

On the Expressiveness and Complexity of ATL[†]

François Laroussinie, Nicolas Markey, and Ghassan Oreiby*

LSV, CNRS & ENS Cachan, France
{f1,markey,oreiby}@lsv.ens-cachan.fr

Abstract. ATL is a temporal logic geared towards the specification and verification of properties in multi-agents systems. It allows to reason on the existence of strategies for coalitions of agents in order to enforce a given property. We prove that the standard definition of ATL (built on modalities “Next”, “Always” and “Until”) has to be completed in order to express the duals of its modalities: it is necessary to add the modality “Release”. We then precisely characterize the complexity of ATL model-checking when the number of agents is not fixed. We prove that it is Δ_2^P - and Δ_3^P -complete, depending on the underlying multi-agent model (ATS and CGS resp.). We also prove that ATL⁺ model-checking is Δ_3^P -complete over both models, even with a fixed number of agents.

1 Introduction

Model checking. Temporal logics were proposed for the specification of reactive systems almost thirty years ago [16]. They have been widely studied and successfully used in many situations, especially for model checking —the automatic verification that a finite-state model of a system satisfies a temporal logic specification. Two flavors of temporal logics have mainly been studied: *linear-time temporal logics*, e.g. LTL [16], which expresses properties on the possible *executions* of the model; and *branching-time temporal logics*, such as CTL [7,17], which can express requirements on *states* (which may have several possible futures) of the model.

Alternating-time temporal logic. Over the last ten years, a new flavor of temporal logics has been defined: *alternating-time temporal logics*, e.g. ATL [2,3]. ATL is a fundamental logic for verifying properties in *synchronous multi-agent systems*, in which several agents can concurrently influence the behavior of the system. This is particularly interesting for modeling control problems. In that setting, it is not only interesting to know if something *can arrive* or *will arrive*, as can be expressed in CTL or LTL, but rather if some agent(s) can *control* the evolution of the system in order to enforce a given property.

The logic ATL can precisely express this kind of properties, and can for instance state that “there is a strategy for a coalition A of agents in order to

[†] Work partially supported by ACI “Sécurité & Informatique” CORTOS, a program of the French Ministry of Research.

* This author is supported by a PhD grant from Région Ile-de-France.

eventually reach an accepting state, whatever the other agents do”. ATL is an extension of CTL, its formulae are built on atomic propositions and boolean combinators, and (following the seminal papers [2,3]) on modalities $\langle\langle A \rangle\rangle \mathbf{X} \varphi$ (coalition A has a strategy to immediately enter a state satisfying φ), $\langle\langle A \rangle\rangle \mathbf{G} \varphi$ (coalition A can force the system to always satisfy φ) and $\langle\langle A \rangle\rangle \varphi \mathbf{U} \psi$ (coalition A has a strategy to enforce $\varphi \mathbf{U} \psi$).

Multi-agent models. While linear- and branching-time temporal logics are interpreted on Kripke structure, alternating-time temporal logics are interpreted on models that incorporate the notion of *multiple agents*. Two kinds of synchronous multi-agent models have been proposed for ATL in the literature. First *Alternating Transition Systems* (ATSS)[2] have been defined: in any location of an ATSS, each agent chooses one *move*, *i.e.*, a subset of locations (the list of possible moves is defined explicitly in the model) in which he would like the execution to go to. When all the agents have made their choice, the intersection of their choices is required to contain one single location, in which the execution enters. In the second family of models, called *Concurrent Game Structures* (CGSs) [3], each of the n agents has a finite number of possible moves (numbered with integers), and, in each location, an n -ary transition function indicates the state to which the execution goes.

Our contributions. While in LTL and CTL, the dual of “Until” modality can be expressed as a disjunction of “Always” and “Until”, we prove that it is not the case in ATL. In other words, ATL, as defined in [2,3], is not as expressive as one could expect (while it is known that adding the dual of “Until” does not increase the complexity of the verification problems [5,9]).

We also precisely characterize the complexity of the model checking problem. The original works about ATL provide model-checking algorithms in time $O(m \cdot l)$, where m is the number of transitions in the model, and l is the size of the formula [2,3], thus in PTIME. However, contrary to Kripke structures, the number of transitions in a CGS or in an ATSS is not quadratic in the number of states [3], and might even be exponential in the number of agents. PTIME-completeness thus only holds for ATSS when the number of agents is bounded, and it is shown in [11,12] that the problem is strictly¹ harder otherwise, namely NP-hard on ATSS and Σ_2^P -hard on CGSs where the transition function is encoded as a boolean function. We prove that it is in fact Δ_2^P -complete and Δ_3^P -complete, resp., correcting wrong algorithms in [11,12] (the problem lies in the way the algorithms handle negations). We also show that ATL^+ is Δ_3^P -complete on both ATSSs and CGSs, even when the number of agents is fixed, extending a result of [18]. Finally we study translations between ATSS and CGS.

Related works. In [2,3] ATL has been proposed and defined over ATSS and CGS. In [10] expressiveness issues are considered for ATL^* and ATL. Complexity of

¹ We adopt the classical hypothesis that the polynomial-time hierarchy does not collapse, and that $\text{PTIME} \neq \text{NP}$. We refer to [15] for the definitions about complexity classes, especially about oracle Turing machines and the polynomial-time hierarchy.

satisfiability is addressed in [9,19]. Complexity results about model checking (for ATL, ATL^+ , ATL^*) can be found in [3,18]. Regarding control and game theory, many papers have focused on this wide area; we refer to [20] for a survey, and to its numerous references for a complete overview.

Plan of the paper. Section 2 contains the formal definitions that are used in the sequel. Section 3 explains our expressiveness result, and Section 4 deals with the model-checking algorithms. Due to lack of space, some proofs are omitted in this article, but can be found in [13].

2 Definitions

2.1 Concurrent Game Structures and Alternating Transition Systems

Definition 1. A Concurrent Game Structure (CGS for short) \mathcal{C} is a 6-tuple $(Agt, Loc, AP, Lab, Mov, Edg)$ s.t.:

- $Agt = \{A_1, \dots, A_k\}$ is a finite set of agents (or players);
- Loc and AP are two finite sets of locations and atomic propositions, resp.;
- $Lab: Loc \rightarrow 2^{AP}$ is a function labeling each location by the set of atomic propositions that hold for that location;
- $Mov: Loc \times Agt \rightarrow \mathcal{P}(\mathbb{N}) \setminus \{\emptyset\}$ defines the (finite) set of possible moves of each agent in each location.
- $Edg: Loc \times \mathbb{N}^k \rightarrow Loc$, where $k = |Agt|$, is a (partial) function defining the transition table. With each location and each set of moves of the agents, it associates the resulting location.

The intended behaviour is as follows [3]: in a given location ℓ , each player A_i chooses one possible move m_{A_i} in $Mov(\ell, A_i)$ and the successor 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 $Next(\ell, A_j, m)$ for the restriction of $Next(\ell)$ to locations reachable from ℓ when player A_j makes the move m .

In the original works about ATL [2], the logic was interpreted on ATSS, which are transition systems slightly different from CGSs:

Definition 2. An Alternating Transition System (ATS for short) \mathcal{A} is a 5-tuple (Agt, Loc, AP, Lab, Mov) where:

- Agt, Loc, AP and Lab have the same meaning as in CGSs;
- $Mov: Loc \times Agt \rightarrow \mathcal{P}(\mathcal{P}(Loc))$ associate with each location ℓ and each agent a the set of possible moves, each move being a subset of Loc . For each location ℓ , it is required that, for any $Q_i \in Mov(\ell, A_i)$, $\bigcap_{i \leq k} Q_i$ be a singleton.

The intuition is as follows: in a location ℓ , once all the agents have chosen their moves (i.e., a subset of locations), the execution goes to the (only) state that belongs to all the sets chosen by the players. Again $Next(\ell)$ (resp. $Next(\ell, A_j, m)$) denotes the set of all possible successor locations (resp. the set of possible successor locations when player A_j chooses the move m).

We prove in Section 4.2 that both models have the same expressiveness (w.r.t. alternating bisimilarity [4]).

2.2 Strategy, Outcomes of a Strategy

Let \mathcal{S} be a CGS or an ATS. A *computation* of \mathcal{S} 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 can use the standard notions of suffix and prefix for these computations; $\rho[i]$ denotes the i -th 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 ². A strategy is *state-based* (or *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 induces a set of computations from ℓ —called the *outcomes* of f_{A_i} from ℓ and denoted³ $\text{Out}_{\mathcal{S}}(\ell, f_{A_i})$ —that player A_i can enforce: $\ell_0 \ell_1 \ell_2 \dots \in \text{Out}_{\mathcal{S}}(\ell, f_{A_i})$ iff $\ell = \ell_0$ and for any i we have $\ell_{i+1} \in \text{Next}(\ell_i, A_i, f_{A_i}(\ell_0 \dots \ell_i))$. Let $A \subseteq \text{Agt}$ be a coalition. A strategy for A is a tuple F_A containing one strategy for every player in A : $F_A = \{f_{A_i} \mid A_i \in A\}$. The outcomes of F_A from a location ℓ contains the computations enforced by the strategies in F_A : $\ell_0 \ell_1 \dots \in \text{Out}_{\mathcal{S}}(\ell, F_A)$ s.t. $\ell = \ell_0$ and for any i , $\ell_{i+1} \in \bigcap_{A_i \in A} \text{Next}(\ell_i, A_i, f_{A_i}(\ell_0 \dots \ell_i))$. The set of strategies for A is denoted³ $\text{Strat}_{\mathcal{S}}(A)$. Finally $\text{Out}_{\mathcal{S}}(\ell, \emptyset)$ represents the set of all computations from ℓ .

2.3 The Logic ATL and Some Extensions

Again, we follow the definitions of [3]:

Definition 3. *The syntax of ATL is defined by the following grammar:*

$$\begin{aligned} \text{ATL} \ni \varphi_s, \psi_s &::= \top \mid p \mid \neg \varphi_s \mid \varphi_s \vee \psi_s \mid \langle\langle A \rangle\rangle \varphi_p \\ \varphi_p &::= \mathbf{X} \varphi_s \mid \mathbf{G} \varphi_s \mid \varphi_s \mathbf{U} \psi_s. \end{aligned}$$

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

In addition, we use standard abbreviations like \perp , \mathbf{F} , etc. ATL formulae are interpreted over states of a game structure \mathcal{S} . The semantics of the main modalities is defined as follows³:

$$\begin{aligned} \ell \models_{\mathcal{S}} \langle\langle A \rangle\rangle \varphi_p &\quad \text{iff} \quad \exists F_A \in \text{Strat}(A). \forall \rho \in \text{Out}(\ell, F_A). \rho \models_{\mathcal{S}} \varphi_p, \\ \rho \models_{\mathcal{S}} \mathbf{X} \varphi_s &\quad \text{iff} \quad \rho[1] \models_{\mathcal{S}} \varphi_s, \\ \rho \models_{\mathcal{S}} \mathbf{G} \varphi_s &\quad \text{iff} \quad \forall i. \rho[i] \models_{\mathcal{S}} \varphi_s, \\ \rho \models_{\mathcal{S}} \varphi_s \mathbf{U} \psi_s &\quad \text{iff} \quad \exists i. \rho[i] \models_{\mathcal{S}} \psi_s \text{ and } \forall 0 \leq j < i. \rho[j] \models_{\mathcal{S}} \varphi_s. \end{aligned}$$

It is well-known that, for the logic ATL, it is sufficient to restrict to state-based strategies (i.e., $\langle\langle A \rangle\rangle \varphi_p$ is satisfied iff there is a state-based strategy all of whose outcomes satisfy φ_p) [3,18].

Note that $\langle\langle \emptyset \rangle\rangle \varphi_p$ corresponds to the CTL formula $\mathbf{A} \varphi_p$ (i.e., universal quantification over all computations issued from the current state), while $\langle\langle \text{Agt} \rangle\rangle \varphi_p$

² I.e., $f_{A_i}(\ell_0 \dots \ell_m) \in \text{Mov}(\ell_m, A_i)$.

³ We might omit to mention \mathcal{S} when it is clear from the context.

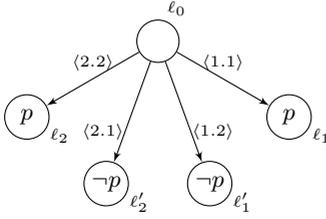


Fig. 1. A CGS that is not determined

$$\text{Loc} = \{\ell_0, \ell_1, \ell_2, \ell'_1, \ell'_2\}$$

$$\text{Mov}(\ell_0, A_1) = \{\{\ell_1, \ell'_1\}, \{\ell_2, \ell'_2\}\}$$

$$\text{Mov}(\ell_0, A_2) = \{\{\ell_1, \ell'_2\}, \{\ell_2, \ell'_1\}\}$$

$$\text{with } \begin{cases} \text{Lab}(\ell_1) = \text{Lab}(\ell_2) = \{p\} \\ \text{Lab}(\ell'_1) = \text{Lab}(\ell'_2) = \emptyset \end{cases}$$

Fig. 2. An ATS that is not determined

corresponds to existential quantification $\mathbf{E}\varphi_p$. Note, however, that $\neg\langle\langle A \rangle\rangle\varphi_p$ is generally *not* equivalent to $\langle\langle \text{Agt} \setminus A \rangle\rangle\neg\varphi_p$ [3,9]. Fig. 1 displays a (graphical representation of a) 2-player CGS for which, in ℓ_0 , both $\neg\langle\langle A_1 \rangle\rangle\mathbf{X}p$ and $\neg\langle\langle A_2 \rangle\rangle\neg\mathbf{X}p$ hold. In such a representation, a transition is labeled with $\langle m_1.m_2 \rangle$ when it correspond to move m_1 of player A_1 and to move m_2 of player A_2 . Fig. 2 represents an (alternating-bisimilar) ATS with the same properties.

Duality is a fundamental concept in modal and temporal logics: for instance, the dual of modality \mathbf{U} , often denoted by \mathbf{R} and read *release*, is defined by $\varphi_s \mathbf{R} \psi_s \stackrel{\text{def}}{=} \neg((\neg\varphi_s) \mathbf{U} (\neg\psi_s))$. Dual modalities allow, for instance, to put negations inner inside the formula, which is often an important property when manipulating formulas. In LTL, modality \mathbf{R} can be expressed using only \mathbf{U} and \mathbf{G} :

$$\varphi \mathbf{R} \psi \equiv \mathbf{G} \psi \vee \psi \mathbf{U} (\varphi \wedge \psi). \tag{1}$$

In the same way, it is well known that CTL can be defined using only modalities \mathbf{EX} , \mathbf{EG} and \mathbf{EU} , and that we have

$$\mathbf{E}\varphi \mathbf{R} \psi \equiv \mathbf{E}\mathbf{G} \psi \vee \mathbf{E}\psi \mathbf{U} (\varphi \wedge \psi) \qquad \mathbf{A}\varphi \mathbf{R} \psi \equiv \neg\mathbf{E}(\neg\varphi) \mathbf{U} (\neg\psi).$$

We prove in the sequel that modality \mathbf{R} cannot be expressed in ATL, as defined in Definition 3. We thus define the following two extensions of ATL:

Definition 4. We define $\text{ATL}_{\mathbf{R}}$ and ATL^+ with the following syntax:

$$\begin{aligned} \text{ATL}_{\mathbf{R}} \ni \varphi_s, \psi_s ::= & \top \mid p \mid \neg\varphi_s \mid \varphi_s \vee \psi_s \mid \langle\langle A \rangle\rangle\varphi_p \\ \varphi_p ::= & \mathbf{X}\varphi_s \mid \varphi_s \mathbf{U}\psi_s \mid \varphi_s \mathbf{R}\psi_s, \end{aligned}$$

$$\begin{aligned} \text{ATL}^+ \ni \varphi_s, \psi_s ::= & \top \mid p \mid \neg\varphi_s \mid \varphi_s \vee \psi_s \mid \langle\langle A \rangle\rangle\varphi_p \\ \varphi_p, \psi_p ::= & \neg\varphi_p \mid \varphi_p \vee \psi_p \mid \mathbf{X}\varphi_s \mid \varphi_s \mathbf{U}\psi_s \mid \varphi_s \mathbf{R}\psi_s. \end{aligned}$$

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

Given a formula φ in one of the logics we have defined, the size of φ , denoted by $|\varphi|$, is the size of the tree representing that formula. The DAG-size of φ is the size of the directed acyclic graph representing that formula (*i.e.*, sharing common subformulas).

3 $\langle\langle A \rangle\rangle(a \mathbf{R} b)$ Cannot Be Expressed in ATL

This section is devoted to the expressiveness of ATL. We prove:

Theorem 5. *There is no ATL formula equivalent to $\Phi = \langle\langle A \rangle\rangle(a \mathbf{R} b)$.*

The proof of Theorem 5 is based on techniques similar to those used for proving expressiveness results for temporal logics like CTL or ECTL [8]: we build two families of models $(s_i)_{i \in \mathbb{N}}$ and $(s'_i)_{i \in \mathbb{N}}$ s.t. (1) $s_i \not\models \Phi$, (2) $s'_i \models \Phi$ for any i , and (3) s_i and s'_i satisfy the same ATL formula of size less than i . Theorem 5 is a direct consequence of the existence of such families of models. In order to simplify the presentation, the theorem is proved for formula⁴ $\Phi = \langle\langle A \rangle\rangle(b \mathbf{R}(a \vee b))$.

The models are described by one single inductive CGS⁵ \mathcal{C} , involving only two players. It is depicted on Fig. 3. A label $\langle\alpha.\beta\rangle$ on a transition indicates that

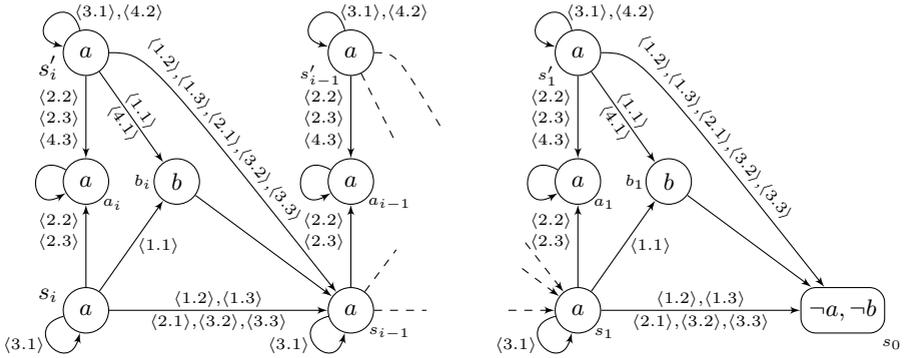


Fig. 3. The CGS \mathcal{C} , with states s_i and s'_i on the left

this transition corresponds to move α of player A_1 and to move β of player A_2 . In that CGS, states s_i and s'_i only differ in that player A_1 has a fourth possible move in s'_i . This ensures that, from state s'_i (for any i), player A_1 has a strategy (namely, he should always play 4) for enforcing $a \mathbf{W} b$. But this is not the case from state s_i : by induction on i , one can prove $s_i \not\models \langle\langle A_1 \rangle\rangle a \mathbf{W} b$. The base case is trivial. Now assume the property holds for i : from s_{i+1} , any strategy for A_1 starts with a move in $\{1, 2, 3\}$ and for any of these choices, player A_2 can choose a move (2, 1, and 2, resp.) that enforces the next state to be s_i where by i.h., A_1 has no strategy for $a \mathbf{W} b$.

We now prove that s_i and s'_i satisfy the same “small” formulae. First, we have the following equivalences:

Lemma 6. *For any $i > 0$, for any $\psi \in \text{ATL}$ with $|\psi| \leq i$:*

$$b_i \models \psi \text{ iff } b_{i+1} \models \psi \quad s_i \models \psi \text{ iff } s_{i+1} \models \psi \quad s'_i \models \psi \text{ iff } s'_{i+1} \models \psi$$

⁴ This formula can also be written $\langle\langle A \rangle\rangle a \mathbf{W} b$, where \mathbf{W} is the “weak until” modality.

⁵ Given the translation from CGS to ATS (see Sec. 4.2), the result also holds for ATSs.

Lemma 7. $\forall i > 0, \forall \psi \in \text{ATL}$ with $|\psi| \leq i: s_i \models \psi$ iff $s'_i \models \psi$.

Proof. The proof proceeds by induction on i , and on the structure of the formula ψ . The case $i = 1$ is trivial, since s_1 and s'_1 carry the same atomic propositions. For the induction step, dealing with CTL modalities ($\langle\langle \emptyset \rangle\rangle$ and $\langle\langle A_1, A_2 \rangle\rangle$) is also straightforward, then we just consider $\langle\langle A_1 \rangle\rangle$ and $\langle\langle A_2 \rangle\rangle$ modalities.

First we consider $\langle\langle A_1 \rangle\rangle$ modalities. It is well-known that we can restrict to state-based strategies in this setting. If player A_1 has a strategy in s_i to enforce something, then he can follow the same strategy from s'_i . Conversely, if player A_1 has a strategy in s'_i to enforce some property, two cases may arise: either the strategy consists in playing move 1, 2 or 3, and it can be mimicked from s_i . Or the strategy consists in playing move 4 and we distinguish three cases:

- $\psi = \langle\langle A_1 \rangle\rangle \mathbf{X} \psi_1$: that move 4 is a winning strategy entails that s'_i, a_i and b_i must satisfy ψ_1 . Then s_i (by i.h. on the formula) and s_{i-1} (by Lemma 6) both satisfy ψ_1 . Playing move 1 (or 3) in s_i ensures that the next state will satisfy ψ_1 .
- $\psi = \langle\langle A_1 \rangle\rangle \mathbf{G} \psi_1$: by playing move 4, the game could end up in s_{i-1} (via b_i), and in a_i and s'_i . Thus $s_{i-1} \models \psi$, and in particular ψ_1 . By i.h., $s_i \models \psi_1$, and playing move 1 (or 3) in s_i , and then mimicking the original strategy (from s'_i), enforces $\mathbf{G} \psi_1$.
- $\psi = \langle\langle A_1 \rangle\rangle \psi_1 \mathbf{U} \psi_2$: a strategy starting with move 4 implies $s'_i \models \psi_2$ (the game could stay in s'_i for ever). Then $s_i \models \psi_2$ by i.h., and the result follows.

We now turn to $\langle\langle A_2 \rangle\rangle$ modalities: clearly if $\langle\langle A_2 \rangle\rangle \psi_1$ holds in s'_i , it also holds in s_i . Conversely, if player A_2 has a (state-based) strategy to enforce some property in s_i : If it consists in playing either 1 or 3, then the same strategy also works in s'_i . Now if the strategy starts with move 2, then playing move 3 in s'_i has the same effect, and thus enforces the same property. □

Remark 1. ATL and $\text{ATL}_{\mathbf{R}}$ have the same distinguishing power as the fragment of ATL involving only the $\langle\langle \cdot \rangle\rangle \mathbf{X}$ modality (see proof of Theorem 6 in [4]). This means that we cannot exhibit two models M and M' s.t. (1) $M \models \Phi$, (2) $M' \not\models \Phi$, and (3) M and M' satisfy the same ATL formula.

Though ATL^+ would not contain the “release” modality in its syntax, it can express it, and is thus strictly more expressive than ATL. However, as for CTL and CTL^+ , it is possible to translate ATL^+ into $\text{ATL}_{\mathbf{R}}$ [10]. Of course, such a translation induces at least an exponential blow-up in the size of the formulae since it is already the case when translating CTL^+ into CTL [21,1]. Finally note that the standard model-checking algorithm for ATL easily extends to $\text{ATL}_{\mathbf{R}}$ (and that MOCHA [5] handles $\text{ATL}_{\mathbf{R}}$ formulae). In the same way, the axiomatization and satisfiability results of [9] can be extended to $\text{ATL}_{\mathbf{R}}$ (as mentioned in the conclusion of [9]).

Turn-based games. In [3], a restriction of CGS —the turn-based considered. In any location of these models (named TB-CGS hereafter), only one player has several moves (the other players have only one possible choice). Such models have

the property of *determinedness* (if the objectives are Borel-definable, which is the case for ATL^+): given a set of players A , either there is a strategy for A to win some objective Φ , or there is a strategy for other players ($\text{Agt} \setminus A$) to enforce $\neg\Phi$. In such systems, modality \mathbf{R} can be expressed as follows: $\langle\langle A \rangle\rangle \varphi \mathbf{R} \psi \equiv_{\text{TB-CGS}} \neg \langle\langle \text{Agt} \setminus A \rangle\rangle (\neg\varphi) \mathbf{U} (\neg\psi)$.

4 Complexity of ATL Model-Checking

In this section, we establish the precise complexity of ATL model-checking. All the complexity results below are stated for ATL but they are also true for $\text{ATL}_{\mathbf{R}}$.

Model-checking issues have been addressed in the seminal papers about ATL, on both ATSS [2] and CGSs [3]. The time complexity is shown to be in $O(m \cdot l)$, where m is the size of the transition table and l is the size of the formula. The authors then claim that the model-checking problem is in PTIME (and obviously, PTIME-complete). However, it is well-known (and already explained in [2,3]) that the size m of the transition table may be exponential in the number of agents. Thus, when the transition table is not given explicitly (as is the case for ATS), the algorithm requires in fact exponential time.

Before proving that this problem is indeed not in PTIME, we define the model of *implicit* CGSs, with a succinct representation of the transition table [11]. Besides the theoretical aspect, it may be quite useful in practice since it allows to not explicitly describe the full transition table.

4.1 Explicit and Implicit CGSs

We distinguish between two classes of CGSs:

Definition 8. • An implicit CGS is a CGS where, in each location ℓ , the transition function is defined by a finite sequence $((\varphi_0, \ell_0), \dots, (\varphi_n, \ell_n))$, where $\ell_i \in \text{Loc}$ is a location, and φ_i is a boolean combination of propositions $A_j = c$ that evaluate to true iff agent A_i chooses move c .

The transition table is then defined as follows: $\text{Edg}(\ell, m_{A_1}, \dots, m_{A_k}) = \ell_j$ iff j is the lowest index s.t. φ_j evaluates to true when players A_1 to A_k choose moves m_{A_1} to m_{A_k} . We require that the last boolean formula φ_i be \top , so that no agent can enforce a deadlock.

• An explicit CGS is a CGS where the transition table is defined explicitly.

The size $|\mathcal{C}|$ of a CGS \mathcal{C} is defined as $|\text{Loc}| + |\text{Edg}|$. For explicit CGSs, $|\text{Edg}|$ is the size of the transition table. For implicit CGSs, $|\text{Edg}|$ is the sum $\sum |\varphi|$ used for the definition of Edg .

The size of an ATS is $|\text{Loc}| + |\text{Mov}|$ where $|\text{Mov}|$ is the sum of the number of locations in each possible move of each agent in each location.

4.2 Expressiveness of CGSs and ATSS

We prove in this section that CGSs and ATSS are closely related: they can model the same concurrent games. In order to make this statement formal, we use the following definition, which extends bisimulation to strategies of coalitions:

Definition 9 ([4]). Let \mathcal{A} and \mathcal{B} be two models of concurrent games (either ATSS or CGSSs) over the same set Agt of agents. Let $R \subseteq \text{Loc}_{\mathcal{A}} \times \text{Loc}_{\mathcal{B}}$ be a (non-empty) relation between states of \mathcal{A} and states of \mathcal{B} . That relation is an alternating bisimulation when, for any $(\ell, \ell') \in R$, the following conditions hold:

- $\text{Lab}_{\mathcal{A}}(\ell) = \text{Lab}_{\mathcal{B}}(\ell')$;
- for any coalition $A \subseteq \text{Agt}$, we have

$$\forall m: A \rightarrow \text{Mov}_{\mathcal{A}}(\ell, A). \exists m': A \rightarrow \text{Mov}_{\mathcal{B}}(\ell', A).$$

$$\forall q' \in \text{Next}(\ell', A, m'). \exists q \in \text{Next}(\ell, A, m). (q, q') \in R.$$
- symmetrically, for any coalition $A \subseteq \text{Agt}$, we have

$$\forall m': A \rightarrow \text{Mov}_{\mathcal{B}}(\ell', A). \exists m: A \rightarrow \text{Mov}_{\mathcal{A}}(\ell, A).$$

$$\forall q \in \text{Next}(\ell, A, m). \exists q' \in \text{Next}(\ell', A, m'). (q, q') \in R.$$

where $\text{Next}(\ell, A, m)$ is the set of locations that are reachable from ℓ when each player $A_i \in A$ plays $m(A_i)$.

Two models are said to be alternating-bisimilar if there exists an alternating bisimulation involving all of their locations.

It turns out that ATSSs and CGSSs (both implicit and explicit ones) have the same expressive power w.r.t. this equivalence:

Theorem 10. 1. Any explicit CGS can be translated into an alternating-bisimilar implicit one in linear time; 2. Any implicit CGS can be translated into an alternating-bisimilar explicit one in exponential time; 3. Any explicit CGS can be translated into an alternating-bisimilar ATS in cubic time; 4. Any ATS can be translated into an alternating-bisimilar explicit CGS in exponential time; 5. Any implicit CGS can be translated into an alternating-bisimilar ATS in exponential time; 6. Any ATS can be translated into an alternating-bisimilar implicit CGS in quadratic time;

Figure 4 summarizes the above results. From our complexity results (and the assumption that the polynomial-time hierarchy does not collapse), the costs of the above translations is optimal.

4.3 Model Checking ATL on Implicit CGSSs

Basically, the algorithm for model-checking ATL [2,3] is similar to that for CTL: it consists in recursively computing fixpoints, *e.g.* based on the following equivalence:

$$\langle\langle A \rangle\rangle p \mathbf{U} q \equiv \mu Z. (q \vee (p \wedge \langle\langle A \rangle\rangle \mathbf{X} Z))$$

The difference with CTL is that we have to compute the pre-image of a set of states *for some coalition*.

It has been remarked in [11] that computing the pre-images is not in PTIME anymore when considering implicit CGSSs: the algorithm has to non-deterministically guess the moves of players in A in each location, and for each pre-image, to

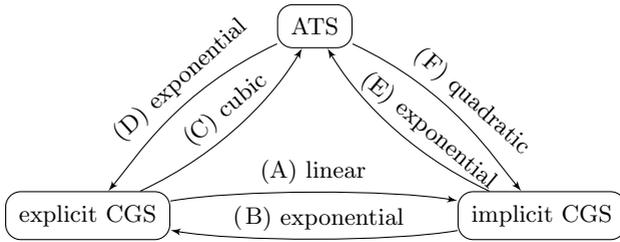


Fig. 4. Costs of translations between the three models

solve the resulting SAT queries derived from those choices and from the transition table. As a consequence, model-checking ATL on implicit CGSs is Σ_2^P -hard [11]. However (see below), the Σ_2^P -hardness proof can very easily be adapted to prove Π_2^P -hardness. It follows that the Σ_2^P algorithm proposed in [11] cannot be correct. The flaw is in the way it handles negation: games played on CGSs (and ATSs) are generally not determined, and the fact that a player has no strategy to enforce φ does not imply that the other players have a strategy to enforce $\neg\varphi$. It rather means that the other players have a *co-strategy* to enforce $\neg\varphi$ (see [9] for precise explanations about co-strategies).

Still, the Σ_2^P -algorithm is correct for formulas whose main operator is not a negation. As a consequence:

Proposition 11. *Model checking ATL on implicit CGSs is in Δ_3^P .*

Since the algorithm consists in labeling the locations with the subformulae it satisfies, the above result holds even if we consider the DAG-size of the formula.

Before proving optimality, we briefly recall the Σ_2^P -hardness proof of [11]. It relies on the following Σ_2^P -complete problem:

EQSAT₂:

Input: two families of variables $X = \{x^1, \dots, x^n\}$ and $Y = \{y^1, \dots, y^n\}$, a boolean formula φ on the set of variables $X \cup Y$.

Output: True iff $\exists X. \forall Y. \varphi$.

This problem can be encoded in an ATL model-checking problem on an implicit CGS: the CGS has three states q_1, q_\top and q_\perp , and $2n$ agents $A^1, \dots, A^n, B^1, \dots, B^n$, each having two possible choices in q_1 and only one choice in q_\top and q_\perp . The transitions out of q_\top and q_\perp are self loops. The transitions from q_1 are given by: $\delta(q_1) = ((\varphi[x^j \leftarrow (A^j \stackrel{?}{=} 1), y^j \leftarrow (B^j \stackrel{?}{=} 1)], q_\top)(\top, q_\perp)$.

Then clearly, the coalition A^1, \dots, A^n has a strategy for reaching q_\top , i.e., $q_1 \models \langle\langle A^1, \dots, A^n \rangle\rangle \mathbf{X} q_\top$, iff there exists a valuation for variables in X s.t. φ is true whatever B -agents choose for Y .

Now, this encoding can easily be adapted to the dual (thus Π_2^P -complete) problem AQSAT₂, in which, with the same input, the output is the value of $\forall X. \exists Y. \varphi$. It suffices to consider the same implicit CGS, and the formula

$\neg \langle\langle A^1, \dots, A^n \rangle\rangle \mathbf{X} \neg q_{\top}$. It states that there is no strategy for players A^1 to A^n to avoid q_{\top} : whatever their choice, players B^1 to B^n can enforce φ .

Following the same idea, we prove the following result:

Proposition 12. *Model checking ATL on implicit CGSs is Δ_3^P -hard.*

Proof. We consider the following Δ_3^P -complete problem[14,18].

SNSAT₂:

Input: m families of variables $X_i = \{x_i^1, \dots, x_i^{n_i}\}$, m families of variables $Y_i = \{y_i^1, \dots, y_i^{n_i}\}$, m variables z_i , m boolean formulae φ_i , with φ_i involving variables in $X_i \cup Y_i \cup \{z_1, \dots, z_{i-1}\}$.

Output: The value of z_m , defined by

$$\begin{cases} z_1 \stackrel{\text{def}}{=} \exists X_1. \forall Y_1. \varphi_1(X_1, Y_1) \\ z_2 \stackrel{\text{def}}{=} \exists X_2. \forall Y_2. \varphi_2(z_1, X_2, Y_2) \\ \dots \\ z_m \stackrel{\text{def}}{=} \exists X_m. \forall Y_m. \varphi_m(z_1, \dots, z_{m-1}, X_m, Y_m) \end{cases}$$

We pick an instance \mathcal{I} of this problem, and reduce it to an instance of the ATL model-checking problem. Note that such an instance uniquely defines the values of variables z_i . We write $v_{\mathcal{I}}: \{z_1, \dots, z_m\} \rightarrow \{\top, \perp\}$ for this valuation. Also, when $v_{\mathcal{I}}(z_i) = \top$, there exists a witnessing valuation for variables in X_i . We extend $v_{\mathcal{I}}$ to $\{z_1, \dots, z_m\} \cup \bigcup_i X_i$, with $v_{\mathcal{I}}(x_i^j)$ being a witnessing valuation if $v_{\mathcal{I}}(z_i) = \top$.

We now define an implicit CGS \mathcal{C} as follows: it contains mn agents A_i^j (one for each x_i^j), mn agents B_i^j (one for each y_i^j), m agents C_i (one for each z_i), and one extra agent D . The structure is made of m states q_i , m states \bar{q}_i , m states s_i , and two states q_{\top} and q_{\perp} . There are three atomic propositions: s_{\top} and s_{\perp} , that label the states q_{\top} and q_{\perp} resp., and an atomic proposition s labeling states s_i . The other states carry no label.

Except for D , the agents represent booleans, and thus always have two possible choices (0 and 1). Agent D always has m choices (0 to $m - 1$). The transition relation is defined as follows: for each i ,

$$\begin{aligned} \delta(\bar{q}_i) &= ((\top, s_i)); \\ \delta(s_i) &= ((\top, q_i)); \\ \delta(q_{\top}) &= ((\top, q_{\top})); \\ \delta(q_{\perp}) &= ((\top, q_{\perp})); \end{aligned} \quad \delta(q_i) = \left(\begin{array}{l} ((D \stackrel{?}{=} 0) \wedge \varphi_i[x_i^j \leftarrow (A_i^j \stackrel{?}{=} 1), \\ y_i^j \leftarrow (B_i^j \stackrel{?}{=} 1), z_k \leftarrow (C_k \stackrel{?}{=} 1)], q_{\top}) \\ ((D \stackrel{?}{=} 0), q_{\perp}) \\ ((D \stackrel{?}{=} k) \wedge (C_k \stackrel{?}{=} 1), q_k) \text{ for each } k < i \\ ((D \stackrel{?}{=} k) \wedge (C_k \stackrel{?}{=} 0), \bar{q}_k) \text{ for each } k < i \\ (\top, q_{\top}) \end{array} \right)$$

Intuitively, from state q_i , the boolean agents choose a valuation for the variable they represent, and agent D can either choose to check if the valuation really witnesses φ_i (by choosing move 0), or “challenge” player C_k , with move $k < i$.

The ATL formula is built recursively by $\psi_0 = \top$ and, writing \mathbf{AC} for the coalition $\{A_1^1, \dots, A_m^n, C_1, \dots, C_m\}$: $\psi_{k+1} \stackrel{\text{def}}{=} \langle\langle \mathbf{AC} \rangle\rangle (\neg s) \mathbf{U} (q_{\top} \vee \mathbf{EX} (s \wedge \mathbf{EX} \neg \psi_k))$.

Let $f_{\mathcal{I}}(A)$ be the state-based strategy for agent $A \in \text{AC}$ that consists in playing according to the valuation $v_{\mathcal{I}}$ (i.e. move 0 if the variable associated with A evaluates to 0 in $v_{\mathcal{I}}$, and move 1 otherwise). The following lemma completes the proof of Proposition 12:

Lemma 13. *For any $i \leq m$ and $k \geq i$, the following three statements are equivalent: (a) $\mathcal{C}, q_i \models \psi_k$; (b) the strategies $f_{\mathcal{I}}$ witness the fact that $\mathcal{C}, q_i \models \psi_k$; (c) variable z_i evaluates to \top in $v_{\mathcal{I}}$. \square*

Finally, Lemma 13 and Proposition 11, this implies:

Theorem 14. *Model checking ATL on implicit CGSs is Δ_3^P -complete.*

4.4 Model Checking ATL on ATSSs

Also for ATSSs, the PTIME upper bound holds only when the number of agents is fixed. As in the previous section, the NP algorithm proposed in [11] for ATL model-checking on ATSSs does not handle negation correctly. Again, the algorithm involves the fixpoint computation with pre-images, and the pre-images are now computed in NP [11]. This yields a Δ_2^P algorithm for full ATL.

Proposition 15. *Model checking ATL over ATSSs is in Δ_2^P .*

The NP-hardness proof of [11] can be adapted in order to give a direct reduction of 3SAT, and then extended to SNSAT:

Proposition 16. *Model checking ATL on ATSSs is Δ_2^P -hard.*

Proof. Let us first recall the definition of the SNSAT problem [14]:

SNSAT:

Input: p families of variables $X_r = \{x_r^1, \dots, x_r^m\}$, p variables z_r , p boolean formulae φ_r in 3-CNF, with φ_r involving variables in $X_r \cup \{z_1, \dots, z_{r-1}\}$.

Output: The value of z_p , defined by

$$\left\{ \begin{array}{l} z_1 \stackrel{\text{def}}{=} \exists X_1. \varphi_1(X_1) \\ z_2 \stackrel{\text{def}}{=} \exists X_2. \varphi_2(z_1, X_2) \\ z_3 \stackrel{\text{def}}{=} \exists X_3. \varphi_3(z_1, z_2, X_3) \\ \dots \\ z_p \stackrel{\text{def}}{=} \exists X_p. \varphi_p(z_1, \dots, z_{p-1}, X_p) \end{array} \right.$$

Let \mathcal{I} be an instance of SNSAT, where we assume that each φ_r is made of n clauses S_r^1 to S_r^n , with $S_r^j = \alpha_r^{j,1} s_r^{j,1} \vee \alpha_r^{j,2} s_r^{j,2} \vee \alpha_r^{j,3} s_r^{j,3}$. Again, such an instance uniquely defines a valuation $v_{\mathcal{I}}$ for variables z_1 to z_r , that can be extended to the whole set of variables by choosing a witnessing valuation for x_r^1 to x_r^n when z_r evaluates to true.

We now describe the ATS \mathcal{A} : it contains $(8n + 3)p$ states: p states $\overline{q_r}$ and p states q_r , p states s_r , and for each formula φ_r , for each clause S_r^j of φ_r , eight states $q_r^{j,0}, \dots, q_r^{j,7}$, as in the previous reduction.

States s_r are labeled with the atomic proposition s , and states $q_r^{j,k}$ that do not correspond to clause S_r^j are labeled with α .

There is one player A_r^j for each variable x_r^j , one player C_r for each z_r , plus one extra player D . As regards transitions, there are self-loops on each state $q_r^{j,k}$, single transitions from each \overline{q}_r to the corresponding s_r , and from each s_r to the corresponding q_r . From state q_r ,

- player A_r^j will choose the value of variable x_r^j , by selecting one of the following two sets of states:

$$\begin{aligned} & \{q_r^{g,k} \mid \forall l \leq 3. s_r^{g,l} \neq x_r^j \text{ or } \alpha_r^{g,l} = 0\} \cup \{q_t, \overline{q}_t \mid t < r\} && \text{if } x_r^j = \top \\ & \{q_r^{g,k} \mid \forall l \leq 3. s_r^{g,l} \neq x_r^j \text{ or } \alpha_r^{g,l} = 1\} \cup \{q_t, \overline{q}_t \mid t < r\} && \text{if } x_r^j = \perp \end{aligned}$$

Both choices also allow to go to one of the states q_t or \overline{q}_t . In q_r , players A_t^j with $t \neq r$ have one single choice, which is the whole set of states.

- player C_t also chooses for the value of the variable it represents. As for players A_t^j , this choice will be expressed by choosing between two sets of states corresponding to clauses that are not made true. But as in the proof of Prop. 12, players C_t will also offer the possibility to “verify” their choice, by going either to state q_t or \overline{q}_t . Formally, this yields two sets of states:

$$\begin{aligned} & \{q_r^{g,k} \mid \forall l \leq 3. s_r^{g,l} \neq z_t \text{ or } \alpha_r^{g,l} = 0\} \cup \{q_u, \overline{q}_u \mid u \neq t\} \cup \{q_t\} && \text{if } z_t = \top \\ & \{q_r^{g,k} \mid \forall l \leq 3. s_r^{g,l} \neq z_t \text{ or } \alpha_r^{g,l} = 1\} \cup \{q_u, \overline{q}_u \mid u \neq t\} \cup \{\overline{q}_t\} && \text{if } z_t = \perp \end{aligned}$$

- Last, player D chooses either to challenge a player C_t , with $t < r$, by choosing the set $\{q_t, \overline{q}_t\}$, or to check that a clause S_r^j is fulfilled, by choosing $\{q_r^{j,0}, \dots, q_r^{j,7}\}$.

Let us first prove that any choices of all the players yields exactly one state. It is obvious except for states q_r . For a state q_r , let us first restrict to the choices of all the players A_r^j and C_r , then:

- if we only consider states $q_r^{1,0}$ to $q_r^{n,7}$, the same argument as in the previous proof ensures that precisely one state per clause is chosen,
- if we consider states q_t and \overline{q}_t , the choices of players B_t ensure that exactly one state has been chosen in each pair $\{q_t, \overline{q}_t\}$, for each $t < r$.

Clearly, the choice of player D will select exactly one of the remaining states.

Now, we build the ATL formula. It is a recursive formula (very similar to the one used in the proof of Prop. 12), defined by $\psi_0 = \top$ and (again writing AC for the set of players $\{A_1^1, \dots, A_p^m, C_1, \dots, C_p\}$):

$$\psi_{r+1} \stackrel{\text{def}}{=} \langle\langle \text{AC} \rangle\rangle (-s) \mathbf{U} (\alpha \vee \mathbf{EX} (s \wedge \mathbf{EX} \neg \psi_r)).$$

Then, writing $f_{\mathcal{I}}$ for the state-based strategy associated to $v_{\mathcal{I}}$:

Lemma 17. *For any $r \leq p$ and $t \geq r$, the following statements are equivalent: (a) $q_r \models \psi_t$; (b) the strategies $f_{\mathcal{I}}$ witness the fact that $q_r \models \psi_t$; (c) variable z_r evaluates to true in $v_{\mathcal{I}}$. \square*

Theorem 18. *Model checking ATL on ATSS is Δ_2^P -complete.*

4.5 Model Checking ATL^+

The complexity of model checking ATL^+ over ATSSs has been settled Δ_3^P -complete in [18]. But Δ_3^P -hardness proof of [18] is in LOGSPACE only w.r.t. the DAG-size of the formula. We prove that model checking ATL^+ is in fact Δ_3^P -complete (with the standard definition of the size of a formula) for our three kinds of game structures.

Theorem 19. *Model checking ATL^+ is Δ_3^P -complete on ATSSs as well as on explicit CGSs and implicit CGSs.*

5 Conclusion

In this paper, we considered the basic questions of expressiveness and complexity of ATL. We showed that ATL, as originally defined in [2,3], is not as expressive as it could be expected, and we argue that the modality “Release” should be added in its definition [12].

We also precisely characterized the complexity of ATL and ATL^+ model-checking, on both ATSSs and CGSs, when the number of agents is not fixed. These results complete the previously known results about these formalisms and it is interesting to see that their complexity classes (Δ_2^P or Δ_3^P) are unusual in the model-checking area.

As future works, we plan to focus on the extensions EATL (extending ATL with a modality $\langle\langle \cdot \rangle\rangle_{\mathbf{F}}^{\infty}$ expressing a Büchi-like winning condition, and for which state-based strategies are still sufficient) and $EATL^+$ (the obvious association of both extensions, but for which state-based strategies are not sufficient anymore).

Acknowledgement. We thank Wojtek Jamroga for pointing out that formulas in $SNSAT_2$ cannot be restricted to CNF [6].

References

1. M. Adler and N. Immerman. An $n!$ lower bound on formula size. In *Proceedings of the 16th Annual Symposium on Logic in Computer Science (LICS'01)*, p. 197–206. IEEE CSP, 2001.
2. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. In *Proc. 38th Annual Symp. Foundations of Computer Science (FOCS'97)*, p. 100–109. IEEE CSP, 1997.
3. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002.
4. R. Alur, T. A. Henzinger, O. Kupferman, and M. Y. Vardi. Alternating refinement relations. In *Proc. 9th Intl Conf. Concurrency Theory (CONCUR'98)*, vol. 1466 of LNCS, p. 163–178. Springer, 1998.
5. R. Alur, T. A. Henzinger, F. Y. C. Mang, S. Qadeer, S. K. Rajamani, and S. Tasiran. Mocha: Modularity in model checking. In *Proc. 10th Intl Conf. Computer Aided Verification (CAV'98)*, vol. 1427 of LNCS, p. 521–525. Springer, 1998.

6. M. Cadoli, A. Giovanardi, and M. Schaerf. An algorithm to evaluate quantified boolean formulae. *Journal of Automated Reasoning*, 28(2):101–142, 2002.
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)*, vol. 131 of *LNCS*, p. 52–71. Springer, 1981.
8. E. A. Emerson. Temporal and modal logic. In *Handbook of Theoretical Computer Science*, vol. B, chapter 16, p. 995–1072. Elsevier, 1990.
9. V. Goranko and G. van Drimmelen. Complete axiomatization and decidability of alternating-time temporal logic. *Theoretical Computer Science*, 353(1-3):93–117, Mar. 2006.
10. A. Harding, M. Ryan, and P.-Y. Schobbens. Approximating ATL* in ATL. In *Revised Papers 3rd Intl Workshop Verification, Model Checking, and Abstract Interpretation (VMCAI'02)*, vol. 2294 of *LNCS*, p. 289–301. Springer, 2002.
11. W. Jamroga and J. Dix. Do agents make model checking explode (computationally)? In *Proc. 4th Intl Centr. and East. Europ. Conf. Multi-Agent Systems (CEEMAS'05)*, vol. 3690 of *LNCS*. Springer, 2005.
12. W. Jamroga and J. Dix. Model checking abilities of agents: A closer look. Technical Report IfI-06-02, Institut für Informatik, Technische Universität Clausthal, Germany, 2006.
13. F. Laroussinie, N. Markey, and G. Oreiby. Expressiveness and complexity of ATL. Technical Report LSV-06-03, Lab. Spécification & Vérification, ENS Cachan, France, 2006.
14. F. Laroussinie, N. Markey, and Ph. Schnoebelen. Model checking CTL⁺ and FCTL is hard. In *Proc. 4th Intl Conf. Foundations of Software Science and Computation Structure (FoSSaCS'01)*, vol. 2030 of *LNCS*, p. 318–331. Springer, 2001.
15. Ch. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
16. A. Pnueli. The temporal logic of programs. In *Proc. 18th Ann. Symp. Foundations of Computer Science (FOCS'77)*, p. 46–57. IEEE Comp. Soc. Press, 1977.
17. J.-P. Queille and J. Sifakis. Specification and verification of concurrent systems in CESAR. In *Proc. 5th Intl Symp. on Programming (SOP'82)*, vol. 137 of *LNCS*, p. 337–351. Springer, 1982.
18. P.-Y. Schobbens. Alternating-time logic with imperfect recall. In *Proc. 1st Workshop Logic and Communication in Multi-Agent Systems (LCMAS'03)*, vol. 85 of *ENTCS*. Elsevier, 2004.
19. D. Walther, C. Lutz, F. Wolter, and M. Wooldridge. ATL satisfiability is indeed EXPTIME-complete. *Journal of Logic and Computation*, 2006. To appear.
20. I. Walukiewicz. A landscape with games in the background. In *Proc. 19th Ann. Symp. Logic in Computer Science (LICS'04)*, p. 356–366. IEEE CSP, 2004.
21. Th. Wilke. CTL⁺ is exponentially more succinct than CTL. In *Proc. 19th Conf. Foundations of Software Technology and Theoretical Computer Science (FSTTCS'99)*, vol. 1738 of *LNCS*, p. 110–121. Springer, 1999.