

Approximating Perfect Recall when Model Checking Strategic Abilities

Francesco Belardinelli

Laboratoire IBISC, Université d'Evry
France

Alessio Lomuscio

Imperial College London
United Kingdom

Vadim Malvone

Laboratoire IBISC, Université d'Evry
France

Abstract

We investigate the notion of bounded recall in the context of model checking ATL^* and ATL specifications in multi-agent systems with imperfect information. We present a novel three-valued semantics for ATL^* , respectively ATL , under bounded recall and imperfect information, and study the corresponding model checking problems. Most importantly, we show that the three-valued semantics constitutes an approximation with respect to the traditional two-valued semantics. In the light of this we construct a sound, albeit partial, algorithm for model checking two-valued perfect recall via its approximation as three-valued bounded recall.

1 Introduction

Alternating-time Temporal Logic (ATL) and its extension ATL^* are widely used formalisms to reason about strategic abilities in multi-agent systems (Alur, Henzinger, and Kupferman 2002). At the heart of ATL is the notion of the temporal sequence of events that a coalition of agents can bring about in a system, irrespective of the actions of other agents outside the coalition. ATL has been extended in various directions giving raise to even more expressive formalisms, e.g., by taking into account time (André et al. 2017), bounded resources (Alechina et al. 2015), epistemic concepts (Ågotnes et al. 2015), and beyond.

A key consideration when using expressive specification languages, including ATL , is the computational complexity of the resulting model checking problem. In the case of ATL , this was shown to be $PTIME$ -complete under perfect information (Alur, Henzinger, and Kupferman 2002). Agents in a multi-agent system (MAS), however, typically operate under imperfect information about the other agents and the environment. Once imperfect information is assumed, the resulting model checking problem becomes Δ_2^P -complete under observational semantics (Jamroga and Dix 2006) and it is undecidable under perfect recall (Dima and Tiplea 2011). The latter case is particularly problematic as (i) it is a traditionally natural set-up to consider (Meyden and Shilov 1999), and (ii) it hinders the development of any verification toolkit.

Recent approaches have attempted to overcome such problems. For instance, recent works show that if agents

can only communicate via broadcasting, decidability can be retained although with high complexity (Belardinelli et al. 2017). Further, hierarchical systems where information is shared in a strictly pre-determined manner have also been shown to provide decidable fragments (Berthon et al. 2017). These and other contributions analyse the verification problem under perfect recall and imperfect information by restricting the class of MAS.

While we share the objectives that have previously been pursued in this area, here we make a significant departure from the approaches above. In particular, rather than assuming perfect recall, we explore the verification problem against the weaker notion of *bounded recall*. In a nutshell, under n -bounded recall an agent's strategy does not depend on her whole history, but only on her last n steps. This is a reasonable assumption on the abilities of agents.

Similar notions of resource-bounded strategies have been previously investigated in the literature. In particular, as we highlight when discussing the related work in Section 5, (Vester 2013) introduced strategies as finite-memory transducers. While related, our treatment is nonetheless different. Indeed, as we show in Section 5, some finite-memory transducers cannot be translated polynomially into our bounded recall strategies and some bounded recall strategies cannot be polynomially recast as transducers. A further key point of departure is that our notion of bounded recall is intended to provide a basis for an iterative verification procedure for MAS based on a novel three-valued semantics.

Our treatment is also motivated by practical considerations. Firstly, bounded recall is a useful concept when modelling concrete systems, as no real-life MAS can have unbounded memory. Secondly, as we show in Section 3, in cases of interest bounded recall can provide a provably sound approximation of perfect recall during verification. A key result that we prove (Corollary 1) states that MAS properties under perfect recall can be determined by analysing their bounded recall approximations.

Structure of the paper. In Section 2 we introduce the notion of bounded recall in the context of interpreted systems and ATL^* , and compare bounded and perfect recall from the perspective of verification. In Section 3 we develop a novel three-valued semantics for bounded as well as perfect recall and study the corresponding model checking problems. Then, we analyse its formal properties against the

classic version in Section 2. These theoretical results lay the foundations for a verification procedure, presented in Section 4, for model checking MAS under imperfect information and perfect recall via iterative checking of bounded recall versions of the same MAS in the three-valued semantics with increasing amount of memory. While the algorithm is incomplete in general, we show that if a bound on recall is assumed, it terminates in *PSPACE*. We discuss the related work in depth in Section 5 and conclude in Section 6.

2 Classic Bounded Recall

In this section we introduce a two-valued semantics for ATL^* under imperfect information and bounded recall. Then, we study the complexity of the corresponding model checking problem, and compare it with perfect recall. Hereafter we assume sets $Ag = \{1, \dots, m\}$ of indices for agents and AP of atomic propositions. Given a set U , \bar{U} denotes its complement. We denote the length of a tuple v as $|v|$, and its i th element either as v_i or $v.i$. Let $last(v) = v_{|v|}$ be the last element in v . For $i \leq |v|$, let $v_{\geq i}$ be the suffix $v_i, \dots, v_{|v|}$ of v starting at v_i and $v_{< i}$ the prefix v_1, \dots, v_i of v starting at v_1 . Finally, $\mathbb{N}^+ = \mathbb{N} \setminus \{0\}$ is the set of positive naturals.

2.1 Interpreted Systems

We follow the presentation of interpreted systems as given in (Fagin et al. 1995). We will use these as a semantics for ATL^* as originally done in (Lomuscio and Raimondi 2006), rather than concurrent game structures.

Definition 1 (Agent). Given a set Ag of indices for agents, an *agent* is a tuple $i = \langle L_i, Act_i, P_i, t_i \rangle$ such that

- L_i is the finite set of *local states*;
- Act_i is the finite set of *individual actions*;
- $P_i : L_i \rightarrow (2^{Act_i} \setminus \emptyset)$ is the *protocol function*;
- $t_i : L_i \times ACT \rightarrow L_i$ is the *local transition function*, where $ACT = Act_1 \times \dots \times Act_{|Ag|}$ is the set of *joint actions*, s.t. for $l \in L_i$, $a \in ACT$, $t_i(l, a)$ is defined iff $a_i \in P_i(l)$.

By Def. 1 an agent i is situated in some local state $l \in L_i$, which represents the information she has about the system. At any state she can perform the actions in Act_i according to protocol P_i . A joint action brings about a change in the state of the agent, according to transition function t_i . Hereafter we identify an agent index i with the corresponding agent.

Given set Ag of agents, a *global state* $s \in \mathcal{G}$ is a tuple $\langle l_1, \dots, l_{|Ag|} \rangle$ of local states, one for each agent in Ag . Notice that an agent's protocol and transition function depend only on his local state, which might contain strictly less information than the global state. In this sense agents have *imperfect information* about the system. A history $h \in \mathcal{G}^+$ is a finite (non-empty) sequence of global states. For $n \geq 1$, \mathcal{G}^n denotes the set of histories of length n , and $\mathcal{G}^{<1+n} = \bigcup_{1 \leq m \leq n} \mathcal{G}^m$ is the set of histories of length *at most* n . Hereafter, $\mathcal{G}^{<\omega}$ denotes the set of all finite histories, i.e., $\mathcal{G}^{<\omega} = \mathcal{G}^+$.

For every agent $i \in Ag$, we define an *indistinguishability relation* \sim_i between global states based on the identity of

local states, that is, $s \sim_i s'$ iff $s_i = s'_i$ (Fagin et al. 1995). This indistinguishability relation is extended to histories in a synchronous, pointwise way, i.e., histories $h, h' \in \mathcal{G}^+$ are *indistinguishable* for agent $i \in Ag$, or $h \sim_i h'$, iff (i) $|h| = |h'|$ and (ii) for all $j \leq |h|$, $h_j \sim_i h'_j$.

Definition 2. An *interpreted system* (IS) is a tuple $M = \langle Ag, s_0, T, \Pi \rangle$, where

- Ag is the set of *agents*;
- $s_0 \in \mathcal{G}$ is the (*global*) *initial state*;
- $T : \mathcal{G} \times ACT \rightarrow \mathcal{G}$ is the *global transition function* such that $s' = T(s, a)$ iff for every $i \in Ag$, $s'_i = t_i(s_i, a)$;
- $\Pi : \mathcal{G} \times AP \rightarrow \{\text{tt}, \text{ff}\}$ is the *labelling function*.

An interpreted system describes the interactions of a group Ag of agents, starting from initial state s_0 , according to transition function T . Notice that T is defined on state s for joint action a iff $a_i \in P_i(s_i)$ for every $i \in Ag$.

2.2 ATL on Bounded Recall

We make use of the Alternating-time Temporal Logic ATL^* (Alur, Henzinger, and Kupferman 2002) to reason about the strategic abilities of agents in interpreted systems.

Definition 3 (ATL^*). State (φ) and path (ψ) formulas in ATL^* are defined as follows, for $q \in AP$ and $\Gamma \subseteq Ag$:

$$\begin{aligned} \varphi &::= q \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle\Gamma\rangle\rangle\psi \\ \psi &::= \varphi \mid \neg\psi \mid \psi \wedge \psi \mid X\psi \mid (\psi U \psi) \end{aligned}$$

Formulas in ATL^* are all and only the state formulas.

As customary, a formula $\langle\langle\Gamma\rangle\rangle\Phi$ is read as ‘the agents in coalition Γ have a strategy to achieve Φ ’. The meaning of LTL operators ‘next’ X and ‘until’ U is standard (Baier and Katoen 2008). Operators $\llbracket\Gamma\rrbracket$, F and G can be introduced as usual.

Formulas in the ATL fragment of ATL^* are obtained from Def. 3 by restricting path formulas ψ as follows, where φ is a state formula and R is the release operator¹:

$$\psi ::= X\varphi \mid (\varphi U \varphi) \mid (\varphi R \varphi)$$

In the rest of the paper we consider two other relevant fragments of ATL^* : the existential and universal fragments.

Definition 4. Let $q \in AP$ and $\Gamma \subseteq Ag$. State (φ) and path (ψ) formulas in the *existential fragment* $EATL^*$ of ATL^* are defined as follows:

$$\begin{aligned} \varphi &::= q \mid \neg q \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle\langle\Gamma\rangle\rangle\psi \\ \psi &::= \varphi \mid \psi \vee \psi \mid \psi \wedge \psi \mid X\psi \mid (\psi U \psi) \mid (\psi R \psi) \end{aligned}$$

State formulas (φ) in the *universal fragment* $AATL^*$ of ATL^* are defined as follows:

$$\varphi ::= q \mid \neg q \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \llbracket\Gamma\rrbracket\psi$$

whereas path formulas (ψ) are defined as for $EATL^*$.

¹Notice that the release operator R can be defined in ATL^* as the dual of until U (indeed, it does not appear in the syntax of Def. 3), while it must be assumed as a primitive operator in ATL .

By Def. 4 in the existential (resp. universal) fragment formulas are only of the form $\langle\langle\Gamma\rangle\rangle\psi$ (resp. $\llbracket\Gamma\rrbracket\psi$) or boolean combinations thereof.

Since the behaviour of agents in interpreted systems depends only on their local state, we assume agents employ *uniform strategies* (Jamroga and van der Hoek 2004; Lomuscio and Raimondi 2006). That is, they perform the same action whenever they have the same information. Moreover, we assume that agents have some bounded recall of the local states visited during an execution. This is formalised as follows.

Definition 5 (Uniform Strategy with Bounded Recall). For $n \in \mathbb{N}^+ \cup \{\omega\}$, a *uniform strategy with n -bounded recall* for agent $i \in Ag$ is a function $f_i^n : \mathcal{G}^{<1+n} \rightarrow Act_i$ such that for all histories $h, h' \in \mathcal{G}^{<1+n}$, (i) $f_i^n(h) \in P_i(last(h), i)$; and (ii) $h \sim_i h'$ implies $f_i^n(h) = f_i^n(h')$.

By Def. 5 any strategy for agent i has to return actions that are enabled for i . Also, whenever two histories are indistinguishable for i , then the same action is returned. Notice that for $n = 1$ we obtain memoryless (or imperfect recall) strategies; whereas for $n = \omega$ we have memoryfull (or perfect recall) strategies.

Given an IS M , a *path* p is an infinite sequence $s_1 s_2 \dots$ of states. For a set $F_\Gamma^n = \{f_i^n \mid i \in \Gamma\}$ of strategies, one for each agent in coalition Γ , a path p is F_Γ^n -*compatible* iff for every $j > 0$, $p_{j+1} = T(p_j, a)$ for some joint action $a \in ACT$ such that for every $i \in \Gamma$, $a_i = f_i^n(p_1, \dots, p_j)$ for $j \leq n$, $a_i = f_i^n(p_{j-n}, \dots, p_j)$ otherwise. Hence, for $n \in \mathbb{N}^+$, n -bounded recall strategies take into account at most the n previously visited states. This modelling choice is meant to account for agents with finite memory of past events (Ågotnes and Walther 2009; Vester 2013). In particular, any actual implementation of MAS with some sort of recall can only employ bounded recall, for some bound determined by the system's memory capacity. Finally, let $out(s, F_\Gamma^n)$ be the set of all F_Γ^n -compatible paths starting from s .

We can now assign a meaning to ATL^* formulas on IS based on a semantics with two truth values: ff and tt.

Definition 6 (Satisfaction). Let $n \in \mathbb{N}^+ \cup \{\omega\}$. The two-valued satisfaction relation \models_n^2 for an IS M , state s , path p , and ATL^* formula ϕ is defined as follows:

$$\begin{aligned} (M, s) \models_n^2 q & \quad \text{iff } \Pi(s, q) = \text{tt} \\ (M, s) \models_n^2 \neg\phi & \quad \text{iff } (M, s) \not\models_n^2 \phi \\ (M, s) \models_n^2 \phi \wedge \phi' & \quad \text{iff } (M, s) \models_n^2 \phi \text{ and } (M, s) \models_n^2 \phi' \\ (M, s) \models_n^2 \langle\langle\Gamma\rangle\rangle\psi & \quad \text{iff for some } F_\Gamma^n, \text{ for all } p \in out(s, F_\Gamma^n), \\ & \quad (M, p) \models_n^2 \psi \\ (M, p) \models_n^2 \phi & \quad \text{iff } (M, p_1) \models_n^2 \phi \\ (M, p) \models_n^2 \psi \wedge \psi' & \quad \text{iff } (M, p) \models_n^2 \psi \text{ and } (M, p) \models_n^2 \psi' \\ (M, p) \models_n^2 X\psi & \quad \text{iff } (M, p_{\geq 2}) \models_n^2 \psi \\ (M, p) \models_n^2 \psi U \psi' & \quad \text{iff for some } k \geq 1, (M, p_{\geq k}) \models_n^2 \psi', \\ & \quad \text{and for all } j, \\ & \quad 1 \leq j < k \text{ implies } (M, p_{\geq j}) \models_n^2 \psi \end{aligned}$$

Def. 6 is parametric in n for bounded recall. For $n = 1$ and $n = \omega$ we obtain respectively the logics ATL_{ir}^* and ATL_{iR}^* for imperfect information and imperfect (resp. perfect) recall (Schobbens 2004). We say that formula ϕ is

true in an IS M (for n -bounded recall), or $M \models_n^2 \phi$, iff $(M, s_0) \models_n^2 \phi$.

Further, we observe that Def. 6 corresponds to the *objective* interpretation of ATL^* , whereby formulas $\langle\langle\Gamma\rangle\rangle\psi$ are evaluated against all paths $p \in out(s, F_\Gamma^n)$ compatible with current state s and joint strategy F_Γ^n . This is a well-established semantical account in logics for strategies (Jamroga and van der Hoek 2004), which has found applications in MAS verification (Busard et al. 2015). Our results can be extended with minor modification to the *subjective* interpretation of strategy operators, according to which in state s one considers all paths starting from some state s' , epistemically consistent with s .

2.3 Model Checking Bounded Recall

We now analyse the model checking problem for bounded recall within the two-valued semantics, defined as follows.

Definition 7 (Model Checking). Given an IS M , a formula ϕ , and a bound $n \in \mathbb{N}^+ \cup \{\omega\}$, the model checking problem (for n -bounded recall) concerns determining whether $M \models_n^2 \phi$.

We immediately state that model checking ATL with perfect recall (i.e., for $n = \omega$) and imperfect information is undecidable.

Theorem 1. The model checking problem for ATL on two-valued semantics with perfect recall and imperfect information is undecidable.

Proof sketch. In (Dima and Tiplea 2011) it is proved the undecidability of model checking ATL with perfect recall and imperfect information over concurrent game structure (CGS). Then, the result follows from the fact that IS and CGS can be translated polynomially one into the other. Specifically, every IS $M = \langle Ag, s_0, T, \Pi \rangle$ induces a CGS $G_M = \langle Ag, AP, S, s_0, \{Act_i\}_{i \in Ag}, T, \Pi \rangle$ that satisfies exactly the same formulas in ATL^* . For the other direction, given a CGS G satisfying some mild conditions (which are fulfilled by the CGS used in the undecidability proof of (Dima and Tiplea 2011)), we can extract a set Ag_G of agents, inducing an IS M_G such that G and M_G satisfy the same formulas in ATL^* . \square

As an immediate consequence of Theorem 1, model checking ATL^* with perfect recall and imperfect information is also undecidable.

In contrast we show that model checking ATL^* with bounded recall and imperfect information is decidable.

Theorem 2. For $n \in \mathbb{N}^+$, the model checking problem for ATL^* (resp. ATL) under n -bounded recall and imperfect information is in $EXPTIME$. Moreover, for a fixed $n \in \mathbb{N}^+$, the same model checking problem is $PSPACE$ -complete (resp. Δ_2^P -complete).

Proof sketch. First, we sketch the decision procedure in $EXPTIME$ for formulas of type $\langle\langle\Gamma\rangle\rangle\psi$, where in particular ψ is an LTL formula. To decide $\langle\langle\Gamma\rangle\rangle\psi$, we first inflate the IS M to a model M' whose states are sequences of states in M of length n . This procedure determines an exponential blow-up. Then, we guess an n -bounded strategy F_Γ^n and prune

Algorithm $MC(M, \langle\langle\Gamma\rangle\rangle\psi, n)$:

```

1  $M' = Inflate(M, n)$ ;
2  $s_\Gamma = GuessStrat(M', \Gamma)$ ;
3 Prune  $M'$  according to  $s_\Gamma$ , obtaining  $M''$ ;
4 return  $MC\_CTL^*(M'', A\psi)$ ;

```

Figure 1: Algorithm to decide ATL^* Model checking.

model M' of all transitions that cannot occur by following F_Γ^n . Finally, in the pruned model thus obtained we check the CTL^* formula $A\psi$. The whole procedure is outlined in Figure 1. Notice that, since we use a polynomial oracle to guess F_Γ^n , we need to consider the inflated model M' exponential in n . Further, model checking CTL^* is in $PSPACE$ (Clarke, Emerson, and Sistla 1986). Then, the whole procedure is in $EXPTIME$. In the case of a general formula φ in ATL^* , we call the procedure above a polynomial number of times in the size of φ , so the procedure is again in $EXPTIME$.

On the other hand, if we consider ATL^* with bound n as fixed, we obtain a $PSPACE$ upper bound. This is because an n -bounded strategy, with n fixed, can be guessed in polynomially many steps, and therefore by using only polynomially many memory cells. As for the lower bound with n fixed, since every formula ψ in LTL is equivalent to $\langle\langle\emptyset\rangle\rangle\psi$ in ATL^* and model checking LTL is $PSPACE$ -hard, the result follows.

As regards the upper bound for ATL with a fixed n , we can adapt the proof of ATL^* described above. Specifically, we substitute the subroutine for model checking CTL^* with the one for CTL , which is in $PTIME$ (Clarke, Emerson, and Sistla 1986). Then, the whole procedure is polynomial with calls to an NP oracle. As a result, the overall complexity is in Δ_2^P . As for the lower bound with n fixed, we can use the same reduction to $SNSAT_2$ as in (Jamroga and Dix 2006). \square

We remark that for a fixed $n \in \mathbb{N}^+$, the complexity of model checking ATL and ATL^* with n -bounded recall (and imperfect information) is the same as for the imperfect recall case, i.e., for $n = 1$ (Schobbens 2004; Jamroga and Dix 2006). Moreover, we provided tight complexity results only for a fixed n . Indeed, here we are mainly interested in the fact that, differently from the perfect information case, the model checking problem for bounded recall is decidable, irrespectively of its complexity.

The decidability results above can be the basis of a partial model checking procedure for perfect recall consisting in increasing the bound n on the memory of agents. However, as the following demonstrates, increasing memory only preserves rather limited fragments of ATL^* and may, therefore, only be of limited interest.

Lemma 1. Let $m, n \in \mathbb{N}^+ \cup \{\omega\}$ be such that $m \leq n$; let ψ be an existential and ϕ an universal formula in ATL^* . Then,

$$(M, p) \models_m^2 \psi \Rightarrow (M, p) \models_n^2 \psi \quad (1)$$

$$(M, p) \not\models_m^2 \phi \Rightarrow (M, p) \not\models_n^2 \phi \quad (2)$$

Proof sketch. The proofs for (1) and (2) are both by induction on the structure of the formula. We only consider the

case for (1) where the main operator is the strategic modality. The other cases are immediate and thus omitted.

(1) By Def. 6 $(M, s) \models_m^2 \langle\langle\Gamma\rangle\rangle\psi$ iff for some F_Γ^m , for all $p \in out(s, F_\Gamma^m)$, $(M, p) \models_m^2 \psi$. Given F_Γ^m we construct a set F_Γ^n of strategies as follows: for every agent $i \in \Gamma$ and history $h \in \mathcal{G}^{<1+n}$, define $f_i^n(h) = f_i^m(h_{(|h|-m)}, \dots, h_{|h|})$ for $m < |h|$, $f_i^n(h) = f_i^m(h)$ otherwise. Notice that each f_i^n so defined is uniform, provided that f_i^m is. Given such F_Γ^n , we obtain $out(s, F_\Gamma^n) = out(s, F_\Gamma^m)$. In particular, for all $p \in out(s, F_\Gamma^n)$, $(M, p) \models_n^2 \psi$ by induction hypothesis, and therefore $(M, s) \models_n^2 \langle\langle\Gamma\rangle\rangle\psi$. \square

By Lemma 1 adding memory preserves the truth of existential formulas as well as falsehood of universal formulas. However, it is not difficult to find counterexamples to the extensions of (1) and (2) even in ATL .

Lemma 2. Let $m, n \in \mathbb{N}^+ \cup \{\omega\}$ be such that $m < n$. There exists formulas φ and $\varphi' = \neg\varphi$ in ATL such that

$$(M, p) \models_m^2 \varphi \quad \text{and} \quad (M, p) \not\models_n^2 \varphi \quad (3)$$

$$(M, p) \not\models_m^2 \varphi' \quad \text{and} \quad (M, p) \models_n^2 \varphi' \quad (4)$$

Proof sketch. We only provide a sketch of proof for (4) with $m = 1$ and $n = 2$. Consider a revisited version of a game of matching pennies in which Player 1 chooses first head or tail and Player 2 can see the choice, i.e., they end up in two different and distinguishable states. But after that, there is another step in which the coin is hidden from Player 2, i.e., there are two indistinguishable states for Player 2. Let $\varphi' = \langle\langle\{2\}\rangle\rangle Fw$. Player 2 has no 1-bounded recall strategy to win the game, but she has a 2-bounded recall, and therefore an n -bounded recall strategy, for every $n \in \mathbb{N}^+ \cup \{\omega\}$. This line of reasoning can be generalised to all $m, n \in \mathbb{N}^+ \cup \{\omega\}$ with $m < n$, as depicted in Figure 2. \square

By Lemmas 1 and 2 any naive attempt to approximate perfect recall by increasing bounded recall is severely restricted in two ways. Firstly, Lemma 1 holds only for the existential and universal fragments of ATL^* . Secondly, only the truth of existential formulas is preserved by adding memory, whereas negative results can only be lifted for the universal fragment. In the next section we present a three-valued semantics to overcome these difficulties. Before this, we exemplify the formal machinery introduced so far on an example.

Example 1. Alice wants to buy a gift for her son's birthday, but she is unable to buy it herself. So, she asks her friend Bob to buy a toy (t) and some wrapping paper (p). But Bob is a playful guy. He buys t and p, but then suggests they play a game. First, Bob puts t and p in two different rooms (r_1 and r_2), one in front of the other. Then, Bob puts Alice blindfolded in between the doors. Alice can go either forward (f) or backward (b); her goal (g) is to enter both rooms to get both t and p. This scenario is modeled as the interpreted system M in Figure 3. More formally, we have an IS $M = \langle Ag, s_0, T, \Pi \rangle$, where $Ag = \{Alice, Bob\}$, $\mathcal{G} = \{s_0, s_1, s_2, s_3, s_4, s_5, s_6, s_7\}$, $Act_{Alice} = \{b, f, I\}$, $Act_{Bob} = \{1, 2, I\}$, the global transition function is given as in Figure 3, $\Pi(s_3, p) = tt$, $\Pi(s_4, t) = tt$, $\Pi(s_5, t) = tt$,

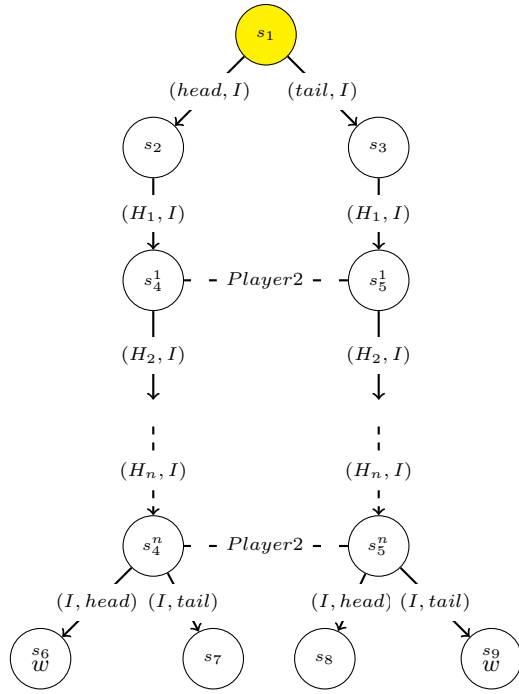


Figure 2: The IS M for a revisited version of matching pennies game. Here, we consider the general setting in which there are n steps of hiding before the choice of Player2.

$\Pi(s_6, p) = tt$, $\Pi(s_7, g) = tt$, and all atoms are false in any other state. We remark informally that if Alice only does f or b , she cannot get both objects. But if she chooses first one action and then the other, then she gets both t and p . The IS M is useful to describe the power of bounded-recall strategies. In particular, Alice has no memoryless strategy to achieve her goal, as she is supposed to perform two different actions in her initial state. However, she has a 2-bounded recall strategy to win the game. More formally, we have that $(M, s_0) \models_n^2 \langle\langle Alice \rangle\rangle (Fp \wedge Ft \wedge Fg)$ is true for $n \geq 2$ and false otherwise.

3 Three-valued Bounded Recall

In Section 2 we remarked that model checking ATL^* under imperfect information and perfect recall is undecidable in general (Dima and Tiplea 2011). Moreover, naively approximating perfect recall through bounded recall leads to the weak results in Lemmas 1 and 2.

To tackle these issues, in this section we lay the theoretical foundations of a partial model checking procedure based on a three-valued semantics. The procedure is partial as in some cases it returns “undefined” (uu) as truth value. On the other hand, differently from Lemma 1, the satisfaction of all ATL^* formulas is preserved by adding memory.

3.1 Three-valued ATL on Bounded Recall

We start by providing the three-valued satisfaction relation for ATL^* .




Figure 3: The IS M for the purchase gift game, where the actions 1 and 2 represents the choice of Bob to put Alice in front of room r_1 or r_2 , respectively. Moreover, s_1 and s_2 are indistinguishable to Alice and I stands for the idle action.

Definition 8 (Satisfaction). Let $n \in \mathbb{N}^+ \cup \{\omega\}$. The three-valued satisfaction relation \models_n^3 for an IS M , state s , path p , ATL^* formula ϕ , and $v \in \{tt, ff\}$ is defined as in Fig. 4. In all other cases the value of ϕ is undefined (uu).

Notice that all clauses for the three-valued semantics mirror the corresponding two-valued clauses, with a notable exception: for $\langle\langle \Gamma \rangle\rangle \psi$ to be false we require the existence of a joint strategy for the complement coalition $\bar{\Gamma}$ that enforces ψ to be false. Similar conditions have previously been proposed (Lomuscio and Michaliszyn 2014). It is a stronger requirement than the usual clause on the coalition Γ not being able to enforce ψ . However, it has the advantage of being preserved when adding memory, as it will become apparent in Lemma 5. Further, as for the two-valued semantics, we normally refer to the cases for $n = 1$ and $n = \omega$ as imperfect, resp. perfect, recall. Also notice that, as regards the Boolean operators, our semantics correspond to Kleene’s three-valued logic.

We say that formula φ is *true* (resp. *false*) in an IS M (for n -bounded recall), or $(M \models_n^3 \varphi) = tt$ (resp. ff), iff $((M, s_0) \models_n^3 \varphi) = tt$ (resp. ff); otherwise φ is *undefined*. Again, we observe that Def. 8 is an objective, three-valued interpretation of ATL^* . The corresponding subjective semantics can be obtained with minor modification, but it is beyond the scope of the present contribution.

We immediately prove that the three-valued notion of satisfaction in Def. 8 is an extension of the two-valued relation in Def. 6.

Lemma 3. For every $n \in \mathbb{N}^+ \cup \{\omega\}$, formula ϕ in ATL^* ,

$$((M, s) \models_n^3 \phi) = tt \Rightarrow (M, s) \models_n^2 \phi \quad (5)$$

$$((M, s) \models_n^3 \phi) = ff \Rightarrow (M, s) \not\models_n^2 \phi \quad (6)$$

Proof sketch. The proofs for both (5) and (6) are by simultaneous induction on the structure of the formula. We present the case for (6) where the main operator is the strategic modality. The cases for the other operators are immediate.

(6) By Def. 8 $((M, s) \models_n^3 \langle\langle \Gamma \rangle\rangle \psi) = ff$ iff for some $F_{\bar{\Gamma}}^n$, for all $p \in out(s, F_{\bar{\Gamma}}^n)$, $((M, p) \models_n^3 \psi) = ff$. Fix such a $F_{\bar{\Gamma}}^n$. By induction hypothesis we obtain that for all $p \in out(s, F_{\bar{\Gamma}}^n)$, $(M, p) \not\models_n^2 \psi$. In particular, for every

$((M, s) \models_n^3 q) = v$	iff	$\Pi(s, q) = v$
$((M, s) \models_n^3 \neg\varphi) = v$	iff	$((M, s) \models_n^3 \varphi) = \neg v$
$((M, s) \models_n^3 \varphi \wedge \varphi') = \text{tt}$	iff	$((M, s) \models_n^3 \varphi) = \text{tt}$ and $((M, s) \models_n^3 \varphi') = \text{tt}$
$((M, s) \models_n^3 \varphi \wedge \varphi') = \text{ff}$	iff	$((M, s) \models_n^3 \varphi) = \text{ff}$ or $((M, s) \models_n^3 \varphi') = \text{ff}$
$((M, s) \models_n^3 \langle\langle \Gamma \rangle\rangle \psi) = \text{tt}$	iff	for some F_Γ^n , for all $p \in \text{out}(s, F_\Gamma^n)$, $((M, p) \models_n^3 \psi) = \text{tt}$
$((M, s) \models_n^3 \langle\langle \Gamma \rangle\rangle \psi) = \text{ff}$	iff	for some F_Γ^n , for all $p \in \text{out}(s, F_\Gamma^n)$, $((M, p) \models_n^3 \psi) = \text{ff}$
$((M, p) \models_n^3 \varphi) = v$	iff	$((M, p_1) \models_n^3 \varphi) = v$
$((M, p) \models_n^3 \neg\psi) = v$	iff	$((M, p) \models_n^3 \psi) = \neg v$
$((M, p) \models_n^3 \psi \wedge \psi') = \text{tt}$	iff	$((M, p) \models_n^3 \psi) = \text{tt}$ and $((M, p) \models_n^3 \psi') = \text{tt}$
$((M, p) \models_n^3 \psi \wedge \psi') = \text{ff}$	iff	$((M, p) \models_n^3 \psi) = \text{ff}$ or $((M, p) \models_n^3 \psi') = \text{ff}$
$((M, p) \models_n^3 X\psi) = v$	iff	$((M, p_{\geq 2}) \models_n^3 \psi) = v$
$((M, p) \models_n^3 \psi U \psi') = \text{tt}$	iff	for some $k \geq 1$, $((M, p_{\geq k}) \models_n^3 \psi') = \text{tt}$, and for all j , $1 \leq j < k$ implies $((M, p_{\geq j}) \models_n^3 \psi) = \text{tt}$
$((M, p) \models_n^3 \psi U \psi') = \text{ff}$	iff	for all $k \geq 1$, $((M, p_{\geq k}) \models_n^3 \psi') = \text{ff}$, or for some $j \geq 1$, $((M, p_{\geq j}) \models_n^3 \psi) = \text{ff}$

Figure 4: The three-valued satisfaction relation \models_n^3 for an IS M , state s , path p , ATL^* formula ϕ , and $v \in \{\text{tt}, \text{ff}\}$.

joint strategy F_Γ^n there exists some path $p' \in \text{out}(s, F_\Gamma^n)$ (which is obtained when coalition $\bar{\Gamma}$ plays according to F_Γ^n) such that $(M, p') \not\models_n^2 \psi$ by hypothesis. As a result, $(M, s) \not\models_n^2 \langle\langle \Gamma \rangle\rangle \psi$. \square

On the other hand, the three-valued semantics is not a *conservative* extension of the two-valued one. Specifically, the following lemma provides counterexamples to the converse of (5) and (6).

Lemma 4. For $n \in \mathbb{N}^+ \cup \{\omega\}$, there exists an IS M with state s , and ATL formulas φ and $\varphi' = \neg\varphi$ such that

$$(M, s) \models_n^2 \varphi \quad \text{and} \quad ((M, s) \models_n^3 \varphi) \neq \text{tt} \quad (7)$$

$$(M, s) \not\models_n^2 \varphi' \quad \text{and} \quad ((M, s) \models_n^3 \varphi') \neq \text{ff} \quad (8)$$

Proof sketch. As regards (8) consider again the game of matching pennies presented in the proof of Lemma 2. We remarked therein that $\langle\langle \{2\} \rangle\rangle Fw$ is false in the (positional) two-valued semantics. However, in the same game Player 1 has no 1-bounded strategy to enforce Player 2 to lose, i.e., in the three-valued semantics the value of $\langle\langle \{2\} \rangle\rangle Fw$ is different from ff; actually, it is uu. \square

3.2 Model Checking Three-valued Bounded Recall

We now analyse the model checking problem for the three-valued semantics.

Definition 9 (Model Checking). Given an IS M , a formula ϕ , bound $n \in \mathbb{N}^+ \cup \{\omega\}$, and truth value $v \in \{\text{tt}, \text{ff}, \text{uu}\}$, the model checking problem (for n -bounded recall) amounts to determining whether $(M \models_n^3 \phi) = v$.

Similarly as in the two-valued semantics, we obtain the following undecidability result.

Theorem 3. The model checking problem for ATL in the three-valued semantics with perfect recall and imperfect information is undecidable.

Proof sketch. The proof follows by adapting the undecidability result in (Dima and Tiplea 2011), which makes use of the ATL formula $\varphi = \langle\langle \{1, 2\} \rangle\rangle Gok$ to express that

a Turing machine does not halt on the empty word. Specifically, we observe that the two- and three-valued interpretations coincide for this particular formula φ . That is, for any bound $n \in \mathbb{N}^+ \cup \{\omega\}$, we have that $(M \models_n^3 \varphi) = \text{tt}$ iff $M \models_n^2 \varphi$. Indeed, the value of atom ok is always defined, and the structure of the clauses for operator $\langle\langle \{1, 2\} \rangle\rangle$ is the same in the two- and three-valued semantics. As a consequence, we obtain that a Turing machine T does not halt on the empty word iff $(M_T \models_n^3 \varphi) = \text{tt}$, where M_T is obtained from T as described in (Dima and Tiplea 2011). \square

By Theorem 3 model checking ATL^* in the same setting is also undecidable. Again, by assuming bounded recall we retrieve decidability.

Theorem 4. For $n \in \mathbb{N}^+$ the model checking problem for ATL^* in the three-valued semantics with n -bounded recall and imperfect information is in $EXPTIME$. Moreover, for a fixed $n \in \mathbb{N}^+$, the same model checking problem is $PSPACE$ -complete.

Proof sketch. As regards the general case with n as a parameter, we adapt the model checking procedure for Theorem 2. Again, the case of interest is for strategic formulas $\varphi = \langle\langle \Gamma \rangle\rangle \psi$. We consider values tt, ff, and uu separately. Since the clause for checking $((M, p) \models_n^3 \varphi) = \text{tt}$ is the same as for $(M, p) \models_n^2 \varphi$, we can use the same procedure as in Theorem 2, which is in $EXPTIME$. To check whether $((M, p) \models_n^3 \varphi) = \text{ff}$ we observe that this is tantamount to the two-valued clause for $\langle\langle \bar{\Gamma} \rangle\rangle \neg\psi$. Therefore, we use the procedure in Fig. 1 with input formula $\langle\langle \bar{\Gamma} \rangle\rangle \neg\psi$. Again, its complexity is in $EXPTIME$. Finally, if both cases (tt) and (ff) return false, then the result is undefined (uu). Since to determine uu we use two procedures in $EXPTIME$, this is also in $EXPTIME$. In Fig. 5 we outline the procedure above. As for Theorem 2, the procedure above is called a polynomial number of times in the size of φ , so the overall complexity is still in $EXPTIME$.

For a fixed $n \in \mathbb{N}^+$ the procedure above is in $PSPACE$, as we recall that guessing a strategy on the inflated model can be done in polynomial space. As regards the lower bound, we make use of the same reduction as in Theorem 2. In particular, we can reduce model checking an LTL formula ψ to the verification of the truth of the ATL^* formula $\langle\langle \emptyset \rangle\rangle \psi$ in

Algorithm $MC3(M, \langle\langle\Gamma\rangle\rangle\psi, n, v)$:	
1	if $v = \text{tt}$ then return $MC(M, \langle\langle\Gamma\rangle\rangle\psi, n)$;
2	else if $v = \text{ff}$ then return $MC(M, \langle\langle\bar{\Gamma}\rangle\rangle\neg\psi, n)$;
3	else if $v = \text{uu}$ and
4	$MC(M, \langle\langle\Gamma\rangle\rangle\psi, n) \vee MC(M, \langle\langle\bar{\Gamma}\rangle\rangle\neg\psi, n)$ then
5	return ff ;
6	else return tt ;

Figure 5: Algorithm to decide ATL^* three-valued model checking.

the three-valued semantics. \square

By Theorems 2 and 4 model checking ATL^* on the two- and three-valued semantics has the same complexity. This is also the case for ATL .

Theorem 5. For $n \in \mathbb{N}^+$ the model checking problem for ATL in the three-valued semantics with n -bounded recall and imperfect information is in $EXPTIME$. Moreover, for a fixed $n \in \mathbb{N}^+$, the same model checking problem is Δ_2^P -complete.

Proof sketch. Clearly, the $EXPTIME$ upper bound for the general case still holds.

As for a fixed $n \in \mathbb{N}^+$, we adapt the proof of Theorem 4. In particular, in the procedure to decide strategy formulas we change the model checking subroutine as we did in the proof of Theorem 2, by using the one for CTL (which we recall to be in Δ_2^P) instead of the one for CTL^* . \square

Again, for a fixed $n \in \mathbb{N}^+$, the complexity of model checking three-valued ATL and ATL^* with n -bounded recall (and imperfect information) is the same as for imperfect recall. Also, as in Sec. 2 we are primarily interested in the decidability of the model checking problem for bounded recall, irrespectively of tight complexity results.

Our aim in the rest of this section is to lay the theoretical foundations of a (partial) model checking procedure that is able to deal with the whole of ATL^* . To this end, the next result, which is akin to Lemma 1, details the preservation of ATL^* formulas when adding memory. However, differently from Lemma 1, this result holds for all ATL^* formulas.

Lemma 5. Let $m, n \in \mathbb{N}^+ \cup \{\omega\}$ be such that $m \leq n$; let ψ be a formula in ATL^* . Then,

$$((M, s) \models_m^3 \psi) = \text{tt} \Rightarrow ((M, s) \models_n^3 \psi) = \text{tt} \quad (9)$$

$$((M, s) \models_m^3 \psi) = \text{ff} \Rightarrow ((M, s) \models_n^3 \psi) = \text{ff} \quad (10)$$

Proof sketch. The proofs for both (9) and (10) are by simultaneous induction on the structure of the formula. We only present the case for (10) where the main operator is the strategic modality, the other operators being immediate.

(10) By Def. 8 $((M, s) \models_m^3 \langle\langle\Gamma\rangle\rangle\psi) = \text{ff}$ if for some $F_{\bar{\Gamma}}^m$, for all $p \in \text{out}(s, F_{\bar{\Gamma}}^m)$, $((M, p) \models_m^3 \psi) = \text{ff}$. Given $F_{\bar{\Gamma}}^m$ we construct a set $F_{\bar{\Gamma}}^n$ of strategies as follows: for every agent $i \in \bar{\Gamma}$ and history $h \in \mathcal{G}^{<1+n}$, define $f_i^n(h) = f_i^m(h_{(|h|-m)}, \dots, h_{|h|})$ for $m < |h|$, $f_i^n(h) = f_i^m(h)$ otherwise. Notice that f_i^n so defined is uniform, provided that f_i^m is. Given such $F_{\bar{\Gamma}}^n$, we obtain $\text{out}(s, F_{\bar{\Gamma}}^n) = \text{out}(s, F_{\bar{\Gamma}}^m)$.

Algorithm $Iterative_MC(M, \psi, n)$:	
1	$j = 1, k = \text{uu}$;
2	while $j \leq n$ and $k = \text{uu}$
3	if $MC3(M, \psi, j, \text{tt})$ then $k = \text{tt}$
4	else if $MC3(M, \psi, j, \text{ff})$ then $k = \text{ff}$
5	$j = j + 1$;
6	end while ;
7	if $k \neq \text{uu}$ then return (j, k) ;
8	else return -1 ;

Figure 6: The procedure $Iterative_MC$ to decide ATL^* iteratively.

In particular, for all $p \in \text{out}(s, F_{\bar{\Gamma}}^n)$, $((M, p) \models_n^3 \psi) = \text{ff}$ by induction hypothesis, and therefore $((M, s) \models_n^3 \langle\langle\Gamma\rangle\rangle\psi) = \text{ff}$. \square

By Lemma 5 adding memory preserves defined truth values for all formulas in ATL^* . This is in marked contrast with Lemma 1. Indeed, even though in some cases the value of an ATL^* formula may be undefined in the three-valued semantics, whenever it is defined, it does not change if memory is added.

By combining together Lemmas 3 and 5 we prove our main result on the relationship between bounded recall and the two- and three-valued semantics.

Corollary 1. Let $m, n \in \mathbb{N}^+ \cup \{\omega\}$ be such that $m \leq n$; let ψ be a formula in ATL^* . Then,

$$((M, p) \models_m^3 \psi) = \text{tt} \Rightarrow ((M, p) \models_n^2 \psi) \quad (11)$$

$$((M, p) \models_m^3 \psi) = \text{ff} \Rightarrow ((M, p) \not\models_n^2 \psi) \quad (12)$$

Of particular interest is the case for $m \in \mathbb{N}^+$ and $n = \omega$. By Corollary 1 we can outline a verification procedure for perfect recall, whereby ATL^* formulas are checked in the three-valued semantics iteratively. If either true or false are returned, by Corollary 1 this is also the truth value for the two-valued semantics under perfect recall. We provide a detailed presentation and formal analysis of this verification procedure in the next section.

4 Approximating Perfect Recall

In this section we provide a partial decision procedure for model checking ATL^* with imperfect information and n -bounded recall. It is partial, as it is not guaranteed to terminate for the case of perfect recall, i.e., $n = \omega$. This procedure is described in an algorithmic way in Fig. 6. It takes as input an IS M , an ATL^* formula ψ , and a bound $n \in \mathbb{N}^+ \cup \{\omega\}$. Procedure $Iterative_MC()$ includes a loop, whose guard checks whether the bound has not yet been attained ($j \leq n$) and ψ has not yet been decided ($k = \text{uu}$). Within the loop ψ is model-checked in M according to the three-valued semantics by subroutine $MC3()$, and variable k stores the result. On exiting the loop, variable k is tested. If $k \neq \text{uu}$, the loop was exited because of a defined answer for the three-valued model checking problem with j -bounded recall. By Corollary 1 we can then transfer the value returned to the corresponding model checking problem in the two-valued

semantics. On the other hand, if $k = uu$ then the bound has been attained in the loop and -1 is given as a result. We now prove the termination of the algorithm in Fig. 6 for $n \in \mathbb{N}^+$, as well as its soundness.

Theorem 6. For $n \in \mathbb{N}^+$, $Iterative_MC()$ terminates in $EXPTIME$. Moreover, $Iterative_MC()$ is sound: if the value returned is different from -1, then $M \models_n^2 \phi$ iff $k = tt$.

Proof sketch. As regards termination in $EXPTIME$, notice that for $n \in \mathbb{N}^+$ the algorithm in Fig. 6 calls procedure $MC3()$, which is in $EXPTIME$ (Theorem 4), a bounded number of times. Then, the overall complexity is in $EXPTIME$.

As for soundness, suppose that the value returned is different from -1. In particular, either $k = tt$ or $k = ff$. If $M \models_n^2 \phi$ and $k = ff$, then $((M, s) \models_j^3 \psi) = ff$ for $j \leq n$, and by Corollary 1.(12) we have $M \not\models_n^2 \phi$, a contradiction. Hence, $k = tt$ as required. On the other hand, if $k = tt$ then by Corollary 1.(11) we obtain $M \models_n^2 \phi$. \square

Incidentally, for a fixed $n \in \mathbb{N}^+$, algorithm $Iterative_MC()$ actually runs in $PSPACE$.

An important application of $Iterative_MC()$ is for $n = \omega$, namely model checking perfect recall. In such a case termination is no longer guaranteed, but soundness is.

Theorem 7. For $n = \omega$, $Iterative_MC()$ does not necessarily terminate. However, $Iterative_MC()$ is sound: if the value returned is different from -1, then $M \models_n^2 \phi$ iff $k = tt$.

Proof sketch. We have remarked that in several games, such as matching pennies, neither player has a strategy to win the game, no matter how much recall we assume on our players. Soundness follows again by Corollary 1. \square

As a result, by Theorem 7 we have a sound, albeit incomplete, decision procedure for model checking ATL^* with perfect recall and imperfect information. Observe that no complete procedure is attainable as the problem is undecidable in general (Dima and Tiplea 2011).

Example 2. We conclude with an example for illustrative purposes. In connection with the IS M for the purchase gift game in Example 1, consider the specification $\phi = \llbracket Bob \rrbracket X \langle\langle Alice \rangle\rangle (Fp \wedge Ft)$ which intuitively states that no matter what Bob does, at the next step Alice has a strategy whereby all outcomes eventually satisfy atoms p and t (possibly at different times), that is, she will be able to collect the gift and the wrapping paper. This specification is neither existential nor universal, and therefore does not fall within the hypothesis of Lemma 1. Nevertheless, ϕ is amenable to algorithm $Iterative_MC()$ in Fig. 6. Specifically, given the IS M in Fig. 3, formula ϕ , and bound $n \geq 2$, the algorithm $Iterative_MC(M, \phi, n)$ initializes the bound on recall to 1 and the value of k to undefined uu . Then, proceeds with the first iteration. Both subroutines $MC3(M, \phi, 1, tt)$ and $MC3(M, \phi, 1, ff)$ return false because, according to the three-valued semantics, Alice does not have a memoryless strategy to enforce $\langle\langle Alice \rangle\rangle (Fp \wedge Ft)$ at the next step, nor Bob has a (memoryless) strategy to prevent $\langle\langle Alice \rangle\rangle (Fp \wedge Ft)$ at the next step. On the other hand, in the second iteration the call $MC3(M, \phi, 2, tt)$ returns true, as Alice has

a 2-bounded recall strategy to enforce $\langle\langle Alice \rangle\rangle (Fp \wedge Ft)$ at the next step, and therefore ϕ holds. Thus, we conclude that the IS M in Example 1 satisfy specification ϕ under the assumptions of imperfect information and perfect recall.

5 Related Work

We now discuss our work in the context of recent contributions on logic-based languages for the specification of strategic abilities of agents in multi-agent systems. In particular, inspired by (Ball and Kupferman 2006; Shoham and Grumberg 2004) a stream of papers has recently appeared on three-valued semantics for ATL . In (Lomuscio and Michaliszyn 2014; 2015) three-valued abstractions for interpreted systems are introduced to tackle the complexity of MAS verification. These investigations were developed further in (Belardinelli, Lomuscio, and Michaliszyn 2016; Lomuscio and Michaliszyn 2016) by means of predicate abstraction. While we take our inspiration from this line of works, our present contribution differs significantly.

First of all, the semantics and the underlying classes of systems we study here are different from those in (Lomuscio and Michaliszyn 2014; 2015; 2016). Specifically, these works assume *non-uniform* strategies (Lomuscio and Raimondi 2006), with crucial implications on the decidability and complexity of the model checking problem. In particular, in the semantics of non-uniform strategies model checking ATL with *imperfect information* is decidable in PTIME both for the memoryfull and memoryless case. Hence, approximating perfect recall is not even an issue in the setting of (Lomuscio and Michaliszyn 2014; 2015; 2016). On the other hand, we here insist on considering uniform strategies, as this is the framework commonly used when analysing strategic abilities of agents in MAS and game-theoretical contexts (Jamroga and van der Hoek 2004).

Further, in the related (Belardinelli and Lomuscio 2017) imperfect recall is assumed, whereas here we tackle the case of perfect recall. Their three-valued semantics of ATL operators is also different from what we propose here. Formally, in (Belardinelli and Lomuscio 2017) the falsehood of a formula of type $\langle\langle \Gamma \rangle\rangle \psi$ is given in terms of *may-strategies* of coalition Γ , whereas we define it in terms of the strategic abilities of the complement coalition $\bar{\Gamma}$. This is a key feature of our semantics, as it allows us to preserve defined truth values when adding recall (Lemma 1).

Finally, our objective in the present contribution is different from the one addressed in (Lomuscio and Michaliszyn 2014; 2015; 2016): we aim at approximating undecidable perfect recall via decidable bounded recall. As remarked above, this is not an objective in (Lomuscio and Michaliszyn 2014; 2015; 2016).

Classic, two-valued bounded recall and bounded strategies have been studied quite extensively in the literature. In (Ågotnes and Walther 2009) the authors consider strategies according to two different bounds: over the set of histories and over the length of histories. In this framework, they show that ATL with bounded memory is strictly more expressive than standard ATL . In (Brihaye et al. 2009) ATL is extended in two directions: strategy contexts and bounded memory. Then, the model checking problem is proved to be

in *EXPSpace*. In (Jamroga, Malvone, and Murano 2017) strategies are defined as a list of condition-action rules. Then, the authors present a variant of *ATL* that makes use of strategy operators with a bound on the size of this list. To the best of our knowledge, in most contributions cited above the emphasis on boundedness always concerns expressiveness issues and complexity of the related model checking problem, but it has not been considered as approximation of perfect recall, which is the main focus of the present work. In some cases, including (Ågotnes and Walther 2009; Brihaye et al. 2009), the semantics are incomparable.

A further representation of finite-memory strategies was introduced in (Vester 2013), where strategies are defined as *deterministic finite-state transducers* (DFST). Since this work is the most closely related to the present contribution, we discuss it in detail and provide some comparison results. We show below that bounded strategies and DFST cannot always be translated (polynomially) one into the other. Hence, the two formalisms are orthogonal. We begin by introducing the definition of DFST, but refer to (Vester 2013) for more details.

Definition 10 (DFST). A *deterministic finite-state transducer* is a tuple $D = \langle V, v_0, In, Out, F_{in}, F_{out} \rangle$, where (i) V is a finite non-empty set of states; (ii) v_0 is the initial state; (iii) In is the input alphabet; (iv) Out is the output alphabet; (v) $F_{in} : V \times In \rightarrow V$ is the transition function; and (vi) $F_{out} : V \times In \rightarrow Out$ is the output function.

Intuitively, the set V of states are the possible values of the internal memory of the strategy. The initial state v_0 corresponds to the initial memory value. The input symbols in In are the states of the game, and set Out of output symbols is the set of actions in the game. In each round of the game the DFST reads the current state. Then, it updates its memory based on the current memory value and the input state according to F_{in} , and performs an action based on the current memory value and the input state according to F_{out} .

A function $\sigma : \mathcal{G}^+ \rightarrow Act$ is a *finite-memory strategy* if there exists a DFST such that for all histories $h \in \mathcal{G}^+$:

$$\sigma(h) = F_{out}(G(v_0, h_{\leq |h|-1}), last(h))$$

where for every state v and history h , function G is defined recursively as follows:

$$G(v, h) = \begin{cases} F_{in}(v, h_1) & \text{if } |h| = 1; \\ F_{in}(G(v, h_{\leq |h|-1}), last(h)) & \text{otherwise.} \end{cases}$$

That is, G is the function that repeatedly applies the transition function F_{in} on a sequence of inputs to calculate the state after a given history.

We now compare formally our definition of bounded strategy with finite-memory strategies given via DFST. Hereafter we say that two strategies are *equivalent* if they correspond to the same function $\sigma : \mathcal{G}^+ \rightarrow Act$.

Proposition 1. For every bound $n \in \mathbb{N}$, there exists some DFST D for which there is no equivalent strategy with $g(n)$ -bounded memory, for any polynomial function g .

Proof sketch. Notice that we can construct a DFST D such that for some history h of size exponential in n , and different states s, s' , it is the case that $F_{out}(G(v_0, s \cdot$

$h_{\leq |h|-1}), last(h)) \neq F_{out}(G(v_0, s' \cdot h_{\leq |h|-1}), last(h))$. For instance, consider a DFST D that loops between two different states in V a number of times exponential in n . This kind of behaviour cannot be captured by any strategy whose recall is bounded by some polynomial $g(n)$. \square

As for the translation from bounded recall strategies to DFST we have the following result.

Proposition 2. For every bound $n \in \mathbb{N}$, there exists some n -bounded recall strategy f for which there is no equivalent DFST with $g(n)$ states, for any polynomial function g .

Proof sketch. We provide a hint of a proof by contradiction. Given a n -bounded recall strategy f , suppose that we can construct a DFST D with m states, where $m < |S|^{n-1}$. In particular, this means that for two different histories h and h' of length n , at some points k and j , the function G returns the same state of memory, i.e. $G(v_0, h_{\leq j}) = G(v_0, h'_{\leq k})$. Suppose that $h_j = h'_k$, then we have that $F_{out}(G(v_0, h_{\leq j}), h_j) = F_{out}(G(v_0, h'_{\leq k}), h'_k)$, i.e., the same action is returned by histories $h_{\leq j}$ and $h'_{\leq k}$. Since we suppose that our strategies have n -bounded recall, then w.l.o.g. we can assume that f assigns different actions to $h_{\leq j}$ and $h'_{\leq k}$. But this contradicts the fact that $m < |S|^{n-1}$ states of memory in a DFST are sufficient to describe a n -bounded strategy. \square

Intuitively, Propositions 1 and 2 can be summarised as follows: while DFST can be seen as representing finite memory in strategies, the bounded strategies here introduced express recall. Memory and recall are related, but orthogonal notions. Moreover, to the best of our knowledge, the one here presented is the first three-valued semantics for bounded recall. This modelling choice is the key feature as to why Theorem 7 applies unrestricted.

6 Conclusions

Model checking MAS against Alternating-time Temporal Logic is known to be undecidable under perfect recall and imperfect information. In this paper we put forward a sound, albeit incomplete, verification procedure for perfect recall based on a novel notion of bounded recall. To do so, we introduced bounded recall on interpreted systems by giving both a two- and a three-valued semantics. By using the three-valued semantics for bounded recall we were able to prove a preservation result of defined truth values from the bounded to the perfect recall case for all *ATL** specifications. As shown, this is not possible in the classic two-valued semantics. This laid the foundation for an iterative procedure, which can, in some cases, solve the model checking problem under perfect recall by considering a bounded amount of memory. Since model checking perfect recall is undecidable in general, the procedure discussed is naturally incomplete. It is, however, the first procedure that we are aware of, which can give a solution in cases of interest.

In future work we plan to explore implementing the procedure above on a symbolic model checker for *ATL** with imperfect information, such as Verics (Knapik et al. 2010) or MCMAS (Lomuscio, Qu, and Raimondi 2017). Given the

complexity bounds of Theorem 6, we do not expect such an implementation to provide an efficient tool for verifying concrete systems for large values of the recall parameter. Moreover, it should be noted that such an implementation is likely to require a major effort. Indeed, to our knowledge no efficient model checker for ATL^* on uniform strategies is available, nor is any model checker on this class of systems for a 3-valued semantics. So, this task would be a significant endeavour in itself.

Indeed, the present contribution is primarily focused on exploring the theoretical underpinnings upon which approximations of perfect recall could be realised. With this in mind, it may be worth exploring further approximations for perfect recall or classes of formulas and models for which the present method is complete. We leave both of these issues to further work.

Acknowledgements. F. Belardinelli acknowledges the support of ANR JCJC Project SVeDaS (ANR-16-CE40-0021). A. Lomuscio is supported by a Royal Academy of Engineering Chair in Emerging Technology.

References

- Ågotnes, T., and Walther, D. 2009. A Logic of Strategic Ability Under Bounded Memory. *Journal of Logic, Language and Information* 18(1):55–77.
- Ågotnes, T.; Goranko, V.; Jamroga, W.; and Wooldridge, M. 2015. Knowledge and ability. In *Handbook of Logics for Knowledge and Belief*. Elsevier. 543–589.
- Alechina, N.; Logan, B.; Nguyen, H. N.; Raimondi, F.; and Mostarda, L. 2015. Symbolic model-checking for resource-bounded ATL. In *Proceedings of AAMAS 2015*, 1809–1810.
- Alur, R.; Henzinger, T.; and Kupferman, O. 2002. Alternating-time temporal logic. *Journal of the ACM* 49(5):672–713.
- André, É.; Petrucci, L.; Jamroga, W.; Knapik, M.; and Penczek, W. 2017. Timed ATL: forget memory, just count. In *Proceedings of AAMAS 2017*, 1460–1462.
- Baier, C., and Katoen, J. P. 2008. *Principles of Model Checking*. MIT press.
- Ball, T., and Kupferman, O. 2006. An abstraction-refinement framework for multi-agent systems. In *Proceedings of LICS06*, 379–388.
- Belardinelli, F., and Lomuscio, A. 2017. Agent-based abstractions for verifying alternating-time temporal logic with imperfect information. In *Proceedings of AAMAS 2017*, 1259–1267.
- Belardinelli, F.; Lomuscio, A.; Murano, A.; and Rubin, S. 2017. Verification of multi-agent systems with imperfect information and public actions. In *Proceedings of AAMAS 2017*, 1268–1276.
- Belardinelli, F.; Lomuscio, A.; and Michaliszyn, J. 2016. Agent-based refinement for predicate abstraction of multi-agent systems. In *Proceedings of ECAI16*, 286–294.
- Berthon, R.; Maubert, B.; Murano, A.; Rubin, S.; and Vardi, M. Y. 2017. Strategy logic with imperfect information. In *Proceedings of LICS 2017*, 1–12.
- Brihaye, T.; Lopes, A.; Laroussinie, F.; and Markey, N. 2009. ATL with Strategy Contexts and Bounded Memory. In *Proceedings of LFCS'09*, 92–106.
- Busard, S.; Pecheur, C.; Qu, H.; and Raimondi, F. 2015. Reasoning about memoryless strategies under partial observability and unconditional fairness constraints. *Information and Computation* 242:128–156.
- Clarke, E.; Emerson, E.; and Sistla, A. 1986. Automatic Verification of Finite-State Concurrent Systems Using Temporal Logic Specifications. *Proceedings of TOPLAS* 8(2):244–263.
- Dima, C., and Tiplea, F. 2011. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. *CoRR* abs/1102.4225.
- Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning about Knowledge*. MIT Press.
- Jamroga, W., and Dix, J. 2006. Model checking abilities under incomplete information is indeed Δ_P^2 -complete. In *Proceedings of EUMAS'06*, 14–15.
- Jamroga, W., and van der Hoek, W. 2004. Agents that know how to play. *Fundamenta Informaticae* 62:1–35.
- Jamroga, W.; Malvone, V.; and Murano, A. 2017. Reasoning about natural strategic ability. In *Proceedings of AAMAS 2017*, 714–722.
- Knapik, M.; Niewiadomski, A.; Penczek, W.; Pólrola, A.; Szreter, M.; and Zbrzezny, A. 2010. Parametric model checking with verics. *Proceedings of TPNOMC* 4:98–120.
- Lomuscio, A., and Michaliszyn, J. 2014. An abstraction technique for the verification of multi-agent systems against ATL specifications. In *Proceedings of KR14*, 428–437.
- Lomuscio, A., and Michaliszyn, J. 2015. Verifying multi-agent systems by model checking three-valued abstractions. In *Proceedings of AAMAS 2015*, 189–198.
- Lomuscio, A., and Michaliszyn, J. 2016. Verification of multi-agent systems via predicate abstraction against ATLK specifications. In *Proceedings of AAMAS 2016*, 662–670.
- Lomuscio, A., and Raimondi, F. 2006. Model checking knowledge, strategies, and games in multi-agent systems. In *Proceedings of AAMAS 2006*, 161–168. ACM Press.
- Lomuscio, A.; Qu, H.; and Raimondi, F. 2017. MCMAS: an open-source model checker for the verification of multi-agent systems. *Proceedings of STTT* 19(1):9–30.
- Meyden, R. v., and Shilov, H. 1999. Model checking knowledge and time in systems with perfect recall. In *Proceedings of FST&TCS'99*, 432–445.
- Schobbens, P. 2004. Alternating-Time Logic with Imperfect Recall. *Proceedings of ENTCS* 85(2):82–93.
- Shoham, S., and Grumberg, O. 2004. Monotonic abstraction-refinement for CTL. In *Proceedings of TACAS04*, 546–560.
- Vester, S. 2013. Alternating-time temporal logic with finite-memory strategies. In *Proceedings of GandALF 2013*, 194–207.