ATL with contexts: agency and explicit strategies

Nicolas Troquard[1] and Dirk Walther[2]

[1] Laboratory for Applied Ontology (ISTC-CNR), Trento, Italy
[2] Technical University of Madrid (UPM), Spain

Abstract. This paper reports on first results of our study of the recently proposed family of ATL_{sc} logics. We provide the intuitions for that it could serve as an (already well-studied) unifying framework for the verification of properties about time, ability, strategies and agency in societies of agents. We relate it with STIT logics of actual agency, and with ATL with explicit strategies. We establish that the problem of satisfiability checking for ATL_{sc} over general concurrent game structures (with possibly infinitely many moves and states) is undecidable.

1 Introduction

The first aim of this paper is to consider the recent extension of ATL with *strategy contexts* [6, 7, 11, 10] and reveal its relevance for the general discussion of strategies, ability and agency, and ground it in the topic of logics for multi-agent systems as well.[1] We explain how ATL_{sc} and ATL^*_{sc} can capture a variety of notions of *strategic actual agency* that lie beyond the mere ability of coalitions as captured by ATL.

The second goal is to contrast the use of strategy contexts with *explicit strategies*, to point out their similarities and differences in expressivity and flexibility. ATL_{sc} and ATLES capture both notions of commitment to and release of strategies. We relate the two logics and discuss how they capture these notions as well as notions like *irrevocable strategies*, *forgetting forever* and *recall* of strategies.

The third contribution is technical. The focus of ATL_{sc} has been on model checking so far, and not satisfiability. To make as few semantic assumptions as possible, we consider a generalisation of concurrent game structures with possibly infinitely many states and possibly infinitely many choices. We establish that, over such structures, the satisfiability problem for the logic ATL_{sc} is undecidable. In this more general setting, it is then not fit for *reasoning* about multi-agent systems. However, this is the price to pay as even apparently much simpler logics present the same drawback (e.g., Chellas' STIT logic of group agency [14]). Nevertheless, when interpreted over finite models, we identify a positive fragment

[1] For the moment, we leave aside the variants of ATL_{sc} designed to specify perfect-recall and bounded-memory strategies, although it is also a prominent aspect of societies of agents. We do not consider strategies in imperfect information either.

of ATL_{sc} that can be translated into ATLES, for which a decision procedure is known.

The paper is organized as follows. We present syntax and semantics of ATL_{sc} in Section 2. In Section 3, we explain how ATL_{sc} and ATL^*_{sc} are adequate for describing notions of actual agency and ATL-like ability. We try to illustrate the richness of the languages by proposing several variants and we point out a difficulty with how the ATL modality was defined in ATL_{sc}. We introduce ATLES on concurrent game structures in Section 4 and compare ATL_{sc} with ATLES. Moreover, we determine a fragment of ATL_{sc} that corresponds to ATLES. In Section 5, we show that ATL^*_{sc} is undecidable over general concurrent game structures (with possibly infinitely many states and choices).

2 ATL with strategy contexts

Any language in this paper is defined over a signature containing an infinite supply of ingredients. While it is typically accepted to have infinitely many propositions available to use in formulas, languages for multiple agents often assume the set of agents to be finite. Using a finite set of agents in the signature, instead of an infinite set, gives rise to a different language that, although of same cardinality, may lead to a different computational complexity for their respective reasoning problems. A finite bound on the number of agents limits the modelling capability of the language and, thus, restricts its generality. We fix $\mathbf{\Pi}$ and $\mathbf{\Sigma}$ to be countable infinite sets of, respectively, *atomic propositions* and *agents* (or *players*). All languages here are defined using $\mathbf{\Pi}$ and $\mathbf{\Sigma}$ as their signature.

The following grammar was given for ATL^*_{sc} in [11]. In this paper, however, we differ from the original definition in that we do not assume the set of agent symbols in the signature to be finite.

Definition 1 (ATL^*_{sc} syntax). *The following grammar defines state formulas φ and path formulas ψ, where p ranges over $\mathbf{\Pi}$ and A over finite subsets of $\mathbf{\Sigma}$. The language of ATL^*_{sc} consists of the state formulas.*

$$\begin{aligned}\varphi &::= \ p \ \mid \ \neg\varphi \ \mid \ \varphi \vee \varphi \ \mid \ \cdot\rangle A\langle\cdot\varphi \ \mid \ \langle\cdot A\cdot\rangle\psi \\ \psi &::= \ \varphi \ \mid \ \neg\psi \ \mid \ \psi \vee \psi \ \mid \ \bigcirc\varphi \ \mid \ \varphi\,\mathcal{U}\,\varphi\end{aligned}$$

The language is enumerable. To see this, verify that there are countably many coalitions, i.e. finite subsets of a countable set.

The remaining Boolean operators $\wedge$, $\rightarrow$ and $\leftrightarrow$ as well as the logical constants $\top$ and $\bot$ can be defined as usual in terms of the operators given. The linear temporal logic operators 'sometime' and 'forever' can be defined as path formulas $\Diamond\varphi = (\top\,\mathcal{U}\,\varphi)$ and $\Box\varphi = \neg(\top\,\mathcal{U}\,\neg\varphi)$. Informally, the formula $\langle\cdot A\cdot\rangle\psi$ states that A has a strategy to ensure the temporal property ψ. The modality $\langle\cdot A\cdot\rangle$ *commits* the members of A to their selected strategy, while the operator $\cdot\rangle A\langle\cdot$ *releases* this commitment.

The language of ATL_{sc} consists only of some formulas from ATL^*_{sc}. The syntax of the path formulas ψ is restricted as follows (where φ refers to the state formulas in Def. 1):

$$\psi ::= \neg\psi \ \mid \ \bigcirc\varphi \ \mid \ \varphi\,\mathcal{U}\,\varphi$$

Notice that $\Box\varphi$ is still definable in ATL_{sc} as this grammar allows for path formulas of the form $\neg(\varphi_1 \,\mathcal{U}\, \varphi_2)$. In contrast, the syntax of ATL [4] is restricted to not allow the application of negation to the next-time and until operator. It does not matter in the case of $\neg\bigcirc\varphi$, because the negation can be pushed inside the next-time operator yielding the equivalent ATL path formula $\bigcirc\neg\varphi$. But it does matter for path formulas of the form $\neg(\varphi_1 \,\mathcal{U}\, \varphi_2)$. Their absence in the ATL syntax is compensated by including the $\Box$ operator explicitly. However, the compensation is only partial, because the dual of until[2] cannot be expressed in ATL, cf. [17].

ATL has been defined using Alternating Transition Systems (ATSs) [2, 3] and Concurrent Game Structures (CGSs) [4]. It is readily seen that both types of structures can be used interchangeably for logics that do no address the names for moves in the object language. In the case of ATL, this has been shown in [13]. In terms of computational complexity of model checking, however, it makes a difference when we use ATSs or CGSs as was studied in [16, 17].

In this paper, we evaluate formulas on Concurrent Game Structures (CGSs), which are defined as follows.

Definition 2 (Concurrent Game Structure). *Let $\Sigma = \{1, \ldots, n\} \subset \mathbf{\Sigma}$, with $n \geq 1$, be a finite set of agents, and $\Pi \subset \mathbf{\Pi}$ be a finite set of atomic propositions. A* Concurrent Game Structure *(CGS) $\mathcal{C}$ for $\langle \Sigma, \Pi \rangle$ is a tuple $\mathcal{C} = \langle W, V, \Sigma, M, Mov, E \rangle$, where:*

- *W is a finite, non-empty set of worlds (or game positions);*
- *$V : W \to 2^{\Pi}$ is a valuation function;*
- *M is a finite, non-empty set of moves;*
- *$Mov : W \times \Sigma \to 2^M \setminus \emptyset$ specifies for every world w and agent a a set $Mov(w, a)$ of moves available to a at w;*
- *$E : W \times M^{\Sigma} \to W$ is a transition function mapping a world w and a move profile $\boldsymbol{m} = \langle m_1, \ldots, m_n \rangle$ (one move for each agent) to the world $E(w, \boldsymbol{m})$.*

Let $\mathcal{C}$ be a CGS. The component Mov determines which of the moves from M are available for an agent at a world w. Let $\mathsf{prof}(w)$ be the set of available move profiles at world w, i.e.,

$$\mathsf{prof}(w) = \{\langle m_1, \ldots, m_n \rangle \mid m_i \in Mov(w, i)\}.$$

A move profile is used to determine a successor of a world using the transition function E. Let $\mathsf{succ}(w)$ be the set of possible successors of w, formally

$$\mathsf{succ}(w) = \{E(w, \boldsymbol{m}) \mid \boldsymbol{m} \in \mathsf{prof}(w)\}.$$

An infinite sequence $\lambda = x_0 x_1 x_2 \cdots \in W^{\omega}$ of worlds is called a *play* or *computation* if $x_{i+1} \in \mathsf{succ}(x_i)$ for all positions $i \geq 0$. Denote with $\lambda[i]$ the i-th component x_i in λ, and with $\lambda[0, i]$ the initial sequence $x_0 \cdots x_i$ of λ.

---

[2] The dual of the temporal logic operator until $\mathcal{U}$ is called *release* $\mathcal{R}$, and it is defined as $(\varphi \mathcal{R} \psi) \stackrel{\text{def}}{=} (\neg\varphi \,\mathcal{U}\, \neg\psi)$. In LTL we have the equivalence $(\varphi \mathcal{R} \psi) \equiv \Box\psi \vee (\psi \,\mathcal{U}\, (\varphi \wedge \psi))$. In CTL, it is defined as $(\varphi \mathcal{R} \psi) \stackrel{\text{def}}{=} E\Box\psi \vee E(\psi \,\mathcal{U}\, (\varphi \wedge \psi))$.

A *strategy for an agent* $a \in \Sigma$ is a function f_a that maps a world w from W to a move profile $f_a(w) \in Mov(w, a)$ available to a at w. A *strategy for a coalition* $A \subseteq \Sigma$ is a set F_A of strategies with $F_A = \{\sigma_a \mid a \in A\}$ containing one strategy for every agent in A. We refer to a strategy also as *strategy context*. We denote with $\mathsf{strat}(A)$ the set of strategies available to coalition A. The strategies considered here are *memoryless* as they are functions from worlds to move profiles and, thus, do not take previously visited states into account.

We define two operations on strategies: upgrade and release of strategies. Let F_A and F be strategies for sets of agents, where F_A contains strategies for the agents in A. The *upgrade* of F with the strategies in F_A is the result of *overwriting* F with strategies for the agents in $A \cap \mathsf{dom}(F)$ and *supplementing* F with strategies for agents for which F does not already provide a strategy (i.e., for agents in $A \setminus \mathsf{dom}(F)$). We will use $\circ$ as a strategy upgrade operator. Formally,

$$F_A \circ F = F_A \cup \{f_a \in F \mid a \notin A\}.$$

The *release* of the strategies for the agents in B from F is the *restriction* of F to strategies for agents that do not occur in B (i.e., for agents in $\Sigma \setminus B$). Formally, for $C = \Sigma \setminus B$,

$$F|_C = \{f_a \in F \mid a \in C\}.$$

The set $\mathsf{out}(w, F_A)$ of *outcomes* of a strategy F_A for the agents in A starting at a world w is the set of all plays $\lambda = x_0x_1x_2 \cdots \in W^\omega$ such that $x_0 = w$ and, for every $i \geq 0$, there is a move profile $\boldsymbol{m} = \langle m_1, \ldots, m_n \rangle \in \mathsf{prof}(x_i)$ such that

(i) $m_a = f_a(x_i)$, for all $a \in A$, and
(ii) $x_{i+1} = E(x_i, \boldsymbol{m})$.

The semantics of ATL^*_{sc} over CGSs is given as follows, where state formulas are evaluated at worlds (or game positions) and path formulas over infinite paths in a CGS.

Definition 3 (ATL^*_{sc} Semantics). *Given a CGS $\mathcal{C} = \langle W, R, V, \Sigma, M, Mov, E \rangle$ for $\langle \Sigma, \Pi \rangle$ and a strategy context F, the consequence relation $\models$ is inductively defined as follows.*

- $\mathcal{C}, w \models_F p$ *iff* $p \in V(w)$, *for all atomic propositions* $p \in \Pi$;
- $\mathcal{C}, w \models_F \neg\varphi$ *iff* $\mathcal{C}, w \not\models_F \varphi$;
- $\mathcal{C}, w \models_F \varphi_1 \vee \varphi_2$ *iff* $\mathcal{C}, w \models_F \varphi_1$ *or* $\mathcal{C}, w \models_F \varphi_2$;
- $\mathcal{C}, w \models_F \rangle\!\cdot A \cdot\!\langle \varphi$ *iff* $\mathcal{C}, w \models_S \varphi$, *where* $S = F|_{\Sigma \setminus A}$;
- $\mathcal{C}, w \models_F \langle\!\cdot A \cdot\!\rangle \psi$ *iff there is* $F_A \in \mathsf{strat}(A)$ *such that for all plays* $\lambda \in \mathsf{out}(w, S)$, *it holds that* $\mathcal{C}, \lambda \models_S \psi$, *where* $S = F_A \circ F$;
- $\mathcal{C}, \lambda \models_F \varphi$ *iff* $\mathcal{C}, \lambda[0] \models_F \varphi$, *when* φ *is a state formula;*
- $\mathcal{C}, \lambda \models_F \neg\psi$ *iff* $\mathcal{C}, \lambda \not\models_F \psi$;
- $\mathcal{C}, \lambda \models_F \psi_1 \vee \psi_2$ *iff* $\mathcal{C}, \lambda \models_F \psi_1 \vee \psi_2$;
- $\mathcal{C}, \lambda \models_F \bigcirc\varphi$ *iff* $\mathcal{C}, \lambda[1] \models_F \varphi$;
- $\mathcal{C}, \lambda \models_F (\varphi_1 \mathcal{U} \varphi_2)$ *iff there is an* $i \geq 0$ *such that* $\mathcal{C}, \lambda[i] \models_F \varphi_2$ *and* $\mathcal{C}, \lambda[j] \models_F \varphi_1$ *for all* j *with* $0 \leq j < i$.

A formula φ is satisfiable *if $\mathcal{C}, w \models_F \varphi$ for some CGS $\mathcal{C}$, some world w in $\mathcal{C}$ and some strategy context F in $\mathcal{C}$; a formula is called* valid *if $\mathcal{C}, w \models_F \varphi$ for all $\mathcal{C}$, all w and all F.*

In this paper, we do not assume agents being capable of perfect recall. In fact, we use a semantics for ATL_{sc} and ATL^*_{sc} that is based on memoryless strategies. This means that agents use strategies that prescribe for every world which move to take. The history of previously visited worlds is not taken into account. This differs from the original definition in [11] that introduces the logics with a perfect recall semantics.

The language seems rather rich. Sometimes, different formulations of the same simple property will seem natural. We shall illustrate this in the next section by defining a few modalities that the community of logics in MAS has become familiar with.

3 Common modalities of agency

We now turn to the definition in the object language of ATL_{sc} and ATL^*_{sc} of a few modalities often discussed in the literature: $\mathsf{ATL}^{(*)}$ modality of ability (Section 3.2) and the modality of *strategic* actual agency (Section 3.3).[3] In order to express those modalities in the language of ATL^*_{sc}, it requires to write formulas referring explicilty to all agents. For this purpose we have to consider a fragment of the language defined in the previous section containing only a finite number of agents. We leave for future work a study of properties that can be expressed in the full language or variants of it.

3.1 Whatever A do

Brihaye *et al.* [6] define a modality that is going to be useful later:

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

They provide the reading: "for any strategy of coalition A, *every* run in the corresponding outcome satisfies a formula ψ". Notice that it is defined in the language of ATL^*_{sc} and not in the one of ATL_{sc}. Its semantics is:

– $\mathcal{C}, w \models_{F_B} [\cdot A \cdot]\psi$ iff $\forall S_A \in \mathsf{strat}(A), \forall \lambda \in \mathsf{out}(w, S_A \circ F_B)$. $\mathcal{C}, \lambda \models_{S_A \circ F_B} \psi$

The modality $[\cdot A \cdot]$ is *not* the dual of $\langle\cdot A \cdot\rangle$. It is also important to observe that the truth of the path formula ψ is in the context of $S_A \circ F_B$. A more precise reading of $[\cdot A \cdot]\psi$ is therefore: "for any strategy of coalition A, *every* run in the corresponding outcome satisfies a formula ψ *in the current strategy context updated by the new strategies of A*."

[3] 'Strategic' is not to be understood in the sense of game theory, where agents strategize to maximize their utility. It is to be opposed to actual agency that considers *only the current move*. Strategic actual agency is a property of agents or coalitions currently doing something by planning *more than one move ahead*.

3.2 Simulating the $\mathsf{ATL}^{(*)}$ path quantifier

Brihaye *et al.* [7, 11], propose to simulate the $\mathsf{ATL}^{(*)}$ formula $\langle\langle A\rangle\rangle\psi$ as follows:

$$\langle\langle A\rangle\rangle^1\psi \stackrel{\text{def}}{=} \rangle\Sigma\langle\cdot\langle\cdot A\cdot\rangle\psi.$$

That is, one first releases the current strategies of all agents, then we find a strategy for A that only yields runs that satisfy ψ. Its truth condition is:

– $\mathcal{C}, w \models_{F_B} \langle\langle A\rangle\rangle^1\psi$ iff $\exists S_A \in \mathsf{strat}(A), \forall\lambda \in \mathsf{out}(w, S_A).\ \mathcal{C}, \lambda \models_{S_A} \psi$

Notice that ψ must hold on each elected run, in the context of the current strategy of the members of A.

When the signature contains a finite set Σ of agents, $\langle\langle\Sigma\rangle\rangle$ and $\langle\langle\emptyset\rangle\rangle$ are dual: we have that $\langle\langle\Sigma\rangle\rangle\psi \leftrightarrow \neg\langle\langle\emptyset\rangle\rangle\neg\psi$ is a valid schema in ATL^*. Now, take the ATL^*_{sc} path formula $\Psi = \bigcirc[\cdot b\cdot]\bigcirc p$. We can see that $\langle\langle\Sigma\rangle\rangle^1\Psi \rightarrow \neg\langle\langle\emptyset\rangle\rangle^1\neg\Psi$ is not an ATL^*_{sc}-validity. It is falsified at the world w_0 of the model in Fig. 1.

We have that $\langle\langle\Sigma\rangle\rangle^1\bigcirc p \rightarrow \neg\langle\langle\emptyset\rangle\rangle^1\bigcirc\neg p$, with p a propositional variable from $\mathbf{\Pi}$ is indeed a valid formula in ATL^*_{sc}. But we have just seen that the uniform substitution of p with $[\cdot b\cdot]\bigcirc p$ yields a non-validity of ATL^*_{sc}. It means that ATL^*_{sc} does not obey the rule of uniform substitution.

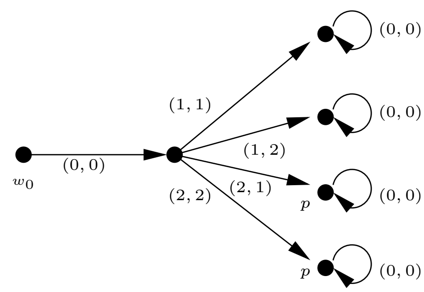

Fig. 1. A CGS for two agents.

A translation tr from the language of ATL^* into the language of ATL^*_{sc} such that $tr(\langle\langle A\rangle\rangle\psi) \stackrel{\text{def}}{=} \langle\langle A\rangle\rangle^1 tr(\psi)$ and homomorphic for the propositional connectives is indeed satisfiability preserving. But the definition does not interact well with the richer language of ATL^*_{sc}. A more fitting definition of the $\mathsf{ATL}^{(*)}$ modality would be:

$$\langle\langle A\rangle\rangle^2\psi \stackrel{\text{def}}{=} \rangle\Sigma\langle\cdot\langle\cdot A\cdot\rangle\rangle\Sigma\langle\cdot\psi.$$

That is, one first releases the current strategies of all agents, then we find a strategy for A, and one finally releases again all the current strategies to evaluate the path formula ψ. Its truth condition is:

– $\mathcal{C}, w \models_{F_B} \langle\langle A\rangle\rangle^2\psi$ iff $\exists S_A \in \mathsf{strat}(A), \forall\lambda \in \mathsf{out}(w, S_A).\ \mathcal{C}, \lambda \models_{\emptyset} \psi$

This seems more adequate with the notion of non-committed ability that we are familiar in $\mathsf{ATL}^{(*)}$. At least we regain duality in the sense that $\langle\langle\Sigma\rangle\rangle^2\bigcirc\varphi \leftrightarrow \neg\langle\langle\emptyset\rangle\rangle^2\bigcirc\neg\varphi$ is a valid axiom schema in ATL_{sc} and $\langle\langle\Sigma\rangle\rangle^2\psi \leftrightarrow \neg\langle\langle\emptyset\rangle\rangle^2\neg\psi$ is a valid axiom schema in ATL^*_{sc}.

In ATL^*_{sc}, there is at least one more way to capture the $\mathsf{ATL}^{(*)}$ path quantifier. It is sometimes interpreted as "coalition A has a strategy to enforce ψ whatever the choices of the other agents." This is actually the reading given in [7, p. 97].

It would then seem natural to express it as

$$\langle\langle A\rangle\rangle^3\psi \stackrel{\text{def}}{=} \langle\cdot A\cdot\rangle[\cdot\Sigma\setminus A\cdot]\psi.$$

(This definition does not fall into the language of ATL_{sc}, but of ATL^*_{sc}.) Its truth condition is:

– $\mathcal{C}, w \models_{F_B} \langle\langle A\rangle\rangle^3\psi$ iff $\exists S_A \in \mathsf{strat}(A), \forall S_{\overline{A}} \in \mathsf{strat}(\Sigma\setminus A)$, $\forall\lambda \in \mathsf{out}(w, S_{\overline{A}} \circ S_A)$. $\mathcal{C}, \lambda \models_{S_{\overline{A}}\circ S_A} \psi$

The path formula ψ is then evaluated with respect to a complete context of strategies, one for each member of the counter-coalition.

To conclude, we have now three sensible notions of ATL-like ability:

$$\begin{array}{ll}
\mathcal{C}, w \models_{F_B} \langle\cdot A\cdot\rangle[\cdot\Sigma\setminus A\cdot]\psi & \psi \text{ eval. wrt. a } \Sigma\text{-commitment}\\
\mathcal{C}, w \models_{F_B} \cdot\rangle\Sigma\langle\cdot\langle\cdot A\cdot\rangle\psi & \psi \text{ eval. wrt. an } A\text{-commitment}\\
\mathcal{C}, w \models_{F_B} \cdot\rangle\Sigma\langle\cdot\langle\cdot A\cdot\rangle\cdot\rangle\Sigma\langle\cdot\psi & \psi \text{ eval. wrt. an } \emptyset\text{-commitment}
\end{array}$$

The successive definitions involve an ever decreasing commitment for the evaluation of the path formula in its scope. Interestingly however, their sets of outcomes are identical, and correspond to these sets of runs that a coalition can enforce (in the sense of $\mathsf{ATL}^{(*)}$). They are distinct in ATL^*_{sc} because of the different commitments, but all are sufficient for an embedding of $\mathsf{ATL}^{(*)}$.

3.3 Strategic actual agency

The modality of actual agency has been widely studied, and is most prominently known for its treatment in the STIT theories (STIT for "seeing to it that"). It is a large family of logics with each their own semantics and modalities [5, 15]. Nonetheless, they all share an Ockhamist view of time [19]. Formulas are evaluated in tree models, with respect to a state and a play. The most basic modality is the Chellas' STIT operator (somewhat a misnomer) of actual (one-shot strategy) agency, proposed by Horty. Integrated in discrete time it allows to embed Coalition Logic ([9]).

A challenge in formal philosophy of action is to devise a modality similar to the Chellas' STIT but for long-term strategies. There is a truth-value gap of strategic statements, analogous to the truth-value gap for future-tense statements addressed, e.g., in [20] and [19]. In a nutshell, a state and a play are not enough to evaluate a statement reading that the coalition A strategically see to it that ψ is true. This is because in general, the context of only a state and a play does not determine a unique strategy of an arbitrary coalition. See [15, p. 149] for an illustration. (In a CGS, a play does determine a unique strategy for the coalition Σ, though.)

Horty observes that two lines of resolution are possible [15, Sec. 7.2]. The Peircean-like one is to consider all strategies that could determine the current play. The Ockhamist-like one, that we adopt here, is that a modality of strategic actual agency should additionally be evaluated wrt. to a strategy context. (This has been investigated by Müller [18] for the individual case and by Broersen *et al.* [8] for the case of coalitions.)

We can say here that a coalition A see to it that ψ in a context F_B iff the strategies of A in F_B are enough to make all the plays satisfy ψ. The truth value of such a modality would then be:

– $\mathcal{C}, w \models_{F_B} [A\ \mathsf{sstit}^1]\psi$ iff $\forall \lambda \in \mathsf{out}(w, F_B|_A).\ \mathcal{C}, \lambda \models_{F_B|_A} \psi$

In fact, the modality $[A\ \mathsf{sstit}^1]$ can readily be captured in the language of ATL_{sc} as follows:

$$[A\ \mathsf{sstit}^1]\psi \stackrel{\text{def}}{=} \cdot\rangle \Sigma \setminus A \langle \cdot \langle \cdot \emptyset \cdot \rangle \psi.$$

The notion of strategic actual agency captured by $[A\ \mathsf{sstit}^1]$ is the one that appears the most immediate in the CGSs with contexts. It does capture perfectly that the current strategy of a coalition ensures that something happens, *independently of the commitment of the counter-coalition*, and *independently of the currently non-committed members of A*. We postpone for future research a thorough comparison, but this will turn out rather different from the solutions in the more traditional STIT literature, e.g., the proposal in [8]. A striking difference is that so far we did not feel compelled to explicitly put into the semantics of actual agency the fact that a coalition see to something *whatever the other agents do*. It might be a blunt conceptual error. But like in the simulation of the $\mathsf{ATL}^{(*)}$ path quantifier in Section 3.2, it might as well reveal interesting differences on the assumptions about agents' commitment to strategies between the two frameworks.

Trying to emulate *whatever the other agents do*, we can employ the modality $[\cdot A \cdot]$ defined in Section 3.1. A direct translation of "the coalition A see to it that ψ whatever the other agents do" would then be:

$$[A\ \mathsf{sstit}^2]\psi \stackrel{\text{def}}{=} [A\ \mathsf{sstit}^1][\cdot \Sigma \setminus A \cdot]\psi.$$

It is clearly equivalent to $[\cdot \Sigma \setminus A \cdot]\psi$. We would then have:

– $\mathcal{C}, w \models_{F_B} [A\ \mathsf{sstit}^2]\psi$ iff $\forall S_{\overline{A}} \in \mathsf{strat}(\overline{A}), \forall \lambda \in \mathsf{out}(w, S_{\overline{A}} \circ F_B).\ \mathcal{C}, \lambda \models_{S_{\overline{A}} \circ F_B} \psi$

The modalities $[A\ \mathsf{sstit}^1]$ and $[A\ \mathsf{sstit}^2]$ are nevertheless very significantly different in that the evaluation of the path formula in the scope of $[A\ \mathsf{sstit}^2]$ is within the context of a commitment of the counter-coalition. (The evaluation of the path formula is still independent from the currently non-committed members of A.)

Variants of these modalities can be defined semantically, where instead of being *independent* of the strategies of the non-committed members of A, they reflect a type of actual agency that remains true *in whatever context of strategies for the non-committed members of A.*

– $\mathcal{C}, w \models_{F_B} [A\ \mathsf{sstit}^3]\psi$ iff $\mathcal{C}, w \models_{F_B} [A\ \mathsf{sstit}^1][\cdot A \setminus B \cdot]\psi$
– $\mathcal{C}, w \models_{F_B} [A\ \mathsf{sstit}^4]\psi$ iff $\mathcal{C}, w \models_{F_B} [A\ \mathsf{sstit}^2][\cdot A \setminus B \cdot]\psi$

Their truth condition is straightforward. However, note that there is no syntactic reference in $[A\ \mathsf{sstit}^3]$ and $[A\ \mathsf{sstit}^4]$ to the committed agents B. Hence, it is doubtful that they are expressible syntactically in the language of ATL_{sc} or ATL^*_{sc} as they require some sort of reification of who are the committed agents in the context.

To sum up, when giving an interpretation to a formula representing a statement about actual agency in the context of a long-term strategy, we are offered again more than one distinct possibility, depending on what commitments we wish to consider when evaluating the path formulas.

$$\begin{array}{ll} \mathcal{C}, w \models_{F_B} \cdot\rangle \Sigma \setminus A \langle\cdot \langle\cdot \emptyset \cdot\rangle \psi & \psi \text{ eval. wrt. a } (B \cap A)\text{-commitment} \\ \mathcal{C}, w \models_{F_B} [\cdot \Sigma \setminus A \cdot] \psi & \psi \text{ eval. wrt. a } (B \cup (\Sigma \setminus A))\text{-commitment} \\ \mathcal{C}, w \models_{F_B} \cdot\rangle \Sigma \setminus A \langle\cdot \langle\cdot \emptyset \cdot\rangle [\cdot A \setminus B \cdot] \psi & \psi \text{ eval. wrt. an } A\text{-commitment} \\ \mathcal{C}, w \models_{F_B} [\cdot \Sigma \setminus A \cdot][\cdot A \setminus B \cdot] \psi & \psi \text{ eval. wrt. a } \Sigma\text{-commitment} \end{array}$$

Of course, we did not exhaust the seemingly sensible characterizations of a modality of strategic actual agency that can be directly expressed in the language of ATL^*_{sc}. One could also have the very simple variants where we release the commitment of all the agents when we evaluate the path formula. It is readily seen that for any $1 \leq i, j \leq 4$, we have that $[A \ \mathsf{sstit}^{\mathsf{i}}] \cdot\rangle \Sigma \langle\cdot \psi \leftrightarrow [A \ \mathsf{sstit}^{\mathsf{j}}] \cdot\rangle \Sigma \langle\cdot \psi$ is valid.

4 Strategy contexts and explicit strategies

We now turn to the second contribution of the paper. Here we contrast the notion of strategy contexts with explicit strategies. We first present ATLES, the extension of ATL with *explicit strategies* from [22] (Section 4.1), and then we translate a fragment of ATL_{sc} into ATLES (Section 4.2).

4.1 ATLES

The language of ATL is enriched with symbols for strategies and commitment functions that assign agents to strategies they are committed to play. Thus ATLES allows to reason explicitly about strategies. This is not possible with any of ATL and ATL_{sc} (and their respective LTL-extensions) as strategies are pure semantic constructs and they do not occur in the object language.

Formally, the signature of the language is extended by a set $\mathbf{\Upsilon}$ of *strategy terms*, where $\mathbf{\Upsilon} = \bigcup_{a \in \mathbf{\Sigma}} \mathbf{\Upsilon}_a$ and $\mathbf{\Upsilon}_a$ is a countable infinite set of *strategy terms* σ_a for agent a in $\mathbf{\Sigma}$. A *commitment function* is a partial function $\rho : \mathbf{\Sigma} \to \mathbf{\Upsilon}$ with a finite domain mapping an agent $a \in \mathbf{\Sigma}$ to a strategy term $\rho(a) \in \mathbf{\Upsilon}_a$ for a. Note that a commitment function ρ is a finite object and as such it is used to additionally parameterise path-quantifiers as $\langle\!\langle A \rangle\!\rangle_\rho$. The set $\mathsf{dom}(\rho)$ consists of the committed agents. If $\rho(a)$ is defined, then ρ contains a mapping of the form $a \mapsto \sigma_a$ which is called a *commitment* of agent a (or *a commits*) to play the strategy denoted by the strategy term σ_a. On the other hand, if $\rho(a)$ is undefined, then a does not commit to any strategy and, thus, a can quantify freely over the strategies available to a. The reading of an ATL-path quantifier $\langle\!\langle A \rangle\!\rangle$ with commitment function ρ is as follows:

> $\langle\!\langle A \rangle\!\rangle_\rho \varphi$ states that, given the commitment of any agent b in $\mathsf{dom}(\rho)$ to use the strategy denoted by $\rho(b)$, the agents in $A \setminus \mathsf{dom}(\rho)$ have a strategy to ensure the temporal property φ, no matter what the agents in $\Sigma \setminus (\mathsf{dom}(\rho) \cup A)$ do.

Notice that the committed agents in $\mathsf{dom}(\rho)$ do not take part in the quantification over strategies in $\langle\!\langle A \rangle\!\rangle_\rho$.

We remark that $\langle\!\langle A\rangle\!\rangle_\rho$ is not how the path quantifier really looks like when used in a formula. The symbol ρ is merely a meta-logical reference to an actual commitment function, which is a collection of mappings of the form $a \mapsto \sigma_a$, where σ_a is a strategy term for agent a. This should be considered when analysing the length of a formula. For instance, take $\rho = \{a \mapsto \sigma_a, b \mapsto \sigma_b\}$. Then we write $\langle\!\langle A\rangle\!\rangle_\rho$ for convenience in order to abbreviate $\langle\!\langle A\rangle\!\rangle_{\{a\mapsto\sigma_a, b\mapsto\sigma_b\}}$. For modelling purposes, one may modify the syntax and write $\langle\!\langle A : a \mapsto \sigma_a, b \mapsto \sigma_b\rangle\!\rangle$ instead.

The notion of commitment to strategies requires the same strategies to be played again later stage. This means, in formulas of the form $\langle\!\langle A\rangle\!\rangle_\rho \Psi$, the same commitment $a \mapsto \sigma_a$ from ρ occurs in a commitment function ξ of a nested path quantifier $\langle\!\langle B\rangle\!\rangle_\xi$ in Ψ. That is, both, ρ and ξ, prescribe the strategy term σ_a for agent a (or, in both cases, a commits to σ_a). We have that $\rho(a) = \xi(a)$. Release of commitment to σ_a is modelled as easy as committing to it in the first place. This is achieved by having a commitment function χ of a nested path quantifier not include the commitment $a \mapsto \sigma_a$, i.e., either $\chi(a) \neq \sigma_a$ or χ is undefined for a. In case release of commitment is not desired, the notion of irrevocable strategies is used. It can be modelled explicitly in ATLES by only allowing commitment functions ρ to extend conservatively the commitment functions ξ under whose range they occur, i.e., ρ and ξ agree for all agents in $\mathsf{dom}(\xi)$. Thus, IATL can be defined in ATLES while avoiding the update semantics employed in [1].

The language of ATLES is defined over the extended signature $\langle \mathbf{\Pi}, \mathbf{\Sigma}, \mathbf{\Upsilon}\rangle$.

Definition 4 (ATLES **Syntax).** *The following grammar defines state formulas φ and path formulas ψ, where p ranges over $\mathbf{\Pi}$, A ranges over finite subsets of $\mathbf{\Sigma}$ and ρ over commitment functions. The language of* ATLES *consists of state formulas.*

$$\begin{aligned} \varphi &::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle\!\langle A\rangle\!\rangle_\rho \psi \\ \psi &::= \bigcirc\varphi \mid \Box\varphi \mid \varphi\,\mathcal{U}\,\varphi \end{aligned}$$

The language of ATLES can easily be extended to allow for negation of the temporal operators next-time and until. But we refrain from extending the syntax in this paper as we use the established complexity result of the satisfiability problem for ATLES from [22] in order to use ATLES to determine a decidable fragment of ATL_{sc} whose satisfiability can be solved in ExpTime.

We define ATLES using concurrent game structures which differs from its original definition in [22]. The logic was introduced using alternating transition systems from [4] extended with strategy terms and a denotation function mapping the strategy terms to actual strategies. Another suitable extension of alternating transition systems was introduced in [21], so-called action-based alternating transition systems, which explicitly accounts for actions and action pre-conditions. It can readily be seen that these transition systems give rise to the same logics.

Strategy terms in $\mathbf{\Upsilon}$ are interpreted as strategies in a CGS via assignments. An *assignment* $\mathfrak{a}$ in $\mathcal{C}$ is a function mapping strategy terms σ_a in $\mathbf{\Upsilon}_a$ for any agent a in $\mathbf{\Sigma}$ to a strategy $\mathfrak{a}(\sigma_a)$ in $\mathsf{strat}(a)$ for a in $\mathcal{C}$. Note that the assignment $\mathfrak{a}$ in a CGS acts like an assignment in First-order Logic with the difference that in ATLES strategy terms are mapped to actual strategies in the CGS instead of domain elements as in FOL. In [22] an assignment is called denotation function, which comes as a component of an ATS.

To define the semantics of ATLES, we use the notions of a strategy and outcome as in Section 2. We lift the notion of assignment to commitment functions as follows. The application of an assignment $\mathfrak{a}$ to a commitment function ρ is the set $\mathfrak{a}(\rho)$ of strategies for the agents in $\mathsf{dom}(\rho)$. Formally,

$$\mathfrak{a}(\rho) = \{f_a \in \mathsf{strat}(a) \mid f_a = \mathfrak{a}(\rho(a)), a \in \mathsf{dom}(\rho)\}.$$

It is readily checked that $\mathfrak{a}(\rho)$ is indeed a set of strategies, one for each agent in $\mathsf{dom}(\rho)$. To see this, recall that ρ is functional, i.e., it yields exactly one strategy term $\rho(a)$ for every agent for which ρ is defined.

An assignment $\mathfrak{a}$ acts as an interpretation of the commitment function ρ (i.e. the strategy terms in ρ). We can view a strategy term $\sigma_a = \rho(a)$, for any a in $\mathsf{dom}(\rho)$, as a constant rather than a variable. As we will see below in the semantics of ATLES, the assignment $\mathfrak{a}$ does not change during the evaluation of a formula and, thus, the strategy $\mathfrak{a}(\sigma_a)$ is fixed. We can think of the strategy term σ_a as being existentially quantified in the sense that there exists a strategy for a that is referenced by σ_a and provided by $\mathfrak{a}$. ATLES does not provide references to universally quantified strategies.

Using the notion of assignments, we can now define how to interpret the formulas of ATLES over CGSs.

Definition 5 **(**ATLES **Semantics).** *Given a CGS* $\mathcal{C} = \langle W, R, V, \Sigma, M, Mov, E\rangle$ *for* $\langle \Sigma, \Pi\rangle$ *and an assignment* $\mathfrak{a}$*, the consequence relation* $\models$ *is inductively defined as follows. The notions of validity and satisfiability are defined as usual.*

- $\mathcal{C}, w \models^{\mathfrak{a}} p$ *iff* $w \in V(p)$*, for all atomic propositions* $p \in \Pi$*;*
- $\mathcal{C}, w \models^{\mathfrak{a}} \neg\varphi$ *iff* $\mathcal{C}, w \not\models^{\mathfrak{a}} \varphi$*;*
- $\mathcal{C}, w \models^{\mathfrak{a}} \varphi_1 \vee \varphi_2$ *iff* $\mathcal{C}, w \models^{\mathfrak{a}} \varphi_1$ *or* $\mathcal{C}, w \models^{\mathfrak{a}} \varphi_2$*;*
- $\mathcal{C}, w \models^{\mathfrak{a}} \langle\!\langle A\rangle\!\rangle_\rho \psi$ *iff there is a strategy* F_A *in* $\mathsf{strat}(A)$ *such that for all plays* $\lambda \in \mathsf{out}(w, S)$*, it holds that* $\mathcal{C}, \lambda \models^{\mathfrak{a}} \psi$*, where* $S = \mathfrak{a}(\rho) \circ F_A$*;*
- $\mathcal{C}, \lambda \models^{\mathfrak{a}} \bigcirc\varphi$ *iff* $\mathcal{C}, \lambda[1] \models^{\mathfrak{a}} \varphi$*;*
- $\mathcal{C}, \lambda \models^{\mathfrak{a}} \Box\varphi$ *iff* $\mathcal{C}, \lambda[i] \models^{\mathfrak{a}} \varphi$ *for all positions* $i \geq 0$*;*
- $\mathcal{C}, \lambda \models^{\mathfrak{a}} (\varphi_1 \mathcal{U} \varphi_2)$ *iff there is an* $i \geq 0$ *such that* $\mathcal{C}, a, \lambda[i] \models^{\mathfrak{a}} \varphi_2$ *and* $\mathcal{C}, \lambda[j] \models^{\mathfrak{a}} \varphi_1$ *for all positions* j *with* $0 \leq j < i$*.*

The ATLES semantics of $\langle\!\langle A\rangle\!\rangle_\rho$ is similar to the semantics of $\langle\cdot A\cdot\rangle$ in ATL_{sc}, which facilitates comparison. We recall that the operator $\circ$ from Section 2 yields $\mathfrak{a}(\rho) \circ F_A = \mathfrak{a}(\rho) \cup \{f_a \in F_A \mid a \notin \mathsf{dom}(\rho)\}$. Intuitively, $\mathfrak{a}(\rho) \circ F_A$ states that commitments of agents are respected as prescribed in ρ, all other agents in A play their just selected strategies.

4.2 Comparing ATL_{sc} and ATLES

Obvious differences between ATL_{sc} and ATLES are that, while the former includes a separate release modality $\cdot\rangle A\langle\cdot$ and strategy contexts in the semantics, the latter allows for commitments of the form $a \mapsto \sigma_a$ in the syntax which are interpreted using assignments. However, commitments and assignments play the role of strategy contexts in ATL_{sc}. A crucial difference between the logics is the semantics of the path quantifiers $\langle\!\cdot A \cdot\!\rangle$ and $\langle\!\langle A \rangle\!\rangle_\rho$; cf. Def. 3 and Def. 5, respectively. For $\langle\!\cdot A \cdot\!\rangle$, the strategies F_A selected by A upgrade (overwrite) the strategy context F_{context}, whereas, for $\langle\!\langle A \rangle\!\rangle_\rho$, the strategy commitments $\mathfrak{a}(\rho)$ are supplemented by F_A. Consequently, due to how context or commitments are respected in $\langle\!\cdot A \cdot\!\rangle$ and $\langle\!\langle A \rangle\!\rangle_\rho$, different agents end up quantifying over strategies in general. The following proposition states under which conditions $\langle\!\cdot A \cdot\!\rangle$ and $\langle\!\langle A \rangle\!\rangle_\rho$ determine the same set $\mathsf{out}(x, S)$ of plays, where S is defined as $S = F_A \circ F_{\mathsf{context}}$ in the former case, and $S = \mathfrak{a}(\rho) \circ F_A$ in the latter.

Proposition 1. *It holds that $F_A \circ F_{\mathit{context}} = \mathfrak{a}(\rho) \circ F_A$ if one of the following conditions is satisfied:*

(i) $F_{\mathit{context}} = \mathfrak{a}(\rho) = \emptyset$;
(ii) $F_A = \emptyset$ *and* $F_{\mathit{context}} = \mathfrak{a}(\rho)$; *or*
(iii) $F_A = F_{\mathit{context}} = \mathfrak{a}(\rho)$.

The proposition can be shown by using the fact that the strategy upgrade operator $\circ$ forms an idempotent semigroup on the set strat of strategies, and that $\circ$ is not commutative.[4]

Proposition 1 makes clear that a strategy context F_{context} in ATL_{sc} corresponds to the strategy commitment $\mathfrak{a}(\rho)$ in ATLES with the difference that F_{context} is a purely semantic object, whereas $\mathfrak{a}(\rho)$ consists of a syntactic component ρ and a semantic component $\mathfrak{a}$. This means we can explicitly describe strategy contexts in the language of ATLES, whereas in ATL_{sc} we have to make use of $\langle\!\cdot A \cdot\!\rangle$ and $\cdot\rangle A \langle\cdot$ that describe that strategies for A are either pushed into the context or released from it. Notice how using strategy commitments in the syntax is more flexible than the strategy context model as every path quantifier in ATLES can be parameterised with a different commitment function, which describes explicitly which agent is using what strategy. In particular, this does not require a dedicated release operator.

The notion of *irrevocable strategies* is captured in ATL_{sc} by carefully avoiding quantification over strategies of committed agents. In ATLES, irrevocability can be made explicit in the syntax.

Once a strategy in the strategy context is overwritten with a new strategy or released, it cannot be recovered in ATL_{sc}, because any reference to it is lost. This could be described with the notion of *forgetting forever*. Not so in ATLES, where 'forgetting forever' can be modelled explicitly in the language, but it is no restriction of the logic as in ATL_{sc}. In fact, an agent in ATLES may *resume* a commitment after releasing it, which also captures a notion of agents having a *strategy memory*.

A strength of ATL_{sc} is to push *any* strategy that is available to an agent into the context. This is achieved with formulas of the form $\neg\langle\!\cdot A \cdot\!\rangle\psi$, where the agents

[4] The operation $\circ$ is a binary function on strat, it is associative (i.e., $(F_A \circ F_B) \circ F_C = F_A \circ (F_B \circ F_C)$), the empty strategy $\emptyset$ forms the identity element (i.e., $F \circ \emptyset = \emptyset \circ F = F$), and $\circ$ is idempotent (i.e., $F \circ F = F$).

in A quantify universally over their strategies F_A. In the semantics, before we continue with the evaluation of the path formula ψ, the strategies F_A are used to upgrade the strategy context (cf. Def. 3). This is another crucial difference to ATLES, which is restricted to existential quantification over commitments. To make more precise the relationship between ATL_{sc} and ATLES, we present an equivalence preserving mapping from a fragment of ATL_{sc} into ATLES. We define a translation $tr(\cdot,\cdot)$ as a partial function that maps an ATL_{sc}-formula, in which every occurrence of a path quantifier $\langle\!\cdot A\cdot\!\rangle$ is under the scope of an even number of negations, and a commitment function to formulas of ATLES as follows:

$$
\begin{aligned}
tr(p,\xi) &\stackrel{\text{def}}{=} p;\\
tr(\neg\varphi,\xi) &\stackrel{\text{def}}{=} \neg tr(\varphi,\xi);\\
tr(\varphi_1 \vee \varphi_2,\xi) &\stackrel{\text{def}}{=} tr(\varphi_1,\xi) \vee tr(\varphi_2,\xi);\\
tr(\cdot\rangle A\langle\cdot\varphi,\xi) &\stackrel{\text{def}}{=} tr(\varphi,\chi),\ \text{where } \chi = \xi|_{\Sigma\setminus A};\\
tr(\neg\cdot\rangle A\langle\cdot\neg\varphi,\xi) &\stackrel{\text{def}}{=} tr(\varphi,\chi),\ \text{where } \chi = \xi|_{\Sigma\setminus A};\\
tr(\langle\cdot A\cdot\rangle\bigcirc\varphi,\xi) &\stackrel{\text{def}}{=} \langle\!\langle A\rangle\!\rangle_\rho \bigcirc tr(\varphi,\rho);\\
tr(\langle\cdot A\cdot\rangle\Box\varphi,\xi) &\stackrel{\text{def}}{=} \langle\!\langle A\rangle\!\rangle_\rho \Box\, tr(\varphi,\rho);\\
tr(\langle\cdot A\cdot\rangle(\varphi_1\,\mathcal{U}\,\varphi_2),\xi) &\stackrel{\text{def}}{=} \langle\!\langle A\rangle\!\rangle_\rho (tr(\varphi_1,\rho)\,\mathcal{U}\,tr(\varphi_2,\rho));
\end{aligned}
$$

where the commitment function ρ overwrites/updates ξ at A with fresh strategy terms. Formally,

$$\rho = \xi|_{\mathsf{dom}(\xi)\setminus A} \cup \{a \mapsto \sigma_a \mid a \in A, \sigma_a \text{ is fresh}\}.$$

The following proposition states that $tr(\cdot,\cdot)$ is indeed equivalence preserving. The proof works by induction on the structure of ATL_{sc}-formulas that are translated.

Proposition 2. *Let φ be an ATL_{sc}-formula, $\mathcal{C}$ a CGS, x a world in $\mathcal{C}$ and F a strategy in $\mathcal{C}$. The following are equivalent:*

(a) $\mathcal{C}, x \models_F \varphi$*;*
(b) $\mathcal{C}, x \models^{\mathfrak{a}} tr(\varphi, \rho_F)$*, for some $\langle \rho_F, F\rangle$-compatible assignment $\mathfrak{a}$,*

*where $\rho_F = \{a \mapsto \sigma_a \mid f_a \in F, \sigma_a$ is fresh$\}$ and an assignment $\mathfrak{a}$ is $\langle\rho_F, F\rangle$-*compatible *if $\mathfrak{a}(\rho_F(a)) = f_a$, for every $a \in \mathsf{dom}(\rho_F)$ and $f_a \in F$.*

The ATL_{sc}-fragment determined by $tr(\cdot,\cdot)$ is the fragment that does not allow for universal quantification over strategy commitments. The latter is expressed by formulas of the form $\neg\langle\!\cdot A\cdot\!\rangle\psi$ or, in general, by the modality $\langle\!\cdot A\cdot\!\rangle$ under the scope of an odd number of negations. The satisfiability checking problem for this fragment can be solved in ExpTime by Proposition 2 and the fact that ATLES is in ExpTime [22]. This is in contrast with the complexity of full ATL_{sc}, which we establish in the following section.

5 Complexity

This section is devoted to investigate the computational complexity of ATL_{sc} and ATL^*_{sc} over *general* CGSs: we relax CGSs from Def. 2 by allowing infinite number of states and infinite number of moves.

Generally, high expressiveness tends to come with the price of high computational complexity of reasoning problems. While the model checking problem was already considered in [11, 7] (and shown to be between 2ExpTime-hard and nonelementary for ATL_{sc}, while it is 2ExpTime-complete for ATL^* [4]), we focus here on the satisfiability problem. Clearly, the lower complexity bounds carry over to ATL_{sc} and ATL^*_{sc} from their respective fragments ATL and ATL^*. It turns out, however, that extending ATL with strategy contexts comes with a much higher price. In the following we show that ATL_{sc} is undecidable. To show this, we use a reduction of the satisfiability problem for the product logic $\mathsf{S5}^n$, which is known to be undecidable. In Section 3, we demonstrated that ATL_{sc} can capture some notion of actual group agency (cf. operator $[A\ \mathsf{sstit}^1]$ in Section 3.3). Thus the undecidability of ATL_{sc} may not come as a surprise considering the undecidability of Chellas' STIT logic of group agency [14].

We obtain the following lower complexity bounds. It remains to be shown that Thm. 1 holds for finite CGSs (as defined in Def. 2), which amounts to showing that $\mathsf{S5}^n$ over finite frames is undecidable. We also leave the matching upper bounds as an open problem.

Theorem 1. *The satisfiability problem for* ATL_{sc} *(over general CGSs) is*

(i) NP-hard for formulas with $n = 1$ *agent;*
(ii) NExpTime-hard for formulas with $n = 2$ *agents; and*
(iii) undecidable for formuals with $n \geq 3$ *agents.*

The lower bounds in Theorem 1 can be shown by the following reduction of the satisfiability problem for $\mathsf{S5}^n$ to the problem for ATL_{sc}.[5] For a formal definition of $\mathsf{S5}^n$, we refer to, e.g., [12]. Define a translation $tr(\cdot)$ mapping $\mathsf{S5}^n$-formulas to formulas of ATL_{sc} as follows:

$$\begin{aligned}
tr(p) &\stackrel{\text{def}}{=} \langle\!\cdot\emptyset\cdot\!\rangle \bigcirc p;\\
tr(\neg\varphi) &\stackrel{\text{def}}{=} \neg tr(\varphi);\\
tr(\varphi \vee \psi) &\stackrel{\text{def}}{=} tr(\varphi) \vee tr(\psi);\\
tr(\Diamond_i \varphi) &\stackrel{\text{def}}{=} \langle\!\cdot i\cdot\!\rangle (\bot\, \mathcal{U}\, tr(\varphi)).
\end{aligned}$$

We can show the following lemma.

Lemma 1. *Let* φ *be an* $S5^n$*-formula and let* Σ_φ *be the set of agents that occur in* φ*. The following are equivalent:*

(i) φ *is satisfiable wrt.* $\models_{S5^n}$*;*
(ii) $\langle\!\cdot \Sigma_\varphi \cdot\!\rangle (\bot\, \mathcal{U}\, tr(\varphi))$ *is satisfiable wrt.* $\models_{ATLsc}$.

Acknowledgements

We are indebted to the reviewers of LAMAS 2012 and the workshop "Modeling Strategic Reasoning", Lorentz Center, Leiden, The Netherlands, 20-24 February 2012. Nicolas is supported by a Marie Curie COFUND fellowship, and Dirk by a Juan de la Cierva fellowship of Spain and the project Agreement Technologies (Grant CONSOLIDER CSD2007-0022, INGENIO 2010), and the MICINN projects TIN2006-15455 and TIN2009-14562-CO5.

[5] Note that the lower bound of Theorem 1(i) already follows from propositional logic.

References

1. T. Ågotnes, V. Goranko, and W. Jamroga. Alternating-time temporal logics with irrevocable strategies. In *Proceedings of TARK XI*, pages 15–24. Presses Universitaires de Louvain, 2007.
2. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. In *Proceedings of FOCS'97*, pages 100–109. IEEE Computer Society, 1997
3. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Lecture Notes in Computer Science*, 1536:23–60, 1998.
4. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
5. N. Belnap, M. Perloff, and M. Xu. *Facing the future: agents and choices in our indeterminist world*. Oxford, 2001.
6. T. Brihaye, A. Da Costa, F. Laroussinie, and N. Markey. ATL with strategy contexts and bounded memory. Technical Report LSV-08-14, ENS Cachan, 2008.
7. T. Brihaye, A. Da Costa, F. Laroussinie, and N. Markey. ATL with strategy contexts and bounded memory. In *Proceedings of LFCS'09*, LNCS, vol. 5407, pages 92–106. Springer, 2009.
8. J. Broersen, A. Herzig, and N. Troquard. A STIT-extension of ATL. In *Proceedings of JELIA'06*, LNAI, vol. 4160, pages 69–81. Springer, 2006.
9. J. Broersen, A. Herzig, and N. Troquard. From Coalition Logic to STIT. In *Proceedings of LCMAS'05*, ENTCS, vol. 157:4, pages 23–35. Elsevier, 2006.
10. A. Da Costa. *Propriétés de jeux multi-agents*. Thèse de doctorat, Laboratoire Spécification et Vérification, ENS Cachan, France, Sept. 2011.
11. A. Da Costa, F. Laroussinie, and N. Markey. ATL with strategy contexts: Expressiveness and model checking. In *Proceedings of FSTTCS'10*, LIPIsc, vol. 8, pages 120–132. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2010.
12. D. Gabbay, A. Kurucz, F. Wolter, and M. Zakharyaschev. *Many-dimensional modal logics: theory and applications*. Studies in Logic, 148. Elsevier Science, 2003.
13. V. Goranko and W. Jamroga. Comparing semantics of logics for multi-agent systems. *Synthese*, 139(2):241–280, 2004.
14. A. Herzig and F. Schwarzentruber. Properties of logics of individual and group agency. In *Proceedings of AiML*, pages 133–149. College Publications, 2008.
15. J. F. Horty. *Agency and Deontic Logic*. Oxford University Press, Oxford, 2001.
16. W. Jamroga and J. Dix. Do agents make model checking explode (computationally)? In *Proceedings of CEEMAS'05*, pages 398–407. Springer-Verlag, 2005.
17. F. Laroussinie, N. Markey, and G. Oreiby. On the expressiveness and complexity of ATL. *Logical Methods in Computer Science*, 4(2), 2008.
18. T. Müller. On the formal structure of continuous action. In *Proceedings of AiML 5*, pages 191–209. King's College Publications, 2005.
19. A. N. Prior. *Past, Present, and Future*. Clarendon Press, 1967.
20. R. H. Thomason. Indeterminist time and truth-value gaps. *Theoria*, 36:264–81, 1970.
21. W. van der Hoek, W. Jamroga, and M. Wooldridge. A logic for strategic reasoning. In *Proceedings of AAMAS'05*, pages 157–164. ACM Press, 2005.
22. D. Walther, W. van der Hoek, and M. Wooldridge. Alternating-time temporal logic with explicit strategies. In *Proceedings of TARK XI*, pages 269–278. Presses Universitaires de Louvain, 2007.