

Some Things are Easier for the Dumb and the Bright Ones (Beware of the Average!)

Model Checking for Strategies in Extreme Epistemic Structures

Wojciech Jamroga^{1,2} and Michał Knapik¹

¹ Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland

² Interdisciplinary Centre for Security, Reliability and Trust, SnT, University of Luxembourg
{w.jamroga, michal.knapik}@ipipan.waw.pl

Abstract

Model checking strategic abilities in multi-agent systems is hard, especially for agents with partial observability of the state of the system. In that case, it ranges from NP-complete to undecidable, depending on the precise syntax and the semantic variant. That, however, is the *worst case complexity*, and the problem might as well be easier when restricted to particular subclasses of inputs. In this paper, we look at the verification of models with “extreme” epistemic structure, and identify several special cases for which model checking is easier than in general. We also prove that, in the other cases, no gain is possible even if the agents have almost full (or almost nil) observability. To prove the latter kind of results, we develop generic techniques that may be useful also outside of this study.

1 Introduction

Many relevant properties of multi-agent systems (MAS) refer to *strategic abilities* of agents and their groups. Such properties can be neatly specified in alternating-time temporal logic (ATL) [Alur *et al.*, 2002]. In its basic version, the logic allows to specify strategic properties of agents and their coalitions under the assumption of perfect information about the current state of affairs. As the assumption is rather unrealistic, there is a growing number of works that study the syntactic and semantic variants of ATL for agents with imperfect information, cf. [Ågotnes *et al.*, 2015] for an overview.

Unfortunately, verification of strategic properties of agents with imperfect information is difficult. More precisely, model checking of ATL variants with imperfect information is Δ_2^P - to PSPACE-complete for agents playing memoryless (a.k.a. positional) strategies [Bulling *et al.*, 2010; Jamroga and Dix, 2006; Schobbens, 2004] and EXPTIME-complete to undecidable for agents with perfect recall of the past [Dima and Tiplea, 2011; Guelev *et al.*, 2011]. This concurs with the results for solving imperfect information games and synthesis of winning strategies, which are also known to be hard [Doyen and Raskin, 2011; Chatterjee *et al.*, 2007; Peterson and Reif, 1979]. Note, however, that theoretical complexity results refer to the *worst case complexity*. The

problem might as well be easier when restricted to a particular subclass of inputs. Indeed, many hard problems have relatively small “hardness cores,” and are fairly easy elsewhere. In the context of model checking for strategies, such results are especially known for strategies with perfect recall [Schewe and Finkbeiner, 2007; Berwanger and Kaiser, 2010; Belardinelli *et al.*, 2017; Berwanger *et al.*, 2018; Maubert and Murano, 2018].

In this paper, we study some natural restrictions on models, that might lead to cheaper verification. More specifically, we look at models with “extreme” epistemic structure, arising when the agents have almost nil, or, symmetrically, almost perfect observability. A sensor observing only one variable, with a fixed number of possible values, provides a natural example of the former type. For the latter class, consider a central controller monitoring a team of robots, with only a fixed number of units being unavailable at a time. It turns out that, when we consistently pair those restrictions with the assumptions about agents’ memory (i.e., assume almost perfect observability and perfect recall, or almost nil observability and no recall), model checking can become easier than in general. This applies especially to the verification of abilities of singleton coalitions. We also show that no gain is possible for the other combinations. To prove the latter kind of results, we develop general reduction techniques which may be relevant also for other formal problems in AI.

2 Model Checking Strategic Abilities

2.1 ATL: What Agents Can Achieve

Alternating-time temporal logic ATL [Alur *et al.*, 2002] generalizes branching time logic CTL by replacing path quantifiers with *cooperation modalities* $\langle\langle A \rangle\rangle$. Informally, $\langle\langle A \rangle\rangle\gamma$ expresses that the group of agents A has a collective strategy to enforce temporal property γ . ATL formulae include temporal operators: “X” (“in the next state”), “G” (“always from now on”) and U (“until”). The additional operator “F” (“now or sometime in the future”) is defined as $F\gamma \equiv \top U \gamma$.

The language of ATL is given by the grammar below, where A is a set of agents, and p is an atomic proposition:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A \rangle\rangle X\varphi \mid \langle\langle A \rangle\rangle G\varphi \mid \langle\langle A \rangle\rangle \varphi U \varphi.$$

2.2 Models of Multi-Agent Interaction

The semantics of **ATL** is defined over a variant of transition systems where transitions are labeled with combinations of actions, one per agent. Formally, a *concurrent game structure (CGS)* [Alur *et al.*, 2002] is a tuple $M = \langle \text{Agt}, St, \Pi, \pi, Act, d, o \rangle$ which includes a nonempty finite set of all agents $\text{Agt} = \{1, \dots, k\}$, a nonempty set of states St , a set of atomic propositions Π and their valuation $\pi: \Pi \rightarrow 2^{St}$, and a nonempty finite set of (atomic) actions Act . Function $d: \text{Agt} \times St \rightarrow 2^{Act} \setminus \emptyset$ defines the sets of actions available to agents at each state. We will often write $d_i(q)$ instead of $d(i, q)$, and denote the set of collective choices of group A at state q by $d_A(q) = \prod_{i \in A} d_i(q)$. Finally, o is a transition function that assigns the outcome state $q' = o(q, \alpha_1, \dots, \alpha_k)$ to state q and a tuple of actions $\alpha_i \in d_{\text{Agt}}(q)$.

Concurrent epistemic game models (CEGM) [Schobbens, 2004], are **CGS**'s augmented with a family of equivalence relations $\sim_a \subseteq St \times St$, one per agent $a \in \text{Agt}$. The relations describe agents' uncertainty: $q \sim_a q'$ means that agent a cannot distinguish between states q and q' . It is also required that agents have the same choices in indistinguishable states: if $q \sim_a q'$ then $d_a(q) = d_a(q')$. The abstraction classes of \sim_a are sometimes called *information sets*. We use $\#is$ to denote the maximum number of information sets per agent, and $|is|$ for the size of the largest information set in the CEGM.

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 $i = 0$). The set of paths starting in q is denoted by $Paths[M](q)$, and the set of their finite prefixes by $Paths[M]^{fin}(q)$.

A *history* h is a finite sequence of states. We use h_F to denote its final state. Two histories $h = q_0 q_1 \dots q_n$ and $h' = q'_0 q'_1 \dots q'_n$ are *indistinguishable for agent a* ($h \approx_a h'$) iff $n = n'$ and $q_i \sim_a q'_i$ for $i = 1, \dots, n$. Additionally, for any equivalence relation \mathcal{R} over a set X we use $[x]_{\mathcal{R}}$ to denote the equivalence class of x . Moreover, we use the abbreviations $\sim_A := \bigcup_{a \in A} \sim_a$ and $\approx_A := \bigcup_{a \in A} \approx_a$. Note that relations \sim_A and \approx_A implement the ‘‘everybody knows’’ type of collective knowledge.

2.3 Semantic Variants of Strategic Ability

A number of semantic variations have been proposed for **ATL**, cf. e.g. [Jamroga, 2003; Schobbens, 2004; Jamroga and van der Hoek, 2004; Ågotnes *et al.*, 2007; Ågotnes and Walther, 2009]. In this paper, we are concerned with the semantic variants **ATL_{ir}** and **ATL_{iR}** [Schobbens, 2004].

The following types of strategies are used in the respective semantic variants:

- **ir** (imperfect information + imperfect recall): $s_a: St \rightarrow Act$ s.t. $s_a(q) \in d_a(q)$ for all q , with the constraint that $q \sim_a q'$ implies $s_a(q) = s_a(q')$;
- **iR** (imperfect information + perfect recall): $s_a: St^+ \rightarrow Act$ s.t. $s_a(q_0 \dots q_n) \in d_a(q_n)$ for all q_0, \dots, q_n , with the constraint that $h \approx_a h'$ implies $s_a(h) = s_a(h')$.

That is, strategy s_a is a conditional plan specifying a 's actions in each state of the system (for memoryless agents) or for every possible history of the system evolution (for agents with perfect recall). Moreover, strategies specify the same choices

	Single agents	Coalitions
Memoryless	Δ_2^P -complete	Δ_2^P -complete
Perfect recall	EXPTIME -complete	undecidable

Figure 1: Existing complexity results

for indistinguishable states (resp. histories). Collective strategies s_A are tuples of individual strategies s_a , one per $a \in A$.

The ‘‘objective outcome’’ function $out^o(q, s_A)$ returns the set of all paths that may occur when agents A execute strategy s_A from state q onward. The set of ‘‘subjectively possible outcomes’’ is defined as $out^i(q, s_A) = \bigcup_{q \sim_A q'} out^o(q', s_A)$.

The semantics of **ATL_{xy}**, parameterized by the notion of outcome ($x \in \{o, i\}$) and the type of recall ($y \in \{r, R\}$), can be given by the following clauses:

$M, q \models_{xy} p$ iff $q \in \pi(p)$, where $p \in \Pi$;

$M, q \models_{xy} \neg\varphi$ iff $M, q \not\models_{xy} \varphi$;

$M, q \models_{xy} \varphi \wedge \psi$ iff $M, q \models_{xy} \varphi$ and $M, q \models_{xy} \psi$;

$M, q \models_{xy} \langle\langle A \rangle\rangle X\varphi$ iff there is a collective *iy*-strategy s_A such that, for each path $\lambda \in out^x(q, s_A)$, we have $M, \lambda[1] \models_{xy} \varphi$;

$M, q \models_{xy} \langle\langle A \rangle\rangle G\varphi$ iff there exists s_A such that, for each $\lambda \in out^x(q, s_A)$, we have $M, \lambda[i] \models_{xy} \varphi$ for every $i \geq 0$;

$M, q \models_{xy} \langle\langle A \rangle\rangle \varphi \cup \psi$ iff there exists s_A such that, for each $\lambda \in out^x(q, s_A)$, there is $i \geq 0$ for which $M, \lambda[i] \models_{xy} \psi$, and $M, \lambda[j] \models_{xy} \varphi$ for each $0 \leq j < i$.

2.4 Known Complexity Results

In this paper, we focus on verifying MAS with imperfect information, i.e., on model checking **ATL_{ir}** and **ATL_{iR}**. The former problem is known to be Δ_2^P -complete [Schobbens, 2004; Jamroga and Dix, 2006].¹ The latter problem is undecidable in general [Dima and Tiplea, 2011], but it becomes **EXPTIME**-complete when only singleton coalitions are allowed in the formula (the upper bound follows from [Guelev *et al.*, 2011, Prop. 33], the lower bound from [Reif, 1984]). A brief summary of the results is presented in Figure 1; a more comprehensive overview can be found in [Bulling *et al.*, 2010]. All the complexity results in this paper are given w.r.t. the number of transitions in the model and the length of the formula.

In contrast, model checking for perfect information strategies is much cheaper, namely **P**-complete [Alur *et al.*, 2002].

3 Abilities of Single Agents: Imperfect Recall

Model checking agents with imperfect information is significantly harder than ones with perfect information. But what if the agents have *almost* perfect information, e.g., their information sets are of size at most 2? Or, symmetrically, they have almost no incoming information (say, all the states are split between only 2 information sets)? In this paper, we systematically study the subproblems generated by such assumptions. In the next two sections we look at the simpler case of

¹ Where $\Delta_2^P = \mathbf{P}^{\text{NP}}$ is the class of problems solvable in polynomial time by a deterministic Turing machine sending adaptive queries to an oracle for **NP**.

Single agents	Small info sets ($ is \leq const$)	Few info sets ($\#is \leq const$)
Memoryless	Δ_2^P -complete	P-complete
Perfect recall	P-complete	in PSPACE for $\#is = 1$ EXPTIME-c. for $\#is > 1$

Figure 2: Model checking complexity for abilities of single agents

individual abilities, i.e., when only singleton coalitions are allowed in the formulae. We refer to the fragment of **ATL** containing only such formulae as **1ATL**. Later, in Section 5, we consider arbitrary coalitional strategies.

To help the reader navigate through the maze of formal arguments, we summarize our findings now. An outline of the main results is presented in Figure 2. On the one hand, we distinguish between agents playing memoryless strategies (i.e., **1ATL_{ir}**) and agents with perfect recall (i.e., **1ATL_{iR}**). On the other hand, we look at models of almost perfect information (information sets of constant size, or bounded by a constant) and models of almost nil observability (constant number of information sets per agent). The cases with complexity lower than for the general problem are highlighted. As it turns out, if we consistently pair *weak* observability with *weak* recall, or almost perfect observability with perfect recall, model checking becomes easy. Interestingly, the complexity decreases also in the case of blindfold memoryful agents (essentially, agents who can only count).

3.1 Agents that Don't Miss Much (Small Info Sets)

ATL_{ir} is appealing in practice because it avoids overestimating the agents' epistemic capabilities, and at the same time yields the lowest verification complexity. We already mentioned that model checking for perfect information strategies is tractable. One would hope that, for agents with *almost perfect information*, the complexity is similarly low. Unfortunately, it turns out to be as hard as in the general case.

Theorem 1. *Model checking **1ATL_{ir}** over CEGMs with information sets of size at most 2 is Δ_2^P -complete.*

The proof relies on a translation \mathcal{T} of models and formulae to be model-checked, so that each transition executing an action for agent 1 is split into a tree of transitions, with only doubleton information sets on the way. Moreover, there is a one-to-one correspondence between potentially successful uniform strategies for agent 1 in the new and the original model. In consequence, we get that $M, q \models_{ir} \langle\langle 1 \rangle\rangle \phi$ iff $\mathcal{T}(M), q \models_{ir} \mathcal{T}(\langle\langle 1 \rangle\rangle \phi)$. We explain the translation now, and use it to complete the proof at the end of the subsection.

Formula Translation

Let **ATL_U¹** be the subset of **ATL** using only agent 1 in strategic operators and only the Until modality. We will modify the model by adding new states, hence we introduce a fresh proposition *real* to label the original states. Now, for each $\phi, \phi' \in \mathbf{ATL}_U^1$ and $p \in \Pi$:

- $\mathcal{T}(p) = p$, $\mathcal{T}(\phi \wedge \phi') = \mathcal{T}(\phi) \wedge \mathcal{T}(\phi')$, $\mathcal{T}(\neg \phi) = \neg \mathcal{T}(\phi)$,
- $\mathcal{T}(\langle\langle 1 \rangle\rangle \phi \cup \phi') = \langle\langle 1 \rangle\rangle (\text{real} \implies \mathcal{T}(\phi)) \cup (\text{real} \wedge \mathcal{T}(\phi'))$.

Model Translation: Idea

The transformation of models is more involved. Let $M = \langle \text{Agt}, St, \Pi, \pi, Act, d, o \rangle$ be a CEGM with at least two agents, such that $\text{real} \notin \Pi$. Let $q_0 \in St$ and $\mathcal{Q} = \{q_0\}_{\sim_1} = \{q_0, q_1, \dots, q_k\}$, where $k > 2$. We build a model $M_{\mathcal{Q}}$ that “simulates” the outgoing transitions in \mathcal{Q} by a forest where $|is| \leq 2$ and the strategic abilities of agent 1 remain the same.

For convenience, denote $Acts = d_1(q_0)$ and introduce a new dummy action *nop* of agent 2. We also define a magic number $H = \binom{|\mathcal{Q}|}{2} \times |Acts| \cdot (|Acts| - 1)$, later used as the “height” of the structure that replaces \mathcal{Q} after transformation. Now, for each $q_i \in \mathcal{Q}$ and $\alpha \in Acts$, define the set of new states $q_i, q_i^\alpha, q_i^{\alpha\alpha}, \dots, q_i^{\alpha^H}$, and denote $\mathcal{Q}' = \{q_i^{\alpha^n} \mid q_i \in \mathcal{Q} \text{ and } 0 \leq n \leq H\}$ (by convention, $a^0 = \epsilon$). We also introduce transitions $q_i^{\alpha^n} \xrightarrow{(\alpha, \text{nop})} q_i^{\alpha^{n+1}}$ for all $0 \leq n < H$. Moreover, we introduce a fresh state *sink* and put $q_i^{\alpha^n} \xrightarrow{(\beta, \gamma)} \text{sink}$, for all $0 < n \leq H$ and $\gamma \in d_2(q_i)$, where $\alpha \neq \beta$. Intuitively, for a given $\alpha \in Acts$, once the transition labeled with α is selected in q_i , the same action α needs to be executed until reaching $q_i^{\alpha^H}$ if *sink* is to be avoided.

We now define the indistinguishability relation \sim^* on \mathcal{Q}' for agent 1 as any equivalence relation on \mathcal{Q}' s.t. for each $q \in \mathcal{Q}'$ we have $|\{q\}_{\sim^*}| \leq 2$ and for all $q_i, q_j \in \mathcal{Q}$:

$$\forall \alpha, \beta \in Acts ((q_i \neq q_j \wedge \alpha \neq \beta) \implies \exists n q_i^{\alpha^n} \sim^* q_j^{\alpha^n}) \quad (\clubsuit)$$

So far we have created a temporal-epistemic structure over the set $\mathcal{Q}' \cup \{\text{sink}\}$. While this construction may seem involved, it serves a simple purpose. Observe that a uniform strategy for agent 1 can enforce a path from q_i to $q_i^{\alpha^H}$ only by repeatedly executing the action $\alpha \in Acts$; any deviation from choosing α is “punished” by an immediate transition to *sink*. Thus, the requirement of uniformity together with Condition (\clubsuit) yield that if $q_i^{\alpha^H}$ is reached from q_i and $q_j^{\alpha^H}$ is reached from q_j by the same strategy, then the strategy repeatedly executes the same action over both paths.

The selected value of H enables a construction that satisfies Condition (\clubsuit) . An instance of the construction is shown in Fig. 3, where $\mathcal{Q} = \{q_0, q_1, q_2\}$ and $d_1(q_0) = \{A, B\}$. The transitions to the *sink* state are omitted. The key to understanding the magic formula H is to notice that for each pair of states (hence the Newton symbol) we discern each pair of different actions by adding a new level to the tower in Fig. 3.

Model Translation: Formal Construction

The CEGM $M_{\mathcal{Q}} = \langle \text{Agt}, St', \Pi', \pi', Act \cup \{\text{nop}\}, d', o' \rangle$ is defined by:

- $St' = (St \setminus \mathcal{Q}) \cup \mathcal{Q}' \cup \{\text{sink}\}$ and $\Pi' = \Pi \cup \{\text{real}\}$;
- $\pi'(q) = \pi(q) \cup \{\text{real}\}$ for all $q \in St$ and $\pi'(q) = \emptyset$ for the remaining states;
- the new protocol:
 - $d'_i(q) = d_i(q)$, for all $q \in St \setminus \mathcal{Q}$ and $i \in \{1, 2\}$ (the inherited protocol),
 - $d'_1(q^{\alpha^n}) = d_1(q)$ for all $q^{\alpha^n} \in \mathcal{Q}'$,
 - $d'_2(q^{\alpha^H}) = d_2(q)$ for each $q^{\alpha^H} \in \mathcal{Q}'$,
 - $d'_i(q) = \{\text{nop}\}$ otherwise, where $i \in \{1, 2\}$;

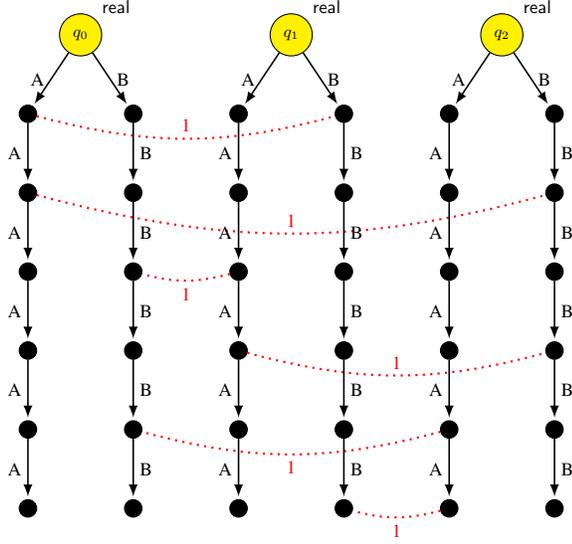


Figure 3: Enforcing uniformity for action A.

- the new transition function:

$$o'(q, \alpha, \beta) = \begin{cases} o(q, \alpha, \beta) & \text{if } q \in St' \setminus \mathcal{Q}' \text{ or } q = q_i^{\alpha^H} \in \mathcal{Q}' \\ \text{as defined above for the remaining cases.} & \end{cases}$$

Note that all the states copied from M are labeled with *real*. The new transition function o' behaves as follows: (1) if an action is executed in a state q outside of \mathcal{Q}' , then the outcome is the same as for o ; (2) if $q = q_i \in \mathcal{Q}'$ and $q \neq q_i^{\alpha^H}$, then agent 1 is in control and can decide to either execute α and move towards $q_i^{\alpha^H}$ or dive in *sink*; (3) if $q = q_i^{\alpha^H}$, then agent 2 regains its part of control.

Finally, we define the indistinguishability relation \sim'_1 of agent 1 over $M_{\mathcal{Q}}$ by requesting that $q \sim'_1 q'$ iff $q, q' \in St \setminus \mathcal{Q}$ and $q \sim_1 q'$ or $q, q' \in \mathcal{Q}'$ and $q \sim^* q'$.

We can now define the final translation $\mathcal{T}(M)$ of CEGM M . Namely, $\mathcal{T}(M)$ is obtained by an iterative reduction of all information sets of size greater than 2 until there are none.

The construction preserves enforceability for agent 1 on the set of *objective* outcome paths:

Theorem 2. For each $q \in St$ and $\phi \in \text{ATL}_{\text{U}}^1$ that does not contain *real*: $M, q \models_{or} \phi \iff \mathcal{T}(M), q \models_{or} \mathcal{T}(\phi)$.

Proof sketch. The proof follows by induction on structure of ϕ . It is sufficient to prove the thesis for a single step of reduction, i.e., $M_{\mathcal{Q}}$ instead $\mathcal{T}(M)$. We omit the details of a rather tedious but not difficult proof due to lack of space. \square

We can now proceed to the proof of Theorem 1.

Proof sketch. We only need to show Δ_2^{P} -hardness. The method used to this end in [Jamroga and Dix, 2006] is based on a reduction of SNSAT_2 [Laroussinie *et al.*, 2001] to model checking of certain ATL_{ir} formulae over two-player CEGMs. Namely, a set F of propositional formulae in CNF is given as an instance of SNSAT_2 , and each of them is translated into a CEGM component in a satisfiability-encoding manner, see [Jamroga and Dix, 2006, Sec. 3.1 and Fig. 2].

The resulting model is denoted by M_{Δ} . In [Jamroga and Dix, 2006, Theorem 4], a formula $\Phi_{\Delta} \in \text{ATL}_{\text{ir}}$ is proposed such that F is satisfiable iff $M_{\Delta} \models_{ir} \Phi_{\Delta}$. The formula contains the “Next” operator, but it can be easily replaced by “Until” to obtain a satisfiability-preserving formula $\Phi'_{\Delta} \in \text{ATL}_{\text{U}}^1$. Moreover, on M_{Δ} , the *or* and *ir* semantics of Φ'_{Δ} coincide. The same applies to their translations $\mathcal{T}(M_{\Delta})$ and $\mathcal{T}(\Phi'_{\Delta})$. Thus, by Theorem 2 we get that $\text{SNSAT}_2(F)$ iff $M_{\Delta} \models_{ir} \Phi'_{\Delta}$ iff $M_{\Delta} \models_{or} \Phi'_{\Delta}$ iff $\mathcal{T}(M_{\Delta}) \models_{or} \mathcal{T}(\Phi')$ iff $\mathcal{T}(M_{\Delta}) \models_{ir} \mathcal{T}(\Phi')$. All the transformations can be done in polynomial time w.r.t. the size of the inputs. \square

3.2 Agents that Don’t See Much (Few Info Sets)

For memoryless agents with limited observational capabilities, model checking becomes easy.

Theorem 3. Let k be a constant. Model checking $\mathbf{1ATL}_{\text{ir}}$ over the class of CEGMs with $\#is \leq k$ is **P**-complete.

Proof. The lower bound follows from **P**-completeness of $\mathbf{1ATL}_{\text{ir}}$ [Alur *et al.*, 2002]. For the upper bound, observe that each agent has only $O(|Act|^k)$ available strategies, and one can determine if a given strategy is winning in linear time by **CTL** model checking. Thus, we can check the strategies one by one in deterministic polynomial time. \square

4 Abilities of Single Agents: Perfect Recall

We continue the analysis from the previous section, now turning to specifications in $\mathbf{1ATL}_{\text{ir}}$.

4.1 Good Memory, Agents that Don’t Miss Much

Model checking of agents with perfect recall and almost perfect information also becomes easy.

Theorem 4. Let k be a constant. Model checking $\mathbf{1ATL}_{\text{ir}}$ over CEGMs with $|is| \leq k$ is **P**-complete.

Proof. The lower bound follows from **P**-completeness of $\mathbf{1ATL}_{\text{ir}}$ [Alur *et al.*, 2002]. For the upper bound, we use the construction in [Guelev *et al.*, 2011, Section 6] that translates model checking of $\mathbf{1ATL}_{\text{ir}}$ in CEGM M to verification of perfect information strategies in a CGS M' . Note that the number of transitions in the new model is $|M'| = O(|M| \cdot 2^{|is|}) \leq O(|M| \cdot 2^k) = O(|M|)$. Moreover, model checking for perfect information can be done in polynomial time w.r.t. the size of the model and the formula. \square

4.2 Good Memory, Agents that Don’t See Much

Consider now the case of models with few information sets.

Theorem 5. Model checking $\mathbf{1ATL}_{\text{ir}}$ over CEGMs with at most 2 information sets per agent is **EXPTIME**-complete.

The inclusion is straightforward from [Guelev *et al.*, 2011, Prop. 33]. To obtain the lower bound, we develop a reduction from model checking CEGMs with arbitrarily many observations to CEGMs with only two available observations, modeled by the “white” and the “dark” states. The idea is that, for a transition $q \longrightarrow q'$, the observation provided by the target state q' is simulated by a sequence of intermediate “dark” and “white” states, signalling in binary the index of the original information set $[q']$.

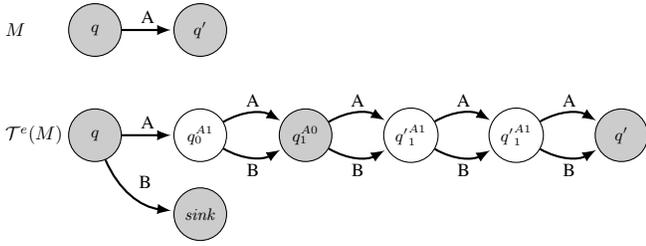


Figure 4: Reducing the number of information sets for a single agent

Model Translation

Let $M = \langle \text{Agt}, St, \Pi, \pi, Act, d, o \rangle$ be a CEGM and $n = |St/\sim_1|$ be the number of information sets for agent $1 \in \text{Agt}$. We label the information sets of \sim_1 with ordinals from 0 to $n - 1$. Let $ctr(q)$ be the index of $[q]_{\sim_1}$. Moreover, $ctr_i(q)$ denotes the i th bit of binary representation of $ctr(q)$, for $0 \leq i \leq \lceil \log n \rceil$.

For any transition $q \xrightarrow{\gamma} q'$ in M , we introduce $2^{\lceil \log n \rceil}$ fresh states $F = \{q_i^{\gamma,0}, q_i^{\gamma,1}\}_{i=0}^{\lceil \log n \rceil}$ and the usual *sink* state. We use the states in F to encode $ctr(q)$ and $ctr(q')$. Namely, we replace $q \xrightarrow{\gamma} q'$ with the following sequence:

$$q \xrightarrow{\gamma} \quad (1)$$

$$q_0^{\gamma, ctr_0(q)} \xrightarrow{\star} \dots \xrightarrow{\star} q_{\lceil \log n \rceil}^{\gamma, ctr_{\lceil \log n \rceil}(q)} \xrightarrow{\star} \quad (2)$$

$$q_0^{\gamma', ctr_0(q')} \xrightarrow{\star} \dots \xrightarrow{\star} q_{\lceil \log n \rceil}^{\gamma', ctr_{\lceil \log n \rceil}(q')} \quad (3)$$

$$\xrightarrow{\star} q' \quad (4)$$

where \star should be replaced with the bundle of all possible actions for the grand coalition, i.e., $Act^{\lceil \text{Agt} \rceil}$. This process is repeated for each transition in M . We will denote part (2) of the sequence by $enc(q)$, and part (3) by $enc'(q')$.

In the next stage we unify the protocol for agent 1 by firstly adding the usual fresh *sink* state and adding for each action $\alpha \in Act \setminus d_1(q)$ and $\beta \in d_{\text{Agt} \setminus \{1\}}(q)$ a new transition $q \xrightarrow{(\alpha, \beta)} \text{sink}$. Intuitively, agent 1 is punished for not following the original protocol by a transition to *sink*. We also copy the labeling of q to all the intermediate states $\hat{q} \in F$, i.e., $q \in \pi(p)$ iff $\hat{q} \in \pi(p)$ for all $p \in \Pi$.

Finally, we add the indistinguishability relation for agent 1 as follows: (1) the original states from St , *sink*, and $\{q_i^{\gamma,0}\}_{i=0}^{\lceil \log n \rceil}$ are colored grey; (2) the states from $\{q_i^{\gamma,1}\}_{i=0}^{\lceil \log n \rceil}$ are colored white; (3) we assume that agent 1 can observe only the color of a state. The resulting CEGM is denoted by $\mathcal{T}^e(M)$.

We illustrate the construction in Fig. 4. In the example, there is only one agent, and states q, q' belong to two information sets with $ctr(q) = 10_{bin}$ and $ctr(q') = 11_{bin}$. Moreover, $Act = \{A, B\}$ and $d_1(q) = \{A\}$.

Theorem 6. For ATL_{iR} formulae $\langle\langle A \rangle\rangle \gamma$ containing no nested strategic modalities and no operator X , we have: $M, q \models_{iR} \langle\langle A \rangle\rangle \gamma \iff \mathcal{T}^e(M), q \models_{iR} \langle\langle A \rangle\rangle \gamma$.

Proof sketch. For simplicity let $A = \{1\}$. The proof of the general case differs only in the number of the fresh states used

to encode the observations of the protagonist coalition. A possible scheme of such encoding can be based on enumerating the information sets of each $i \in A$ with functions $ctr^i(\cdot)$. Then, each transition $q \xrightarrow{\gamma} q'$ is swapped with its replacement, where the agents of the coalition take turns in a fixed order to enforce the encoding of source and target information sets using separate fresh locations. These locations are created and labeled as in the single-agent above.

We establish a correspondence between iR -strategies over M and $\mathcal{T}^e(M)$ that preserves U and G. Let $\lambda = q_0 q_1 \dots q_k$ be a history in M . By $lft(\lambda) = q_0 enc(q_0) enc'(q_1) q_1 \dots enc(q_{k-1}) enc'(q_k) q_k$ we denote the lifting of λ to $\mathcal{T}^e(M)$. Note that by construction each finite history in $\mathcal{T}^e(M)$ is of form $\lambda' = q_0 enc(q_0) enc'(q_1) q_1 \dots enc(q_{k-1}) enc'(q_k) q_k \mathcal{R}$, where \mathcal{R} is a sequence containing only fresh states. We can thus define the casting of λ' to M as $cst(\lambda') = q_0 q_1 \dots q_k$.

Now, let s_1 be an iR -strategy for agent 1, over M . We define the lifting $lft(s_1)$ of s_1 to $\mathcal{T}^e(M)$ as a function s.t. $lft(s_1)(\lambda') = s_1(cst(\lambda'))$ for all the histories λ' over $\mathcal{T}^e(M)$ such that $\lambda'[0], \lambda'_F \notin F$ and $lft(s_1)(\lambda') = B$ for a fixed $B \in Act$ for all the remaining histories. Intuitively, $lft(s_1)$ makes the same choices as s_1 , unless the path reaches a fresh state where the fixed action is used.

Let s'_1 be an iR -strategy for agent 1, over $\mathcal{T}^e(M)$. The casting $cst(s'_1)$ of s'_1 to M is a function s.t. $cst(s'_1)(\lambda) = s'_1(lft(\lambda))$ for each history λ over M . Intuitively, $cst(s'_1)$ makes the same choices as s'_1 while ignoring the fresh states.

To conclude the proof, it is easy to see that s_1 enforces p along each path starting from q iff $lft(s_1)$ does so. Moreover, it is routine to show that the uniformity is preserved, i.e., $lft(s_1)$ and $cst(s'_1)$ as defined above are iR -strategies. \square

The lower bound in Theorem 5 follows from Theorem 6 and the **EXPTIME**-hardness of solving reachability games with imperfect information [Reif, 1984].

4.3 Special Case: Blindfold Agents with Recall

Two information sets are enough to make model checking ATL_{iR} as hard as in arbitrary models. What happens if there is only a single information set, comprising of the whole state space? Such CEGMs are called *blindfold*. Note that, in blindfold CEGMs, memoryful strategies for any $A \subseteq \text{Agt}$ can be interpreted as functions $\sigma_A: \mathbb{N} \rightarrow Act$, where $\sigma_A(i)$ is the joint action selected by A in the i th step. Moreover, the set of the outcome paths of σ_A does not depend on the initial state, so one can as well write $out(\sigma_A)$ instead of $out(q, \sigma_A)$.

Interestingly, model checking ATL_{iR} over blindfold CEGMs is easier than in the general case.

Theorem 7. Model checking ATL_{iR} over the class of blindfold CEGMs is in **PSPACE**.

Proof. It suffices to show that $\langle\langle A \rangle\rangle p \text{ U } r$ and $\langle\langle A \rangle\rangle G p$ can be verified in **PSPACE**, where $A \subseteq \text{Agt}$ and $p, r \in \Pi$. The idea is as follows: firstly we show that if these properties are true, then they can be attained by using finite strategies with memory that needs to cover only all the possible subsets of the state space. Secondly, we use this limit to build a model checker in a form of a non-deterministic Turing machine.

Let us start with $q \models \langle\langle A \rangle\rangle \text{pUr}$ and let $\sigma_A: \mathbb{N} \rightarrow \text{Act}^A$ be a joint strategy for A s.t. $\lambda \models \text{pUr}$ for each $\lambda \in \text{out}(\sigma_A)$. For each $i \in \mathbb{N}$ we inductively define the set A_i as follows: $A_0 = St \setminus \llbracket r \rrbracket$ and $A_{i+1} = \{\text{states reachable in one step from } A_i \text{ via } \sigma_A\} \setminus \llbracket r \rrbracket$ for $i > 0$.

It follows from the definition of the Until modality that there exists the smallest index $k_{fin} \in \mathbb{N}$ s.t. $A_i = \emptyset$ for all $i \geq k_{fin}$. Now, let us select any $B \in \{A_i\}_{i=0}^{k_{fin}}$ and let $k_{min}, k_{max} \in \mathbb{N}$ be the minimal and maximal, resp., indices s.t. $A_{k_{min}} = B = A_{k_{max}}$. If $k_{min} < k_{max}$, then we can transform σ_A into σ_A^B as follows: $\sigma_A^B(i) = \sigma_A(i)$ for all $0 \leq i < k_{min}$ and $\sigma_A^B(i) = \sigma_A(i + k_{max} - k_{min})$ for all $i \geq k_{min}$. It is easy to see that $\lambda \models \text{pUr}$ for each $\lambda \in \text{out}(\sigma_A^B)$. The process of recomputing the sets $\{A_i\}_{i \in I}$ and further reducing the working strategy can be repeated until no reduction is possible, i.e., the family $\{A_i\}_{i \in I}$ contains no repetitions. This in turn means that $I \leq 2^{|St|}$ and $k_{fin} \leq 2^{|St|}$. Therefore, there exists a joint strategy for A that enforces pUr in less than $2^{|St|}$ steps.

The construction for $\langle\langle A \rangle\rangle \text{Gp}$ follows analogously.

To build a non-deterministic Turing machine for $\langle\langle A \rangle\rangle \text{pUr}$ and $\langle\langle A \rangle\rangle \text{Gp}$, we equip it with a deterministic $|St|$ -bit counter that enables to track the progress of execution up to $2^{|St|}$ steps. The machine consecutively guesses joint actions for A for the current step indicated by the counter, executes them and then increments the counter. Only recently reached states are preserved. The machine rejects if the counter exceeds $2^{|St|}$ or a state violating the verified property has been reached. It accepts if while traversing along p -labeled states a state labeled with r has been reached along each path (the case of pUr) or a loop has been detected (the case of Gp). \square

5 Abilities of Coalitions

In Sections 3 and 4, we focused on formulae containing only singleton coalitions. We now briefly wrap up the study, presenting analogous results for multi-player teams. A summary is shown in Figure 5; again, the cases with lower complexity than for the general problem are highlighted.

We observe that, for models with small information sets, Δ_2^P -completeness follows from Theorem 1 and the complexity of the general problem [Schobbens, 2004; Jamroga and Dix, 2006]. Moreover, Theorem 7 (inclusion in PSPACE for blindfold agents) is formulated and proved for the whole language of ATL_{IR} . Finally, our reduction in Theorem 6 works also for coalitional abilities. Thus, model checking ATL_{IR} for $\#is = k$ and $k \geq 2$ is undecidable.²

The last case that we address is that of coalitional abilities for memoryless agents with weak observational capabilities.

Theorem 8. *Model checking ATL_{IR} over CEGMs with at most 2 information sets per agent is Δ_2^P -complete.*

Proof sketch. We adapt the Δ_2^P -hardness proof in [Jamroga and Dix, 2006] by using a team of verifiers $V = \{v_0^1, \dots, v_0^{n*m}, v_1, \dots, v_k\}$ where n is the number of nested queries, m is the maximal number of clauses per query and

² Alternatively, one can observe that the undecidability proof in [Dima and Tiplea, 2011] actually uses a model with $\#is = 2$.

Coalitions	Small info sets ($ is \leq \text{const}$)	Few info sets ($\#is \leq \text{const}$)
Memoryless	Δ_2^P -complete	P-compl./Δ_2^P-compl.
Perfect recall	?	in PSPACE for $\#is = 1$ undecidable for $\#is > 1$

Figure 5: Model checking complexity for abilities of coalitions

k is the number of propositional variables in the instance of SNSAT_2 . Each agent v_0^i controls the choice of the literal in a particular clause, and each agent v_j controls the valuation of the Boolean variable underlying “her” literal. Every agent can only distinguish between the states she controls and the rest of the state space. Finally, we replace each occurrence of $\langle\langle v \rangle\rangle$ with $\langle\langle V \rangle\rangle$ in the formula from [Jamroga and Dix, 2006], and the reduction goes through. \square

Note that the above proof requires that the number of agents in the class of models is variable (and is a parameter of the model checking problem). As it turns out, the requirement is essential for Theorem 8 to hold. This is especially important, as a fixed finite set of agents is often assumed beforehand, when defining the syntax of the agent logic. In those cases, model checking agents with limited epistemic capabilities is easy even in the coalitional case.

Theorem 9. *Let k, n be constants. Model checking ATL_{IR} over the class of CEGMs with at most n agents and at most k information sets per agent is **P**-complete.*

Proof. Straightforward extension of the proof of Theorem 3 (the number of coalitional strategies is now polynomial). \square

6 Conclusions

Verification of autonomous agents in multi-agent systems is an important research path. Despite some recent advances [Huang and van der Meyden, 2014; Pilecki *et al.*, 2014; Busard, 2017; Jamroga *et al.*, 2017], the problem is still open due to its inherent computational complexity. In this paper, we show that the complexity is in fact lower than expected in some borderline cases, in particular for agents with consistently good (resp. weak) observational and mental capabilities. We also show that, for agents whose capabilities are “in between,” the problem is as hard as in the general case.

An interesting question for the future concerns the corresponding complexity results for ATL^* . We would also like to establish the model checking complexity for coalitions with perfect recall and almost perfect information.

Acknowledgements

The authors acknowledge the support of the National Centre for Research and Development, Poland (NCBR), and the Luxembourg National Research Fund (FNR), under the PolLux/FNR-INTER project VoteVerif (POLLUX-IV/1/2016).

References

- [Ågotnes and Walther, 2009] T. Ågotnes and D. Walther. A logic of strategic ability under bounded memory. *Journal of Logic, Language and Information*, 18(1):55–77, 2009.
- [Ågotnes et al., 2007] T. Ågotnes, V. Goranko, and W. Jamroga. Alternating-time temporal logics with irrevocable strategies. In *Proceedings of TARK XI*, pages 15–24, 2007.
- [Ågotnes et al., 2015] T. Ågotnes, V. Goranko, W. Jamroga, and M. Wooldridge. Knowledge and ability. In H.P. van Ditmarsch, J.Y. Halpern, W. van der Hoek, and B.P. Kooi, editors, *Handbook of Epistemic Logic*, pages 543–589. College Publications, 2015.
- [Alur et al., 2002] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.
- [Belardinelli et al., 2017] F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin. Verification of multi-agent systems with imperfect information and public actions. In *Proceedings of AAMAS*, pages 1268–1276, 2017.
- [Berwanger and Kaiser, 2010] D. Berwanger and L. Kaiser. Information tracking in games on graphs. *Journal of Logic, Language and Information*, 19(4):395–412, 2010.
- [Berwanger et al., 2018] D. Berwanger, A. B. Mathew, and M. van den Bogaard. Hierarchical information and the synthesis of distributed strategies. *Acta Inf.*, 55(8):669–701, 2018.
- [Bulling et al., 2010] N. Bulling, J. Dix, and W. Jamroga. Model checking logics of strategic ability: Complexity. In M. Dastani, K. Hindriks, and J.-J. Meyer, editors, *Specification and Verification of Multi-Agent Systems*, pages 125–159. Springer, 2010.
- [Busard, 2017] S. Busard. *Symbolic Model Checking of Multi-Modal Logics: Uniform Strategies and Rich Explanations*. PhD thesis, Universite Catholique de Louvain, 2017.
- [Chatterjee et al., 2007] K. Chatterjee, L. Doyen, T.A. Henzinger, and J.-F. Raskin. Algorithms for omega-regular games of incomplete information. *Logical Methods in Computer Science*, 3(3), 2007.
- [Dima and Tiplea, 2011] C. Dima and F.L. Tiplea. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. *CoRR*, abs/1102.4225, 2011.
- [Doyen and Raskin, 2011] L. Doyen and J.-F. Raskin. Games with imperfect information: Theory and algorithms. In *Lecture Notes in Game Theory for Computer Scientists*, pages 185–212. Cambridge University Press, 2011.
- [Guelev et al., 2011] D.P. Guelev, C. Dima, and C. Enea. An alternating-time temporal logic with knowledge, perfect recall and past: axiomatisation and model-checking. *Journal of Applied Non-Classical Logics*, 21(1):93–131, 2011.
- [Huang and van der Meyden, 2014] X. Huang and R. van der Meyden. Symbolic model checking epistemic strategy logic. In *Proceedings of AAAI Conference on Artificial Intelligence*, pages 1426–1432, 2014.
- [Jamroga and Dix, 2006] W. Jamroga and J. Dix. Model checking ATL_{ir} is indeed Δ_2^P -complete. In *Proceedings of EUMAS*, volume 223 of *CEUR Workshop Proceedings*, 2006.
- [Jamroga and van der Hoek, 2004] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2–3):185–219, 2004.
- [Jamroga et al., 2017] W. Jamroga, M. Knapik, and D. Kurpiewski. Fixpoint approximation of strategic abilities under imperfect information. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1241–1249. IFAAMAS, 2017.
- [Jamroga, 2003] W. Jamroga. Some remarks on alternating temporal epistemic logic. In B. Dunin-Keplicz and R. Verbrugge, editors, *Proceedings of Formal Approaches to Multi-Agent Systems (FAMAS 2003)*, pages 133–140, 2003.
- [Laroussinie et al., 2001] F. Laroussinie, N. Markey, and Ph. Schnoebelen. Model checking CTL+ and FCTL is hard. In *Proceedings of FoSSaCS’01*, volume 2030 of *Lecture Notes in Computer Science*, pages 318–331. Springer, 2001.
- [Maubert and Murano, 2018] Bastien Maubert and Aniello Murano. Reasoning about knowledge and strategies under hierarchical information. In *Proceedings of KR*, pages 530–540, 2018.
- [Peterson and Reif, 1979] G. Peterson and J. Reif. Multiple-person alternation. In *Proceedings of the 20th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 348–363. IEEE Computer Society Press, 1979.
- [Pilecki et al., 2014] J. Pilecki, M.A. Bednarczyk, and W. Jamroga. Synthesis and verification of uniform strategies for multi-agent systems. In *Proceedings of CLIMA XV*, volume 8624 of *Lecture Notes in Computer Science*, pages 166–182. Springer, 2014.
- [Reif, 1984] John H. Reif. The complexity of two-player games of incomplete information. *J. Comput. Syst. Sci.*, 29(2):274–301, 1984.
- [Schewe and Finkbeiner, 2007] S. Schewe and B. Finkbeiner. Distributed synthesis for alternating-time logics. In *Proceedings of ATVA*, volume 4762 of *Lecture Notes in Computer Science*, pages 268–283. Springer, 2007.
- [Schobbens, 2004] P. Y. Schobbens. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science*, 85(2):82–93, 2004.
- [van der Hoek et al., 2006] W. van der Hoek, A. Lomuscio, and M. Wooldridge. On the complexity of practical ATL model checking. In *Proceedings of International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 201–208. ACM, 2006.