

Centre Fédéré en Vérification

Technical Report number 2009.114

ATL with strategy contexts and bounded memory

Thomas Brihaye, Arnaud Da Costa, François Laroussinie, Nicolas Markey



This work was partially supported by a FRFC grant: 2.4530.02 and by the MoVES project. MoVES (P6/39) is part of the IAP-Phase VI Interuniversity Attraction Poles Programme funded by the Belgian State, Belgian Science Policy

<http://www.ulb.ac.be/di/ssd/cfv>

ATL with strategy contexts and bounded memory

Thomas Brihaye¹, Arnaud Da Costa²,
François Laroussinie³, and Nicolas Markey²

¹ Institut de mathématiques, Université de Mons-Hainaut, Belgium

² Lab. Spécification & Vérification, ENS Cachan – CNRS UMR 8643, France

³ LIAFA, Univ. Paris 7 – CNRS UMR 7089, France

thomas.brihaye@umh.ac.be, dacosta@lsv.ens-cachan.fr,
francoisl@liafa.jussieu.fr, markey@lsv.ens-cachan.fr

Abstract. We extend the alternating-time temporal logics ATL and ATL* with *strategy contexts* and *memory constraints*: the first extension make strategy quantifiers to not “forget” the strategies being executed by the other players. The second extension allows strategy quantifiers to restrict to memoryless or bounded-memory strategies.

We first consider expressiveness issues. We show that our logics can express important properties such as equilibria, and we formally compare them with other similar formalisms (ATL, ATL*, Game Logic, Strategy Logic, ...). We then address the problem of model-checking for our logics, especially we provide a PSPACE algorithm for the sublogics involving only memoryless strategies and an EXPSPACE algorithm for the bounded-memory case.

1 Introduction

Temporal logics and model checking. Temporal logics (LTL, CTL) have been proposed for the specification of reactive systems almost thirty years ago [Pnu77,CE81,QS82]. Since then, they have been widely studied and successfully used in many situations, especially for model checking—the automatic verification that a model of a system satisfies a temporal logic specification.

Alternating-time temporal logic (ATL). Over the last ten years, ATL has been proposed as a new flavor of temporal logics for specifying and verifying properties in *multi-agent systems* (modeled as *Concurrent Game Structures* (CGS) [AHK02]), in which several agents can concurrently act upon the behaviour of the system. In these models, it is not only interesting to know if something *can* or *will* happen, as is expressed in CTL or LTL, but also if some agent(s) can *control* the evolution of the system in order to enforce a given property, whatever the other agents do. ATL can express this kind of properties thanks to its quantifier over strategies, denoted $\langle\langle A \rangle\rangle$ (where A is a coalition of agents). That coalition A has a strategy for reaching a winning location is then written $\langle\langle A \rangle\rangle \mathbf{F} \text{win}$ (where \mathbf{F} is the LTL modality for “eventually”).

Our contributions. In this paper, we extend ATL and ATL* in two directions first, while ATL strategy quantifier drops strategies introduced by earlier quantifiers in the evaluation of the formula, our logics keep executing those strategies. To achieve this idea, we naturally adapt the semantics of ATL* in order to interpret a formula within a *strategy context*. Our new modality, written $\langle A \rangle$, can for instance express that “A has a strategy s.t. (1) Player B always has a strategy (given that of A) to enforce Φ and (2) Player C always has a strategy (given the *same* strategy of A) to enforce Ψ ”: $\langle A \rangle \mathbf{G} (\langle B \rangle \Phi \wedge \langle C \rangle \Psi)$. Naive attempts to express this property in standard ATL fail: in the ATL formula $\langle\langle A \rangle\rangle \mathbf{G} (\langle\langle B \rangle\rangle \Phi \wedge \langle\langle C \rangle\rangle \Psi)$, the coalitions do not cooperate anymore; in $\langle\langle A \rangle\rangle \mathbf{G} (\langle\langle A, B \rangle\rangle \Phi \wedge \langle\langle A, C \rangle\rangle \Psi)$, coalition A is allowed to use different strategies when playing with B and C.

Our second extension consists in parameterising strategy quantifiers with the *resources* (in terms of memory) allowed for strategies: we define the quantifier $\langle A^s \rangle$ with $s \in (\mathbb{N} \cup \{\infty\})$, which restricts the quantification to strategies using memory of size s (called s -memory strategies) for Player A. It is well-known that memoryless strategies are enough to enforce ATL properties, but this is not the case for ATL* formulae, nor for our extension of ATL (and ATL*) with strategy contexts.

Our results are twofold: on the one hand, we study the increase in expressiveness brought by our extensions, comparing our logics to ATL and ATL* and several related logics such as Game Logic [AHK02], Strategy Logic [CHP07] and QD_μ [Pin07], ... We also illustrate their convenience with some sample formulas expressing *e.g.* equilibrium properties.

On the other hand, we study the model-checking problem for our extensions: while we only have a non-elementary algorithm for the most general logic, we propose a polynomial-space algorithm for model-checking our logic in the memoryless case, and extend it to an exponential-space algorithm for the bounded-memory setting.

Related works. Recently, several works have focused on the same kind of extensions of ATL, and come up with different solutions which we list below. Generally speaking, this leads to very expressive logics, able to express *e.g.* equilibrium properties, and drastically increases model-checking complexity.

- IATL [ÅGJ07] extends ATL with strategy contexts, with a similar definition as ours, but it requires players to commit to a strategy, which they are not allowed to modify in the sequel. This logic is then studied in the memoryless case (which is proven to be a strict restriction to memory-based strategies).
- SL [CHP07] extends temporal logics with first-order quantification over strategies. This extension has been defined and studied only in the two-player turn-based setting, where a non-elementary algorithm is proposed.
- QD_μ [Pin07] considers strategies as labellings of the computation tree of the game structure with fresh atomic propositions. This provides a way of explicitly dealing with strategies. This extension is added on top of the decision μ -calculus $D\mu$, yielding a very expressive, yet decidable framework.

- Stochastic Game Logic [BBGK07] is a similar extension to ours, but for stochastic games. It is undecidable in the general case, but proved decidable when restricting to memoryless strategies.

Instead of defining a completely new formalism, we prefer sticking to an ATL-like syntax, as we believe that our new modality $\langle \cdot \cdot \rangle$ is more intuitive than the standard ATL modality $\langle\langle A \rangle\rangle$. Also, none of the above the extension has the ability to explicitly restrict to bounded-memory strategies, which is of obvious practical relevance and leads to more efficient algorithms.

Plan of the paper. Section 2 contains the definitions of our logics, and of our bounded-memory setting. Section 3 deals with the expressiveness results, and compares our extension with those cited in the related work above. In Section 4, we consider the model-checking problem for our extensions, and provide algorithms for the case of s -memory strategies. Due to lack of space, some of the proofs are postponed to the appendix.

2 Definitions

In this section we introduce classical definitions of concurrent game structures, strategies and outcomes. We then define a notion of s -bounded memory strategies. In the whole paper, AP denotes a finite non-empty set of atomic propositions.

2.1 Concurrent Game Structures

Concurrent game structures are a multi-player extension of classical Kripke structures [AHK02]. Their definition is as follows:

Definition 1. A Concurrent Game Structure (CGS for short) \mathcal{C} is a 8-tuple $(Loc, \ell_0, Lab, \delta, Agt, \mathcal{M}, Mov, Edg)$ where:

- Loc is a finite set of locations, $\ell_0 \in Loc$ is the initial location;
- $Lab: Loc \rightarrow 2^{AP}$ is a labelling function;
- $\delta \subseteq Loc \times Loc$ is the set of transitions;
- $Agt = \{A_1, \dots, A_k\}$ is a finite set of agents (or players);
- \mathcal{M} is a finite, non-empty set of moves;
- $Mov: Loc \times Agt \rightarrow \mathcal{P}(\mathcal{M}) \setminus \{\emptyset\}$ defines the (finite) set of possible moves of each agent in each location.
- $Edg: Loc \times \mathcal{M}^k \rightarrow \delta$, where $k = |Agt|$, is a transition table. With each location and each set of moves of the agents, it associates the resulting transition.

The size $|\mathcal{C}|$ of a CGS \mathcal{C} is defined as $|Loc| + |Edg|$, where $|Edg|$ is the size of the transition table¹. The intended behaviour is as follows [AHK02]: in a

¹ Our results would still hold if we consider symbolic CGSs [LMO07], where the transition table is encoded through boolean formulas.

location ℓ , each player A_i chooses one of his possible moves m_{A_i} and the next transition is given by $\text{Edg}(\ell, m_{A_1}, \dots, m_{A_k})$. We write $\text{Next}(\ell)$ for the set of all transitions corresponding to possible moves from ℓ , and $\text{Next}(\ell, A_j, m)$, with $m \in \text{Mov}(\ell, A_j)$, for the restriction of $\text{Next}(\ell)$ to possible transitions from ℓ when player A_j makes the move m .

2.2 Coalitions, bounded-memory strategies, outcomes.

Coalitions. A coalition is a subset of agents. In multi-agent systems, a coalition A plays against its opponent coalition $\text{Agt} \setminus A$ as if they were two single players. We thus extend Mov and Next to coalitions:

- Given $A \subseteq \text{Agt}$ and $\ell \in \text{Loc}$, $\text{Mov}(\ell, A)$ denotes the set of possible moves for coalition A from ℓ . Those moves m are composed of one single move per agent of the coalition, *i.e.*, $m = (m_a)_{a \in A}$.
- Next is extended to coalitions in a natural way: given $m = (m_a)_{a \in A} \in \text{Mov}(\ell, A)$, we let $\text{Next}(\ell, A, m)$ denote the restriction of $\text{Next}(\ell)$ to locations reachable from ℓ when every player $A_j \in A$ makes the move m_{A_j} .

Strategies and outcomes. Let \mathcal{C} be a CGS. A *computation* of \mathcal{C} is an infinite sequence $\rho = \ell_0 \ell_1 \dots$ of locations such that for any i , $\ell_{i+1} \in \text{Next}(\ell_i)$. We write ρ^i for the i -th suffix of ρ , and $\rho[i \dots j]$ for part of ρ between ℓ_i and ℓ_j . In particular, $\rho[i]$ denotes the $i + 1$ -st location ℓ_i . A *strategy* for a player $A_i \in \text{Agt}$ is a function f_{A_i} that maps any finite prefix of a computation to a possible move for A_i , *i.e.*, satisfying $f_{A_i}(\ell_0 \dots \ell_m) \in \text{Mov}(\ell_m, A_i)$. A strategy is *memoryless* if it only depends on the current state (*i.e.*, $f_{A_i}(\ell_0 \dots \ell_m) = f_{A_i}(\ell_m)$). A strategy for a coalition A of agents is a set of strategies, one for each agent in the coalition. The set of strategies (resp. memoryless strategies) for A is denoted $\text{Strat}(A)$ (resp. $\text{Strat}^0(A)$).

A strategy for A_j induces a set of computations from ℓ , called the *outcomes* of f_{A_j} from ℓ and denoted $\text{Out}(\ell, f_{A_j})$, that player A_j can enforce: $\ell_0 \ell_1 \dots \in \text{Out}(\ell, f_{A_j})$ iff $\ell_0 = \ell$ and $\ell_{i+1} \in \text{Next}(\ell_i, A_j, f_{A_j}(\ell_0 \dots \ell_i))$ for any i . Given a coalition A , a strategy for A is a tuple F_A containing one strategy for each player in A : $F_A = \{f_{A_j} | A_j \in A\}$. The *domain* of F_A ($\text{dom}(F_A)$) is A . The strategy f_{A_j} for A_j is also denoted $(F_A)_{|A_j}$; more generally, $(F_A)_{|B}$ (resp. $(F_A)_{\setminus B}$) denotes the restriction of F_A to the coalition $A \cap B$ (resp. $A \setminus B$). The outcomes of F_A from a location ℓ are the computations enforced by the strategies in F_A : $\ell_0 \ell_1 \dots \in \text{Out}(\ell, F_A)$ iff $\ell_0 = \ell$ and for any i , $\ell_{i+1} \in \text{Next}(\ell_i, A, (f_{A_j}(\ell_0, \dots, \ell_i))_{A_j \in A})$. Note that $\text{Out}(\ell, F_A) \subseteq \text{Out}(\ell, (F_A)_{|B})$ for any coalitions A and B , and in particular that $\text{Out}(\ell, F_\emptyset)$ represents the set of all computations from ℓ .

It is also possible to *combine* two strategies $F \in \text{Strat}(A)$ and $F' \in \text{Strat}(B)$, resulting in a strategy $F \circ F' \in \text{Strat}(A \cup B)$ defined as follows: $(F \circ F')_{|A_j}(\ell_0 \dots \ell_m)$ equals $F_{|A_j}(\ell_0 \dots \ell_m)$ if $A_j \in A$, and it equals $F'_{|A_j}(\ell_0 \dots \ell_m)$ if $A_j \in B \setminus A$.

Finally, given a strategy F , an execution ρ and some integer $i \geq 0$, we define the strategy $F^{\rho, i}$ corresponding to the behaviour of F after prefix $\rho[0 \dots i]$ as follows: $F^{\rho, i}(\pi) = F(\rho[0 \dots i] \cdot \pi)$. Note that if F is memoryless, then $F^{\rho, i} = F$.

Bounded-memory strategies. Between general strategies (without bound over its resources) and *memoryless* strategies, we can consider *bounded-memory strategies*. Let s be a (binary-encoded) integer representing the size of the memory. We define a bounded memory strategy as a memoryless strategy over the locations of the CGS **and** a set of memory cells [Maz02,Tho95]: choosing the move depends on both the location and the current memory cell, and after every move, the player can “update” its memory by moving to another cell. The size of the memory is then defined as the number of cells. Let Cell be the set of $s+1$ memory cells $\{0, \dots, s\}$.

Formally an s -memory strategy F_A for Player A is a 3-tuple $(F^{\text{mov}}, F^{\text{cell}}, c)$ where: F^{mov} is a mapping from $\text{Cell} \times \text{Loc}$ to \mathcal{M} that associates a move with the current memory cell and the current location of the CGS, F^{cell} is a mapping from $\text{Cell} \times \text{Loc}$ to Cell that updates the memory cell, and c is the current memory cell of this strategy. For the sake of readability, given a bounded-memory strategy $F_A = (F^{\text{mov}}, F^{\text{cell}}, c)$, we still write $F_A(\ell)$ for $F^{\text{mov}}(c, \ell)$.

The notions of computations and outcomes are easily extended to this new setting: the set $\text{Next}(\ell, A, F_A(\ell))$ contains the possible successor locations when A plays from ℓ according to F_A . Of course, the memory cell of F_A changes along an execution ρ , and we define $F_A^{\rho, i}$ as the strategy $(F^{\text{mov}}, F^{\text{cell}}, c_i)$ where c_i is defined inductively with: $c_0 = c$ and $c_{j+1} = F^{\text{cell}}(\rho[j], c_j)$. Finally the outcomes $\text{Out}(\ell, F_A)$ are the executions $\rho = \ell_0 \ell_1 \dots$ such that $\ell_{j+1} \in \text{Next}(\ell_j, A, F_A^{\rho, j}(\ell_j))$.

Coalitions are handled the usual way: we use pairs (A, \bar{s}) to represent a coalition $A \subseteq \text{Agt}$ and a memory-bounds vector $\bar{s} \in (\mathbb{N} \cup \{\infty\})^A$ which associates a size $\bar{s}(A_j)$ with the memory that agent $A_j \in A$ can use for its strategy. The set of strategies for A with memory bound \bar{s} is denoted $\text{Strat}^{\bar{s}}(A)$, and we omit to mention the memory bound when none is imposed.

2.3 The logic $\text{ATL}_{sc, \infty}^*$

We now define the logic $\text{ATL}_{sc, \infty}^*$ that extends ATL^* with strategy contexts and bounded-memory strategy quantifiers:

Definition 2. *The syntax of $\text{ATL}_{sc, \infty}^*$ is defined by the following grammar:*

$$\begin{aligned} \text{ATL}_{sc, \infty}^* \ni \varphi_s, \psi_s &::= P \mid \neg \varphi_s \mid \varphi_s \vee \psi_s \mid \langle A, \bar{s} \rangle \varphi_p \mid \cdot \rangle A \langle \cdot \varphi_s \\ \varphi_p, \psi_p &::= \varphi_s \mid \neg \varphi_p \mid \varphi_p \vee \psi_p \mid \mathbf{X} \varphi_p \mid \varphi_p \mathbf{U} \psi_p \end{aligned}$$

with $P \in \text{AP}$, $A \subseteq \text{Agt}$ and $\bar{s} \in (\mathbb{N} \cup \{\infty\})^A$.

An $\text{ATL}_{sc, \infty}^*$ formula Φ is interpreted over a state ℓ of a CGS \mathcal{C} *within a strategy context* $F \in \text{Strat}(B)$ for some coalition B ; this is denoted by $\ell \models_F \Phi$.

The semantics is defined as follows:

$$\begin{aligned}
\ell \models_F \langle A, \bar{s} \rangle \varphi_p & \text{ iff } \exists F_A \in \text{Strat}^{\bar{s}}(A). \forall \rho \in \text{Out}(\ell, F_A \circ F). \rho \models_{F_A \circ F} \varphi_p, \\
\ell \models_F \rangle A \langle \varphi_s & \text{ iff } \ell \models_{F \setminus A} \varphi_s, \\
\rho \models_F \varphi_s & \text{ iff } \rho[0] \models_F \varphi_s, \\
\rho \models_F \mathbf{X} \varphi_p & \text{ iff } \rho^1 \models_{F^{\rho,1}} \varphi_p, \\
\rho \models_F \varphi_p \mathbf{U} \psi_p & \text{ iff } \exists i. \rho^i \models_{F^{\rho,i}} \psi_p \text{ and } \forall 0 \leq j < i. \rho^j \models_{F^{\rho,j}} \varphi_p.
\end{aligned}$$

Given a CGS \mathcal{C} with initial location ℓ_0 , and an $\text{ATL}_{sc,\infty}^*$ formula Φ , the model-checking problem consists in deciding whether² $\ell_0 \models_{\emptyset} \Phi$.

The formula $\langle A, \bar{s} \rangle \varphi$ holds on a location ℓ within a context F for a coalition B iff there exists a \bar{s} -memory strategy for A to enforce φ when B plays according to the strategy F . We use $\langle A \rangle$ to denote the modality with no restriction over the memory allowed for the strategies of A (*i.e.*, the modality $\langle A, \infty^A \rangle$); and we use $\langle A^0 \rangle$ as an abbreviation for $\langle A, 0^A \rangle$ to consider only memoryless strategies.

Conversely the modality $\rangle A \langle$ removes the strategy for A from the current context under which the formula is interpreted. The operator $\rangle \text{Agt} \langle$ allows us to empty the current context, and then we clearly have: $\ell \models_F \rangle \text{Agt} \langle \varphi \Leftrightarrow \ell \models_{F'} \rangle \text{Agt} \langle \varphi$ for any context F and F' .

This entails that $\text{ATL}_{sc,\infty}^*$ contains ATL^* (thus also CTL^*). Indeed the classical strategy quantifier of ATL^* , namely $\langle\langle A \rangle\rangle$, does not handle strategy contexts: $\langle\langle A \rangle\rangle \varphi$ holds for a location ℓ iff A has a strategy to enforce φ whatever the choices of $\text{Agt} \setminus A$. Clearly $\langle\langle A \rangle\rangle \varphi$ is equivalent to $\rangle \text{Agt} \langle \langle A \rangle \varphi$.

Obviously the existence of an \bar{s} -memory strategy for A to enforce φ entails the existence of an \bar{s}' -memory strategy if $\bar{s}' \geq \bar{s}$ (*i.e.*, $\bar{s}'(A_j) \geq \bar{s}(A_j)$ for all $A_j \in A$). Note that the converse is not true except for special cases such as ATL where memoryless strategies are sufficient (see [AHK02,Sch04]).

We will use standard abbreviations such as $\top = P \vee \neg P$, $\perp = \neg \top$, $\mathbf{F} \varphi = \top \mathbf{U} \varphi$, etc.

Now we introduce several fragments of $\text{ATL}_{sc,\infty}^*$:

- $\text{ATL}_{sc,b}^*$ (with $b \in \mathbb{N}$) is the fragment of $\text{ATL}_{sc,\infty}^*$ where the quantifiers $\langle A, \bar{s} \rangle$ only use memory-bounds less than or equal to b . In particular, $\text{ATL}_{sc,0}^*$ only allows memoryless strategies.
- ATL_{sc}^* is the fragment of $\text{ATL}_{sc,\infty}^*$ where no restriction over the memory is allowed (any strategy quantifier deals with infinite-memory strategies).
- $\text{ATL}_{sc,\infty}$ contains the formulae where every temporal modality is in the immediate scope of a strategy quantifier (*i.e.*, the path formulae are restricted to $\varphi_s \mathbf{U} \psi_s$, $\varphi_s \mathbf{R} \psi_s$ — \mathbf{R} is the “dual-until” modality—and $\mathbf{X} \varphi_s$). It follows from the above assertion that $\text{ATL}_{sc,\infty}$ contains ATL and CTL . We also define the fragments $\text{ATL}_{sc,b}$ and ATL_{sc} as above.

² The context can be omitted when it is empty, and we can directly write $\ell \models \Phi$.

3 Expressiveness

In this section, we consider expressiveness issues, first illustrating the ability of $\text{ATL}_{sc,\infty}^*$ to state interesting properties, and then comparing it with related formalisms.

3.1 Some interesting formulas of $\text{ATL}_{sc,\infty}^*$

The new modalities $\langle A \rangle$ allow us to express many interesting properties over the strategies of different players in a game. In [BDCLM08], we show how our logics can express the different properties that motivated the introduction of SL, QD_μ or IATL. Here we just give a few examples.

Nash equilibria. Given two players A_1 and A_2 having their own objectives Φ_1 and Φ_2 , two strategies F_1 and F_2 for players 1 and 2 respectively form a *Nash equilibrium* if there is no “better” strategy F'_1 for A_1 w.r.t. Φ_1 when Player 2 plays according to F_2 , and *vice versa*. Given a strategy context $F = (F_1, F_2)$, the following formula holds in state ℓ under F iff F_1 and F_2 form a Nash equilibrium in ℓ :

$$\left((\langle A_1 \rangle \Phi_1) \Rightarrow \Phi_1 \wedge (\langle A_2 \rangle \Phi_2) \Rightarrow \Phi_2 \right)$$

This provides us with a way of expressing the existence of Nash equilibria having extra properties.

Winning secure equilibria. The *winning secure equilibrium* [CHJ06] (WSE) is a stronger notion of equilibrium: two strategies F_1 and F_2 , for players 1 and 2 with objectives Φ_1 and Φ_2 respectively, form a WSE if each player has no better strategy for himself, and no worse strategy for his opponent. Again, the strategy context F is a winning secure equilibrium in ℓ iff the following formula holds in ℓ within F :

$$\begin{aligned} & (\langle A_1 \rangle \Phi_1) \Rightarrow \Phi_1 \wedge (\langle A_2 \rangle \Phi_2) \Rightarrow \Phi_2 \wedge \\ & \left(\langle A_1 \rangle (\Phi_1 \wedge \neg \Phi_2) \Rightarrow (\Phi_1 \wedge \neg \Phi_2) \right) \wedge \left(\langle A_2 \rangle (\Phi_2 \wedge \neg \Phi_1) \Rightarrow (\Phi_2 \wedge \neg \Phi_1) \right) \end{aligned}$$

Client-server interactions. Given a protocol where a server S has to treat the requests of different agents A_1, \dots, A_n , we can express that S has a strategy to ensure that every agent A_i can act in order to make its requests to be granted. Such a property can be stated as follows:

$$\langle S \rangle \mathbf{G} \left[\bigwedge_{i=1..n} \left(\text{req}_i \Rightarrow \langle A_i \rangle \mathbf{F} \text{grant}_i \right) \right]$$

Clearly this property requires the use of strategy contexts because every agent has to cooperate with the server (but not with other agents).

3.2 Comparison with other formalisms

Figure 1 summarizes the expressiveness results for our logics. An arrow $L \rightarrow L'$ denotes that $L \leq_{ex} L'$, *i.e.*, that L' is at least as expressive as L (*i.e.*, for any formula in L , there exists an equivalent³ formula in L'). Note that in some cases, the relation is strict and we have $L < L'$, see the corresponding theorems for more details. The dotted arrow correspond to results proved in this paper; plain arrows correspond to literature results (they are labeled with bibliographic references) or direct syntactic inclusions.

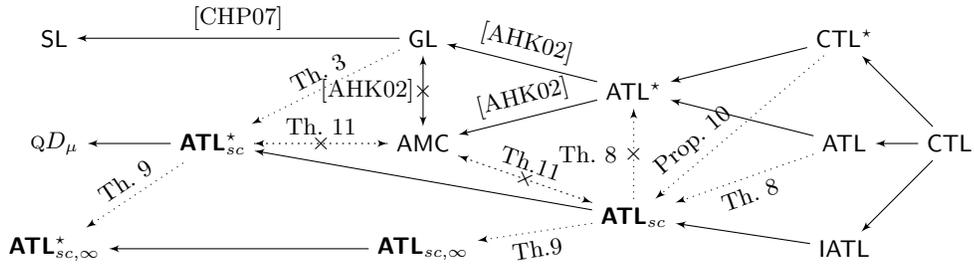


Fig. 1. Expressiveness of $ATL_{sc,\infty}$ and $ATL_{sc,\infty}^*$ compared to classical logics

The results about ATL , ATL^* , CTL^* and AMC are presented in Appendix A.1: most of them are based on the ability of the new modalities with strategy contexts and/or memory bounds to *distinguish* models that are alternating-bisimilar (and thus satisfy the same formulas of the classical AMC fragments). In this appendix, we also show that the $\cdot \langle \cdot \rangle$ modality does not add expressive power to ATL_{sc}^* ; this result is used to compare our logics with QD_μ in Appendix A.2. We also show that adding the quantification over bounded memory increases the expressive power of ATL_{sc} and ATL_{sc}^* (see Th.9 in Appendix).

Here we only develop our results concerning Game Logic, which is a powerful logic to handle properties over strategies, and we discuss the case of Strategy Logic.

Comparison with Game Logic. *Game Logic* was introduced in [AHK02] in order to express the module-checking problem [KVV01]. This logic is an extension of ATL^* where the existence of a strategy for A is written $\exists A$, and where it is possible to deal explicitly with the execution tree induced by a strategy: given such a tree t , it is possible to quantify (with modalities \exists and \forall) over the executions inside t and specify temporal properties. For example, the formula $\exists A.((\exists P \mathbf{U} P') \wedge (\forall F P''))$ specifies the existence of a strategy F_A for A s.t. in the tree induced by F_A , we have: (1) there exists a run along which $P \mathbf{U} P'$

³ That is, having the same truth value in any location of any CGS under any context.

holds and (2) every run satisfies $\mathbf{F}P''$. The formal definition of \mathbf{GL} is given in Appendix A.2. We have the following result:

Theorem 3. $ATL_{sc}^* >_{ex} \mathbf{GL}$

Proof (sketch). First we give a translation from \mathbf{GL} into ATL_{sc}^* ; it is given by the inductively-defined application σ s.t.:

$$\sigma(\exists A\varphi) \stackrel{\text{def}}{=} \langle\langle A \rangle\rangle \bar{\varphi} \quad \bar{\exists\varphi} \stackrel{\text{def}}{=} \neg \langle\emptyset\rangle \neg \bar{\varphi} \quad \bar{P} \stackrel{\text{def}}{=} P.$$

The other inductive rules are defined in the natural way. Note that if φ is a \mathbf{GL} tree-formula, then $\bar{\varphi}$ is an ATL_{sc}^* state-formula. In this translation, we use a strategy context to represent the tree used to interpret \mathbf{GL} path- and tree-formulae. In the following, given a state ℓ of a CGS and a strategy F for some coalition A , we use $\text{ExecTree}(\ell, F)$ to denote the subtree of the computation tree from ℓ whose infinite rooted paths are the elements of $\text{Out}(\ell, F)$. We must show that for any \mathbf{GL} path (resp. tree) formula φ_p (resp. φ_t), any path ρ in some CGS and any strategy F for some coalition A , we have: $(\text{ExecTree}(\rho[0], F), \rho) \models \varphi_p$ iff $\rho \models_F \bar{\varphi}_p$ and $\text{ExecTree}(\rho[0], F) \models \varphi_t$ iff $\rho[0] \models_F \bar{\varphi}_t$. Here we just consider the first equivalence (the second one can be treated in a similar way). The usual induction steps are straightforward, thus we only consider the following two cases :

- $\varphi = \exists A.\psi$. Then $\rho \models_F \bar{\varphi}$ means that there is a strategy F' for coalition A , s.t. any computation ρ' in $\text{Out}(\rho[0], F')$ satisfies $\bar{\psi}$. By i.h., this is equivalent to $\exists F' \in \text{Strat}(A). \forall \rho' \in \text{Out}(\rho[0], F'). (\text{ExecTree}(\rho[0], F'), \rho') \models \psi$, hence to $\rho[0] \models \exists A.\psi$ because ψ is a tree formula. Now $\exists A.\psi$ is a state formula that can be interpreted over any execution tree with root $\rho[0]$, in particular over $\text{ExecTree}(\rho[0], F)$.
- $\varphi = \exists\psi$. Then $\rho \models_F \bar{\varphi}$ means that “not all the computations from $\rho[0]$ and following the strategy context F do not satisfy $\bar{\psi}$ ”, and is then equivalent to $\exists \rho' \in \text{Out}(\rho[0], F). \rho' \models_F \bar{\psi}$. Again from the i.h. we obtain the existence of a path in $\text{ExecTree}(\rho[0], F)$ satisfying ψ , and then $\text{ExecTree}(\rho[0], F) \models \exists\psi$, which is equivalent to $(\text{ExecTree}(\rho[0], F), \rho) \models \varphi$ (as φ is a tree formula).

Finally if we consider the case where φ is a state formula and F is the empty strategy, we get that $\bar{\varphi}$ is an ATL_{sc}^* equivalent formula for φ .

We have $\mathbf{GL} <_{ex} ATL_{sc}^*$ because the ATL_{sc} formula $\langle A_1 \rangle \mathbf{X} (\langle A_2 \rangle \mathbf{X} b \wedge \langle A_3 \rangle \mathbf{X} a)$ has no equivalent in \mathbf{GL} . Indeed consider the CGSs \mathcal{S}_1 and \mathcal{S}_2 in Figure 2. They satisfy the same \mathbf{GL} formulas, since move 3 for Player 1 (in \mathcal{S}_2) does not affect the sets of execution trees induced by all strategies for a fixed coalition: for any coalition A and state q , we have $\text{ExecTree}(q, \text{Strat}_{\mathcal{S}_1}(A)) = \text{ExecTree}(q, \text{Strat}_{\mathcal{S}_2}(A))$. Yet this move ensures that s'_0 satisfies $\langle A_1 \rangle \mathbf{X} (\langle A_2 \rangle \mathbf{X} b \wedge \langle A_3 \rangle \mathbf{X} a)$ (when players 2 and 3 respectively choose moves 2 and 1), while s_0 does not. \square

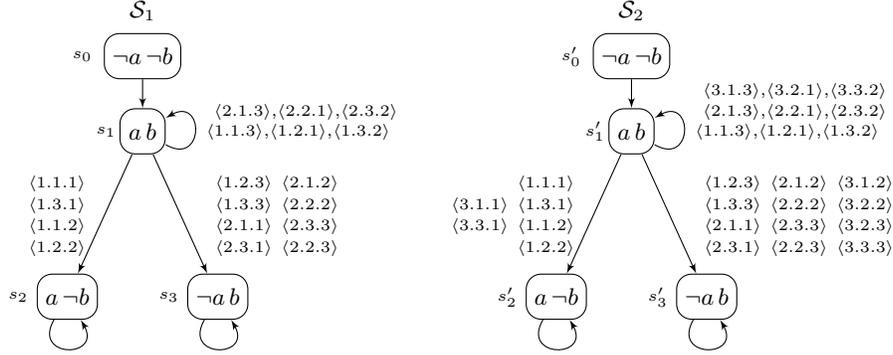


Fig. 2. \mathcal{S}_1 and \mathcal{S}_2 cannot be distinguished by GL

Comparison with Strategy Logic [CHP07]. Strategy Logic has been defined in [CHP07] as an extension of LTL with first-order quantification on strategies. That player A has a strategy to enforce φ is then written $\exists \sigma_A. \forall \sigma_B. \varphi(\sigma_A, \sigma_B)$ where the arguments (*i.e.*, the strategies for the two players) given to φ indicate on which paths φ is evaluated.

While this logic has only been defined on 2-player turn-based games, its definition can easily be extended to our n -player CGS framework. We conjecture that $\text{ATL}_{sc,\infty}$ and SL are incomparable (proving those results seems to be especially challenging due to the particular syntax of SL):

- SL can explicitly manipulate strategies as first-order elements. It can for instance state properties such as $\exists x_1. \exists y_1. \exists x_2. \exists y_2. [\varphi_1(x_1, y_1) \wedge \varphi_2(x_2, y_2) \wedge \varphi_3(x_1, y_2) \wedge \varphi_4(x_2, y_2)]$ which (we conjecture) $\text{ATL}_{sc,\infty}$ cannot express due to the circular constraint.
- on the other hand, SL requires subformulas embedded in modalities to be closed. As a consequence, formula $\exists x_1. \forall y_1. [\mathbf{G}(\exists y_2. [\mathbf{F}p](x_1, y_2))](x_1, y_1)$ is not an SL formula (because $\exists y_2. [\mathbf{F}p](x_1, y_2)$ is not closed), while it is expressed in $\text{ATL}_{sc,\infty}$ as $\langle A \rangle \mathbf{G}(\langle B \rangle \mathbf{F}p)$.

However, it should be noticed that the simple one-alternation fragment of SL (denoted S1ASL) can be translated into $\text{ATL}_{sc,\infty}^*$. Indeed this fragment is built from formulas of the form $\exists x_1. \exists y_1. \forall x_2. \forall y_2. [\varphi_1(x_1, y_2) \wedge \varphi_2(x_2, y_1) \wedge \varphi_3(x_1, y_1)]$ [CHP07] which we can express as $\langle A_1 \rangle [\varphi_1 \wedge \langle A_2 \rangle (\varphi_3 \wedge \cdot) A_1 \langle \varphi_2 \rangle]$.

4 $\text{ATL}_{sc,\infty}$ and $\text{ATL}_{sc,\infty}^*$ model-checking

We begin with proving that model-checking is decidable for our logic. Still, as is the case for Strategy Logic, the resulting algorithm is non-elementary. We thus mainly focus on simpler cases (namely, memoryless and bounded-memory strategies), where more efficient algorithms can be obtained.

Theorem 4. *Model checking $\text{ATL}_{sc,\infty}^*$ formulas over CGS is decidable.*

Proof (sketch). The translation from ATL_{sc}^* to QD_μ (see Section A.2) yields decidability of ATL_{sc}^* . Moreover, as we will see in Section 4.2, it is possible to encode the bounded-memory strategies as memoryless strategies over an extended CGS. Since memorylessness can be expressed with QD_μ , this provides an indirect algorithm for $\text{ATL}_{sc,\infty}^*$ model checking. \square

4.1 Model-checking $\text{ATL}_{sc,0}^*$ and $\text{ATL}_{sc,0}$

Theorem 5. *The model checking problems for $\text{ATL}_{sc,0}^*$ and $\text{ATL}_{sc,0}$ over CGSs are PSPACE-complete.*

Proof. We only address the membership in PSPACE. The hardness proof is given in Lemma 16 in Appendix A.3 (and is similar to that of [BBGK07]).

Let \mathcal{C} be a CGS, ℓ a location and F a memoryless strategy context, assigning a memoryless strategy to each player of some coalition A . Since F contains only memoryless strategies, it associates with each location one move for each agent in A . Dropping the other moves of those agents, we get a CGS, denoted (\mathcal{C}, F) , whose set of executions is exactly the set of outcomes of F in \mathcal{C} .

From this and the fact that a memoryless strategy can be stored in space $O(|Q|)$, we get a simple PSPACE model-checking algorithm for $\text{ATL}_{sc,0}^*$ that relies on a (PSPACE) model-checking algorithm for LTL. The main difficulty is that strategy contexts prevent us from proceeding in a standard bottom-up fashion. As a consequence, our algorithm consists in enumerating strategies starting from outermost strategy quantifiers.

If φ is an $\text{ATL}_{sc,0}^*$ path formula, we denote by $\Phi(\varphi)$ the set of outermost *quantified* φ subformulae (*i.e.* of the form $\langle A \rangle \psi$), and by $\sigma(\varphi)$ the corresponding LTL formula where all subformulae $\psi \in \Phi(\varphi)$ have been replaced by new propositions a_ψ . We enumerate all possible contexts, recursively calling the algorithm at each step of the enumeration, and thus gradually taking care of each labelling a_ψ . Algorithm 1 describes the procedure. \square

Algorithm 1 : MC- $\text{ATL}_{sc,0}^*(\mathcal{C}, F, \ell_0, \varphi)$ – $\text{ATL}_{sc,0}^*$ model checking

Require: a CGS \mathcal{C} , $F \in \text{Strat}^0(A)$, $\ell_0 \in \text{Loc}$ and an $\text{ATL}_{sc,0}^*$ path formula φ

Ensure: YES iff $\forall \lambda \in \text{Out}(\ell_0, F)$, $\lambda \models_F \varphi$

$\mathcal{C}' := (\mathcal{C}, F)$

foreach $\psi \in \Phi(\varphi)$ **do**

case $\psi = \langle B^0 \rangle \psi'$:

for $F_B \in \text{Strat}^0(B)$, $\ell \in \text{Loc}$ **do**

if MC- $\text{ATL}_{sc,0}^*(\mathcal{C}, F_B \circ F, \ell, \psi')$, **then** label ℓ with a_ψ

case $\psi = \rangle B \langle \psi'$:

for $\ell \in \text{Loc}$ **do**

if MC- $\text{ATL}_{sc,0}^*(\mathcal{C}, F \setminus_B, \ell, \psi')$, **then** label ℓ with a_ψ

return MC LTL $(\mathcal{C}', \ell_0, \mathbf{A}\sigma(\varphi))$

Remark 1. Note that PSPACE-completeness straightforwardly extends to “memoryless” extensions (*i.e.*, with quantification over memoryless strategies) of ATL^* and SL. Since ATL objectives do not require memory, ATL_0 is the same as ATL, and its model-checking problem is PTIME-complete. Moreover a similar algorithm would work for *symbolic CGSs*, a succinct encoding of CGS proposed in [JD05,LMO07].

Also notice that both the above algorithm and the PSPACE-hardness proof can be adapted to IATL. This corrects the Δ_2^P -completeness result of [ÅGJ07].

4.2 Bounded-memory strategies

The case of bounded-memory strategies can be handled in a similar way as memoryless strategies. Indeed as explained in Subsection 2.2, we can see an s -bounded strategy for Player A_i as a memoryless strategy over an extended structure containing the original CGS \mathcal{C} and a particular CGS controlled by A_i and describing its memory. Formally, for a player A_i , we define the CGS $\mathbb{M}_{A_i}^s$ as follows: $\mathbb{M}_{A_i}^s = (\text{Agt}, \text{Loc}_s^i, \emptyset, \text{Loc}_s^i \times \text{Loc}_s^i, \emptyset, \mathcal{M}_s^i \cup \{\perp\}, \text{Mov}_s^i, \text{Edg}_s^i)$ where

- $\text{Loc}_s^i = \{0, \dots, s\}$ is the set of (unlabeled) locations;
- \mathcal{M}_s^i is isomorphic to Loc_s^i (and we identify both sets),
- Mov_s^i and Edg_s^i do not depend on the location: Mov_s^i allows only one move \perp to each player, except for player A_i , who is allowed to play any move in \mathcal{M}_s^i . Then Edg_s^i returns the location chosen by A_i .

Let $\bar{s} \in \mathbb{N}^{\text{Agt}}$ be a memory-bounds vector. Now considering the product structure $\mathcal{C}_{\bar{s}} = \prod_{A_i \in \text{Agt}} \mathbb{M}_{A_i}^{\bar{s}(A_i)} \times \mathcal{C}$, for all player A_j we can very simply export $\bar{s}(A_j)$ memory-bounded strategies of \mathcal{C} to *some* memoryless strategies over $\mathcal{C}_{\bar{s}}$. Indeed, given a Player A_j , we do not want to consider all memoryless strategies f over $\mathcal{C}_{\bar{s}}$ but only the ones where A_j exclusively uses the information from $\mathbb{M}_{A_j}^{\bar{s}(A_j)}$ (*i.e.*, such that $f(i_1, \dots, i_j, \dots, i_k, l) = f(0, \dots, i_j, \dots, 0, l)$). Let $\text{RStrat}_{\mathcal{C}_{\bar{s}}}^0(A_j)$ be this restricted set of such strategies; clearly we have $\text{RStrat}_{\mathcal{C}_{\bar{s}}}^0(A_j) \subseteq \text{Strat}_{\mathcal{C}_{\bar{s}}}^0(A_j)$.

Adapting the proof of Theorem 5 to memory-bounded strategies, we get:

Proposition 6. *Let $C = (\text{Agt}, \text{Loc}, \ell_0, AP, \text{Lab}, \mathcal{M}, \text{Mov}, \text{Edg})$ be a CGS. Let $\varphi \in \text{ATL}_{sc,b}^*$ involving only \bar{s} -memory quantifiers. Then φ can be checked in exponential space.*

Proof. We run the algorithm of Theorem 5 over the structure $\mathcal{C}_{\bar{s}}$, restricting the enumerations of $\text{Strat}_{\mathcal{C}_{\bar{s}}}^0(B)$ to those of $\text{RStrat}_{\mathcal{C}_{\bar{s}}}^0(B)$. \square

Remark 2. – If the memory-bounds \bar{s} were given in unary, our algorithm would be PSPACE, since the LTL model-checking over the product structure can be performed on-the-fly.

- Note that this algorithm can deal with formulas containing several subformulas $\langle A, \bar{s}_1 \rangle \varphi_1, \dots, \langle A, \bar{s}_p \rangle \varphi_p$ with different different memory-bounds s_i (for the same coalition A).

- Since our algorithm consists in enumerating the strategies, it could cope with games of incomplete information, where the strategies would be based on (some of) the atomic propositions labeling a location, rather than on the location itself [Sch04].
- Bounded-memory quantification can be defined also for the other formalisms where memory-based strategies are needed, *e.g.* ATL^* or SL. Our EXPSPACE algorithm could easily be adapted to that case.

5 Conclusion

In this paper we propose powerful extensions of ATL and ATL^* logics. These extensions allow us to express many interesting and complex properties that have motivated the definition of new formalisms in the past. An advantage of these extensions is to treat strategies through modalities as in ATL and ATL^* .

As future work, we plan to study the exact complexity of model-checking $ATL_{sc,\infty}$ and $ATL_{sc,\infty}^*$, with the aim of finding reasonably efficient algorithms for fragments of these expressive logics. Finally we think that the ability to deal explicitly with bounded-memory strategies is an interesting approach to develop.

References

- [ÅGJ07] Thomas Ågotnes, Valentin Goranko, and Wojciech Jamroga. Alternating-time temporal logics with irrevocable strategies. In *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK'07)*, pages 15–24, 2007.
- [AHK02] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
- [BBGK07] Christel Baier, Tomáš Brázdil, Marcus Größer, and Antonín Kučera. Stochastic game logic. In *Proceedings of the 4th International Conference on Quantitative Evaluation of Systems (QEST'07)*, pages 227–236. IEEE Comp. Soc. Press, 2007.
- [BDCLM08] Thomas Brihaye, Arnaud Da Costa, François Laroussinie, and Nicolas Markey. Atl with strategy contexts and bounded memory. Technical Report 08-014, Lab. Specification et Verification, February 2008.
- [CE81] Edmund M. Clarke and E. Allen Emerson. Design and synthesis of synchronous skeletons using branching-time temporal logic. In *Proc. 3rd Workshop Logics of Programs (LOP'81)*, volume 131 of *LNCS*, pages 52–71. Springer, 1981.
- [CHJ06] Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdzinski. Games with secure equilibria. *Theoretical Computer Science*, 365(1-2):67–82, 2006.
- [CHP07] Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Strategy logic. In *Proc. of the 18th International Conference on Concurrency Theory (CONCUR'07), Lisbon, Portugal*, volume 47013 of *LNCS*, pages 59–73. Springer, 2007.

- [JD05] Wojciech Jamroga and Jürgen Dix. Do agents make model checking explode (computationally)? In *Proceedings of the 4th International Central and Eastern European Conference on Multi-Agent Systems (CEEMAS'05)*, volume 3690 of *LNCS*. Springer, 2005.
- [KVVW01] O. Kupferman, M. Y. Vardi, and P. Wolper. Module checking. *Information and Computation*, 164(2):322–344, 2001.
- [LMO07] François Laroussinie, Nicolas Markey, and Ghassan Oreiby. On the expressiveness and complexity of ATL. In *Proc. 10th Intl Conf. Foundations of Software Science and Computation Structures (FoSSaCS'07)*, volume 4423 of *LNCS*, pages 243–257, Braga, Portugal, 2007. Springer.
- [Maz02] René Mazala. Infinite games. In *Automata, Logics, and Infinite Games*, volume 2500 of *LNCS*, pages 23–42. Springer-Verlag, 2002.
- [Pin07] Sophie Pinchinat. A generic constructive solution for concurrent games with expressive constraints on strategies. In *Proceedings of the 5th International Symposium on Automated Technology for Verification and Analysis (ATVA'07)*, LNCS. Springer-Verlag, 2007.
- [Pin08] Sophie Pinchinat. Personal communication, 2008.
- [Pnu77] Amir Pnueli. The temporal logic of programs. In *Proc. 18th Ann. Symp. Foundations of Computer Science (FOCS'77)*, pages 46–57. IEEE Comp. Soc. Press, 1977.
- [QS82] Jean-Pierre Queille and Joseph Sifakis. Specification and verification of concurrent systems in CESAR. In *Proc. 5th Intl Symp. on Programming (SOP'82)*, volume 137 of *LNCS*, pages 337–351. Springer, 1982.
- [Sch04] Pierre-Yves Schobbens. Alternating-time logic with imperfect recall. In *Proceedings of the 1st Workshop on Logic and Communication in Multi-Agent Systems (LCMAS'03)*, volume 85 of *ENTCS*. Elsevier, 2004.
- [Tho95] Wolfgang Thomas. On the synthesis of strategies in infinite games. In *Proc. 12th Ann. Symp. Theoretical Aspects of Computer Science (STACS '95), Munich, Germany, Mar. 1995*, volume 900 of *Lecture Notes in Computer Science*, pages 1–13. Springer, 1995.

A Appendix

A.1 Comparison with ATL, ATL^{*}, AMC

Comparison with ATL^{*}.

Proposition 7. – $\langle \text{Agt}^0 \rangle \mathbf{X} (\langle \emptyset^0 \rangle \mathbf{X} P)$ has no equivalent in ATL^{*},
– $\langle A_1 \rangle \mathbf{X} (\langle A_2 \rangle \mathbf{X} P \wedge \langle A_2 \rangle \mathbf{X} P')$ has no equivalent in ATL^{*}.

Proof. – Consider the two one-player CGS \mathcal{C}_1 and \mathcal{C}_2 in Fig. 3. They clearly satisfy the same ATL^{*} formulae since \mathcal{C}_2 corresponds to an unfolding of \mathcal{C}_1 . But we have:

$$s_0 \not\models \langle \text{Agt}^0 \rangle \mathbf{X} (\langle \emptyset^0 \rangle \mathbf{X} a) \quad s'_0 \models \langle \text{Agt}^0 \rangle \mathbf{X} (\langle \emptyset^0 \rangle \mathbf{X} a)$$

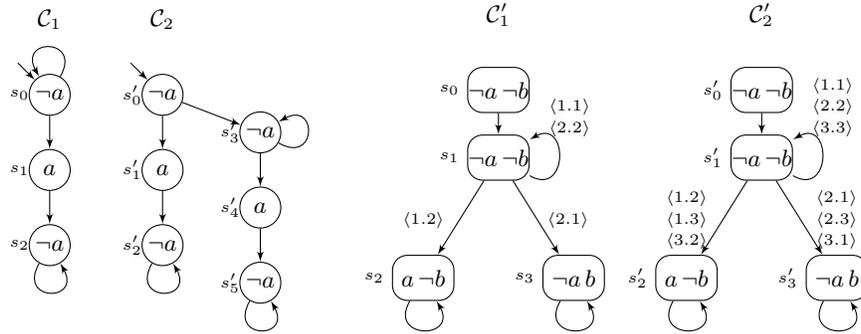


Fig. 3. \mathcal{C}_1 and \mathcal{C}_2 (resp \mathcal{C}'_1 and \mathcal{C}'_2) cannot be distinguished by ATL^{*}

Indeed any memoryless strategy in s_0 will contain either the infinite path $s_0 \rightarrow s_0 \rightarrow s_0 \rightarrow \dots$, or the path $s_0 \rightarrow s_1 \rightarrow s_2 \rightarrow s_2 \rightarrow \dots$. In order to have a after two transitions, it is necessary to first take the transition $s_0 \rightarrow s_0$ and secondly the transition $s_0 \rightarrow s_1$. This can only be achieved by a strategy with memory. On the contrary, s'_0 satisfies clearly the property. This entails that $\langle \text{Agt}^0 \rangle \mathbf{X} (\langle \emptyset^0 \rangle \mathbf{X} a)$ has no equivalent in ATL^{*}.

- Now consider the CGSs \mathcal{C}'_1 and \mathcal{C}'_2 in Fig. 3. They satisfy the same ATL^{*} formulae: in \mathcal{C}'_2 , a third move is possible for both players but it does not give more strategies to enforce particular ATL^{*} properties. But we clearly have that s'_0 satisfies $\langle A_1 \rangle \mathbf{X} (\langle A_2 \rangle \mathbf{X} (a \wedge \neg b) \wedge \langle A_2 \rangle \mathbf{X} (\neg a \wedge b))$ thanks to the third move of Player 1. But this is not the case for s_0 that does not satisfy the ATL_{sc} formula. Thus $\langle A_1 \rangle \mathbf{X} (\langle A_2 \rangle \mathbf{X} P \wedge \langle A_2 \rangle \mathbf{X} P')$ has no equivalent in ATL^{*}.

□

As a direct corollary of the previous results and of the inclusion of ATL (resp. ATL^{*}) in ATL_{sc,∞} (resp. ATL_{sc,∞}^{*}), we obtain the following theorem:

Theorem 8. *For any $b \in \mathbb{N}$, we have:*

- $\{ATL_{sc,\infty}, ATL_{sc,b}, ATL_{sc}\} >_{ex} ATL$
- $\{ATL_{sc,\infty}^*, ATL_{sc,b}^*, ATL_{sc}^*\} >_{ex} ATL^*$
- $ATL^* \not\leq_{ex} \{ATL_{sc,\infty}, ATL_{sc,b}, ATL_{sc}\}$

Remark 3. Note that, for one-player CGS, the use of strategy contexts does not add expressive power to ATL^{*}: a subformula $\langle A \rangle \varphi$ is equivalent to $\langle\langle A \rangle\rangle \varphi$, and $\langle \emptyset \rangle \varphi$ is equivalent to φ (resp. $\langle\langle \emptyset \rangle\rangle \varphi$) if it is interpreted in a strategy context for Agent A (resp. in an empty strategy context). Thus in one-player CGS, any ATL_{sc}^{*} formula is equivalent to an ATL^{*} or CTL^{*} formula (in such structures $\langle\langle A \rangle\rangle$ is equivalent to the **E** quantifier, and $\langle\langle \emptyset \rangle\rangle$ is equivalent to **A**).

We also have the following result showing the expressive power of memory-bounds in modalities $\langle - \rangle$:

Theorem 9. $ATL_{sc,\infty} >_{ex} ATL_{sc}$ and $ATL_{sc,\infty}^* >_{ex} ATL_{sc}^*$.

Proof. Consider the two one-player CGS \mathcal{C}_1 and \mathcal{C}_2 in Fig. 3. As explained in the proof of Proposition 7, they can be distinguished by the ATL_{sc,∞} formula $\langle \text{Agt}^0 \rangle \mathbf{X} (\langle \emptyset^0 \rangle \mathbf{X} P)$. But these structures are one-player CGSs, and thus any ATL_{sc}^{*} formula is equivalent –over \mathcal{C}_1 and \mathcal{C}_2 – to some CTL^{*} formula. As \mathcal{C}_1 and \mathcal{C}_2 are bisimilar, they cannot be distinguished by any CTL^{*} formula. Thus we have $ATL_{sc,\infty} >_{ex} ATL_{sc}$ and $ATL_{sc,\infty}^* >_{ex} ATL_{sc}^*$. \square

Finally note that the use of strategy contexts in ATL allows us to reach the expressive power of CTL^{*}:

Proposition 10. $CTL^* \leq_{ex} ATL_{sc}$

Proof. Let Ψ be a CTL^{*} formula. We first define a translation to build an “almost ATL_{sc}” formula $\bar{\Psi}$ as follows:

$$\begin{aligned} \bar{P} &\stackrel{\text{def}}{=} P & \bar{\mathbf{X}}\varphi &\stackrel{\text{def}}{=} \langle \emptyset \rangle \mathbf{X} \bar{\varphi} \\ \bar{\neg\varphi} &\stackrel{\text{def}}{=} \bar{\varphi} & \bar{\varphi} \mathbf{U} \psi &\stackrel{\text{def}}{=} \langle \emptyset \rangle (\bar{\varphi} \mathbf{U} \bar{\psi}) \\ \bar{\varphi \vee \psi} &\stackrel{\text{def}}{=} \bar{\varphi} \vee \bar{\psi} & \bar{\mathbf{E}}\varphi &\stackrel{\text{def}}{=} \langle \text{Agt} \rangle \bar{\varphi} \end{aligned}$$

We clearly have $q \models \Phi$ iff $q \models_{\emptyset} \bar{\Phi}$. The existential path quantifier is replaced by $\langle \text{Agt} \rangle$. Now given a path formula $\mathbf{X} \varphi$, its corresponding $\bar{\mathbf{X}} \bar{\varphi}$ will be interpreted within a strategy context – fixed by some $\langle \text{Agt} \rangle$ modality – defining the path. Thus inserting a $\langle \emptyset \rangle$ does not change the semantics and allows us to put \mathbf{X} in the immediate scope of a strategy quantifier as required in ATL_{sc,∞}. Nevertheless $\bar{\Psi}$ does not yet belong to ATL_{sc,∞} because it is possible to have a double embedding of strategy quantifiers of the following form: $\langle \text{Agt} \rangle (\langle \emptyset \rangle \varphi \wedge \dots)$. This can be removed *e.g.* by inserting $\perp \mathbf{U}$ between both quantifications. \square

Comparison with AMC. Alternating-time μ -calculus is neither more nor less expressive than our extensions of ATL:

Theorem 11. $AMC \not\leq_{ex} ATL_{sc}^*$, and $ATL_{sc} \not\leq_{ex} AMC$

Proof. – As explained in Remark 3, when considering one-player CGS (*i.e.*, Kripke structures), ATL_{sc}^* is equivalent to CTL^* , which is strictly less expressive than the (classical) μ -calculus.

- Consider the two CGS C'_1 and C'_2 in Figure 3. The extra move allowed from s'_1 does not change the truth value of AMC formulae. Indeed from s_1 , any strategy for A_1 or A_2 to ensure some AMC formula also exists from s'_1 . And from s'_1 , the extra move 3 (both for A_1 and A_2) allows more target nodes for the opponent: any strategy using this move, can be simulated with the move 1 or 2. Thus these two CGS satisfy the same AMC formulae. As they can be distinguished by ATL_{sc} (and ATL_{sc}^*), we have $ATL_{sc} \not\leq_{ex} AMC$. \square

Expressiveness of $\rangle A \langle$ quantifier We have illustrated the use of modality $\rangle A \langle$ by expressing the classical ATL^* modality $\langle\langle A \rangle\rangle$ with $\rangle \text{Agt} \langle \langle A \rangle$: we first forget the current strategy context and then quantify over the existence of a strategy for A : relaxing is necessary because it has to be a real strategy, *i.e.*, correct for any choice for the other agents. In fact, this modality does not add expressive power to $ATL_{sc,\infty}^*$:

Proposition 12. *For any $ATL_{sc,\infty}^*$ formula Φ , there exists a formula Ψ containing no $\rangle - \langle$ modality such that $\Phi \equiv \Psi$.*

Proof. Given a subset of agents $C \subseteq \text{Agt}$ and $\Phi \in ATL_{sc,\infty}^*$, we define formula $\overline{\Phi}^C$ recursively as follows:

$$\begin{aligned} \overline{\langle A, \bar{s} \rangle \Phi}^C &\stackrel{\text{def}}{=} \langle A, \bar{s} \rangle [C \setminus A] \overline{\Phi}^{C \setminus A} & \overline{\rangle A \langle \Phi}^C &\stackrel{\text{def}}{=} \overline{\Phi}^{C \cup A} \\ \overline{\Phi \mathbf{U} \Psi}^C &\stackrel{\text{def}}{=} \overline{\Phi}^C \mathbf{U} \overline{\Psi}^C & \overline{\mathbf{X} \Phi}^C &\stackrel{\text{def}}{=} \mathbf{X} \overline{\Phi}^C \\ \overline{\Phi \wedge \Psi}^C &\stackrel{\text{def}}{=} \overline{\Phi}^C \wedge \overline{\Psi}^C & \overline{\neg \Phi}^C &\stackrel{\text{def}}{=} \neg \overline{\Phi}^C \\ \overline{P}^C &\stackrel{\text{def}}{=} P \end{aligned}$$

Now we have the following lemma:

Lemma 13. *For any strategy context F , any subset $C \subseteq \text{dom}(F)$ and any formula $\Phi \in ATL_{sc,\infty}^*$ and any path formula Φ_p , we have:*

$$\begin{aligned} \ell \models_{F \setminus C} \Phi &\Leftrightarrow \ell \models_F \overline{\Phi}^C \\ \rho \models_{F \setminus C} \Phi_p &\Leftrightarrow \rho \models_F \overline{\Phi}_p^C \end{aligned}$$

Proof. The proof is done by structural induction over the formula. In this proof we will use C' as an abbreviation for the coalition $C \setminus A$. Moreover F_B ranges over strategies for coalition B .

– $\Psi \stackrel{\text{def}}{=} \langle A, \bar{s} \rangle \varphi$.

We have the following equivalences : $\ell \models_F \langle A, \bar{s} \rangle [C'] \bar{\varphi}^{C'}$ means by definition

$$\exists F_A. \forall F_{C'}. \forall \rho \in \text{Out}(q, F_{C'} \circ F_A \circ F). \rho \models_{F_{C'} \circ F_A \circ F} \bar{\varphi}^{C'},$$

Then the induction hypothesis yields

$$\exists F_A. \forall F_{C'}. \forall \rho \in \text{Out}(q, F_{C'} \circ F_A \circ F). \rho \models_{(F_{C'} \circ F_A \circ F) \setminus C'} \varphi,$$

or equivalently, since $(F_{C'} \circ F_A \circ F) \setminus C' = F_A \circ (F \setminus C)$:

$$\exists F_A. \forall F_{C'}. \forall \rho \in \text{Out}(q, F_{C'} \circ F_A \circ F). \rho \models_{F_A \circ (F \setminus C)} \varphi,$$

or also

$$\exists F_A. \forall \rho \in \text{Out}(q, F_A \circ (F \setminus C)). \rho \models_{(F_A \circ (F \setminus C))} \varphi,$$

since we have

$$\bigcup_{F_{C'} \in \text{Strat}(C')} \text{Out}(q, F_{C'} \circ F_A \circ F) = \text{Out}(q, F_A \circ (F \setminus C)).$$

This leads to $\ell \models_{F \setminus C} \langle A, \bar{s} \rangle \varphi$, which is the desired result.

– $\Psi \stackrel{\text{def}}{=} \rangle A \langle \varphi$. On the one hand, by the semantics of $\text{ATL}_{sc, \infty}^*$, we have that:

$$q\ell \models_{F \setminus C} \rangle A \langle \varphi \quad \text{iff} \quad \ell \models_{F \setminus (C \cup A)} \varphi.$$

On the other hand, the induction hypothesis tells us that:

$$\ell \models_{F \setminus (C \cup A)} \varphi \quad \text{iff} \quad \ell \models_F \bar{\varphi}^{C \cup A}.$$

Gathering the two equivalences, we obtain the desired result.

– $\Psi \stackrel{\text{def}}{=} \varphi \mathbf{U} \psi$. The semantics of $\text{ATL}_{sc, \infty}^*$ tells us that $\rho \models_{F \setminus C} \Psi$ if and only if the following formal holds:

$$\exists i. \rho^i \models_{(F \setminus C)^{\rho, i}} \psi \quad \text{and} \quad \forall 0 \leq j < i. \rho^j \models_{(F \setminus C)^{\rho, j}} \varphi.$$

By using the induction hypothesis, the above formula is equivalent to the following one:

$$\exists i. \rho^i \models_{F^{\rho, i}} \bar{\psi}^C \quad \text{and} \quad \forall 0 \leq j < i. \rho^j \models_{F^{\rho, j}} \bar{\varphi}^C,$$

which means that $\rho \models_F \bar{\varphi}^C \mathbf{U} \bar{\psi}^C$. We thus obtain the desired result.

– The remaining cases are straightforward. □

We can now finish the proof by considering $\Psi = \bar{\Phi}^\emptyset$.

A.2 Comparison with other formalisms

Comparison with QD_μ [Pin07]. The logic QD_μ [Pin07] extends the decision μ -calculus with quantification over strategies. In that setting, strategies are thought of as a labelling of the infinite computation tree of the system with extra atomic propositions and this allows to have strategy contexts. There is neither explicit modality $\rangle A \langle$ nor bounded-memory strategy quantification in QD_μ . But as we have seen above, any ATL_{sc}^* formula can be translated into an equivalent $\rangle\langle$ -free formula, which can in turn be expressed in QD_μ [Pin08].

Note also that QD_μ is able to express that a strategy is memoryless (provided that it has access to the names of the locations of the CGS): this is achieved by saying that the strategy always labels the same set of successors in any two copies of the same location in the execution tree.

Comparison with IATL [ÁGJ07]. IATL extends ATL with a context in a similar way as $ATL_{sc,\infty}$ does, but there is no way of replacing a strategy by another one: once a player has committed to applying a strategy, she cannot change her mind. It is thus easy to see that $ATL_{sc,\infty}$ is (strictly) more expressive than IATL.

Game Logic. The formal definition of GL [AHK02] is as follows: Let ℓ be a state and F a strategy for some coalition A , remember that we write $\text{ExecTree}(\ell, F)$ for the subtree of the computation tree from ℓ whose infinite rooted paths are the elements of $\text{Out}(\ell, F)$.

Definition 14. *The syntax of GL is defined by the following grammar :*

$$\begin{aligned} GL \ni \varphi_s, \psi_s &::= P \mid \neg\varphi_s \mid \varphi_s \vee \psi_s \mid \exists A.\varphi_t \\ \varphi_t, \psi_t &::= \varphi_s \mid \neg\varphi_t \mid \varphi_t \vee \psi_t \mid \exists\varphi_p \\ \varphi_p, \psi_p &::= \varphi_t \mid \neg\varphi_p \mid \varphi_p \vee \psi_p \mid \mathbf{X}\varphi_p \mid \varphi_p \mathbf{U}\psi_p \end{aligned}$$

where P ranges over the set AP and A over the subsets of Agt .

Thus in GL, we distinguish state, tree and path formulae. Note also that path formulae are interpreted over a path ρ inside a tree t (denoted (t, ρ)). We give the semantics of the unusual operators :

$$\begin{aligned} \ell \models \exists A.\varphi_t &\quad \text{iff} \quad \exists F_A \in \text{Strat}(A). \text{ExecTree}(\ell, F_A) \models \varphi_t, \\ t \models \varphi_s &\quad \text{iff} \quad \ell \models \varphi_s, \text{ where } \ell \text{ is the root of } t, \\ t \models \exists\varphi_p &\quad \text{iff} \quad \exists \rho \in t. (t, \rho) \models \varphi_p, \\ (t, \rho) \models \varphi_t &\quad \text{iff} \quad t \models \varphi_t. \end{aligned}$$

Proposition 15. *GL can be expressed in ATL_{sc}^* , and $\langle A_1 \rangle \mathbf{X} (\langle A_2 \rangle \mathbf{X} b \wedge \langle A_3 \rangle \mathbf{X} a)$ has no equivalent in GL.*

A.3 PSPACE-hardness of $\text{ATL}_{sc,0}$ model checking

Lemma 16. *The model checking problem for $\text{ATL}_{sc,0}$ over CGSs is PSPACE-hard.*

Proof. We reduce it to the PSPACE-complete problem QBFSAT:

QBFSAT:

Input: A family of variables $X = \{x^1, \dots, x^n\}$, a boolean CNF formula $\varphi = \bigwedge_{j=1, \dots, J} C_j$ on the set of variables X .

Output: The value of $\exists x_1 \forall x_2 \exists x_3 \dots Q_n x_n \varphi(x_1, \dots, x_n)$, where Q_n is a \forall if n is even, and a \exists otherwise

From an instance \mathcal{I} of this problem, we build the turn-based CGS \mathcal{C} on Fig. 4. In \mathcal{C} we have n different agents, one for each variable, who play only once and so give a value to their variable. In fact the players have successively the choice between two moves, for true or false. We then label the states in order to be able to express that every disjunctive clause C_j will eventually be set to true:

- $\forall 1 \leq i \leq n, \text{Lab}(x_i) = \{C_j \mid x_i \text{ appears in } C_j\}$,
- $\forall 1 \leq i \leq n, \text{Lab}(\neg x_i) = \{C_j \mid \neg x_i \text{ appears in } C_j\}$.

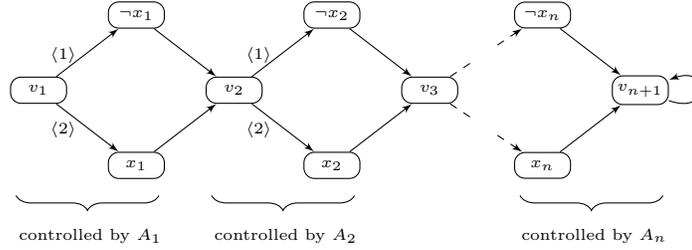


Fig. 4. \mathcal{C} without the labels

Since each player plays once, a strategy for A_i corresponds to a truth value for the variable x_i . Thus, the formula $\langle A_i^0 \rangle \psi_s$ (resp. $\neg \langle A_i^0 \rangle \neg \psi_s$) means that there exists a truth value for x_i such that (resp. no matter what the truth value for x_i is) ψ_s stands. Note that when a player makes his move, every opponent knows the move, because it is written in the context. Therefore, those who play after him can take it into consideration when comes their turn.

Now consider the following $\text{ATL}_{sc,0}^*$ formula :

$$\Psi_{\mathcal{I}} \stackrel{\text{def}}{=} \langle A_1^0 \rangle (\neg \langle A_2^0 \rangle \neg (\langle A_3^0 \rangle (\neg \dots \epsilon_n \bigwedge_{j=1, \dots, J} \langle \emptyset^0 \rangle \mathbf{F} C_j \dots)))$$

where ϵ_n is \neg if n is even, and \emptyset otherwise. Clearly $\Psi_{\mathcal{I}}$ holds for \mathcal{C} iff \mathcal{I} is a positive instance of QBFSAT. Finally it is easy to obtain an equivalent model

checking problem for $\text{ATL}_{sc,0}$. Let $\Psi'_{\mathcal{I}}$ be the $\text{ATL}_{sc,0}$ formula defined from $\Psi_{\mathcal{I}}$ by inserting \mathbf{X} after every $\langle A_i^0 \rangle$ modality. And let \mathcal{C}' be the CGS defined from \mathcal{C} by adding n states q_1, \dots, q_n and $n - 1$ transitions $q_i \rightarrow q_{i+1}$ for $i < n$, and the transition $q_n \rightarrow v_1$. Clearly $q_1 \models \Psi'_{\mathcal{I}}$ iff \mathcal{I} is a positive instance. \square

It should be noticed that $\rangle \cdot \langle$ has not been used in the hardness proof.