

ATL with strategy contexts and bounded memory

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

Research report LSV-08-14

Laboratoire Spécification et Vérification

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, providing 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 [15, 7, 16]. 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. Over the last ten years, a new flavor of temporal logics has been developed: *alternating-time temporal logics* (ATL) [2]. ATL is a fundamental logic for specifying and verifying properties in *multi-agent systems* (modeled as *Concurrent Game Systems* (CGS) [2]), in which several agents can concurrently act upon the behaviour of the system. On 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^{*} to express richer properties: while ATL strategy quantifier drops strategies introduced by earlier quantifiers in the evaluation of the formula, our logics keep executing those strategies. For example, our new modality, which we write $\langle A \rangle$, allows us to express the property “ A has a strategy such that (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 ”. This would be written as follows: $\langle A \rangle \mathbf{G} (\langle B \rangle \Phi \wedge \langle C \rangle \Psi)$. Note that naive attempts to express this property in ATL fail: for instance, 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, while in $\langle\langle A \rangle\rangle \mathbf{G} (\langle\langle A, B \rangle\rangle \Phi \wedge \langle\langle A, C \rangle\rangle \Psi)$, the coalition A is allowed to use different strategies when playing with B and C . In order to achieve this idea, we naturally adapt the semantics of ATL^{*} in order to interpret a formula within a strategy context.

Secondly we parameterise 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: we mainly study the expressiveness of the new formalisms we propose, and compare them with classical formalisms such as ATL, ATL^{*}, Alternating-time- μ -calculus (AMC) [2], the Game Logic (GL) [2], and also with recent formalisms such as Strategy Logic (SL) [6] or QD_μ [13], ... While our logic does not contain any of those two, it can express most of their interesting properties, especially regarding equilibria. We also study the model-checking problem for our logics: while we have a non-elementary algorithm for the general case, 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. Very recently, several works have focused on the same kind of extension of ATL, and come up with different solutions that we list below. Generally speaking, this leads to very expressive logics, able to express equilibrium properties. The counterpart is that those logics have very high complexity.

- IATL [1] extends ATL with strategy contexts, with a similar definition as ours, but it forces 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 [6] 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 [13] 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 [4] 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.

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.

2.1 Concurrent Game Structures

Concurrent game structures are a multi-player extension of classical Kripke structures [2]. Their definition is as follows:

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

- Loc is a finite set of locations, $\ell_0 \in Loc$ is the initial location;
- AP is a finite, non-empty set of atomic propositions;
- $Lab: Loc \rightarrow 2^{AP}$ is a labelling function;
- $Agt = \{A_1, \dots, A_k\}$ 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.
- $Edg: Loc \times \mathcal{M}^k \rightarrow Loc$, where $k = |Agt|$, is a function defining the transition table. With each location and each set of moves of the agents, it associates the resulting location.

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 [2]: in a location ℓ , each player A_i chooses one possible move m_{A_i} in $Mov(\ell, A_i)$ and the next location is given by $Edg(\ell, m_{A_1}, \dots, m_{A_k})$. We write $Next(\ell)$ for the set of all possible successor locations from ℓ , and $Next(\ell, A_j, m)$, with $m \in Mov(\ell, A_j)$, for the restriction of $Next(\ell)$ to locations reachable from ℓ when player A_j makes the move m .

¹ Our results would still hold if we consider symbolic CGSs, where the transition table is encoded through boolean formulas [11].

2.2 Coalition, bounded-memory strategy, outcomes.

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 possible moves for coalition A from ℓ . Such a move m is composed of a single move for every agent of the coalition, that is $m \stackrel{\text{def}}{=} (m_a)_{a \in A}$.
- 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 $\rho[i]$ for 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 (since $\text{Strat}(\emptyset) = \{\emptyset\}$).

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 $F'_{|A_j}(\ell_0 \dots \ell_m)$ otherwise.

Finally, given a strategy F , an execution ρ and some integer $i \geq 0$, we define the restriction $F^{\rho, i}$ of F to the behavior from the future state $\rho[i]$: $F^{\rho, i}(\pi) = F(\rho[0 \dots i] \cdot \pi)$. Note that if F is memoryless, then $F^{\rho, i} = F$.

Bounded-memory strategies. Between the general strategies (without any bound over its resources) and the simple memoryless strategies, we can consider *s-bounded memory strategies*. Let s be a (binary-encoded) integer representing the size of the memory. The most basic notion of memory-based strategies is based on the s -th suffix of the history: in that case, a strategy σ is a function from $Q^{[0, s]}$ to \mathcal{M} . The main drawback of this approach is that the players cannot choose

which part of the history is recorded, they can only remember the s last events (see Appendix A.1 for an example of this very limited kind of memory-bounded strategy).

More interestingly, we can define a bounded memory strategy as a memoryless strategy over the locations of the CGS **and** a set of memory cells [12]: 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 \mathbf{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 $\mathbf{Cell} \times \mathbf{Loc}$ 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 $\mathbf{Cell} \times \mathbf{Loc}$ to \mathbf{Cell} that updates the memory cell, and c is the current memory cell of this strategy. In the following, we also denote $F^{\text{mov}}(c, \ell)$ by $F_A(\ell)$ in order to simplify the notation: $F_A(\ell)$ is the move for A induced by the strategy F_A from the location ℓ and given then current memory cell c .

We can adapt the previous notions to this kind of strategy $F_A = (F^{\text{mov}}, F^{\text{cell}}, c)$. 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 $A \subseteq \mathbf{Agt}$ 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 $\mathbf{Strat}^{\bar{s}}(A)$, and we omit to mention the memory bound when none is imposed.

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

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

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

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

with $P \in \mathbf{AP}$, $A \subseteq \mathbf{Agt}$ and $\bar{s} \in (\mathbb{N} \cup \{\infty\})^A$. For convenience, we write e.g. $\langle A^s, B^t \rangle \varphi$ instead of $\langle \{A, B\}, \{A \mapsto s, B \mapsto t\} \rangle \varphi$.

Given a formula $\varphi \in \mathbf{ATL}_{sc, \infty}^*$, the size of φ , denoted by $|\varphi|$, is the size of the tree representing that formula, assuming binary-encoding of integer constants.

We use standard abbreviations such as $\top \stackrel{\text{def}}{=} P \vee \neg P$, $\perp \stackrel{\text{def}}{=} \neg \top$, $\mathbf{F} \varphi \stackrel{\text{def}}{=} \top \mathbf{U} \varphi$, etc, and omit to mention the memory bound when no memory constraint is imposed.

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 $\ell_0 \models_{\emptyset} \Phi$.

$\text{ATL}_{sc,\infty}^*$ contains several extensions compared to ATL^* : the combination (or nesting) of strategies in a context and the way of handling separately \bar{s} -memory strategies.

Two $\text{ATL}_{sc,\infty}^*$ formulae Φ and Φ' are equivalent, denoted $\Phi \equiv \Phi'$, iff their truth value is the same for any location in any CGS under any strategy context F : $\ell \models_F \Phi \Leftrightarrow \ell \models_F \Phi'$ for any \mathcal{C} , ℓ , and F .

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 context: $\langle\langle A \rangle\rangle \varphi$ 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 $\rangle \text{Agt} \langle \langle A \rangle \varphi$.

Clearly 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 [2, 17]).

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 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 strategies with possibly infinite memory).

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

- $\text{ATL}_{sc,\infty}$ 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 “Release” modality–, and $\mathbf{X} \varphi_s$). It follows from the above discussion that $\text{ATL}_{sc,\infty}$ contains ATL and CTL. We also define the fragments $\text{ATL}_{sc,b}$ and ATL_{sc} as above.

3 Expressiveness

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

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

First we introduce a new strategy quantifier in order to specify that “for any strategy of coalition A , every run in the corresponding outcome satisfies a formula φ ”:

$$[A] \varphi \stackrel{\text{def}}{=} \neg \langle A \rangle \neg \langle \emptyset \rangle \varphi.$$

It translates as follows:

$$\ell \models_F [A] \varphi \quad \text{iff} \quad \forall F_A \in \text{Strat}(A). \forall \rho \in \text{Out}(\ell, F_A \circ F). \rho \models_{F_A \circ F} \varphi.$$

Note that this modality is not the dual of $\langle A \rangle$ because we want an universal quantification over the runs in the outcomes. Thus it is different from the $\llbracket A \rrbracket$ in [2] which specifies the existence of *counter*-strategies for $\text{Agt} \setminus A$ to ensure φ .

Examples of properties using strategy contexts. The new modalities $\langle A \rangle$ allow us to express many interesting properties over the strategies of different players in a game. In particular, our logics can express the different examples that motivated the introduction of SL, QD_μ or IATL:

- We can express the fact that “Player 1 can ensure Φ_1 while Player 2 plays to ensure Φ_2 ” with the following formula:

$$\Psi = \langle A_1 \rangle [A_2] \left(\langle \emptyset \rangle A_1 \langle \emptyset \rangle \Phi_2 \Rightarrow \Phi_1 \right). \quad (1)$$

When we interpret Ψ with the semantics of $\text{ATL}_{sc,\infty}^*$ we get:

$$\ell \models \Psi \quad \text{iff} \quad \exists F_1. \forall F_2. (\forall \rho_2 \in \text{Out}(\ell, \{F_2\}). \rho_2 \models_{F_2} \Phi_2) \Rightarrow (\forall \rho \in \text{Out}(\ell, \{F_1, F_2\}). \rho \models_{\{F_1, F_2\}} \Phi_1) \quad (2)$$

The use of modality $\rangle A_1 \langle$ allows us to specify that Φ_1 has to be ensured by Player 1 only when Player 2 follows a *true* strategy to ensure (whatever Player 1 does) Φ_2 .

- the *winning secure equilibrium* [5] is a special form of Nash equilibrium where both players can cooperate to satisfy an objective $\Phi_1 \wedge \Phi_2$, and if Player 1 (resp. 2) abandons the cooperation to ensure $\neg\Phi_2$ (resp. $\neg\Phi_1$) then Player 2 (resp. 1) can enforce $\neg\Phi_1$ (resp. $\neg\Phi_2$). The existence of such equilibrium can be stated as follows (Φ_1 and Φ_2 are assumed to be LTL formulas):

$$\langle A_1, A_2 \rangle \left(\Phi_1 \wedge \Phi_2 \wedge [A_1] (\neg\Phi_2 \Rightarrow \neg\Phi_1) \wedge [A_2] (\neg\Phi_1 \Rightarrow \neg\Phi_2) \right).$$

Indeed this formula states that Players 1 and 2 can cooperate together to ensure the common objective, but there is no way for one player to enforce its own objective and the negation of the objective of the other player.

- Given two players A_1 and A_2 having their own objectives Φ_1 and Φ_2 , the *Nash equilibrium* holds for two strategies F_1 and F_2 for players 1 and 2 respectively, 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 for Player 2. This property characterises pairs of strategies. Given two strategies F_1 and F_2 in the strategy context F , the following formula holds for states ℓ within F iff F_1 and F_2 correspond to a Nash equilibrium in ℓ :

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

- The notion of *dominating* strategy can also be expressed: if F is a strategy for A_1 , then a state ℓ satisfies the following formula within the strategy context F iff F is a dominating strategy from s w.r.t. the objective Φ : $(\langle \text{Agt} \setminus A_1 \rangle \neg\Phi) \Rightarrow \neg\langle A_1 \rangle \Phi$. Indeed if the opponent can ensure $\neg\Phi$ when A_1 plays F , then it was not possible for A_1 to win.

Bounding the memory of the opponent. Our definition of bounded-memory strategies does not restrict the possible co-strategies for the opponents. However, it might sometimes be interesting to look for strategies that are winning only against *deterministic* bounded-memory opponents. This can be expressed in our logic by dualizing the bounded-memory quantifier: formula $\langle A \rangle [Agt \setminus A, \bar{s}] \Phi$ states that coalition A has a strategy enforcing Φ if the opponent coalition has \bar{s} -bounded memory.

Expressiveness of $\rangle A \langle$ quantifier. In the previous section, we illustrate the use of modality $\rangle A \langle$ by expressing the classical ATL^* modality $\langle\langle A \rangle\rangle$ with $\rangle Agt \langle \langle A \rangle$: 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 3. *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 $\Phi \in \text{ATL}_{sc,\infty}^*$, we define formula $\overline{\Phi}^C$ recursively as follows:

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

Now, for any strategy context F and any $C \subseteq \text{dom}(F)$, and for any state q and any execution ρ , we have the following equivalences:

$$q \models_{F \setminus C} \Phi \Leftrightarrow q \models_F \overline{\Phi}^C \qquad \rho \models_{F \setminus C} \Phi_p \Leftrightarrow \rho \models_F \overline{\Phi}_p^C$$

The proof is done by structural induction over the formula (see Appendix A.2). Letting $\Psi = \overline{\Phi}^\emptyset$ yields the result. \square

3.2 Comparison with other formalisms

Formal comparison of expressiveness. Two comparison criteria can be used to compare the expressive power of logics:

- The *distinguishing power* is the ability of a logic to distinguish models. Two models C_1 and C_2 are distinguished by L iff there exists a formula $\Phi \in L$ s.t. $C_1 \models \Phi$ and $C_2 \not\models \Phi$. We write $L_1 \leq_{dp} L_2$ when L_2 can distinguish any two models that can be distinguished by L_1 .
- The *expressiveness* is the ability of a logic to express properties. We say that L_2 is more expressive than L_1 (written $L_1 \leq_{ex} L_2$) when for every L_1 formula, there exists an equivalent formula in L_2 .

We define $<_{dp}$, \equiv_{dp} , $>_{ex}$ and \equiv_{ex} in a standard way. Note that logics L_1 and L_2 can have the same distinguishing power (*i.e.*, $L_1 \leq_{dp} L_2$ and $L_2 \leq_{dp} L_1$) and differ on the expressiveness criterion. This happens *e.g.* for CTL and CTL* .

Comparison with ATL* . From the distinguishing power point of view, we have the following two results that emphasize the ability of our new modality $\langle A \rangle$:

Proposition 4. *For any $b \in \mathbb{N}$, we have $\{\text{ATL}_{sc,b}, \text{ATL}_{sc}\} >_{dp} \text{ATL}^*$.*

Proof. First note that $\text{ATL}_{sc,0}$ and ATL_{sc} both contain ATL (remember that memoryless quantification is sufficient for ATL) and that ATL and ATL* have the same distinguishing power [3]. This entails $\{\text{ATL}_{sc,0}, \text{ATL}_{sc}\} \geq_{dp} \text{ATL}^*$.

The proof that $\text{ATL}_{sc,0} >_{dp} \text{ATL}^*$ uses the fact formula $\langle \text{Agt}^0 \rangle (\mathbf{X} \langle \emptyset^0 \rangle \mathbf{X} P)$ has no equivalent in ATL* (see Lemma 13 in Appendix A.3). This entails that $\text{ATL}_{sc,b} >_{dp} \text{ATL}^*$ for any b . Finally, we have $\text{ATL}_{sc} >_{dp} \text{ATL}^*$ because $\langle A_1 \rangle \mathbf{X} (\langle A_2 \rangle \mathbf{X} P \wedge \langle A_2 \rangle \mathbf{X} P')$ cannot be expressed in ATL* (see Lemma 14 in Appendix A.3). \square

As a direct corollary of the previous results and of the inclusion of ATL (resp. ATL^*) in $\text{ATL}_{sc,\infty}$ (resp. $\text{ATL}_{sc,\infty}^*$), we obtain the following theorem:

Theorem 5. *For any $b \in \mathbb{N}$, we have:*

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

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

Proposition 6. $\text{CTL}^* \leq_{ex} \text{ATL}_{sc}$

The proof is given in Appendix A.3. Here we just illustrate the underlying translation with two examples:

- $\text{EG}(P \vee \mathbf{X}P)$ is equivalent to $\langle \text{Agt} \rangle \mathbf{G}(P \vee \langle \emptyset \rangle \mathbf{X}P)$;
- $\text{E}(\mathbf{F}P \Rightarrow \mathbf{F}P')$ is equivalent to $\langle \text{Agt} \rangle \mathbf{F}(\langle \emptyset \rangle \mathbf{G}\neg P) \vee \langle \text{Agt} \rangle \mathbf{G}(\langle \emptyset \rangle \mathbf{F}P')$.

Comparison with Game Logic. We now compare our new logics with *Game Logic*, proving that it is strictly less expressive than ATL_{sc}^* . *Game Logic* was introduced in [2] in order to express the module checking problem [10]. This logic is an extension of ATL^* that allows us to deal explicitly with the execution tree induced by a strategy: given such a tree t , it is possible to quantify over the executions inside t and specify temporal properties. For example, the formula $\exists A.((\exists P \mathbf{U} P') \wedge (\forall \mathbf{F} P''))$ specifies that there exists a strategy F_A for A such that we have for the tree induced by F_A : (1) there exists a run verifying $P \mathbf{U} P'$ and (2) every run verifies $\mathbf{F} P''$. The definition of *GL* is given in Appendix A.5. We have the following result:

Theorem 7. $\text{ATL}_{sc} >_{ex} \text{GL}$

Proof (sketch). The proof is based on a translation from *GL* into ATL_{sc} (see Prop. 16 in appendix) and the fact that 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 *GL* (see Prop. 17 in appendix). \square

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

Proposition 8. $\text{ATL}_{sc,\infty} \not\leq_{ex} \text{AMC}$ and $\text{AMC} \not\leq_{ex} \text{ATL}_{sc,\infty}^*$.

Proof. The CGS \mathcal{C}_2 (described in Figure 3 in Appendix A.3) is a simple unfolding of \mathcal{C}_1 , thus they satisfy the same AMC formulae but they can be distinguished by $\text{ATL}_{sc,\infty}$, this gives the first result. When considering one-player CGS (*i.e.*, Kripke structures), our $\text{ATL}_{sc,\infty}^*$ is clearly equivalent to CTL^* , which is strictly less expressive than the (classical) μ -calculus. \square

Comparison with Strategy Logic [6]. Strategy Logic (SL for short) has recently been defined in [6] 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 given to φ indicate on which path 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:

- SL can explicitly manipulate strategies as first-order elements. It can for instance state properties such as

$$\exists x_1. \exists y_1. \exists x_2. \exists y_2. [\varphi_1(x_1, y_1) \wedge \varphi_2(x_2, y_1) \wedge \varphi_3(x_1, y_2) \wedge \varphi_4(x_2, y_2)]$$

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 can be translated into $\text{ATL}_{sc,\infty}^*$.

Comparison with qD_μ [13]. The logic qD_μ [13] 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 [14].

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 [1]. IATL extends ATL with a context in a similar way as $\text{ATL}_{sc,\infty}$ 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 $\text{ATL}_{sc,\infty}$ is (strictly) more expressive than IATL.

Figure 1 summarizes the expressiveness results of our logics.




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

4 $ATL_{sc,\infty}$ and $ATL_{sc,\infty}^*$ 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 has very high complexity. We thus focus on simpler cases (namely, memoryless and bounded-memory strategies), where more efficient algorithms can be obtained.

Theorem 9. *Model checking $ATL_{sc,\infty}^*$ formulas over CGS is decidable.*

Proof. The translation from ATL_{sc}^* to QD_μ 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 $ATL_{sc,\infty}^*$ model checking. \square

4.1 Model-checking $ATL_{sc,0}^*$ and $ATL_{sc,0}$

Theorem 10. *The model checking problems for $ATL_{sc,0}^*$ and $ATL_{sc,0}$ over CGSs are PSPACE-complete.*

Proof. We only address the membership in PSPACE. The hardness proof is given in Lemma 18 in Appendix A.6 (and is similar to that of [4]).

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 $ATL_{sc,0}^*$ 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 $ATL_{sc,0}^*$ 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-ATL $^*_{sc,0}(\mathcal{C}, F, \ell_0, \varphi)$ – ATL $^*_{sc,0}$ model checking

Require: a CGS \mathcal{C} , $F \in \text{Strat}^0(A)$, $l_0 \in \text{Loc}$ and an 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-ATL $^*_{sc,0}(\mathcal{C}, F_B \circ F, \ell, \psi')$ , then label  $l$  with  $a_\psi$ 
  case  $\psi = \rangle B \langle \psi'$  :
    for  $l \in \text{Loc}$  do
      if MC-ATL $^*_{sc,0}(\mathcal{C}, F \setminus_B, l, \psi')$ , then label  $l$  with  $a_\psi$ 
return MC LTL  $(\mathcal{C}', l_0, \mathbf{A}\sigma(\varphi))$ 

```

Note that PSPACE-completeness extends straightforwardly 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 [9].

Remark 1. Both the algorithm and the PSPACE-hardness can be adapted for IATL. This corrects the Δ_2^P -completeness result of [1].

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: $\mathbb{M}_{A_i}^s = (\text{Agt}, \text{Loc}_s^i, 0, \emptyset, \emptyset, \mathcal{M}_s^i \cup \{\perp\}, \text{Mov}_s^i, \text{Edg}_s^i)$ 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 $\text{RStrat}_{\mathcal{C}_{\bar{s}}}^0(A_j) \subset \text{Strat}_{\mathcal{C}_{\bar{s}}}^0(A_j)$.

Adapting the proof of Theorem 10 to memory-bounded strategies, we get:

Proposition 11. *Let $C = (\text{Agt}, \text{Loc}, \ell_0, \text{AP}, \text{Lab}, \mathcal{M}, \text{Mov}, \text{Edg})$ 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 10 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 formula 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 [17].
- 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 $\text{ATL}_{sc,\infty}$ and $\text{ATL}_{sc,\infty}^*$, with the aim of possibly finding reasonably efficient algorithms for our expressive extensions of ATL and ATL^* . Finally we think that the ability to deal explicitly with bounded-memory strategies is an interesting approach to develop.

References

1. T. Ågotnes, V. Goranko, and W. 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.

2. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
3. R. Alur, T. A. Henzinger, O. Kupferman, and M. Y. Vardi. Alternating refinement relations. In *Proc. 9th Intl Conf. Concurrency Theory (CONCUR'98)*, volume 1466 of *LNCS*, pages 163–178. Springer, 1998.
4. Ch. Baier, T. Brázdil, M. Größer, and A. 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.
5. K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Games with secure equilibria. *Theoretical Computer Science*, 365(1-2):67–82, 2006.
6. K. Chatterjee, T. A. Henzinger, and N. 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.
7. E. M. Clarke and E. A. 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.
8. E. A. Emerson. Temporal and modal logic. In *Handbook of Theoretical Computer Science*, volume B, chapter 16, pages 995–1072. Elsevier, 1990.
9. W. Jamroga and J. 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.
10. O. Kupferman, M. Y. Vardi, and P. Wolper. Module checking. *Information and Computation*, 164(2):322–344, 2001.
11. F. Laroussinie, N. Markey, and G. 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.
12. R. Mazala. Infinite games. In *Automata, Logics, and Infinite Games*, volume 2500 of *LNCS*, pages 23–42. Springer-Verlag, 2002.
13. S. 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.
14. S. Pinchinat. Personal communication, 2008.
15. A. 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.
16. J.-P. Queille and J. 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.
17. P.-Y. 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.

A Appendix

A.1 Discussion about the bounded-memory strategy

Here we illustrate the drawback of the basic notion of bounded-memory strategy where we can only consider the s -th suffix of the history.

Consider the game depicted on Fig. 2. It can be proved that there exists no history-based bounded-memory strategy for Player 1 from state s_0 for the LTL objective $(\mathbf{X} a \wedge \mathbf{G} b) \vee (\mathbf{X} \neg a \wedge \mathbf{G} c)$ (stating that only b -states should be visited if a has been visited at the beginning of the execution, and only c -states should be visited otherwise) because the crucial information (visiting s_1) occurs at the very beginning of the execution.




Fig. 2. A game with no history-based bounded-memory strategy

But with the notion of bounded-memory strategy we have adopted, we can easily see that Player 1 has a 1-bounded memory strategy to ensure that the formula $(\mathbf{X} a \wedge \mathbf{G} b) \vee (\mathbf{X} \neg a \wedge \mathbf{G} c)$ holds. Indeed, we can define the strategy $F = (F^{\text{mov}}, F^{\text{cell}}, 0)$ with:

$$\begin{aligned} F^{\text{cell}}(0, s_0) &= F^{\text{cell}}(0, s_2) = F^{\text{cell}}(0, s_3) = F^{\text{cell}}(0, s_4) = F^{\text{cell}}(0, s_5) = 0 \\ F^{\text{cell}}(0, s_1) &= F^{\text{cell}}(1, s_3) = F^{\text{cell}}(1, s_5) = F^{\text{cell}}(1, s_6) = 1 \\ F^{\text{mov}}(0, s_3) &= 2 \quad \text{and} \quad F^{\text{mov}}(1, s_3) = 1 \end{aligned}$$

Memory cell 0 encodes the fact that s_1 has not (or “not yet”) been visited. This strategy is easily seen to be winning.

A.2 Elimination of $\cdot A \langle$

Here we give the proof of the following lemma used in Proposition 3:

Lemma 12. *For any strategy context F , any subset $C \subseteq \text{dom}(F)$ and any formula $\Phi \in \text{ATL}_{sc, \infty}^*$ and any path formula Φ_p , we have:*

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

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 .

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

We have the following equivalences : $q \models_F \langle A \rangle [C'] \bar{\varphi}^{C'}$ 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 else, since we have

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

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

leading to $q \models_{F \setminus C} \langle A \rangle \varphi$, which is the desired result.

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

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

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

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

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:

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

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

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

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.

□

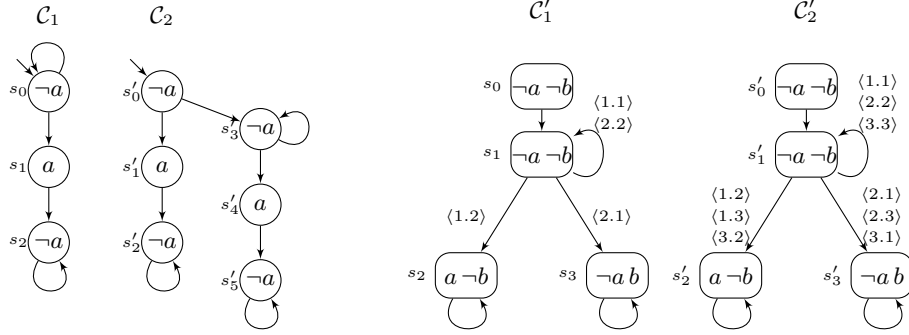


Fig. 3. \mathcal{C}_1 and \mathcal{C}_2 (resp \mathcal{C}'_1 and \mathcal{C}'_2) cannot be distinguished by ATL^*

A.3 Distinguishing power

Lemma 13. $\langle Agt^0 \rangle (X \langle \emptyset^0 \rangle X P)$ 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:

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

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. \square

Note that we could also consider the logic ATL^*_0 , that is the logic ATL^* where every strategy quantifier deals with memoryless strategies. This logic could distinguish the two structures of the previous proof with the formula $\langle\langle Agt \rangle\rangle_0 X X P$.

Thus $ATL_{sc,\infty}$ allows us to distinguish models that cannot be distinguished by ATL^* , and note that it is not necessary to use the modality $\cdot - \langle$ to obtain this result. As a consequence we have:

Lemma 14. $\langle A_1 \rangle X (\langle A_2 \rangle X P \wedge \langle A_2 \rangle X P')$ has no equivalent in ATL^* .

Proof. 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 $\langle A_1 \rangle X (\langle A_2 \rangle X P \wedge \langle A_2 \rangle X P')$ thanks to the third move of Player 1. But this is not the case for s_0 that does not satisfy the $ATL_{sc,\infty}$ formula. \square

A.4 Translation from CTL* to ATL_{sc}

Here we give the proof of Proposition 6:

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

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

We clearly have $q \models \Phi$ iff $q \models_{\emptyset} \overline{\Phi}$. The existential path quantifier is replaced by $\langle \text{Agt} \rangle$. Now given a path formula $\mathbf{X}\varphi$, its corresponding $\overline{\mathbf{X}\varphi}$ 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 $\overline{\Psi}$ 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 by using the method for translating CTL⁺ into CTL [8] (remember that $\langle \text{Agt} \rangle$ is equivalent to E). \square

A.5 Comparison with other formalisms

Game Logic. Game Logic was introduced in [2], it extends of ATL* and allows us to deal explicitly with the execution tree induced by a strategy. Let ℓ be a state and F a strategy for some coalition A , 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 15. *The syntax of GL is defined by the following grammar :*

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

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 16. *Every GL formula can be translated into an equivalent ATL_{sc}^* formula.*

Proof. The translation is given by the inductively-defined application σ s.t.:

$$\overline{\exists A.\varphi} \stackrel{\text{def}}{=} \langle\langle A \rangle\rangle \bar{\varphi}, \quad \overline{\exists \varphi} \stackrel{\text{def}}{=} \neg \langle \emptyset \rangle \neg \bar{\varphi}, \quad \overline{P} \stackrel{\text{def}}{=} P.$$

The other inductive rules are defined the natural way.

Note that if φ is a 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 GL path and tree formulae. We must show that for any 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 two following cases :

- $\varphi = \exists A\psi$. Then $\rho \models_F \bar{\varphi}$ means that there exists a strategy F' for the coalition A , such that any computation ρ' in $\text{Out}(\rho[0], F')$ will satisfy $\bar{\psi}$. With the induction hypothesis, this is equivalent to $\exists F' \in \text{Strat}(A). \forall \rho' \in \text{Out}(\rho[0], F'). (\text{ExecTree}(\rho[0], F'), \rho') \models \psi$, and so to $\rho[0] \models \exists A.\psi$ because ψ is a tree formula. Now $\exists A\psi$ is a state formula which can be interpreted over any execution tree with root $\rho[0]$, most accordingly 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 then this is equivalent to $\exists \rho' \in \text{Out}(\rho[0], F). \rho' \models_F \bar{\psi}$. Again, we obtain by applying the induction hypothesis that there exists a path in $\text{ExecTree}(\rho[0], F)$ that satisfies ψ , and then $\text{ExecTree}(\rho[0], F) \models \exists \psi$ and this is equivalent to $(\text{ExecTree}(\rho[0], F), \rho) \models \varphi$ (as φ is a tree formula).

As a result, if we restrict this equivalence to the cases 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 . \square

Proposition 17. $\langle A_1 \rangle \mathbf{X} (\langle A_2 \rangle \mathbf{X} b \wedge \langle A_3 \rangle \mathbf{X} a)$ has no equivalent in GL.

Proof. Consider the CGSs \mathcal{S}_1 and \mathcal{S}_2 in Figure 4. They satisfy the same GL formulae, since the third move for Player 1 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

A.6 PSPACE-hardness of $ATL_{sc,0}$ model checking

Lemma 18. *The model checking problem for $ATL_{sc,0}$ over CGSs is PSPACE-hard.*




Fig. 4. S_1 and S_2 cannot be distinguished by GL

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. 5. 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. 5. \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 $\cdot \langle$ has not been used in the hardness proof.