
On Epistemic Temporal Strategic Logic

SIEUWERT VAN OTTERLOO

University of Liverpool
sieuwert@csc.liv.ac.uk

GEERT JONKER

Universiteit Utrecht
geertj@cs.uu.nl

ABSTRACT. ATEL is one of the most expressive logics for reasoning about knowledge, time and strategies. Several issues around the interpretation of this logic are still unresolved. This paper contributes to the ongoing discussion by showing that agents do not have to know a specific strategy for doing something in order to have a capability. Furthermore we claim that agents can possess so-called strategic knowledge that is derived from their knowledge of strategies being played. In order to prove these claims we present an alternative interpretation of ATEL over extensive game forms. For the definition of abilities we use strategy domination, and to deal with strategic knowledge we include strategy profiles in the model. We illustrate the interpretation issues mentioned using several small examples. Furthermore we show how perfect recall and perfect memory can be characterized.

1 Introduction

Logic and game theory are important for understanding language and communication. Logic because it can be used to formalise ideas and notions, and game theory because communicating agents often have goals to reach with their communication. In this paper we look at a logic called ATEL (van der Hoek and Wooldridge 2002), which is a logic to reason about knowledge, time and strategies. This logic can be applied to the formal analysis of a wide range of systems, for instance distributed protocols, synchronisation and security, but also to any issue that can be modeled as an extensive game, such as argumentation, auctions or language games. In this paper we are concerned with the interpretation of ATEL, and we use simple abstract games for the issues at hand.

The original interpretation of the language featured attractive computational properties, but suffered from some counterintuitive properties, which led to a number of refinements (van der Hoek and Wooldridge 2003; Jonker 2003; Jamroga and van der Hoek 2003). In this paper we provide a new interpretation for ATEL. In order to keep things simple we do this for turn based systems, but we see no reason why the approach could not be extended to the more general class of concurrent systems.

The history of ATEL started with the definition of Computation Tree Logic, a logic for branching time models (Clarke and Emerson 1981). Using this logic one can express many temporal properties of distributed or concurrent systems, such as ‘all possible computations will reach this state’ or ‘this state eventually occurs in at least one computation’. It was discovered by Alur and others (Alur, Henzinger, and Kupferman 1997) that this logic could be extended to reason about multiagent systems without changing the complexity of model checking. They extended the language with coalition operators and called it Alternating-time Temporal Logic (ATL). Using ATL one could express properties like ‘This coalition can ensure that a state is reached’ or ‘This coalition cannot avoid that p always holds’. Van der Hoek and Wooldridge realized that it would be useful to reason about the knowledge of agents within ATL, and extended the language with a knowledge operator calling it ATEL (van der Hoek and Wooldridge 2002). They gave the knowledge operator a very intuitive interpreted systems semantics (Fagin, Halpern, Moses, and Vardi 1995) but did not alter the notion of a strategy used in the coalition operator. The result was that agents were assumed to always be able to make different choices in different states, even if the agent could not distinguish these states. Assuming that agents can make different choices in states they cannot distinguish is counterintuitive (Jonker 2003) and different semantics along the lines of imperfect recall ATL (Schobbens 2004) were developed by Jonker, Van der Hoek and Jamroga (Jonker 2003; Jamroga and van der Hoek 2003).

In plain English, ATEL originally assumed that a coalition of agents can achieve something if a strategy to achieve it exists. This condition has been strengthened in the sources cited:

- A coalition of agents can achieve X if they have a uniform strategy (or strategy-under-incomplete-knowledge) that achieves X (Jonker 2003)
- A coalition of agents can achieve X if they have a uniform strategy of which they know that it achieves X . (Jonker 2003; Jamroga and van der Hoek 2003)

In section 2 we define the notion of a uniform strategy formally, but it can

be thought of as a strategy with extra restrictions so that it does not use facts that an agent is not supposed to know. One of the main points of this paper is that this latter condition is too strong. A coalition does not have to be able to identify a strategy of which they know it will be successful. A coalition of sensible agents will choose, if they cannot identify a foolproof strategy, the best strategy that they can come up with. More precisely, they will choose a strategy that is not dominated by any other strategy. We say that a coalition of agents can do X if any undominated strategy achieves X . We will define domination later in this paper.

A second point of this paper is that the knowledge of a coalition does not have to depend only on the state of a game or system, but also on the strategies they employ. It seems safe to assume that agents know what strategies they employ, and that this gives them extra information about the future. This phenomenon can be called strategic knowledge (Druiven 2002). The interpretation developed here addresses this issue by assuming all agents in a coalition know the strategy that a coalition uses. We stress that this is intended not as the final answer on this issue, but as a demonstration how one can incorporate some form of strategic knowledge.

The language CTL is in fact a syntactical restriction of the language CTL*. The same is true for ATL and ATL* and also for ATEL and ATEL*. The unstarred versions have received the most attention, because they have a low model checking complexity (van der Hoek and Wooldridge 2003). In this paper however we are more interested in the meaning of interpretations for the languages than in complexity. Therefore we prefer to work with ATEL*, the language without restrictions, rather than to define an interpretation only for ATEL.

In section 2 we present necessary definitions. Section 3 contains examples. In section 4 we define an interpretation for ATEL*. Section 5 uses the logic to analyse the examples of section 3. In section 6 two theorems are proven and section 7 is the conclusion.

2 Extensive Game Forms

Games are models for interaction between agents with different and possibly competing objectives (Osborne and Rubinstein 1994). An extensive game gives a detailed description of such interaction. It shows which decisions are made in order to reach an outcome. It can be represented in a game tree, where each leaf corresponds to an outcome and each node corresponds to a choice between options. The preferences of all agents are part of an extensive game. In many cases one wants to study the structure of a game independent

of the preferences of the agents. In that case one can use the idea of a game form, which is an extensive game without preference function. The notion of an extensive game goes back to Kuhn (Kuhn 1953). We have adapted the next definition from Osborne (Osborne and Rubinstein 1994). Since the structure encodes the fact that agents may not be sure what exactly the current state is when making a decision, we speak of an extensive game with *imperfect information*.

We have adopted notation fashionable in game theory (Osborne and Rubinstein 1994) and notation used in coalition logic (van der Hoek and Wooldridge 2003). The set of all agents is denoted Σ and Γ is used for a subset of agents. Individual agents are denoted X, Y . For game forms and game form interpretations F, G are used. Strategies are called S, T or S_Γ, T_Γ , to indicate for which group the strategy is. A game form and a strategy together form a model, denoted by M . Formulas are denoted ϕ, ψ and atomic propositions p, q . P is a set of propositions. They are interpreted using a function π . Actions are a, b and histories are called h, h' or j . To improve readability, we sometimes abuse notation a little bit and write XY for the set of agents $\{X, Y\}$.

Game Form A *game form* is a tuple (Σ, H, Ow, \sim) , where Σ is a finite set of agents and H is a non-empty set of *histories*. The set H must be prefix-closed, which means that for any sequence $ha \in H$ also $h \in H$. We use the special symbol ϵ to denote the empty sequence. The set of all actions available after h is denoted $A(h) = \{a | ha \in H\}$. A history $h \in H$ is terminal if $A(h) = \emptyset$. The set of all terminal histories of H is denoted $Z(H)$.

The function $Ow(h) : H \setminus Z(H) \rightarrow \Sigma$ defines which player chooses the next action. Intuitively the agent $Ow(h)$ owns the history h , but we can also say that it decides h , controls h or has the initiative in h . For each agent $X \in \Sigma$, the relation \sim_X is an equivalence relation between histories, where $h \sim_X j$ expresses the fact that agent X cannot tell the difference between having gone through history h and history j . One condition applies: if $Ow(h) = X$ and $h' \sim_X h$ then also $Ow(h') = X$ and further $A(h) = A(h')$. This condition ensures that an agent knows when it can select a action and that it knows which actions are available. This definition is taken from Osborne and Rubinstein (Osborne and Rubinstein 1994), definition 200.1, where it is called an extensive game form. We have extended the information sets such that agents also have information when they are not in charge, which is a common extension for logical purposes (van Benthem 2001; Bonanno 2004b).

Game Form Interpretation A *game form interpretation* is defined as a tuple $(\Sigma, H, Ow, \sim, P, \pi)$. The first elements (Σ, H, Ow, \sim) are a game form. The set P contains propositions and $\pi : H \rightarrow 2^P$ is a function that assigns to each history the set of propositions that are true in the final state of that history. The idea is that these propositions can be used to refer to certain histories, for instance to histories where an agent achieves a certain goal.

Strategies A strategy S_Γ for a coalition Γ is a function that takes a history h such that $Ow(h) \in \Gamma$ and returns a non-empty set of actions $S_\Gamma(h)$ such that $S_\Gamma(h) \subseteq A(h)$. This means that strategies can be non-deterministic. We sometimes call the strategy S_Γ a *strategy profile* to indicate that it contains a strategy for every agent in Γ . There is no fundamental difference between strategies and strategy profiles in this paper. A strategy S_Γ is uniform iff for all h, j with $h \sim_{Ow_h} j$ it is the case that $S_\Gamma(h) = S_\Gamma(j)$. A uniform strategy thus prescribes the same actions in histories that an agent cannot distinguish.

For the purposes of this paper, we think of a game form as a description of a commonly known protocol for interaction between autonomous agents. All agents have preferences and it is common knowledge that these are private, thus not known to other agents. It is also commonly known that agents with different objectives cannot communicate outside the structure of the game. If an agent adopts a certain strategy, for instance to reach a certain goal, then the agent itself knows which strategy it is playing, but other agents do not know this. If a coalition of agents has a goal, then it is assumed that agents have the right means of coordination in order to select a group strategy.

Mental capabilities are included in the definition of a game form. We assume that if an agent is for instance forgetful, that this has been encoded in the equivalence relation in the game form. The relations in the game form thus represent what an agent knows from all its sources, not just its observations. Under this assumption it makes sense to relate properties of agents, such as perfect recall and perfect memory in section 6, to properties of the equivalence relations.

3 Scenarios

To illustrate the principles of ATEL covered in this paper we will first have a look at some examples. The first game form we study is game form G_1 in figure 1.1. In this game agent A can first choose between action 1 and action 2. Agent B then chooses for either action 3 or action 4. The dashed line indicated that agent B does not know what agent A has done when B has




Figure 1.1: A simple game form G_1




Figure 1.2: Game form G_2

to choose. For this game form we are interested in the strategic abilities of agent B . In the original ATEL interpretation agent B has a strategy in t_0 to satisfy q . If A chooses action 1 it would choose action 3, if A chooses action 2 it would choose action 4. However, at t_1 agent B doesn't know whether agent A has played action 1 or 2. In terms of Jonker (Jonker 2003), agent B doesn't have a uniform strategy to satisfy q in t_0 . In t_1 things are a bit different. Agent B does have a uniform strategy to achieve q , but doesn't know which one. In terms of Jonker and Van der Hoek and Jamroga, it cannot identify the right strategy. If agent B would want to satisfy p , it would be able to identify a uniform strategy in t_1 : choosing action 3.

But what if agent B would want to satisfy $p \wedge q$? At t_1 it wouldn't be able to identify a uniform strategy, since it doesn't know whether agent A has played action 1 or 2. Nevertheless, it has a 'best bet' strategy: playing action 3. In the view of the agent, this strategy *might* make it reach its goal. There is no other strategy that is better than this one, so we call this strategy *undominated*. It would be dominated if there existed a strategy that performed better in some of the indistinguishable states, and equally in the others. We will define this more formally in section 4. Let us now assume that we are in the left situation, after action 1. We notice that if agent B plays his 'best bet' strategy, it will achieve his goal. We therefore say that agent B is *able* to satisfy $p \wedge q$, or that $p \wedge q$ is *achievable* for agent B .

Game form G_2 in figure 1.2 illustrates that an agent can have several undominated strategies. We assume the current situation is the one after action 2. Agent B wants to make p true in the next states. The agent has two actions to choose from, and it cannot identify which action is best for achieving p . In the left situation, the appropriate strategy would be action 4. In the right situation, the appropriate strategy would be action 5. For the middle situation however it does not matter which strategy the agent chooses. Either action leads to a desired state. Thus, the agent has two strategies that might make him reach his goal and to him it is not clear




Figure 1.3: Imperfect Memory Game Form




Figure 1.4: Block Pushing

whether to use one or the other, they are both undominated. The best he can do is to randomly choose one of the two. Fortunately, since agent A has played action 2, agent B will reach his goal with any of his undominated strategies. Therefore, we say that agent B is *able* to satisfy p , or that p is *achievable* for agent B .

Both in game forms G_1 and G_2 , the ability of agent B to reach his goal was facilitated by agent A 's choice. In G_1 , we saw that agent B was able to satisfy p and q because agent A had chosen to play action 1. We say that agent A has the ability to give agent B the ability to satisfy p and q . The fact that agent A moves before agent B does, does not imply that agent B cannot enable agent A to reach certain goals. If agent B decides beforehand to play action 3, it in a way gives agent A the ability to satisfy (eventually) p and q .

In figure 1.3 another game form is depicted. This game form is unusual since agents lose information, according to the relations indicated. Agent A seems not to remember, after action 1 that it has chosen this action. Similarly, B knows the difference between histories 1 and 2 but one step later it cannot distinguish 13 and 23. These peculiarities are discussed in section 3 and section 6.

The last scenario pictured in figure 1.4 deals with the question whether agents can lose abilities by choosing certain strategies. This is a minor issue, not very related to the previous issues. Nevertheless it illustrates one issue in the interpretation of ATEL. Assume that the leftmost state is the current state. The interesting property of this model is that the agent has the ability to make p true in the next state. It cannot get rid of this ability in one step. If we assume that agents can unconditionally commit to strategies, so that they cannot change later their strategy, than A would be able, by committing to go to the rightmost state, to causing itself not to be able to achieve p . This is not what we think of as intuitive. Therefore we assume that agents can always change their own strategy at a later moment.

Finally, we turn back to game form G_1 again. Consider the scenario where

agent A has decided to play the strategy of choosing action 2. At time t_0 , this influences the knowledge of agent A . It now knows for example that agent B won't be able to satisfy $p \wedge q$ anymore. This is called strategic knowledge (Druiven 2002). It is a kind of provisional knowledge; provided that an agent will follow the strategy it has committed to, it has more knowledge about the future. The same holds for a coalition; provided that they follow the same strategy and that this is common knowledge between them, their knowledge about the future increases. The next section shows how strategic knowledge, along with the other principles explained above, can be formalized using a new interpretation of ATEL^* .

4 Strategic Temporal Epistemic Logic

The language ATEL^* is the smallest language L such that for any formula $\phi, \psi \in L$, any coalition of agents Γ and any agent X it is the case that:

$$\begin{array}{ccccccc} \phi \vee \psi \in L & \neg\phi \in L & \Box\phi \in L & \phi \mathcal{U} \psi \in L & & & \\ \bigcirc\phi \in L & K_X\phi \in L & \ll \Gamma \gg \phi \in L & & & & \end{array}$$

This language is in fact a mixture between ATL^* and epistemic logic (Fagin, Halpern, Moses, and Vardi 1995). The reader will be familiar with disjunction ($\phi \vee \psi$, 'or') and negation ($\neg\phi$, 'not'). The temporal operators \Box , \mathcal{U} and \bigcirc say something about the future. $\Box\phi$ means that ϕ is true in all future states. The formula $\phi \mathcal{U} \psi$ expresses the fact that at a certain point in the future ψ becomes true, and until that time ϕ is true. The next-state operator $\bigcirc\phi$ expresses the fact that in the next state ϕ is true. The epistemic operator $K_X\phi$ indicates that agent X knows that ϕ holds. The coalition operator $\ll \Gamma \gg \phi$ expresses that the set of agents Γ can ensure that ϕ holds.

We interpret formulas ϕ over a model M and a history h and write $M, h \models \phi$ if the formula is true. Unlike previous interpretations we include strategies in the model. A model M is a pair (F, S_Σ) where F is a game form with interpretation and S_Σ a strategy for all agents.

We define the neutral strategy S_Γ^0 as the strategy which returns all available actions: $S_\Gamma^0(h) = A(h)$. Strategies for different coalitions can be combined using the function `combine` into a function for a larger coalition. The strategy `combine`(Γ, S_Γ, T) is equal to S_Γ for agents in Γ and equal to T for other agents:

$$\text{combine}(\Gamma, S_\Gamma, T)(h) = \begin{cases} S_\Gamma(h) & \text{if } Ow(h) \in \Gamma \\ T(h) & \text{if } Ow(h) \notin \Gamma \end{cases}$$

A model $M = (F, S_\Sigma)$ contains information about all strategies that the agents currently use. An agent only knows his own strategy. It knows it will

follow that strategy, but assumes that nothing more is known about others than that they adhere to the neutral strategy S_Σ^0 . Thus, when evaluating the knowledge of an agent about the future, we use the model that is the result of the agent using its strategy, while the others use the neutral strategy. This strategy is in fact the least restrictive strategy indistinguishable to the agent. We define it using the operator $\mathbf{k}(M, X)$ which is defined as

$$\mathbf{k}((F, S_\Sigma), X) = (F, \mathbf{combine}(X, S_\Sigma, S_\Sigma^0))$$

The function k is used in the definition of the knowledge operator.

In order to define the meaning of the strategic operator $\ll \Gamma \gg$, we use the idea of *undominated strategies*. Informally a strategy S dominates T if S is strictly better than T for reaching a certain goal. A coalition of rational agents, we assume, will not play a strategy if that strategy is dominated. To define domination, we first we need to define two other operators. The operator **update** is similar to **k**: it returns the model that represents the view of agents in coalition Γ after the adoption of strategy S_Γ . Agents in Γ adhere to S_Γ but other agents can act in any way they want:

$$\mathbf{update}(M, S_\Gamma) = (F, \mathbf{combine}(\Gamma, S_\Gamma, S_\Sigma^0))$$

Furthermore, we call a strategy S_Γ *successful* (**success**(M, h, S_Γ, ϕ)) in history h (for Γ with respect to ϕ) if and only if **update**(M, S_Γ), $h \models \phi$. A strategy S_Γ *dominates* a strategy T_Γ with respect to a model M , a goal ϕ and a history h if two conditions are met : There is a history $h' \sim_\Gamma h$ such that **success**(M, h, S_Γ, ϕ) but not **success**(M, h, T_Γ, ϕ), and secondly there is no history j such that **success**(M, j, T_Γ, ϕ) but not **success**(M, j, S_Γ, ϕ). This definition of dominance makes the domination relation transitive and asymmetric. These properties ensure that any nonempty, finite set of strategies contains at least one strategy not dominated by another strategy in the set. We say that **achievable**(M, h, Γ, ϕ) if for all undominated strategies S_Γ it is the case that **success**(M, h, S_Γ, ϕ). The reason that we quantify over all undominated strategies is that we imagine that a coalition picks randomly any undominated strategy, since it has no reason to prefer one undominated strategy over the other. Therefore success is only guaranteed if all undominated strategies are successful. Using all these notions we define

the interpretation of ATEL* as follows:

$M, h \models p$	iff $p \in \pi(h)$
$M, h \models \neg\phi$	iff not $M, h \models \phi$
$M, h \models \phi \vee \psi$	iff $M, h \models \phi$ or $M, h \models \psi$
$M, h \models K_X \phi$	iff $\forall h' : h' \sim_X h \implies \mathbf{k}(M, X), h' \models \phi$
$M, h \models \phi \mathcal{U} \psi$	iff $\forall j \in Z(H(S_\Sigma, h)) \exists i \forall k : h \leq k < i \implies$ $M, j_0 \dots j_k \models \phi) \wedge M, j_0 \dots j_i \models \psi$
$M, h \models \Box \phi$	iff $\forall h' \in H(S_\Sigma, h) : M, h' \models \phi$
$M, h \models \bigcirc \phi$	iff $\forall a \in S_\Sigma(h) : M, ha \models \phi$
$M, h \models \ll \Gamma \gg \phi$	iff $\text{achievable}(M, h, \Gamma, \phi)$

The set $H(S, h)$ contains all histories of H that start with h and are consistent (after h) with the strategy S . It can be defined recursively. $H(S, h)$ is the smallest set H' such that $h \in H'$ and $\forall h' \in H', \forall a \in S(h') : h'a \in H'$.

5 Examples

In section 3 we have introduced four game forms. In this section we use these game forms to show the interpretation of example formulas. For all examples the model M_i is defined as (G_i, S_Σ^0) . The empty history is denoted \emptyset . The first examples deal with temporal properties.

$M_1, \emptyset \models t_0 \wedge \bigcirc t_1 \wedge \bigcirc \bigcirc t_2$	Initially t_0 holds, then t_1 and then t_2
$M_1, \emptyset \models \Box(t_2 \rightarrow \neg t_1)$	If t_2 holds, then not t_1
$M_1, \emptyset \models \top \mathcal{U} t_2$	Eventually t_2 holds

We have argued that in the game form G_1 after action 1, the agent B can achieve p , but not q . It can also achieve $p \wedge q$ but it does not know that it can achieve this fact. The translations of these facts are given here. In model G_2 similar properties hold and these are also given.

$M_1, 1 \models \ll B \gg \bigcirc p$	B can make p true in the next state
$M_1, 1 \models \neg \ll B \gg \bigcirc q$	B cannot make q true in the next state
$M_1, 1 \models \ll B \gg \bigcirc (p \wedge q)$	B can make p and q true in the next state
$M_1, 1 \models \neg K_B \ll B \gg \bigcirc (p \wedge q)$	B doesn't know it can make p and q true
$M_2, 1 \models \neg \ll B \gg \bigcirc p$	B cannot make p true in the next state
$M_2, 2 \models \ll B \gg \bigcirc p$	B can make p true in the next state
$M_2, 2 \models \neg K_B \ll B \gg \bigcirc p$	B does not know it can make p true

In game form G_3 , agent A does not remember the choices it has made in the past. Agent B does not always know its previous observations. This is expressed in the next statements. The next model, G_4 , shows that agents cannot in general commit themselves to act against their future preferences.

$M_3, \emptyset \models K_A \ll A \gg \circ p$	A knows it can make p true
$M_3, \emptyset \models \neg \ll A \gg \circ K_{Ap}$	A cannot know p in the next state
$M_3, 1 \models K_B \circ q$	B knows q is true in the next state
$M_3, 1 \models \neg \circ K_B q$	Next, B does not know that q
$M_4, \emptyset \models \ll A \gg \circ \circ p$	A can make p true in 2 steps
$M_4, \emptyset \models \neg \ll A \gg \neg \ll A \gg \circ p$	A cannot ensure it cannot make p true
$M_4, \emptyset \models \neg \ll A \gg \circ \neg \ll A \gg \circ p$	A cannot lose its ability

Turning back to game form G_1 , we give an example of how an agent can have strategic knowledge. Suppose that agent A has committed to the strategy of playing action 2. It then knows that agent B will not be able to achieve the satisfaction of $p \wedge q$ anymore. Let $S_{A,B}$ be the strategy profile consisting of the strategy of playing action 2 for agent A and the neutral strategy for agent B . We can then represent the knowledge described above as:

$$(G_1, S_{A,B}), \emptyset \models K_A \neg \ll B \gg \circ \circ p \wedge q \quad A \text{ knows } B \text{ cannot make } p \wedge q$$

6 Perfect Recall and Perfect Memory

Agents have *perfect recall* if they never forget their previous observations and the actions they have chosen (Osborne and Rubinstein 1994). Similarly they have perfect memory (Koller and Megiddo 1992; Bonanno 2004a) if they do not forget their observations. Traditionally game theory has focused on perfect recall agents, but artificial agents in multiagent systems often do not have these properties. For the complexity of solving games, or model checking temporal formulas, it is relevant whether the systems have perfect recall of perfect memory (Koller and Megiddo 1992; Halpern, van der Meyden, and Vardi 2004). Therefore we present here two theorems that characterize whether a game form interpretation has perfect recall and perfect memory. Especially we want to illustrate the difference between perfect recall and perfect memory, since this difference does not appear in temporal logic without strategies, but does exist in games and strategic logics. Making use of the strategy profiles that we have included in the model is necessary for the proof of the perfect recall property. The authors are not sure whether a different

theorem regarding perfect recall could hold for previous interpretations of ATEL*.

It can be argued that *perfect recall* is not a property of a game, but a property of a player in a game. However we think of the equivalence relations \sim_X in a game form $(\Sigma, H, Ow, [\sim_a]_{a \in \Sigma})$ as representing the knowledge of the agents. We thus assumed that the capabilities of the agents have been included in the equivalence relations. Thus we view perfect recall as a property that a game form can or cannot have. We define perfect recall in terms of observations; an agent has perfect recall if it remembers all its observations, including the actions it has chosen. Let the $O_X(h)$ be the function returning the ordered list of all observations and actions chosen by agent X in history h . Then an agent X has perfect recall if and only if $h \sim_X j \leftrightarrow O_X(h) = O_X(j)$. The function O_X can be defined recursively. The observation function of ha contains all observations of h , plus maybe the action a (if $Ow(h) = X$) and the observation made in ha . Using such recursive definition, it is not hard to show that the property of perfect recall is equivalent to the next two properties.

$$\begin{aligned} ha \sim_X jb &\rightarrow h \sim_X j \\ Ow(h) = X \wedge ha \sim_X jb &\rightarrow a = b \end{aligned}$$

This characterisation of perfect recall is the one we use in the next theorem.

Theorem 1 *Let $F = (\Sigma, H, Ow, [\sim_X]_{X \in \Sigma})$ and $X \in \Sigma$. X has perfect recall in F if and only if for every P, π , every ϕ , every strategy S_Σ and every h :*

$$((F, P, \pi), S_\Sigma), h \models K_X \circ \phi \rightarrow \bigcirc K_X \phi$$

Suppose that X has perfect recall in F and let $P, \pi, \phi, S_\Sigma, h$ be given. Let $G = (F, P, \pi)$ and suppose that $(G, S_\Sigma), h \models K_X \circ \phi$. Define $M' = (G, S'_\Sigma) = k((G, S_\Sigma), X)$. Let $a \in S_\Sigma(h)$ and $jb \sim_X ha$. From the perfect recall properties we know that $h \sim_X j$. From the definition of k one can see that for any h' it is the case that $S_\Sigma(h') \subseteq S'_\Sigma(h')$. Using $(G, S_\Sigma), h \models K_X \circ \phi$ we can conclude that $M', j \models \bigcirc \phi$. If $Ow(h) = X$ then $Ow(j) = X$ and we know that $b = a$ and therefore $b \in S_\Sigma(h) \subseteq S'_\Sigma$. If $Ow(h) \neq X$ then $S'_\Sigma(j) = A(j)$ and thus $b \in S'_\Sigma(j)$. From $M', j \models \bigcirc \phi$ and $b \in S'_\Sigma(j)$ we can conclude that $M', jb \models \phi$. Since we have shown this for an arbitrary $jb \sim_X ha$ we conclude that $(G, S_\Sigma), h \models \bigcirc K_X \phi$.

For the second half, assume that for every P, π , every ϕ , every strategy S_Σ and every h :

$$((F, P, \pi), S_\Sigma), h \models K_X \circ \phi \rightarrow \bigcirc K_X \phi$$

Let $G = (F, P, \pi)$ and let $ha \in H$. Take jb such that $ha \sim_X jb$. Let $P = \{p\}$ and define π such that $\pi(j'b') = \{p\}$ iff $h \sim_X j'$ and $b' \in A(j')$. Let $S = S_\Sigma^0$. This interpretation ensures that $(G, S), h \models K_X \circ p$. We can derive from the assumptions that $(G, S), h \models \circ K_X p$. Thus for every $a \in A(h)$ it is the case that $(G, S), ha \models K_X p$. Since $ha \sim_X jb$, we conclude that $(G, S), jb \models p$. By definition of $\pi(p)$ this gives us $h \sim_X j$.

Let $G = (F, P, \pi)$ and let $ha \in H$ be such that $Ow(h) = X$. take $jb \in H$ such that $ha \sim_X jb$. From the previous part we can already conclude that $h \sim_X j$. Let $P = \{p\}$ and let $S = S_\Sigma$ be a strategy such that $S(h) = \{a\}$. define π such that $\pi(j'b') = \{p\}$ iff $b' = a$ and $j' \sim_X h$. These definitions ensure that $(G, S), h \models K_X \circ p$. We can derive $(G, S), h \models \circ K_X p$. From the definition of next we know that $(G, S), ha \models K_X p$ and thus that $(G, S), jb \models p$. By definition of p we conclude that $a = b$. This concludes the proof.

An agent with perfect memory remembers all its previous observations. Let $M_X(h)$ be the function returning the ordered list of all observations by agent X in history h (excluding the actions it has chosen). We define that an agent X has perfect memory if and only if $h \sim_X j \leftrightarrow M_X(h) = M_X(j)$. Again, one can define the observation function M recursively. Using such recursive definition, it is not hard to show that the property of perfect memory is equivalent to $ha \sim_X jb \rightarrow h \sim_X j$.

Theorem 2 *Let $F = (\Sigma, H, Ow, [\sim_X]_{X \in \Sigma})$ be a game form and $X \in \Sigma$. X has perfect memory in F if and only if for every P, π , every ϕ , and every h :*

$$((F, P, \pi), S_\Sigma^0), h \models K_X \circ \phi \rightarrow \circ K_X \phi$$

In this theorem, instead of just any strategy we use the neutral strategy S_Σ^0 . An important property of this strategy is that $k((F, S_\Sigma^0), X) = (F, S_\Sigma^0)$.

For the first half of the proof, let X have perfect memory in F and assume that P, π, ϕ and h are given. Let $M = ((F, P, \pi), S_\Sigma^0)$ and assume that $M, h \models K_X \circ \phi$. Let $a \in A(h)$ and jb any history with $jb \sim_X ha$. From the perfect memory property we know that $j \sim_X h$. We can derive that $k(M, X), j \models \circ \phi$. Since M contains the neutral strategy, $k(M, X) = M$ and thus $M, j \models \circ \phi$. This means that $M, jb \models \phi$. Since we have shown this for an arbitrary jb with $jb \sim_X ha$ we can conclude that $M, h \models \circ K_X \phi$.

For the second part, assume that for every P, π, ϕ and h : $((F, P, \pi), S_\Sigma^0), h \models K_X \circ \phi \rightarrow \circ K_X \phi$. Let $ha \sim_X jb$ be given. Define $P = \{p\}$ and π such that $\pi(h'a') = \{p\}$ iff $h' \sim_X h$ and $a' \in A(h')$. This definition ensures that $M, h \models K_X \circ p$. We can conclude that $M, h \models \circ K_X p$. From this formula one can derive that $M, ha \models K_X p$ and thus that $M, jb \models p$. this means that $\pi(jb) = \{p\}$ and thus that $j \sim_X h$. Therefore X has perfect memory in F .

7 Conclusions and Further research

We have presented the logic ATEL^* and shown how this logic, because it deals with knowledge, time and strategies, can be used for reasoning about game situations and multi-agent protocols. There are some subtle issues in the interaction between knowing and being able to do something. By giving an interpretation based on the notion of dominated strategies, we have shown that knowing that you can do something is not necessary in order to be able to do something.

This idea has led us to present a new interpretation that considers strategies explicitly, and considers the knowledge that agents have about these strategies. We believe that this is an important issue in reasoning about games, because this is what allows agents to coordinate their actions and to reason more accurately about the future. Within this paper we have presented two technical results, namely that perfect recall and perfect memory of game forms can be characterised in terms of ATEL^* . There is a lot more work to be done on this logic. In order to apply it to linguistic and game-theoretic domains it would be nice to understand which axioms are valid in this logic, and which strategic knowledge assumptions are the most natural for each domain. Another open problem is the computational complexity of this logic. If this problem is tractable, then ATEL^* could be applied to many multi agent verification problems, such as problems similar to the Russian Cards problem (van Ditmarsch 2003) and the Dining Cryptographers problem (van der Meyden and Su 2002). We conjecture that the interpretation presented in this paper may have made the model checking problem harder, but cannot present any results on this issue yet.

Bibliography

- Alur, R., T. A. Henzinger, and O. Kupferman (1997, October). Alternating-time temporal logic. In *Proceedings of the 38th IEEE Symposium on Foundations of Computer Science*, Florida, pp. 100–109.
- Bonanno, G. (2004a). Memory and perfect recall in extensive games. *Games and Economic Behavior* 47, 237–256.
- Bonanno, G. (2004b). Memory implies von neumann-morgenstern games. *Knowledge Rationality and Action to appear*, 1–20.
- Clarke, E. and E. Emerson (1981). Desing and synthesis of synchronisation skeletons using branching time temporal logic. *Lecture Notes in Computer Science* 131, 52–71.
- Druiven, S. (2002). Knowledge development in games of imperfect information. University Maastricht Master Thesis.
- Fagin, R., J. Halpern, Y. Moses, and M. Vardi (1995). *Reasoning about knowledge*. The MIT Press: Cambridge, MA.
- Halpern, J., R. van der Meyden, and M. Vardi (2004). Complete axiomatizations for reasoning about knowledge and time. *SIAM Journal on Computing* 33, 674–703.
- Jamroga, W. and W. van der Hoek (2003). Some remarks on alternating-time temporal epistemic logic. submitted.
- Jonker, G. (2003). Feasible strategies in alternating-time temporal epistemic logic. Universiteit Utrecht Master Thesis.
- Koller, D. and N. Megiddo (1992, October). The complexity of two-person zero-sum games in extensive form. *Games and Economic Behavior* 4 4, 528–552.
- Kuhn, H. (1953). Extensive games and the problem of information. *Contributions to the Theory of Games II*, 193–216.
- Osborne, M. J. and A. Rubinstein (1994). *A Course in Game Theory*. The MIT Press: Cambridge, MA.

- Schobbens, P.-Y. (2004). Alternating-time logic with imperfect recall. In W. van der Hoek, A. Lomuscio, E. de Vink, and M. Wooldridge (Eds.), *Electronic Notes in Theoretical Computer Science*, Volume 85. Elsevier.
- van Benthem, J. (2001). Games in dynamic-epistemic logic. *Bulletin of Economic Research* 53(4), 219–248.
- van der Hoek, W. and M. Wooldridge (2002). Tractable multiagent planning for epistemic goals. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2002)*, Bologna, Italy, pp. 1167–1174.
- van der Hoek, W. and M. Wooldridge (2003). Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications. *Studia Logica* 75(4), 125–157.
- van der Meyden, R. and K. Su (2002). Symbolic model checking the knowledge of the dining cryptographers. submitted.
- van Ditmarsch, H. P. (2003). The russian cards problem. *Studia Logica* 75(4), 31–62.