

Alternating Epistemic Mu-Calculus

Nils Bulling

Department of Informatics,
Clausthal University of Technology, Germany

bulling@in.tu-clausthal.de

Wojciech Jamroga

Computer Science and Communication,
University of Luxembourg

wojtek.jamroga@uni.lu

Abstract

Alternating-time temporal logic (**ATL**) is a well-known logic for reasoning about strategic abilities of agents. An important feature that distinguishes variants of **ATL** for imperfect information scenarios is that the standard fixed point characterizations of temporal modalities do not hold anymore. In this paper, we show that adding explicit fixed point operators to the “next-time” fragment of **ATL** already allows to capture abilities that could not be expressed in **ATL**. We also illustrate that the new language allows to specify important kinds of abilities, namely ones where the agents can always recompute their strategy while executing it. Thus, the agents are not assumed to remember their strategy by definition, like in the existing variants of **ATL**. Last but not least, we show that verification of such abilities can be cheaper than for all the variants of “**ATL** with imperfect information” considered so far.

1 Introduction

Alternating-time temporal logic (**ATL**) [2] is probably the most popular logic that allows reasoning about agents’ abilities in strategic encounters. **ATL** combines features of temporal logic and basic game theory, which are both encapsulated in the main language construct of the logic, $\langle\langle A \rangle\rangle\gamma$, which can be read as “the group of agents A has a strategy to enforce γ ”, where γ is a temporal property including operators \bigcirc (“next”), \square (“always”) and/or \mathcal{U} (“until”). **ATL** was originally proposed for reasoning about agents in perfect information scenarios. Since then, several variations have been studied for imperfect information [9; 7; 1; 5; 8]. One of the features that distinguishes them from the perfect information setting is that the fixed point characterizations of temporal modalities do not hold anymore under imperfect information. On the conceptual level, this means that having a strategy to achieve φ does not mean that the agents will be able to recompute the strategy when they are already executing it. Thus, the semantics of these **ATL** variants includes an implicit assumption that agents *can remember the strategy which they are executing* even if they forget everything else.

This can be a good thing or a bad thing, depending on the notion of ability that one wants to formalize. Nevertheless, the other kind of ability (enforcing φ without resorting to additional memory of the strategy) is at least as important. It captures the idea of agents “*persistently knowing how to play to enforce φ* ”, i.e., so that they can come up with the right strategy not only at the beginning of the game; they will *know* how to recreate the strategy also in any future moment of the play. Moreover, this kind of ability has a minimalistic flavor regarding epistemic prerequisites: agents are supposed to achieve φ while resorting *only* to observations that they can make along the way.

In this paper, we argue that this class of abilities cannot be expressed in the existing variants of **ATL**. We also propose that, instead of considering yet another semantics of strategic operators, it suffices to add explicit fixed point operators to the “next-time” fragment of **ATL**. In terms of technical results, our contribution is threefold: (1) We show that the new logic is incomparable with “**ATL** with imperfect information” regarding expressive power. (2) We propose specifications of natural types of strategic abilities. (3) We point out that verification of such abilities can be cheaper than for all the variants of “**ATL** with imperfect information” considered so far.

Recomputable strategies are ones that do not have to be remembered by the agents in order to execute them. In a sense strategies of this kind impose the weakest requirements on agents’ memory. Moreover, recomputable strategies can be (by definition) synthesized incrementally, which can make the synthesis tractable—at least for small coalitions.

2 Agent Logics: Preliminaries

In this section we provide a brief overview of existing variants of alternating-time temporal logic (**ATL**) in the way they were proposed in [2] and later refined in [9]. Finally, we recall the definition of alternating μ -calculus (**AMC**) which was introduced in [2].

2.1 Concurrent (Epistemic) Game Structures

The semantics for the logics are defined over a variant of transition systems where transitions are labeled with combinations of actions, one per agent. An *imperfect information concurrent game structure* (ICGS) [10; 9] is given by

$M = \langle \text{Agt}, Q, \Pi, \pi, \text{Act}, d, o, \{\sim_a \mid a \in \text{Agt}\} \rangle$ which includes a nonempty finite set of all agents $\text{Agt} = \{1, \dots, k\}$, a nonempty set of states Q , a set of atomic propositions Π and their valuation $\pi : \Pi \rightarrow \mathcal{P}(Q)$, and a nonempty finite set of (atomic) actions Act . Function $d : \text{Agt} \times Q \rightarrow \mathcal{P}(\text{Act})$ defines nonempty sets of actions available to agents at each state, and o is a (deterministic) transition function that assigns the outcome state $q' = o(q, \alpha_1, \dots, \alpha_k)$ to state q and a tuple of actions $\langle \alpha_1, \dots, \alpha_k \rangle$ for $\alpha_i \in d(i, q)$ and $1 \leq i \leq k$, that can be executed by Agt in q . We write $d_a(q)$ instead of $d(a, q)$. Each $\sim_a \subseteq Q \times Q$ is an equivalence relation satisfying $d_a(q) = d_a(q')$ for $q \sim_a q'$. These relations model incomplete information in the way it is done in epistemic logic. Note that the perfect information models from [2] (CGS, concurrent game structures) can be modeled by assuming each \sim_a to be the minimal reflexive relation.

A *strategy* of agent a is a conditional plan that specifies what a is going to do in each situation. A natural taxonomy of four strategy types was introduced in [9] and labeled as follows: i (resp. I) stands for *imperfect* (resp. *perfect*) *information*, and r (resp. R) refers to *imperfect* (resp. *perfect*) *recall*. The following types of strategies can be used in the respective semantic variants:

- Ir : $s_a : Q \rightarrow \text{Act}$ such that $s_a(q) \in d_a(q)$ for all q (agents have perfect information and base their decisions on states only);
- IR : $s_a : Q^+ \rightarrow \text{Act}$ such that $s_a(q_0 \dots q_n) \in d_a(q_n)$ for all $q_0 \dots q_n$ (agents have perfect information and base their decisions on the whole history of events happened);
- ir : like Ir , with the additional constraint that $q \sim_a q'$ implies $s_a(q) = s_a(q')$ (agents have to assign the same choice to states which appear indistinguishable to them);
- iR : like IR , with the additional constraint that $h \approx_a h'$ implies $s_a(h) = s_a(h')$, where \approx_a encodes indistinguishability of histories as follows. For $h = q_0 q_1 \dots q_n \in Q^+$ and $h' = q'_0 q'_1 \dots q'_n \in Q^+$ we have $h \approx_a h'$ iff $n = n'$ and $q_i \sim_a q'_i$ for $i = 1, \dots, n$.

A *collective xy-strategy* s_A is a tuple of xy -strategies, one per agent from A , for $xy = Ir, IR, ir, iR$. Additionally, $s_A|_a$ denotes agent a 's part of collective strategy s_A . The set of A 's collective xy -strategies is defined as Σ_A^{xy} .

A *path* $\lambda = q_0 q_1 q_2 \dots$ is an infinite sequence of states such that there is a transition between each q_i, q_{i+1} . We use $\lambda[i]$ to denote the i th position on path λ (starting from $i = 0$) and $\lambda[i, \infty]$ to denote the subpath of λ starting from i .

Function $out_M(q, s_A)$ returns the set of all paths that can result from the execution of strategy s_A from state q in model M , see [2] for the precise definition. We define $out_M^{xy}(q, s_A) = out_M(q, s_A)$ for $xy = Ir, IR$, and $out_M^{xy}(q, s_A) = \bigcup_{a \in A} \bigcup_{q' \sim_a q} out_M(q', s_A)$ for $xy = ir, iR$.

2.2 Alternating Time Temporal Logic

The *language* $\mathcal{L}_{ATL}(\Pi, \text{Agt})$ is given by the following grammar [2]:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A \rangle\rangle \circ \varphi \mid \langle\langle A \rangle\rangle \square \varphi \mid \langle\langle A \rangle\rangle \varphi \mathcal{U} \varphi$$

where $A \subseteq \text{Agt}$ and $p \in \Pi$. Additionally, we define ‘‘some-time in the future’’ as $\diamond\varphi \equiv \top \mathcal{U} \varphi$.

Using xy -strategies defined above we obtain the logics ATL_{xy} for $x \in \{i, I\}$ and $y \in \{r, R\}$. The xy -semantics, \models_{xy} , is defined by the following clauses where we omit the standard cases for propositions, negation, and conjunction:

$$\begin{aligned} M, q \models_{xy} \langle\langle A \rangle\rangle \circ \varphi & \text{ iff there is a collective } xy\text{-strategy } s_A \\ & \text{ such that, for each } \lambda \in out^{xy}(q, s_A), M, \lambda[1] \models_{xy} \varphi; \\ M, q \models_{xy} \langle\langle A \rangle\rangle \square \varphi & \text{ iff there is an } xy\text{-strategy } s_A \text{ s.t., for} \\ & \text{ each } \lambda \in out^{xy}(q, s_A) \text{ and every } i \geq 0, M, \lambda[i] \models_{xy} \varphi; \\ M, q \models_{xy} \langle\langle A \rangle\rangle \varphi \mathcal{U} \psi & \text{ iff there is an } xy\text{-strategy } s_A \text{ s.t.,} \\ & \text{ for each } \lambda \in out^{xy}(q, s_A), \text{ there is } i \geq 0 \text{ for which} \\ & M, \lambda[i] \models_{xy} \psi, \text{ and } M, \lambda[j] \models_{xy} \varphi \text{ for each } 0 \leq j < i. \end{aligned}$$

Informally speaking, $M, q \models \langle\langle A \rangle\rangle \gamma$ iff, there exists a collective strategy s_A such that γ holds on all outcome paths that the agents in A consider possible executions of s_A . Moreover, in the ir - and iR -semantics, we can define *knowledge operators* $K_a \varphi$ as $\langle\langle a \rangle\rangle \varphi \mathcal{U} \varphi$ and we define the ‘‘everybody in A knows’’ modality E_A as $\bigwedge_{a \in A} K_a$. It is easy to check that $M, q \models_{xy} K_a \varphi$ iff for all q' with $q \sim_a q'$, $M, q' \models_{xy} \varphi$, for $xy = ir, iR$.

Note that $M, q \models_{ix} \langle\langle A \rangle\rangle \gamma$ requires A to have a single strategy that is successful in *all* states indistinguishable from q for *any member* of the coalition.

2.3 Alternating Mu-Calculus

Alternating μ -calculus (**AMC**) replaces the temporal-strategic operators $\langle\langle A \rangle\rangle \square, \langle\langle A \rangle\rangle \mathcal{U}$ with the *least fixed point operator* μ . Like the basic version of **ATL**, **AMC** was proposed in the context of perfect information scenarios. The *language* $\mathcal{L}_{AMC}(\Pi, \text{Agt}, \text{Var})$ is given by the following grammar [2]:

$$\varphi ::= p \mid X \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle A \rangle \varphi \mid \mu X(\varphi)$$

where $A \subseteq \text{Agt}$, $p \in \Pi$, $X \in \text{Var}$ and each φ in $\mu X(\varphi)$ is X -*positive*, i.e. each free occurrence of X in φ is under the scope of an even number of negations in φ . We define $\nu X(\varphi(X))$ as $\neg \mu X(\neg \varphi(\neg X))$ where $\varphi(\neg X)$ is equivalent to $\varphi(X)$ but each free occurrence of X in φ is replaced by $\neg X$. We also take $[A]\varphi$ as $\neg \langle A \rangle \neg \varphi$ and define conjunction in the standard way. Var is the set of fixed point variables, i.e., second order variables ranging over sets of states.

Here we consider only the *alternation-free* fragment of **AMC**. The corresponding language $\mathcal{L}_{af-AMC}(\Pi, \text{Agt}, \text{Var})$ consists of all alternation-free formulae of $\mathcal{L}_{AMC}(\Pi, \text{Agt}, \text{Var})$; i.e. all formulae of which their positive normal form (i.e. negations occur only in front of propositions) contains no occurrences of ν (resp. μ) on any syntactic path from an occurrence of μX (resp. νX) to a bound occurrence of X (cf. [2]).

A (Var, Q) -*valuation* \mathcal{V} is a mapping $\mathcal{V} : \text{Var} \rightarrow \mathcal{P}(Q)$. Given a variable X and a set $Z \subseteq Q$ of states we define

$$\mathcal{V}[X := Z](Y) = \begin{cases} \mathcal{V}(Y) & \text{if } X \neq Y; \\ Z & \text{else.} \end{cases} \quad \text{That is, the valuation } \mathcal{V}[X := Z] \text{ equals } \mathcal{V} \text{ for all variables different from } X \text{ and assigns } Z \text{ to } X.$$

The semantics for \mathcal{L}_{AMC} is given by the denotation function $\llbracket \cdot \rrbracket_{\mathcal{V}}^M$ that maps \mathcal{L}_{AMC} -formulae to sets of states (i.e. $\llbracket \varphi \rrbracket_{\mathcal{V}}^M \subseteq Q$) where M is a CGS and

\mathcal{V} is a valuation. The denotation function is defined as follows: $\llbracket p \rrbracket_{\mathcal{V}}^M = \pi(p)$, $\llbracket X \rrbracket_{\mathcal{V}}^M = \mathcal{V}(X)$, $\llbracket \neg\varphi \rrbracket_{\mathcal{V}}^M = Q \setminus \llbracket \varphi \rrbracket_{\mathcal{V}}^M$, $\llbracket \varphi \wedge \psi \rrbracket_{\mathcal{V}}^M = \llbracket \varphi \rrbracket_{\mathcal{V}}^M \cap \llbracket \psi \rrbracket_{\mathcal{V}}^M$, $\llbracket \langle A \rangle \varphi \rrbracket_{\mathcal{V}}^M = \{q \mid \exists \alpha_A \in d_A(q) \forall \alpha_{\text{Agt} \setminus A} \in d_{\text{Agt} \setminus A} : o(q, (\alpha_A, \alpha_{\text{Agt} \setminus A})) \in \llbracket \varphi \rrbracket_{\mathcal{V}}^M\}$, and $\llbracket \mu X(\varphi) \rrbracket_{\mathcal{V}}^M = \bigcap \{Z \subseteq Q \mid \llbracket \varphi \rrbracket_{\mathcal{V}[X:=Z]}^M \subseteq Z\}$.

Finally, we write $M, q \models^{\mathcal{V}} \varphi$ for $q \in \llbracket \varphi \rrbracket_{\mathcal{V}}^M$ and $M, q \models \varphi$ if $M, q \models^{\mathcal{V}} \varphi$ for all (Var, Q) -valuations \mathcal{V} . We omit “ (Var, Q) ” if clear from context.

2.4 Known Expressivity Results

Definition 1 (Comparing expressivity) A logic L_1 is as expressive as L_2 over a class of models \mathcal{M} (written $L_2 \preceq^{\mathcal{M}} L_1$) iff there is a translation TR of L_2 formulae to L_1 -formulae such that for each L_2 -formula φ we have: $M \models_{L_2} \varphi$ iff $M \models_{L_1} \text{TR}(\varphi)$ for all $M \in \mathcal{M}$.

We say that L_1 is more expressive than L_2 ($L_2 \prec^{\mathcal{M}} L_1$) iff $L_2 \preceq^{\mathcal{M}} L_1$ but not $L_1 \preceq^{\mathcal{M}} L_2$. Finally, we call L_1, L_2 incomparable if neither $L_2 \preceq^{\mathcal{M}} L_1$ nor $L_1 \preceq^{\mathcal{M}} L_2$.

We note, that for two classes of models \mathcal{M} and \mathcal{M}' with $\mathcal{M} \subseteq \mathcal{M}'$, if not $L_2 \preceq^{\mathcal{M}} L_1$ then also not $L_2 \preceq^{\mathcal{M}'} L_1$.

The following translation $\text{TR}_{\text{ATL}_{\text{IR}}}(\cdot) : \mathcal{L}_{\text{ATL}} \rightarrow \mathcal{L}_{\text{af-AMC}}$ shows how \mathcal{L}_{ATL} -formulae can be encoded into $\mathcal{L}_{\text{af-AMC}}$ [2]: $\text{TR}_{\text{ATL}_{\text{IR}}}(p) = p$, $\text{TR}_{\text{ATL}_{\text{IR}}}(\neg\varphi) = \neg\text{TR}_{\text{ATL}_{\text{IR}}}(\varphi)$, $\text{TR}_{\text{ATL}_{\text{IR}}}(\langle A \rangle \varphi) = \langle A \rangle \text{TR}_{\text{ATL}_{\text{IR}}}(\varphi)$, $\text{TR}_{\text{ATL}_{\text{IR}}}(\langle A \rangle \square\varphi) = \nu X(\text{TR}_{\text{ATL}_{\text{IR}}}(\varphi) \wedge \langle A \rangle X)$, and $\text{TR}_{\text{ATL}_{\text{IR}}}(\langle A \rangle \mu\varphi) = \mu X(\text{TR}_{\text{ATL}_{\text{IR}}}(\varphi) \vee (\text{TR}_{\text{ATL}_{\text{IR}}}(\varphi_1) \wedge \langle A \rangle X))$.

Proposition 1 ([2]) For all \mathcal{L}_{ATL} -formulae φ , all CGS M and all states $q \in Q_M$ we have: $M, q \models_{\text{ATL}_{\text{IR}}} \varphi$ iff $M, q \models_{\text{af-AMC}} \text{TR}_{\text{ATL}_{\text{IR}}}(\varphi)$.

Theorem 2 ([2]) Alternation-free AMC_{IR} is more expressive than ATL_{IR} .

3 Alternating Epistemic Mu-Calculus

Alternating μ -calculus was introduced for reasoning about systems in which agents always know the current state of affairs. To the best of our knowledge, analogous logics have never been considered for systems with imperfect information. In this section we propose an imperfect information semantics for $\mathcal{L}_{\text{af-AMC}}$ in a straightforward way. The only change concerns the semantic clause for operator $\langle A \rangle$. Now, $\langle A \rangle \varphi$ is only true if agents A have a collective one-step strategy which enforces φ from all the states that are indistinguishable from the current one. Also, since knowledge operators K_a cannot be expressed anymore when the “until” operator is missing, we add them explicitly to the language. This way, we obtain the *alternation-free alternating epistemic μ -calculus*, af-AMC_i in short.

3.1 Syntax and Semantics

Definition 2 ($\mathcal{L}_{\text{af-AMC}_i}$) The language $\mathcal{L}_{\text{AMC}_i}(\Pi, \text{Agt}, \text{Var})$ extends the definition of \mathcal{L}_{AMC} by formulae $K_a\varphi$; that is, it is defined by the grammar:

$$\varphi ::= p \mid X \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle A \rangle \varphi \mid \mu X(\varphi) \mid K_a\varphi$$

where $A \subseteq \text{Agt}$, $a \in \text{Agt}$, $p \in \Pi$, $X \in \text{Var}$ and each φ in $\mu X(\varphi)$ is X -positive, as before. Additionally, to the




Figure 1: Models M and N , each with a single agent 1.

macros defined above we define “everybody in A knows” as $E_A\varphi = \bigwedge_{a \in A} K_a\varphi$ and $\hat{K}_a\varphi = \neg K_a\neg\varphi$. Analogously to the perfect information case, we define the alternation-free sublanguage $\mathcal{L}_{\text{af-AMC}_i}(\Pi, \text{Agt}, \text{Var})$ of $\mathcal{L}_{\text{AMC}_i}(\Pi, \text{Agt}, \text{Var})$.

Definition 3 (Imperfect information semantics) Let M be an ICGS and \mathcal{V} a valuation of fixed point variables. The *i*-semantics for $\mathcal{L}_{\text{af-AMC}_i}(\text{Agt}, \Pi, \text{Var})$ -formulae, denoted $\llbracket \cdot \rrbracket_{i, \mathcal{V}}^M$, is defined as the semantics $\llbracket \cdot \rrbracket_{\mathcal{V}}^M$ (the modification to ICGSS is straightforward) extended with the following semantic clauses:

$$\llbracket \langle A \rangle \varphi \rrbracket_{i, \mathcal{V}}^M = \{q \mid \exists s_A \in \Sigma_A^{ir} \forall \lambda \in \text{out}^{ir}(q, s_A) (\lambda[1] \in \llbracket \varphi \rrbracket_{i, \mathcal{V}}^M)\},$$

$$\llbracket K_a\varphi \rrbracket_{i, \mathcal{V}}^M = \{q \mid \forall q' (q \sim_a q' \text{ implies } q' \in \llbracket \varphi \rrbracket_{i, \mathcal{V}}^M)\}.$$

In other words, $q \in \llbracket \langle A \rangle \varphi \rrbracket_{i, \mathcal{V}}^M$ iff agents A have a collective *ir*-strategy such that everybody in the group knows that it enforces φ in the next step.

We write $M, q \models_i \varphi$ iff $q \in \llbracket \varphi \rrbracket_{i, \mathcal{V}}^M$ for all valuations \mathcal{V} of fixed point variables from Var .

Example 1 Consider the ICGS M shown in Figure 1. The story is as follows: A married man is sitting in a pub drinking with his friends (q_0). In order to get back to his wife (q_2), he needs to finish his drinking session first (q_1), but that will result in a temporary lapse of memory. In particular, the man can have a nagging feeling that something might be wrong with his marriage, e.g. his wife could have left him because of his drinking habits (q'_1), in which case he can only come back to an empty house (q'_2).

Trivially, we have that $M, q_0 \models_{ir} \langle \{1\} \rangle \diamond p$, so the man should rest assured. However, is it really the property he is after? He knows now that everything will be fine, but he also knows that he will get confused on the way, which may prevent him from reaching his goal. The stronger kind of ability is captured by the af-AMC_i formula $\mu X(p \vee \langle 1 \rangle X)$, which is not true in M, q_0 .

The example suggests that ATL_{ir} and af-AMC_i offer means to refer to *different* types of strategic ability.

We proceed with this issue in Section 3.2, where we show that both logics are incomparable, and in Section 4, where we analyse the new kind of strategic ability.

3.2 Expressivity

We start by showing that, like in the perfect information setting, af-AMC_i still allows to capture properties that cannot be expressed in ATL_{ir} .

Proposition 3 ATL_{ir} is not as expressive as $af-AMC_i$ over the class of ICGSSs.

Proof. It is known that the alternation-free modal μ -calculus (**af-MC**) is strictly more expressive than **CTL** over the class \mathcal{M}^K of Kripke structures. That is, for any translation of **af-MC**-formulae to **CTL**-formulae there is a **af-MC**-formula φ such that $\text{not } \mathcal{M}^K \models_{\text{af-MC}} \varphi$ iff $\mathcal{M}^K \models_{\text{CTL}} \mathcal{TR}_{\text{CTL}}(\varphi)$.

Now suppose ATL_{ir} were **af-AMC**_{*i*}-expressive over the class of ICGSSs. Then, ATL_{ir} would also be **af-AMC**_{*i*}-expressive over \mathcal{M}^K as \mathcal{M}^K can be seen as the class of perfect information single agent ICGSSs. Thus, **CTL** would also be **af-MC**-expressive over \mathcal{M}^K . Contradiction. ■

Note that we do *not* use epistemic operators in the proof, so they are *not* the reason for the expressivity gap.

The other direction is more surprising. First, we show that naive ideas for a translation from **ATL** to **af-AMC** in the imperfect information setting do not work.

Proposition 4 There is an \mathcal{L}_{ATL} -formula φ , an ICGS M , and a state $q \in Q_M$ such that: $M, q \models_{ATL_{ir}} \varphi$ and $M, q \not\models_{af-AMC_i} \mathcal{TR}_{ATL_{ir}}(\varphi)$.

Proof. Consider model M from Figure 1 and the **ATL** formula $\varphi \equiv \langle\langle 1 \rangle\rangle \diamond p$. Now, $\mathcal{TR}_{ATL_{ir}}(\varphi) = \mu X(p \vee \langle 1 \rangle X)$. It is easy to see that $M, q_0 \models_{ir} \varphi$ but $M, q_0 \not\models_i \mathcal{TR}_{ATL_{ir}}(\varphi)$. ■

We may try to repair the translation by use of epistemic operators that take into account also states indistinguishable from the current one. Let the *epistemic extension* $\mathcal{TR}_{ATL_{ir}}(\cdot)$ of the standard translation $\mathcal{TR}_{ATL_{ir}}(\cdot)$ be obtained by modifying the standard translation as follows. For every \mathcal{L}_{ATL} -formula $\varphi = \langle\langle A \rangle\rangle \psi$ we set $\mathcal{TR}_{ATL_{ir}}(\varphi) = E_A \mathcal{TR}_{ATL_{ir}}^*(\varphi)$ where $\mathcal{TR}_{ATL_{ir}}^*(\cdot)$ is defined as $\mathcal{TR}_{ATL_{ir}}(\cdot)$ but each nested occurrence of $\mathcal{TR}_{ATL_{ir}}(\cdot)$ is replaced by $\mathcal{TR}_{ATL_{ir}}(\cdot)$. We will refer to $\mathcal{TR}_{ATL_{ir}}(\cdot)$ as *epistemic standard translation*.

Proposition 5 There is an \mathcal{L}_{ATL} -formula φ , an ICGS M , and a state $q \in Q_M$ such that: $M, q \models_{ATL_{ir}} \varphi$ and $M, q \not\models_{af-AMC_i} \mathcal{TR}_{ATL_{ir}}(\varphi)$.

Proof. Again, we use the model M from Figure 1 and $\varphi \equiv \langle\langle 1 \rangle\rangle \diamond p$. $\mathcal{TR}_{ATL_{ir}}(\varphi) = E_1 \mu X(p \vee \langle 1 \rangle X)$. Now, we have again that $M, q_0 \models_{ir} \varphi$ but $M, q_0 \not\models_i \mathcal{TR}_{ATL_{ir}}(\varphi)$. ■

So, is there another translation that gives us the desired characterisation? The next results show that such a translation does not exist. Due to lack of space we omit the full proofs.

Proposition 6 $af-AMC_i$ is not as expressive as ATL_{ir} over the class of ICGSSs.

Proof. (Idea) The basic idea is that the ATL_{ir} -formula $\langle\langle 1 \rangle\rangle \diamond p$ does distinguish the models M and N in state q_0 shown in Fig. 1 but there is no **af-AMC**_{*i*}-formulae which can. ■

The next theorem follows from Propositions 3 and 6.

Theorem 7 $af-AMC_i$ is incomparable to ATL_{ir} in expressive power over the class of ICGSSs.

How does the picture change if agents have memory? That is, we compare ATL_{ir} and **af-AMC**_{*i*}. (Note that it does not

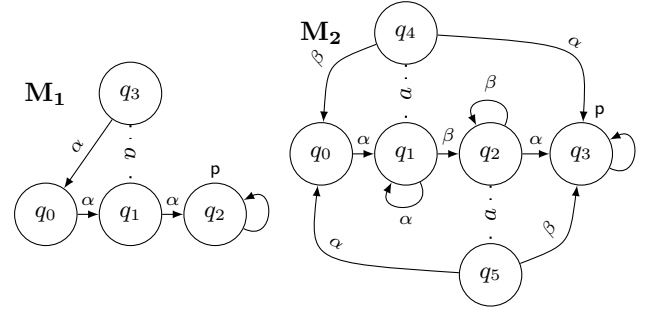


Figure 2: Models M_1 and M_2 with $\text{Agt} = \{a\}$. Arrows represent transitions (labeled by a 's actions). Dotted lines depict indistinguishability of states.

make sense to consider memory in **af-AMC**_{*i*} since the strategic modalities refer only to one-step strategies there.) Clearly, Proposition 6 holds for ATL_{ir} . But also Proposition 3 holds, since ATL_{ir} and ATL_{ir} are *equally expressive* over (pure) CGSSs. Hence, if the analogous result were not true ATL_{ir} would also be at least as expressive as **af-AMC**_{*i*} over single-agent CGSSs. As before this yields a contradiction, for **af-MC** is known to be strictly more expressive than **CTL**.

Theorem 8 $af-AMC_i$ is incomparable to ATL_{ir} in expressive power over the class of ICGSSs.

4 Specification and Verification of Abilities

In this section we argue that **af-AMC**_{*i*} can be useful for specifying and verifying meaningful strategic properties of agents. We begin with a formalization of the intuitions we tried to convey in the introduction. To this end, we define fixed point properties that formally capture those informal intuitions.

4.1 Strategic Fixed Points: Achievement

The main intuition behind our proposal is that we enable expressing that agents A have a strategy to enforce a temporal property γ while knowing how to play *all along the game*. For achievement properties, we have typically $\gamma \equiv \diamond \varphi$. We note that the ability to eventually achieve φ while knowing how to play all along can have at least two meaningful interpretations:

1. A have a strategy which they know to achieve φ , can be recomputed along the execution, and guarantees that they will know when φ has been achieved;
2. The agents have a recomputable strategy that they know to achieve φ , but they will not necessarily know when φ is achieved.

For a single agent the first achievement ability corresponds closely to the **af-AMC**_{*i*}-formula:

$$\mu X(K_a \varphi \vee \langle a \rangle X).$$

We observe that this formula (and formulae of **af-AMC**_{*i*} in general) has a strong “constructive” flavor. For example, the given formula requires not only that a has a uniform strategy to eventually obtain p , but also that

- she can *recompute* the strategy at *any* future state resulting from executing it,
- she *knows* about this; and furthermore,
- with every step, a knows that the goal is closer, and hence the part of the strategy that she needs to recompute becomes smaller.

This is a stronger kind of ability than those typically expressed in ATL_{ir} . For illustration, consider model M_1 in Figure 2. Agent a has a uniform strategy to achieve p eventually: $M_1, q_0 \models \langle\langle a \rangle\rangle \diamond p$ (the strategy is to play α everywhere). It is even the case that a can recompute a successful strategy on the way ($M_1, q_0 \models \langle\langle a \rangle\rangle (\langle\langle a \rangle\rangle \diamond p) \mathcal{U} p$)—which is trivial as only a single action is available. However, the af-AMC_i -formula does not hold: $M_1, q_0 \not\models \mu X(K_{ap} \vee \langle a \rangle X)$. This is because a is never certain that the goal state is really approaching. More precisely, suppose a executes α in state q_0 . In the next state, q_1 , she executes α again. However, as she considers q_3 possible the resulting state could—from her perspective—also be q_0 . Applying this reasoning iteratively a possible execution path could also be $q_0 q_3 (q_0 q_3)^\omega$ which does never satisfy p . Hence, the agent does not know to have a recomputable strategy to achieve p .

Analogously, we can also describe the second interpretation of the achievement property by

$$\mu X(\varphi \vee \langle a \rangle X).$$

We omit the formal characterizations due to lack of space. We note also that the relationships for proper coalitions are more sophisticated; the exact study is left for future work.

4.2 Strategic Fixed Points: Maintenance

In this section we turn to maintenance properties. They are typically expressed by $\Box\varphi$. Like for achievement properties, the ability to maintain φ while knowing how to play all along can have at least two interpretations:

1. A have a strategy to maintain φ and keep knowing that φ , that can be recomputed at every step.
2. The agents have a recomputable strategy to maintain φ (but they may be unsure if φ is true at some steps).

Similar to the case of achievement properties we can describe the first property by

$$\nu X(K_a \varphi \wedge \langle a \rangle X).$$

In this his case the difference between ATL_{ir} and af-AMC_i specifications is perhaps even clearer. As an example we consider model M_2 shown in Figure 2. We have $M_2, q_0 \models \langle\langle a \rangle\rangle \Box \neg p$ which expresses that a can avoid p . For example, take the strategy in which α is played everywhere. We even have $M_2, q_0 \models \langle\langle a \rangle\rangle \Box (\langle\langle a \rangle\rangle \Box \neg p)$ (while avoiding p , she can recompute a successful strategy). However, at q_1 , a can only come up with a strategy that is different from the original strategy (i.e., play β everywhere), and indeed $M_2, q_0 \not\models \nu X(K_a \neg p \wedge \langle a \rangle X)$.

Again, the second property can be described in a similar way by formula

$$\nu X(\varphi \wedge \langle a \rangle X).$$

Combinations of achievement and maintenance, captured by the “until” operator \mathcal{U} , can be treated analogously.

4.3 Verification

In the previous sections we have shown that af-AMC_i allows to specify interesting properties that cannot be expressed in the existing variants of ATL for imperfect information. A natural question arises: How costly is the verification of these properties? It is known that:

- Model checking of strategic logics of perfect information (ATL_{ir} , ATL_{IR} , af-AMC) is in \mathbf{P} [2];
- Verification of ATL_{ir} (imperfect information and memoryless strategies) is Δ_2^P -complete even for turn-based systems consisting of a *single* agent [9; 6];
- Model checking ATL_{ir} (imperfect information, perfect recall) is undecidable [4].

Here, we show that model checking the alternating μ -calculus under imperfect information is in \mathbf{P} for coalitions consisting of at most two agents (in a system including arbitrarily many agents). Moreover the problem is between \mathbf{NP} and Δ_2^P if larger coalitions are involved. Thus, verification with af-AMC_i can be distinctly cheaper than with ATL for abilities of small coalitions, and no harder in the general case.

First, we point out that the μ -operator in af-AMC_i is a standard least fixed point operator. Thus, the crucial part of model checking is the computation of the preimage for $\langle A \rangle$ (i.e., the set of states satisfying $\langle A \rangle p$ for a given proposition p). Then, verification of an arbitrary formula can be done through a polynomial number of calls to preimage computations within the standard iterative algorithm (cf. e.g. [3]). The complexity of the preimage computation turns out to be as follows.

Proposition 9 *Checking if $M, q \models \langle A \rangle p$ for $|A| \leq 2$ can be done in linear time wrt the number of transitions in M .*

Proof. The algorithm for $A = \{a_1, a_2\}$ is given below. It is easy to see that it never processes the same transition twice.

Let $Q = [q]_{\sim_A}$ and $D = d_{a_1}(q) \times d_{a_2}(q)$;

while there is still a collective action (α_1, α_2) in D **do**:

- Fix α_1 for a_1 in $[q]_{\sim_{a_1}}$ and α_2 for a_2 in $[q]_{\sim_{a_2}}$.
- For every state in $[q]_{\sim_A}$ there is at most one agent in A for whom the action has not been fixed. **If** a_i 's action is not fixed for q', q'' such that $q' \sim_{a_i} q''$ **then** collapse q', q'' into a single state (taking the union of the outgoing transitions). **Repeat** iteratively;
- **If** in the resulting perfect information ICGS A have a one-step strategy to enforce p in the next state **then** return *true* **else** remove (α_1, α_2) from D and revert to the original model M ;

If the loop ended with no success, there are no more available actions, so return *false*. ■

Proposition 10 *Model checking $\langle A \rangle p$ in M, q is in \mathbf{NP} wrt. the number of transitions in M .*

Proof. (Sketch) Guess a one-step strategy of A , remove from M the irrelevant transitions and states outside $[q]_{\sim_A}$, and check if p holds for all the remaining “next” states. ■




Figure 3: Model M_Φ for $\Phi \equiv C_1 \wedge C_2$, $C_1 \equiv x_1 \vee x_2$, $C_2 \equiv \neg x_1 \vee x_2$. Only transitions leading to q_\perp are labeled; the other combinations of actions lead to q_\top .

Proposition 11 *Model checking $\langle A \rangle p$ in M, q for $|A| \geq 3$ is NP-hard wrt. the number of transitions in M .*

Proof. (Sketch) We use an adaptation of the SAT reduction from [6]. Given a Boolean formula Φ in CNF, we construct a 3-agent ICGS M_Φ as follows. Each literal l from clause ψ in Φ is associated with a state q_l^ψ . At state q_l^ψ , player 1 indicates a literal from ψ , and player 2 decides on the valuation of the underlying Boolean variable. If 1 indicated a “wrong” literal $l' \neq l$ then the system proceeds to state q_\top where proposition *yes* holds. The same happens if 1 indicated the “right” literal (l) and 2 selected the valuation that makes l true. Otherwise the system proceeds to the “sink” state q_\perp .

Player 1 must select literals uniformly within clauses, so $q_l^\psi \sim_1 q_{l'}^{\psi'}$ iff $\psi = \psi'$. Player 2 is to select uniform valuations of variables, i.e., $q_l^\psi \sim_2 q_{l'}^{\psi'}$ iff $\text{var}(l) = \text{var}(l')$ where $\text{var}(l)$ is the variable contained in l . Finally, all states except q_\top, q_\perp are indistinguishable for 3. An example of the construction is presented in Figure 3.

Then, Φ is satisfiable iff $M_\Phi, q \models \langle 1, 2, 3 \rangle \text{yes}$ where q is an arbitrary “literal” state. ■

The following is a straightforward corollary (We conjecture that it is actually Δ_2^P -complete.).

Theorem 12 *Model checking af-AMC_i for formulae that include only operators $\langle A \rangle$ with $|A| \leq 2$ can be done in polynomial time wrt. the size of the model and the length of the formula. In general, the problem is NP-hard and in Δ_2^P .*

5 Conclusions

In this paper, we have proposed an imperfect information variant of the alternation-free alternating μ -calculus from [2]. The idea is very simple: af-AMC_i is a modification of af-AMC with imperfect information semantics of one-step modalities $\langle A \rangle$, plus explicit modalities for reasoning about agents’ knowledge. We show that the new logic allows to specify properties which cannot be expressed in ATL_{ir} nor ATL_{iR} . Somewhat surprisingly, the converse also holds: there are properties expressible in ATL_{ir} and ATL_{iR} which have no counterparts in af-AMC_i . Thus, unlike af-AMC which strictly extends ATL in perfect information scenarios, af-AMC_i and ATL variants for imperfect information seem

to have different agendas. ATL allows to reason about properties that the agents know how to achieve from the current state of the system. af-AMC_i formulae refer to properties that the agents will know how to achieve *all along the play* because they are able to recompute the right strategy at any moment.

We also point out that model checking af-AMC_i specifications can be done in polynomial time for small coalitions; for larger groups of agents, the verification complexity is the same as for ATL_{ir} . So, the new logic allows to specify a new intuitive class of strategic properties, and at the same time has a tractable model checking problem – albeit in a somewhat limited scope. We consider the result significant. To our knowledge, af-AMC_i is the first strategic logic of imperfect information with tractable verification for a relevant subset of formulae. Even solving 2-player extensive form games with binary payoffs is NP-complete in any reasonable sense (e.g., for surely winning). Model checking of all existing variants of ATL with imperfect information is at least Δ_2^P -hard even for 1-player(!) turn-based ICGS’s. The conceptual and computational aspects hopefully make af-AMC_i appealing in practical contexts.

A deeper study on specification and verification issues wrt. the alternating epistemic μ -calculus is left for future research.

Acknowledgements. Wojciech Jamroga acknowledges the support of the FNR (National Research Fund) Luxembourg under project S-GAMES – C08/IS/03.

References

- [1] T. Ågotnes. Action and knowledge in alternating-time temporal logic. *Synthese*, 149(2):377–409, 2006. Section on Knowledge, Rationality and Action.
- [2] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.
- [3] E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 1999.
- [4] C. Dima and F. Tiplea. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. Submitted. Available at <http://lacl.univ-paris12.fr/dima/work/atlir.html>, 2011.
- [5] W. Jamroga and T. Ågotnes. Constructive knowledge: What agents can achieve under incomplete information. *Journal of Applied Non-Classical Logics*, 17(4):423–475, 2007.
- [6] W. Jamroga and J. Dix. Model checking ATL_{ir} is indeed Δ_2^P -complete. In *Proceedings of EUMAS’06*, 2006.
- [7] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2–3):185–219, 2004.
- [8] H. Schnoor. Strategic planning for probabilistic games with incomplete information. In *Proceedings of AAMAS’10*, pages 1057–1064, 2010.
- [9] P. Y. Schobbens. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science*, 85(2):82–93, 2004.
- [10] W. van der Hoek and M. Wooldridge. Cooperation, knowledge and time: Alternating-time Temporal Epistemic Logic and its applications. *Studia Logica*, 75(1):125–157, 2003.