* (AETS) – add epistemic accessibility relations $\sim_1, ..., \sim_k \subseteq Q \times Q$ for expressing agents' beliefs [4]:

$$S = \langle \Sigma, Q, \Pi, \pi, \sim_1, ..., \sim_k, \delta\rangle$$

Since the definition of ATEL from [4] is based on the previous version of ATL [2], actions are not represented explicitly here, function d is absent, and Σ is an

[1] it should be noted that at least three different versions of ATL have been proposed by Alur and colleagues over the course of the last 6 years, each with a slightly different definition of the semantic structure. The earliest version [1] includes definitions for a synchronous turn-based structure and an asynchronous structure in which every transition is owned by a single agent. [2] offers general games structures with no action labels and more sophisticated transition function. In [3] function d is introduced and δ simplified; moreover, an arbitrary finite set of agents Σ is replaced with set $\{1, ..., k\}$. All of this may lead to some confusion.

arbitrary finite, non-empty set of players. The system transition function $\delta : Q \times \Sigma \to 2^{2^Q}$ is meant to encode all the choices available to agents at each state. Now $\delta(q, a) = \{Q', Q'', ...\}$ $(Q', Q'', ... \subseteq Q)$ defines the possible outcomes of a's decisions at state q, and the choices are identified with the outcomes. The resulting transition is assumed to be the intersection of choices from all the agents: $Q_1 \cap ... \cap Q_k$, $Q_i \in \delta(q, a_i)$. Since the system is required to be deterministic again (given the state and the agents' choices), $Q_1 \cap ... \cap Q_k$ must always be a singleton.

1.3 Agents' Strategies and Semantics of Cooperation Modalities

In a concurrent game structure, a strategy for an agent a is a function $f_a : Q^+ \to \mathbb{N}$ such that $f_a(\lambda q) \leq d_a(q)$, i.e. $f_a(\lambda q)$ is a valid decision in q. The function specifies a's decisions for every possible (finite) history of system transitions. The set of all possible (infinite) computations from q, consistent with a tuple of strategies $F_A : A \to (Q^+ \to \mathbb{N})$ – one strategy for each agent from $A \subseteq \Sigma$ – is denoted with $out(F_A, q)$. Now, informally speaking, $S, q \models \langle\langle A \rangle\rangle \Phi$ iff there exists a collective strategy F_A such that Φ is satisfied for all computations from $out(F_A, q)$. In other words, no matter what the rest of the agents decides to do, the agents from A have a way of enforcing Φ.

In AETS, a strategy for an agent a is a function $f_a : Q^+ \to 2^Q$, giving a choice for every finite history of possible transitions. The strategy must be consistent with the choices available to a, i.e. $f_a(\lambda q) \in \delta(q, a)$. Again (informally) $S, q \models \langle\langle A \rangle\rangle \Phi$ iff there exists a collective strategy F_A such that Φ is satisfied for all computations from $out(F_A, q)$.

2 Problems with AETS Transitions

Something seems to be lacking in the definition of a valid strategy for an agent in AETS. When defining a strategy, the agent can make his choices for every state independently. This is not feasible in a situation of incomplete information if the strategy is supposed to be deterministic: if a can't recognize whether he's in situation s_1 or s_2, he cannot plan to proceed with one action in s_1, and another in s_2. It's very much like with the information sets from von Neumann and Morgenstern [6]: for every state in an information set the same action must be chosen within a strategy.

The following example can be considered: two agents play a very simple card game with the deck consisting of Ace, King and Queen (A, K, Q). It is assumed that A beats K, K beats Q, but Q beats A. First a_1 and a_2 are given a card. Then a_1 can trade his card for the one remaining in the deck, or he can keep the current one. The player with the better card wins the game. A turn-based synchronous AETS for the game is shown on figure 1, with the environment *env* dealing the cards. Right after the cards are given, both a_1 and a_2 don't know what is the hand of the other player; for the rest of the game the players have

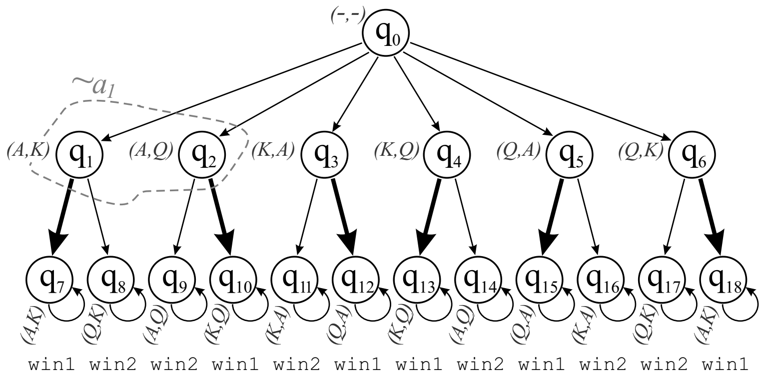

$\Sigma = \{env, a_1, a_2\}$
$\delta(q_0, env) = \{\{q_1\}, ..., \{q_6\}\}$
$\delta(q_0, a_1) = \delta(q_0, a_2) = \{\{q_1, ..., q_6\}\}$
$\delta(q_1, a_1) = \{\{q_7\}, \{q_8\}\}$ etc.
$\delta(q_1, env) = \delta(q_1, a_2) = \{\{q_7, q_8\}\}$ etc.
$\delta(q_7, env) = \delta(q_7, a_1) = \delta(q_7, a_2) = \{\{q_7\}\}$ etc.

$q_0 \sim_{a_1} q_0,\ q_0 \sim_{a_2} q_0$
$q_1 \sim_{a_1} q_1,\ q_1 \sim_{a_1} q_2,\ q_2 \sim_{a_1} q_2$ etc.
$q_1 \sim_{a_2} q_1,\ q_1 \sim_{a_2} q_6,\ q_6 \sim_{a_2} q_6$ etc.
$q_7 \sim_{a_1} q_7,\ q_7 \sim_{a_2} q_7$
$q_7 \sim_{a_1} q_7,\ q_7 \sim_{a_2} q_7$ etc.

Fig. 1. Epistemic transition system for the card game. For every state, the players' hands are described. The thick arrows indicate a_1's winning strategy.

complete information about the state. Two atomic propositions: `win1` and `win2` enable to recognize the final winner.

Note that $q_0 \models \langle\langle a_1 \rangle\rangle \Diamond win1$, although it should definitely be false for this game! Of course, a_1 may *happen* to win, but he doesn't have the *power* to bring about winning because he has no way of recognizing the right decision until it's too late. Even if we ask about whether the player can *know* that he has a winning strategy, it doesn't help: $K_{a_1} \langle\langle a_1 \rangle\rangle \Diamond win1$ is satisfied in q_0, too, because for all $q \in Q$ such that $q_0 \sim_{a_1} q$ we have $q \models \langle\langle a_1 \rangle\rangle \Diamond win1$.

This calls for a constraint like the one from von Neumann and Morgenstern: if $q \sim_a q'$ and the history of previous transitions is the same, then a strategy f_a must specify the same action for both q and q'. Unfortunately, it's impossible to express this constraint because actions are identified with their outcomes in AETS. Note, however, that the same action started in two different states seldom generates the same result: if a_1 trades his Ace in q_1, the system moves to q_8 and a_1 loses the game; if he trades the card in q_2, the system moves to q_{10} and he wins. Still he can't discriminate trading the Ace in both situations. Thus, some relation of 'subjective unrecognizability' is necessary over the agents' choices to tell which decisions will be considered the same in which states. Probably the easiest way to accomplish this is to provide the decisions with explicit labels – the way it has been done in the latest version of ATL – and assume that the choices with the same label represent the same action from the agent's subjective point of view. This kind of solution fits also well in the tradition of game theory.

Identifying actions with their outcomes may make things unnecessarily complicated even for AETS with no incomplete information. Consider a system with a single binary variable `x`. There are two processes: the controller (or server) s can enforce that the variable retains its value in the next step, or let the client change the value. The client c can request the value of `x` to be 0 or 1. Every player has complete information about the current state, i.e. $q \sim_a q'$ iff $q = q'$. The players proceed with their choices simultaneously – they don't know the other player's decision until the transition is done. The states and possible transitions of the system as a whole are shown on figure 2. There are two propositions available to observe the value of `x`: "`x=0`" and "`x=1`" (note: these are just atomic propositions, `=` is not the equality symbol here).

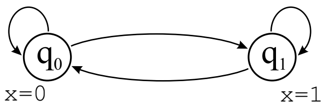

Fig. 2. Transitions of the variable controller/client system

Let's try to have a look at $\delta(q_0, s)$ first. The controller should have a choice to "enforce no change" with deterministic outcome of $\{q_0\}$, so $\{q_0\} \in \delta(q_0, s)$. Now, for all $Q' \in \delta(q_0, c)$, q_0 must be in Q' because $\{q_0\} \cap Q'$ has to be a singleton. Thus $\{q_1\} \notin \delta(q_0, c)$, and if we want to make the transition from q_0 to q_1 possible at all then $\{q_0, q_1\} \in \delta(q_0, c)$. Now $\{q_0, q_1\} \notin \delta(q_0, s)$ because $\{q_0, q_1\} \cap \{q_0, q_1\}$ is no singleton, so $\{q_1\} \in \delta(q_0, s)$ – otherwise the system still never proceeds from q_0 to q_1. In consequence, $\{q_0\} \notin \delta(q_0, c)$, because $\{q_1\} \cap \{q_0\}$ isn't a singleton either. The resulting transition function for q_0 is:

$$\delta(q_0, s) = \{\{q_0\}, \{q_1\}\}$$
$$\delta(q_0, c) = \{\{q_0, q_1\}\}$$

Unfortunately, it is easy to show that $q_0 \models \langle\langle s \rangle\rangle \bigcirc \texttt{x=1}$ for this model, and this is obviously wrong with respect to the original description of the system.

This doesn't necessarily mean that no AETS can be made up for this problem, having added some extra states and transitions. Indeed, for the transition system on figure 3, $q_0 \models \neg\langle\langle s \rangle\rangle \bigcirc \texttt{x=1}$, $q_0 \models \langle\langle s \rangle\rangle \bigcirc \texttt{x=0}$ and so on. The states reflect the value of `x` *and* the last transition made: q_0 is for "x=0 by s's force", q_0' for "x=0 by c's request" etc.[2] The point is, however, that ATEL is aimed mostly for model checking – so we're not going to search for a model in which some formula is satisfied. Actually, it's the other way round: we must come up with

[2] in fact, this shows how we can design a transition system in a general case, too: first we identify the set of 'pure system states' St (tuples of valuations of all the variables in all the involved processes, for instance) and sets of possible actions for all the agents: $Act_1, ..., Act_k$. Then defining Q as a subset of $St \times Act_1 \times ... \times Act_k$ should do the trick.

an AETS which we believe is right (or 'natural') in order to verify the formulae in question. And the first model from figure 2 seems perfectly natural.

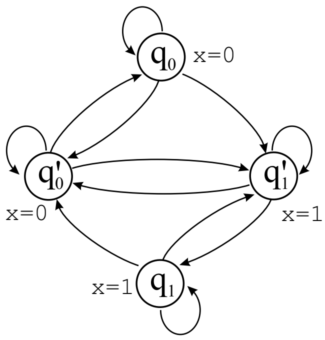

$$\begin{aligned}\delta(q_0, s) &= \delta(q'_0, s) = \{\{q_0\}, \{q'_0, q'_1\}\}\\ \delta(q_1, s) &= \delta(q'_1, s) = \{\{q_1\}, \{q'_0, q'_1\}\}\\ \delta(q_0, c) &= \delta(q'_0, c) = \{\{q_0, q'_0\}, \{q_0, q'_1\}\}\\ \delta(q_1, c) &= \delta(q'_1, c) = \{\{q_1, q'_0\}, \{q_1, q'_1\}\}\end{aligned}$$

Fig. 3. New AETS for the controller/client

3 Improved Definitions for AETS

How to get rid of those problems? Following the evolution of the concurrent game structures and allowing for explicit representation of agents' actions seems the simplest solution. Additionally, it may be useful to allow for symbolic action and agent labels (instead of natural numbers only) – it doesn't change the computational costs, but it can make specifying a system more natural for human designers. Thus an alternating epistemic transition system can be defined as

$$S = \langle \Sigma, Q, \Pi, \pi, \sim_1, ..., \sim_k, Act, d, \delta \rangle$$

where $Act = \{\alpha_1, ..., \alpha_m\}$ are action labels and $d_a(q) \subseteq Act$ defines the actions available to agent a at state q. For every q, q' such that $q \sim_a q'$, it is required that

– $d_a(q) = d_a(q')$

otherwise a can distinguish q from q' by the decisions he can make.[3] Finally, $\delta(q, \alpha^1, ..., \alpha^k) \in Q$ describes the system transition from q for a complete array of decisions $\langle \alpha^1, ..., \alpha^k \rangle \in d_{a_1}(q) \times ... \times d_{a_k}(q)$.

Let $\lambda[i]$ denote the ith item in sequence λ. An *incomplete information strategy* is a function $f_a : Q^+ \to Act$ for which the following constraints hold:

[3] as it turned out during preparation of the final version of this paper, the authors of ATEL have also suggested a similar requirement in a case study submitted independently to another forum [5]. They also considered whether some further constraint on the possible runs of the system should be added, but they dismissed the idea.

- $f_a(\lambda q) \in d_a(q)$,
- if $\lambda[i] \sim_a \lambda'[i]$ for every i, then $f_a(\lambda) = f_a(\lambda')$.

We will require that agents have incomplete information strategies in order to be able to enforce some state of affairs.

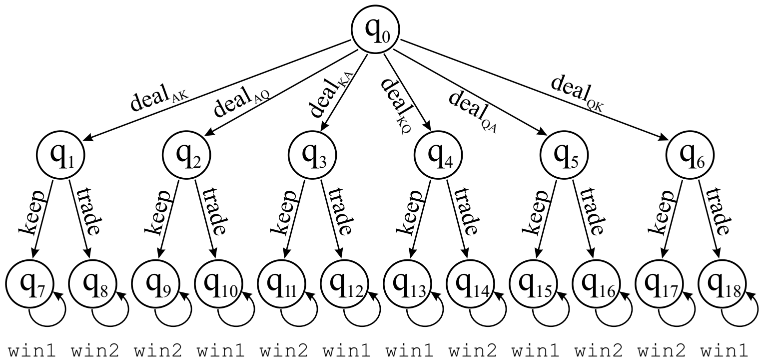

$\Sigma = \{env, a_1, a_2\}$
$Act = \{deal_{AK}, deal_{AQ}, deal_{KA}, deal_{KQ}, deal_{QA}, deal_{QK}, trade, keep, wait\}$

$q_0 \sim_{a_1} q_0,\ q_0 \sim_{a_2} q_0$
$q_1 \sim_{a_1} q_1,\ q_1 \sim_{a_1} q_2,\ q_2 \sim_{a_1} q_2$ etc.
$q_1 \sim_{a_2} q_1,\ q_1 \sim_{a_2} q_6,\ q_6 \sim_{a_2} q_6$ etc.
$q_7 \sim_{a_1} q_7,\ q_7 \sim_{a_2} q_7$
$q_7 \sim_{a_1} q_7,\ q_7 \sim_{a_2} q_7$ etc.

Fig. 4. New AETS for the card game. The transitions are labeled with decisions from the player who takes turn.

A new transition system for the card game is shown on figure 4. Now a_1 can be proven unable to bring about winning on his own: $q_0 \models \neg\langle\langle a_1\rangle\rangle\Diamond win1$. Like in the real game, he can win only with some 'help' from the environment: $q_0 \models \langle\langle a_1, env\rangle\rangle\Diamond win1$.

4 Final Remarks

ATEL and ATEL* are interesting formalisms to describe and verify properties of autonomous processes in situations of incomplete information. A few small changes are suggested in this paper, mostly to make ATEL/ATEL* models consistent with the incomplete information assumption. The notion of a strategy – the way it is defined in [4] – makes formula $\langle\langle A\rangle\rangle\Phi$ describe what coalition A may *happen* to bring about against the most efficient enemies (i.e. when the enemies know the current state and even the A's collective strategy beforehand), whereas the original idea from ATL was rather about A being able to *enforce* Φ.

Explicit representation of actions makes specifying the incomplete information strategies constraint easy. It also enables to express some interesting features of multiagent systems. For instance, we can define an *epistemically monotonic agent* as a player who uses new observations only to narrow down his beliefs:

if $q'_1 \sim_a q'_2$ and there exist $\alpha^1, ..., \alpha^k$ such that $\delta(q_1, \alpha^1, ..., \alpha^k) = q'_1$ and $\delta(q_2, \alpha^1, ..., \alpha^k) = q'_2$
then $q_1 \sim_a q_2$.

i.e. who does only monotonic reasoning with no belief revision.

The incomplete information constraint on strategies, proposed in this paper, is far from being perfect. Consider the last game structure and state q_1, for example. It is easy to show that $q_1 \models \langle\langle a_1 \rangle\rangle \Diamond win1$. Moreover, $q_0 \models \langle\langle\rangle\rangle \bigcirc \langle\langle a_1 \rangle\rangle \Diamond win1$, although still $q_0 \not\models \langle\langle a_1 \rangle\rangle \Diamond win1$. These and other issues about ATEL will be hopefully made clear in the next paper, which is actually in preparation.

The author would like to thank the anonymous reviewers and Wiebe van der Hoek for their valuable remarks.

References

1. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 100–109. IEEE Computer Society Press, 1997. citeseer.nj.nec.com/170985.html
2. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Lecture Notes in Computer Science*, 1536:23–60, 1998. citeseer.nj.nec.com/174802.html
3. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002. Updated, improved, and extended text. www-cad.eecs.berkeley.edu/~tah/Publications/alternating-time_temporal_logic.html
4. W. van der Hoek and M. Wooldridge. Tractable multiagent planning for epistemic goals. In *Proceedings of the First International Conference on Autonomous Agents and Multiagent Systems (AAMAS-02)*, 2002.
5. W. van der Hoek and M. Wooldridge. Model checking cooperation, knowledge and time – a case study. 2003. Submitted.
6. J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behaviour*. Princeton University Press: Princeton, NJ, 1944.