

Augmenting ATL with strategy contexts

François Laroussinie and Nicolas Markey

March 2014

Research report LSV-14-05 (Version 2)



Laboratoire Spécification & Vérification

École Normale Supérieure de Cachan
61, avenue du Président Wilson
94235 Cachan Cedex France

Augmenting ATL with strategy contexts

François Laroussinie^a, Nicolas Markey^b

^a*LIAFA – CNRS & U. Paris-Diderot – Case 7014 – F-75205 Paris cedex 13 – France*

^b*LSV – CNRS & ENS Cachan – 61 av. Pdt Wilson – F-94230 Cachan – France*

Abstract

We study the extension of the alternating-time temporal logic (ATL) with strategy contexts: contrary to the original semantics, in this semantics the strategy quantifiers do not reset the previously selected strategies.

We show that our extension ATL_{sc} is very expressive, but that its decision problems are quite hard: model checking is k -EXPTIME-complete when the formula has k nested strategy quantifiers; satisfiability is undecidable, but we prove that it is decidable when restricting to turn-based games. Our algorithms are obtained through a very convenient translation to QCTL (the computation-tree logic CTL extended with atomic quantification), which we show also applies to Strategy Logic, as well as when strategy quantification ranges over memoryless strategies.

Keywords: temporal logics, games for synthesis, model checking, satisfiability

1. Introduction

The alternating-time temporal logic (ATL) is a convenient extension of CTL for expressing properties of multi-agent systems. For this, it includes quantification over strategies of the agents, instead of the sole quantifiers over paths of CTL. However, in order to retain the nice algorithmic properties of CTL, the quantification over strategies in ATL is *forgetful*, in the sense that each quantifier deletes the previously selected strategies. Under this

[☆]This paper is a long version of [15], whose focus was mainly on satisfiability of ATL_{sc} via QCTL. We added several results from [8, 9], with new proofs and a uniform presentation using QCTL.

^{☆☆}Work partly sponsored by ERC Stg Grant EQualIS and FP7 FET project Cassting.

semantics, ATL model checking is PTIME-complete [2], while satisfiability is EXPTIME-complete [26].

ATL with strategy contexts (ATL_{sc} for short) lifts this restriction, and stores the previously selected strategies within a context. The players keep on following their strategies until the formula says otherwise. This makes the logic very expressive, and very convenient for expressing properties of multi-agent systems, in setting where the agents are neither collaborative nor completely antagonist. In a setting of a client/server interaction, ATL_{sc} can express the existence of a policy for the server in order to enforce mutual exclusion to the resource, and to make each client have a way of eventually accessing the resource. Such a property mixes collaboration between the server and each client, and possibly antagonism between clients. Such an expressiveness comes with a cost: ATL_{sc} model checking was proved to be k -EXPTIME-complete, where k is the number of nested strategy quantifiers in the formula being checked [9]. It follows that ATL_{sc} model-checking problem is Tower-complete (Tower contains the classes k -EXPTIME for all k ; see [22] for more details about the complexity class Tower). Satisfiability of ATL_{sc} is undecidable [24].

In this paper, we revisit these problems with a uniform technique: we establish a very tight link between ATL_{sc} and QCTL (the extension of CTL with quantification over atomic propositions [23, 11, 10, 14]). We prove that the model-checking problems for both logics are essentially the same problem (there are reductions both ways). Moreover, while satisfiability is undecidable for ATL_{sc} , we use our translation to QCTL to prove that satisfiability is decidable in restricted cases, and in particular when restricted to turn-based games.

The convenience of QCTL is not particular to ATL_{sc} , and we prove that our technique also applies to Strategy Logic (SL) [7, 19], another extension of ATL with explicit handling of strategies. Also, using a different semantics for quantification in QCTL, we use our translation to study the variant of ATL_{sc} where strategy quantifiers range over memoryless strategies. In this case, model checking is PSPACE-complete, but satisfiability is undecidable, even for finitely branching two-player turn-based games.

Related work. Over the last decade, several formalisms have been proposed as improvements to ATL for specifying properties of multi-agent models. Strategy Logic (SL) [7, 19] is the most closely related to ATL_{sc} . It builds on a different approach, where strategies are explicit: first-order quantification

can be used to select strategies, which can then be assigned to players. LTL is then used to impose constraints on the resulting paths. The logic was defined (and its model-checking problem was proved decidable) in [7] for turn-based games. This was extended to concurrent games in [19, 16], and satisfiability was proved undecidable. In Section 6.2, we show that our QCTL technique also applies to SL, and prove that our results for ATL_{sc} also hold for SL.

Other related approaches include stochastic game logic [4] and the quantified decision μ -calculus [21], as well as various fragments of Strategy Logic [17, 18] and ATL_{sc} [1, 27].

2. Definitions

2.1. Preliminaries

Given two mappings $f: A \rightarrow B$ and $g: A' \rightarrow B'$, and a subset $C \subseteq A \cap A'$, we say that f and g coincide on C , written $f \simeq_C g$, when $f(c) = g(c)$ for all $c \in C$. In case $A = A'$ and $B = 2^S$ and $B' = 2^{S'}$, given $T \subseteq S \cap S'$, we say that f and g agree in T , written $f \equiv_T g$, if for all $a \in A$, it holds $f(a) \cap T = g(a) \cap T$. If A is a subset of a larger set \mathbb{A} , then f is said to be a *partial function* (w.r.t \mathbb{A}), and A is called its *domain* (written $\text{dom}(f)$).

For $k \in \mathbb{N}$, we define $[k] = \{i \in \mathbb{N} \mid 0 \leq i < k\}$. We also let $[\infty] = \mathbb{N}$. Let Σ be a set. A *word* over Σ is a mapping $w: [k] \rightarrow \Sigma$, for some $k \in \mathbb{N} \cup \{\infty\}$; for $n \in [k]$, we usually write w_n for $w(n)$. The word w is finite when $k \in \mathbb{N}$ (then k is the *length* of w , denoted $|w|$), otherwise it is infinite (then $|w| = +\infty$). When w is finite, we write $\text{last}(w)$ for its last element $w(|w| - 1)$. The (only) word of length zero is denoted with ϵ . Given a finite word v and a word w , their *concatenation* $v \cdot w$ is the sequence u s.t. $u(n) = v(n)$ when $n < |v|$ and $u(n) = w(n - |v|)$ when $n \geq |v|$. When v is a one-letter word, we sometimes write $v_0 \cdot w$ to denote $v \cdot w$. A *prefix* of a word w is a finite word p such that there exists a word s such that $w = p \cdot s$. For any $n \leq |w|$, w has a unique prefix of length n , which we denote $w_{<n}$ (or sometimes $w_{\leq n-1}$).

Let D be a set. A *D-tree* is a non-empty set T of finite words over D such that for any $t \in T$, all the prefixes of t are in T . The elements of T are called *nodes*, and the special node ϵ (the empty word) is the *root* of T . A Σ -labeled *D-tree* is a pair $\langle T, \ell \rangle$ where T is a *D-tree* and $\ell: T \rightarrow \Sigma$ labels the nodes of T with a letter in Σ .

2.2. Kripke structures

Let AP be a set of atomic propositions, and $\Sigma = 2^{\text{AP}}$.

Definition 1. A Kripke structure over AP is a 3-tuple $\mathcal{S} = \langle Q, R, \ell \rangle$ where Q is a countable set of states, $R \subseteq Q^2$ is a binary relation and $\ell: Q \rightarrow \Sigma$ is a labeling function. We always assume that the relation R be left-total, i.e., for any $q \in Q$, there is a $q' \in Q$ such that $(q, q') \in R$.

Let $\mathcal{S} = \langle Q, R, \ell \rangle$ be a Kripke structure, and $q \in Q$. A path in \mathcal{S} from q is a non-empty word π over Q such that $\pi_0 = q$ and for all $n \in [|\pi| - 1]$, it holds $(\pi_n, \pi_{n+1}) \in R$. Given a finite path π and a path ρ such that $\text{last}(\pi) = \rho_0$, the join of π and ρ , denoted $\pi : \rho$ is defined as the concatenation $\pi_{<|\pi|-1} \cdot \rho$. Notice that since each state in a Kripke structure must have at least one successor, any finite path can be enlarged.

Given a subset $I \subseteq Q$, we write $IQ_{\mathcal{S}}^*$ (resp. $IQ_{\mathcal{S}}^{\omega}$) for the set of finite (resp. infinite) paths in \mathcal{S} from some $q \in I$; we write $qQ_{\mathcal{S}}^*$ (resp. $qQ_{\mathcal{S}}^{\omega}$) in case $I = \{q\}$, and $Q_{\mathcal{S}}^+$ (resp. $Q_{\mathcal{S}}^{\omega}$) in case $I = Q$. With a path π , we associate a trace $\ell \circ \pi: [|\pi|] \rightarrow \Sigma$, which is a word over Σ .

The *unwinding* of \mathcal{S} from q is the Q -tree $T = \{w \in Q^* \mid q \cdot w \in qQ_{\mathcal{S}}^*\}$. The *computation tree* of \mathcal{S} from q is the Σ -labeled Q -tree $\langle T, l \rangle$ such that $l(w) = \ell(\text{last}(q \cdot w))$ for all $w \in T$. Notice that from our assumption that each state in \mathcal{S} has at least one outgoing transition, any node in the computation tree has at least one successor. A *branch* in T is an infinite word $w \in Q^{\omega}$ such that $q \cdot w \in qQ_{\mathcal{S}}^{\omega}$. Any finite prefix of a branch is a node of T .

2.3. Quantified CTL

The temporal logics CTL* and CTL were defined in the 1980s. Let AP be a set of atomic propositions. The syntax of CTL* is as follows:

$$\begin{aligned} \text{CTL}^* \ni \varphi_s, \psi_s &::= p \mid \neg \varphi_s \mid \varphi_s \vee \psi_s \mid \mathbf{E}\varphi_p \mid \mathbf{A}\varphi_p \\ \varphi_p, \psi_p &::= \varphi_s \mid \neg \varphi_p \mid \varphi_p \vee \psi_p \mid \mathbf{X}\varphi_p \mid \varphi_p \mathbf{U}\psi_p \end{aligned}$$

where p ranges over AP. Formulas of CTL* are interpreted in the computation tree (hence the name) of a given Kripke structure. Let $\mathcal{S} = \langle Q, R, \ell \rangle$ be such a structure, let ρ be an infinite path in \mathcal{S} , and $n \in \mathbb{N}$. That formula φ holds

true at position n along ρ in \mathcal{S} is defined inductively as follows:

$$\begin{aligned}
\mathcal{S}, \rho, n \models p & \text{ iff } p \in \ell(\rho_n) \\
\mathcal{S}, \rho, n \models \neg \varphi & \text{ iff } \mathcal{S}, \rho, n \not\models \varphi \\
\mathcal{S}, \rho, n \models \varphi \vee \psi & \text{ iff } \mathcal{S}, \rho, n \models \varphi \text{ or } \mathcal{S}, \rho, n \models \psi \\
\mathcal{S}, \rho, n \models \mathbf{E}\varphi_p & \text{ iff } \exists \rho' \in \rho_n Q_{\mathcal{S}}^{\omega}. \mathcal{S}, \rho', 0 \models \varphi_p \\
\mathcal{S}, \rho, n \models \mathbf{A}\varphi_p & \text{ iff } \forall \rho' \in \rho_n Q_{\mathcal{S}}^{\omega}. \mathcal{S}, \rho', 0 \models \varphi_p \\
\mathcal{S}, \rho, n \models \mathbf{X}\varphi_p & \text{ iff } \mathcal{S}, \rho', n+1 \models \varphi_p \\
\mathcal{S}, \rho, n \models \varphi_p \mathbf{U} \psi_p & \text{ iff } \exists l \geq 0. \mathcal{S}, \rho, n+l \models \psi_p \text{ and } \forall 0 \leq m < l. \\
& \quad \mathcal{S}, \rho, n+m \models \varphi_p
\end{aligned}$$

Several useful abbreviations can be defined: besides the classical $\top = p \vee \neg p$ (which always evaluates to true), $\perp = \neg \top$ and $\varphi \wedge \psi = \neg(\neg \varphi \vee \neg \psi)$, the following modalities will be used throughout this paper:

$$\mathbf{F}\varphi_p = \top \mathbf{U}\varphi_p \qquad \mathbf{G}\varphi_p = \neg \mathbf{F}\neg \varphi_p.$$

The former states that φ_p will eventually hold true along the current path, while the latter states that it holds true at any position along that path.

An important remark about this semantics is that the evaluation of state formulas (of the form φ_s in the grammar defining CTL*) at position n along ρ only depends on ρ_n . In other terms, for any two paths ρ and ρ' and any two positions n and n' such that $\rho_n = \rho'_{n'}$, and for any state formula φ_s , it holds

$$\mathcal{S}, \rho, n \models \varphi_s \quad \text{iff} \quad \mathcal{S}, \rho', n' \models \varphi_s.$$

As a consequence, for a state formula φ_s , we often replace $\mathcal{S}, \rho, 0 \models \varphi_s$ with $\mathcal{S}, \rho_0 \models \varphi_s$.

Another remark is that CTL* is invariant under bisimulation: two structures that are bisimilar satisfy the same CTL* formulas. In particular, evaluating a CTL* formula over a Kripke structure and over its computation tree are equivalent. Formally, let $q \in Q$, and w be a branch in the computation tree \mathcal{T} of \mathcal{S} from q (so that $q \cdot w$ is an infinite path in \mathcal{S} ; we abusively see w as a path in \mathcal{T} starting from the root, when \mathcal{T} is seen as an infinite-state Kripke structure). Let $n \in \mathbb{N}$. Then for any $\varphi \in \text{CTL}^*$, it holds

$$\mathcal{S}, q \cdot w, n \models \varphi \quad \text{iff} \quad \mathcal{T}_{\mathcal{S}}, w, n \models \varphi.$$

The fragment CTL of CTL* is obtained by restricting the form of path formulas to the following grammar:

$$\varphi_p, \psi_p ::= \neg \varphi_p \mid \mathbf{X} \varphi_s \mid \varphi_s \mathbf{U} \psi_s.$$

In other terms, the modalities \mathbf{X} and \mathbf{U} (and negations thereof) can only appear in the immediate scope of a path quantifier \mathbf{E} or \mathbf{A} .

We now present an extension of CTL* with quantification over atomic propositions, which will be our main technical tool in the sequel.

For $P \subseteq \mathbf{AP}$, two Kripke structures $\mathcal{S} = \langle Q, R, \ell \rangle$ and $\mathcal{S}' = \langle Q', R', \ell' \rangle$ are *P-equivalent* (denoted $\mathcal{S} \equiv_P \mathcal{S}'$) whenever $Q = Q'$, $R = R'$, and $\ell \equiv_P \ell'$ (i.e., $\ell(q) \cap P = \ell'(q) \cap P$ for any $q \in Q$). In other terms, $\mathcal{S} \equiv_P \mathcal{S}'$ if \mathcal{S}' can be obtained from \mathcal{S} by modifying the labeling function of \mathcal{S} for propositions not in P .

Definition 2. *The syntax of QCTL* is defined by the following grammar:*

$$\begin{aligned} \text{QCTL}^* \ni \varphi_s, \psi_s &::= p \mid \neg \varphi_s \mid \varphi_s \vee \psi_s \mid \mathbf{E} \varphi_p \mid \mathbf{A} \varphi_p \mid \exists p. \varphi_s \\ \varphi_p, \psi_p &::= \varphi_s \mid \neg \varphi_p \mid \varphi_p \vee \psi_p \mid \mathbf{X} \varphi_p \mid \varphi_p \mathbf{U} \psi_p. \end{aligned}$$

The *structure semantics* of QCTL* is derived from the semantics of CTL* by adding the following rule:

$$\mathcal{S}, \rho, n \models \exists p. \varphi_s \quad \text{iff} \quad \exists \mathcal{S}'. \mathcal{S}' \equiv_{\mathbf{AP} \setminus \{p\}} \mathcal{S} \text{ and } \mathcal{S}', \rho, n \models \varphi_s.$$

In other terms, $\exists p. \varphi_s$ means that it is possible to (re)label the Kripke structure with p in order to make φ_s hold.

While CTL* is invariant under bisimulation, it is not true for QCTL*. While evaluating CTL* on a Kripke structure and on its execution tree (seen as an infinite-state Kripke structure) are equivalent, this is not true of QCTL*. As a consequence, we define another semantics for QCTL*, called the *tree semantics* of QCTL*. This semantics is obtained by applying the structure semantics to the execution tree of the Kripke structure. Evaluating $\exists p. \varphi_s$ then amounts to (re)labeling the execution tree with p in such a way that φ_s holds true.

We refer to [14] for a detailed study of QCTL* and QCTL. Here we just recall the following important properties of these logics. First note that QCTL is actually as expressive as QCTL* (with an effective translation) [10, 9]. Secondly model checking and satisfiability are decidable but non elementary

More precisely the complexity depends on the number of alternations of propositional quantifications in the formulas. In the following we will refer to the fragments EQ^kCTL and Q^kCTL of QCTL and to EQ^kCTL^* and Q^kCTL^* the corresponding fragments of QCTL^* : EQ^kCTL contains the QCTL formulas in prenex normal form (where quantifications are external to the CTL formula), starting with an existential quantification \exists and where the number of alternations is k . In Q^kCTL formulas are not supposed to be in prenex normal form but the alternation of quantifier is bounded by k : formally, Q^1CTL is $\text{CTL}[\text{EQ}^1\text{CTL}]$, and Q^{k+1}CTL is $\text{Q}^1\text{CTL}[\text{Q}^k\text{CTL}]$. The decision procedures for these logics are based on automata construction: given a QCTL formula φ and a (finite) set $D \subset \mathbb{N}$, one can build a tree automaton $\mathcal{A}_{\varphi, D}$ recognizing the D -trees satisfying φ . This provides a decision procedure for model checking as the Kripke structure \mathcal{S} fixes the set D , and it remains to check whether the execution tree of \mathcal{S} is accepted by $\mathcal{A}_{\varphi, D}$. For satisfiability the decision procedure is obtained by building a formula φ_2 from φ such that φ_2 is satisfied by some $[2]$ -tree if, and only if, φ is satisfied by some finitely-branching tree. Finally, it remains to notice that a QCTL formula is satisfiable if, and only if, it is satisfiable in a finitely-branching tree (since QCTL is as expressive as MSO) to get the decision procedure for QCTL satisfiability. As a consequence, a QCTL formula is satisfiable if, and only if, it is satisfied by a regular tree (corresponding to the execution tree of some finite Kripke structure).

2.4. Concurrent Game Structures

Game structures extend Kripke structures with several agents acting on the evolution of the system.

Definition 3 ([2]). *A concurrent game structure (CGS) is a 7-tuple $\mathcal{C} = \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge} \rangle$ where: $\langle Q, R, \ell \rangle$ is a Kripke structure, $\text{Agt} = \{a_1, \dots, a_p\}$ is a finite set of agents, \mathcal{M} is a non-empty countable set of moves, $\text{Mov}: Q \times \text{Agt} \rightarrow \mathcal{P}(\mathcal{M}) \setminus \{\emptyset\}$ defines the set of available moves of each agent in each state, and $\text{Edge}: Q \times \mathcal{M}^{\text{Agt}} \rightarrow R$ is a transition table associating, with each state q and each set of moves of the agents, the resulting transition departing from q .*

Let $\mathcal{C} = \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge} \rangle$ be a CGS. The notions of paths and execution trees of \mathcal{C} are inherited from the underlying Kripke structure $\langle Q, R, \ell \rangle$. The size of \mathcal{C} , denoted $|\mathcal{C}|$, is $|Q| + |\text{Edge}|$. Notice that for

all $q \in Q$, $\text{Edge}(q)$ is a $|\text{Agt}|$ -dimensional table, whose size is then exponential in $|\text{Agt}|$ (provided that $|\mathcal{M}| > 1$). A *move vector* in \mathcal{C} is a mapping $m: \text{Agt} \rightarrow \mathcal{M}$. For a state $q \in Q$, we define

$$\text{Next}(q) = \{q' \in Q \mid \exists m \in \mathcal{M}^{\text{Agt}}. \forall a_i \in \text{Agt}. m(a_i) \in \text{Mov}(q, a_i) \\ \text{and } \text{Edge}(q, m) = (q, q')\}$$

and, for a coalition $C \subseteq \text{Agt}$ and a partial move vector m with $\text{dom}(m) = C$,

$$\text{Next}(q, C, m) = \{q' \in Q \mid \exists m' \in \mathcal{M}^{\text{Agt}}. m' \simeq_C m \text{ and } \text{Edge}(q, m') = (q, q')\}.$$

A *turn-based* game structure (TBGS for short) is a CGS where each state q is controlled by a single agent, called the *owner* of q (and denoted $\text{Own}(q)$). Formally, for every $q \in Q$, for any two move vectors m and m' with $m =_{\text{Own}(q)} m'$, it holds $\text{Edge}(q, m) = \text{Edge}(q, m')$.

A *strategy* for a player $a_i \in \text{Agt}$ in a CGS \mathcal{C} is a function $f: Q_S^+ \rightarrow \mathcal{M}$ that maps any finite path to a possible move for a_i , *i.e.*, satisfying $f(\pi) \in \text{Mov}(\text{last}(\pi), a_i)$. A strategy f is *memoryless* if $f(\pi) = f(\pi')$ whenever $\text{last}(\pi) = \text{last}(\pi')$. A strategy for a coalition A is a mapping assigning a strategy to each agent in A . The set of strategies for A is denoted $\text{Strat}(A)$ (and $\text{Strat}_0(A)$ is the subset of memoryless strategies). In the sequel, when no ambiguity arises, we subscript the strategies with their domains, writing f_A for a strategy of $A \subseteq \text{Agt}$ and f_i for a strategy of player a_i (hence $f_i = f_A(a_i)$ when $a_i \in A$). Given a strategy $f_A \in \text{Strat}(A)$ and a coalition B , the strategy $(f_A)|_B$ (resp. $(f_A)_{\setminus B}$) denotes the restriction of f_A to the coalition $A \cap B$ (resp. $A \setminus B$). Given two strategies $f_A \in \text{Strat}(A)$ and $g_B \in \text{Strat}(B)$, we define $f_A \circ g_B \in \text{Strat}(A \cup B)$ as $(f_A \circ g_B)_j = f_j$ (resp. g_j) if $a_j \in A$ (resp. $a_j \in B \setminus A$).

Let ρ be a finite path in \mathcal{C} , and $f_A \in \text{Strat}(A)$ for some coalition A . A path π is *compatible with f_A after ρ* if it contains ρ as a prefix and, for all $n \in [|\rho|; |\pi| - 1]$, letting $m: a_i \in A \mapsto f_i(\pi_{<n})$, it holds $\pi_{n+1} \in \text{Next}(\pi_n, A, m)$. The set of *outcomes of f_A after ρ* , denoted $\text{Out}(\rho, f_A)$, is the set of infinite paths that are compatible with f_A after ρ .

Example 4. *Fig. 1 represents two three-state two-player CGSs. The transitions are decorated with their corresponding move vectors. In \mathcal{C} , each player has only two allowed moves, while in \mathcal{C}' , they have three. One can easily check e.g. that $\text{Next}(\ell_0, a_1, 1)$ is $\{\ell_0, \ell_1\}$ in \mathcal{C} , and $\text{Next}(\ell'_0, a_2, 3)$ is $\{\ell'_0, \ell'_1, \ell'_2\}$ in \mathcal{C}' .*




Fig. 1: Two concurrent game structures

2.5. ATL with strategy contexts

We now introduce our logic, which extends the alternating-time temporal logic of [2] with strategy contexts. We assume a fixed set of atomic propositions AP and a fixed set of agents Agt .

Definition 5. *The formulas of ATL_{sc}^* are defined by the following grammar:*

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

where p ranges over AP and A ranges over 2^{Agt} .

ATL_{sc}^* formulas are interpreted over CGSs, within a context (*i.e.*, a preselected strategy): *state formulas* of the form φ_s in the grammar above are evaluated over states, while *path formulas* of the form φ_p are evaluated along paths. In order to have a uniform definition, we evaluate all formulas at a given position along a path.

The semantics is quite similar as that of CTL^* , but ATL_{sc}^* now has *strategy quantifiers* in place of path quantifiers. Informally, formula $\langle A \rangle \varphi_p$ holds at position n along ρ under the context F if it is possible to extend F with a strategy for the coalition A such that the outcomes of the resulting strategy after $\rho_{\leq n}$ all satisfy φ_p . Similarly, $\langle \bar{A} \rangle \varphi_p$ is true when we can complete the current context with a strategy for coalition $\text{Agt} \setminus A$ in such a way that φ_p holds along every outcome. The use of complementary coalition in the syntax is discussed in Remark 8.

We now define the semantics formally. Let \mathcal{C} be a CGS, ρ be an infinite path of \mathcal{C} , and $n \in \mathbb{N}$ point to a position along ρ . Let $B \subseteq \text{Agt}$ be a coalition, and $f_B \in \text{Strat}(B)$. That a (state or path) formula φ holds at a position n along ρ in \mathcal{C} under strategy context f_B , denoted $\mathcal{C}, \rho, n \models_{f_B} \varphi$, is defined inductively as follows (omitting atomic propositions and boolean operators):

$$\begin{aligned}
\mathcal{C}, \rho, n \models_{f_B} (\downarrow A) \varphi_s & \text{ iff } \mathcal{C}, \rho, n \models_{(f_B) \setminus A} \varphi_s \\
\mathcal{C}, \rho, n \models_{f_B} (\overline{\downarrow A}) \varphi_s & \text{ iff } \mathcal{C}, \rho, n \models_{(f_B) \upharpoonright A} \varphi_s \\
\mathcal{C}, \rho, n \models_{f_B} \langle \downarrow A \rangle \varphi_p & \text{ iff } \exists g_A \in \text{Strat}(A). \forall \rho' \in \text{Out}(\rho_{\leq n}, g_A \circ f_B). \\
& \quad \mathcal{C}, \rho', n \models_{g_A \circ f_B} \varphi_p \\
\mathcal{C}, \rho, n \models_{f_B} \langle \overline{\downarrow A} \rangle \varphi_p & \text{ iff } \exists g_{\overline{A}} \in \text{Strat}(\text{Agt} \setminus A). \forall \rho' \in \text{Out}(\rho_{\leq n}, g_{\overline{A}} \circ f_B). \\
& \quad \mathcal{C}, \rho', n \models_{g_{\overline{A}} \circ f_B} \varphi_p \\
\mathcal{C}, \rho, n \models_{f_B} \mathbf{X} \varphi_p & \text{ iff } \mathcal{C}, \rho, n + 1 \models_{f_B} \varphi_p \\
\mathcal{C}, \rho, n \models_{f_B} \varphi_p \mathbf{U} \psi_p & \text{ iff } \exists l \geq 0. \mathcal{C}, \rho, n + l \models_{f_B} \psi_p \text{ and } \forall 0 \leq m < l. \\
& \quad \mathcal{C}, \rho, n + m \models_{f_B} \varphi_p
\end{aligned}$$

Notice how the existential strategy quantifier embeds an implicit universal quantification over the set of outcomes of the selected strategy. Also notice that, given a state formula φ_s , two paths ρ and ρ' , a position n such that $\rho_{\leq n} = \rho'_{\leq n}$, and a context f_B , it holds

$$\mathcal{C}, \rho, n \models_{f_B} \varphi_s \quad \text{iff} \quad \mathcal{C}, \rho', n \models_{f_B} \varphi_s$$

In particular, when $n = 0$, that $\mathcal{C}, \rho, 0 \models_{f_B} \varphi_s$ does not depend on the whole path ρ but only on its first state ρ_0 . In the sequel we equivalently write $\mathcal{C}, \rho_0 \models_{f_B} \varphi_s$ in place of $\mathcal{C}, \rho, 0 \models_{f_B} \varphi_s$.

The usual shorthands such as F and \mathbf{G} are defined as for CTL*. It will also be convenient to use the constructs $\downarrow A \varphi_p$ as a shorthand for $\neg \langle \downarrow A \rangle \neg \varphi_p$, and $\langle \downarrow A \rangle \varphi_s$ as a shorthand for $\langle \downarrow A \rangle \perp \mathbf{U} \varphi_s$.

The fragment ATL_{sc} of ATL_{sc}^* is defined as usual, by restricting the set of path formulas to

$$\varphi_p, \psi_p ::= \neg \varphi_p \mid \mathbf{X} \varphi_s \mid \varphi_s \mathbf{U} \psi_s.$$

Example 6. *We consider again the CGSs of Fig. 1. One can check that in both CGSs, player a_1 has a strategy to avoid visiting ℓ_2 (resp. ℓ'_2), but he has no strategy for leaving state ℓ_0 (resp. ℓ'_0). But what distinguishes these two CGSs is the following: in \mathcal{C} , for any move m_1 of player a_1 , it holds $|\text{Next}(\ell_0, a_1, m_1)| = 2$. On the other hand, for move 3 of a_1 in \mathcal{C}' , we have $|\text{Next}(\ell'_0, a_1, 3)| = 3$. In particular, under move 3 of a_1 in ℓ'_0 , player a_2 still has a move to reach ℓ'_1 and another move to reach ℓ'_2 , so that formula $\langle a_1 \rangle (\langle a_2 \rangle \mathbf{X} a \wedge \langle a_2 \rangle \mathbf{X} b)$ holds true in ℓ'_0 . No such move exists in \mathcal{C} , and the same formula is false in ℓ_0 .*

The classical semantics of ATL^* and ATL , as defined in [2], did not involve a strategy context. Syntactically, the logic was defined as

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

The semantics was similar to that of ATL_{sc}^* , but without the strategy context. The semantics of the strategy quantifier $\langle\langle A \rangle\rangle \varphi_p$ is then defined as follows:

$$\mathcal{C}, \rho, n \models \langle\langle A \rangle\rangle \varphi_p \quad \text{iff} \quad \exists g_A \in \text{Strat}(A). \forall \rho' \in \text{Out}(\rho_{\leq n}, g_A). \mathcal{C}, \rho', n \models \varphi_p.$$

Example 7. Consider the ATL formula $\langle\langle a_1 \rangle\rangle (\langle\langle a_2 \rangle\rangle \mathbf{X} a \wedge \langle\langle a_2 \rangle\rangle \mathbf{X} b)$. The first quantification $\langle\langle a_1 \rangle\rangle$ in this formula is useless, since the selected strategy is lost when quantifying over strategies of a_2 . As a consequence, this formula is equivalent to $\langle\langle a_2 \rangle\rangle \mathbf{X} a \wedge \langle\langle a_2 \rangle\rangle \mathbf{X} b$, which is false both in ℓ_0 and in ℓ'_0 . It can be proved that ℓ_0 and ℓ'_0 are alternating bisimilar [3], so that they cannot be distinguished by ATL^* .

Remark 8. Previous definitions of ATL_{sc}^* (see [5, 8]) did not allow complementary coalitions in the syntax for $\langle _ \rangle$ and $(_)$. We add them here for two reasons. The first one is convenience: given a coalition A , it is sometimes useful to consider the other players. The second is mathematical correctness: our expressiveness result of Theorem 9, in the way we state it in this paper, requires talking about complementary coalitions.

3. Expressiveness of ATL_{sc} and ATL_{sc}^*

We devote this section to expressiveness issues. We begin with some examples of ATL_{sc} formulas, witnessing how useful our new formalism can be. We then give some theoretical results about the expressiveness of ATL_{sc} , showing for instance that ATL_{sc} and ATL_{sc}^* have the same expressive power. We finish the section with a comparison (w.r.t. expressiveness) to plain ATL . Comparisons with other formalisms are deferred to Section 6, where we review several logics related to ATL_{sc} .

3.1. Examples of formulas

For the reader to get acquainted with ATL_{sc} , we give in this section several examples of ATL_{sc} formulas. This will also witness the expressive power and usefulness of our logic.

Client-server interactions. Consider a situation where a server S controls the access to a resource shared among several clients a_1 to a_n . The server has a double role: it has to ensure that at most one client uses the resource at any point in time (*mutual exclusion*), but also to provide a way for each client to be able to access the resource; the latter mixes collaboration between the server and individual agents, and antagonism between agents. In ATL_{sc}^* , we can express this requirement as follows:

$$\langle S \rangle \mathbf{G} \left[\left(\bigwedge_{i,j \in [1,n], i \neq j} \neg(\text{access}_i \wedge \text{access}_j) \right) \wedge \left(\bigwedge_{i \in [1,n]} \langle a_i \rangle \mathbf{F} \text{access}_i \right) \right]$$

Nash equilibria. In n -player games (with $n > 2$), players will usually have non-zero-sum objectives, and the notion of *winning strategy* is not relevant anymore. Instead, equilibria positions are sought, in which all players has *optimal* outcome. Here *optimal* may have many different meanings. One of them, corresponding to *Nash equilibria* [20], is related to the other players strategies: a Nash equilibrium is a strategy profile in which each individual strategy is the best response to the others' strategies.

ATL_{sc} can express the existence of a Nash equilibrium (here in a setting where the individual objectives of the players are boolean, and only considering pure strategies): this would be written as

$$\langle \text{Agt} \rangle \left[\bigwedge_{a \in \text{Agt}} (\neg \varphi_a \Rightarrow \neg \langle a \rangle \varphi_a) \right]$$

In this formula, we write that there is a strategy profile (the one witnessing the outermost quantifier) such that no player can improve their payoff (*i.e.*, if they are not winning in the equilibrium, they don't have a way of achieving their goal in this situation).

In ATL_{sc} , we can additionally impose extra requirements to Nash equilibria. Indeed, several Nash equilibria might coexist, and some might be better than others (for instance, Nash equilibria where all the players fail to achieve their objectives might coexist with Nash equilibria where all the objectives are met). In ATL_{sc} we can additionally impose constraints on the equilibria strategies.

Winning secure equilibria. A *winning secure equilibrium* [6] is a winning (for all players) Nash equilibria with the additional requirement that if a player deviates and worsens the payoff of some player, then she also worsens her own payoff. In other terms, no player can harm another player without

harming herself. The existence of a winning secure equilibrium can be written as

$$\langle \text{Agt} \rangle \left[\bigwedge_{a \in \text{Agt}} \varphi_a \wedge \bigwedge_{a, b \in \text{Agt}} [a] (\neg \varphi_b \Rightarrow \neg \varphi_a) \right]$$

Dominant strategy. A strategy is said *dominant* if it is a best response to any strategies of the other players. The existence of a dominant strategy for player a_i can be expressed as

$$\langle a_i \rangle [\overline{a_i}] (\langle a_i \rangle \varphi_i \Rightarrow \varphi_i).$$

3.2. ATL_{sc} vs ATL_{sc}^*

From ATL_{sc}^* to ATL_{sc} . Surprisingly, strategy contexts bring ATL_{sc} to the same expressiveness as ATL_{sc}^* : any ATL_{sc}^* formula can be translated into an equivalent ATL_{sc} formula. The main idea is to replace the (implicit) universal quantification over outcomes with explicit universal quantification over strategies. This way, all players are assigned a strategy in the context. In that case, there is only one outcome (because our CGSs are deterministic), so that we can insert empty strategy quantifier $\langle \emptyset \rangle$ in front of any temporal modality.

Given an ATL_{sc}^* formula Φ and two coalitions $(B, B') \in 2^{\text{Agt}(\Phi)} \times (2^{\text{Agt}(\Phi)} \cup \text{Agt}_c)$ (with the intuitive idea that $B \cup (\text{Agt}_c \setminus B')$ is the set of players that have been artificially assigned a strategy in the context); here Agt_c should be seen as a symbol representing all the agents appearing in the underlying CGS. Notice that our resulting formula should not depend on the set Agt_c , as our translation will hold for any CGS \mathcal{C} , we define $\widehat{\Phi}^{[B, B']}$ inductively as follows:

$$\begin{aligned} \widehat{P}^{[B, B']} &= P \\ \widehat{\neg \varphi}^{[B, B']} &= \neg \widehat{\varphi}^{[B, B']} & \widehat{\varphi \wedge \psi}^{[B, B']} &= \widehat{\varphi}^{[B, B']} \wedge \widehat{\psi}^{[B, B']} \\ \widehat{\mathbf{X} \varphi}^{[B, B']} &= \langle \emptyset \rangle \mathbf{X} \widehat{\varphi}^{[B, B']} & \widehat{\varphi \mathbf{U} \psi}^{[B, B']} &= \langle \emptyset \rangle (\widehat{\varphi}^{[B, B']} \mathbf{U} \widehat{\psi}^{[B, B']}) \\ \widehat{(\overline{A}) \varphi}^{[B, B']} &= \widehat{\varphi}^{[B \cup A, B']} & \widehat{(\overline{A}) \varphi}^{[B, B']} &= \widehat{\varphi}^{[B, B' \cap A]} \\ \widehat{\langle A \rangle \varphi}^{[B, B']} &= \langle A \rangle [B \setminus A] [\overline{B' \cup A}] \widehat{\varphi}^{[B \setminus A, B' \cup A]} \\ \widehat{\langle \overline{A} \rangle \varphi}^{[B, B']} &= \langle \overline{A} \rangle [(B \cap A) \cup (A \setminus B')] \widehat{\varphi}^{[(B \cap A) \cup (A \setminus B'), \text{Agt}_c]} \end{aligned}$$

Clearly, $\widehat{\Phi}^{[B, B']}$ is an ATL_{sc} formula, thanks to the $\langle \emptyset \rangle$ inserted in front of the temporal modalities. Notice that the resulting formula only involves coalitions

that appear in the original formula, and does not depend on $\text{Agt}_{\mathcal{C}}$ (because $\text{Agt}_{\mathcal{C}} \cup A = \text{Agt}_{\mathcal{C}}$, $A \setminus \text{Agt}_{\mathcal{C}} = \emptyset$, $\text{Agt}_{\mathcal{C}} \cap A = A$, and $[\overline{\text{Agt}_{\mathcal{C}}}] \psi$ is equivalent to $[\emptyset] \psi$). This transformation achieves the following result (whose full proof is rather technical, so that we moved it to Appendix Appendix A).

Theorem 9. *Given a formula $\varphi \in \text{ATL}_{sc}^*$ and a coalition B' , there exists an ATL_{sc} formula $\hat{\varphi}^{[\emptyset, B']}$, involving only players in $\text{Agt}_{\varphi} \cup B'$, such that for any strategy context f with $\text{dom}(f) = B'$, φ and $\hat{\varphi}^{[\emptyset, B']}$ are equivalent under context f .*

Remark 10. *One can notice that the resulting formula does not make use of $(_)$, even when applied to an ATL_{sc} formula. Hence, as a side result, we obtain that the $(_)$ operator does not increase the expressive power of ATL_{sc} .*

3.3. Comparison with ATL

Clearly enough, ATL^* properties can be expressed in ATL_{sc}^* : indeed, the ATL strategy quantifier $\langle\langle A \rangle\rangle$ is equivalent to $(\overline{\emptyset}) \langle A \rangle$. Notice that following Examples 6 and 7, ATL_{sc} is actually strictly more expressive than ATL^* .

Similarly, CTL^* is translated in ATL_{sc}^* by rewriting $\mathbf{E}\varphi$ as $\langle \overline{\emptyset} \rangle \varphi$ and $\mathbf{A}\varphi$ as $(\overline{\emptyset}) \langle \emptyset \rangle \varphi$. Notice that these transformations do not preserve the strategy context, and that they are different from the path quantifiers used in Game Logic (GL) [2]. There, the path quantifiers range over the set of outcomes of the strategies already in use. More precisely, in GL, we have

$$\mathcal{C}, \rho, i \models_f \mathbf{E}_{\text{GL}} \varphi_p \quad \text{iff} \quad \exists \rho' \in \text{Out}(\rho_{\leq i}, f). \mathcal{C}, \rho', i \models_f \varphi_p$$

The universal path quantifier is dual. Both path quantifiers can be expressed in ATL_{sc}^* as follows:

$$\mathbf{E}_{\text{GL}} \varphi_p \equiv \neg \langle \overline{\emptyset} \rangle \neg \varphi_p \quad \mathbf{A}_{\text{GL}} \varphi_p \equiv \langle \emptyset \rangle \varphi_p$$

4. Model checking

In this section, we present an algorithm for model checking ATL_{sc} and ATL_{sc}^* , and study its complexity. Our algorithm is based on a translation of the model-checking problem from ATL_{sc}^* into the model-checking problem for QCTL^* . Using a translation in the other direction, we prove that ATL_{sc}^* model checking is complete for k -EXPTIME (where k is the quantifier height of the formula).

4.1. From ATL_{sc}^* to QCTL^*

Let $\mathcal{C} = \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge} \rangle$ be a finite-state CGS, with a finite set of moves $\mathcal{M} = \{m_1, \dots, m_k\}$ and $\text{Agt} = \{a_1, \dots, a_n\}$. We consider the following sets of fresh atomic propositions: $\mathbf{P}_Q = \{\mathbf{p}_q \mid q \in Q\}$, $\mathbf{P}_{\mathcal{M}}^j = \{\mathbf{m}_1^j, \dots, \mathbf{m}_k^j\}$ for every $a_j \in \text{Agt}$, and write $\mathbf{P}_{\mathcal{M}} = \bigcup_{a_j \in \text{Agt}} \mathbf{P}_{\mathcal{M}}^j$. Let $\mathcal{S}_{\mathcal{C}}$ be the Kripke structure $\langle Q, R, \ell_+ \rangle$ where for any state q , we have: $\ell_+(q) = \ell(q) \cup \{\mathbf{p}_q\}$. A strategy for an agent a_j can be seen as a function $f_j: qQ_S^* \rightarrow \mathbf{P}_{\mathcal{M}}^j$ labeling the execution tree of $\mathcal{S}_{\mathcal{C}}$ with propositions in $\mathbf{P}_{\mathcal{M}}^j$.

Let C be a coalition in Agt , $f_C \in \text{Strat}(C)$ be a strategy context, and $\Phi \in \text{ATL}_{sc}^*$. We transform the question whether $\mathcal{C}, q \models_{f_C} \Phi$ into an instance of QCTL^* model checking over $\mathcal{S}_{\mathcal{C}}$ (assuming the tree semantics). For this, we define a QCTL^* formula $\overline{\Phi}^C$ inductively. Except for strategy quantifiers, the translation is straightforward:

$$\begin{aligned} \overline{\varphi \wedge \psi}^C &= \overline{\varphi}^C \wedge \overline{\psi}^C & \overline{\neg \varphi}^C &= \neg \overline{\varphi}^C & \overline{p}^C &= p \\ \overline{\langle A \rangle \varphi_s}^C &= \overline{\varphi_s}^{C \setminus A} & \overline{\varphi_p \mathbf{U} \psi_p}^C &= \overline{\varphi_p}^C \mathbf{U} \overline{\psi_p}^C & \overline{\mathbf{X} \varphi_p}^C &= \mathbf{X} \overline{\varphi_p}^C \end{aligned}$$

For a formula $\langle A \rangle \varphi_p$ with $A = \{a_{j_1}, \dots, a_{j_l}\}$, we let:

$$\begin{aligned} \overline{\langle A \rangle \varphi_p}^C &= \exists \mathbf{m}_1^{a_{j_1}} \dots \mathbf{m}_k^{a_{j_1}} \dots \mathbf{m}_1^{a_{j_l}} \dots \mathbf{m}_k^{a_{j_l}} \cdot \mathbf{p}_{\text{out}} \cdot \\ &\quad \left(\Phi_{\text{strat}}(A) \wedge \Phi_{\text{out}}(A \cup C) \wedge \mathbf{A} \left(\mathbf{G} \mathbf{p}_{\text{out}} \Rightarrow \overline{\varphi_p}^{C \cup A} \right) \right) \end{aligned}$$

where:

$$\begin{aligned} \Phi_{\text{strat}}(A) &= \bigwedge_{a \in A} \bigwedge_{q \in Q} \mathbf{AG} \left(\mathbf{p}_q \Rightarrow \bigvee_{m_i \in \text{Mov}(q, a)} (\mathbf{m}_i^a \wedge \bigwedge_{j \neq i} \neg \mathbf{m}_j^a) \right) \\ \Phi_{\text{out}}(A) &= \mathbf{p}_{\text{out}} \wedge \mathbf{AG} \left[\neg \mathbf{p}_{\text{out}} \Rightarrow \mathbf{AX} \neg \mathbf{p}_{\text{out}} \right] \wedge \mathbf{AG} \left[\mathbf{p}_{\text{out}} \Rightarrow \right. \\ &\quad \left. \bigvee_{q \in Q} \left(\mathbf{p}_q \wedge \bigvee_{\overline{m} \in \text{Mov}(q, A)} (\mathbf{p}_{\overline{m}} \wedge \mathbf{AX} ((\bigvee_{q' \in \text{Next}(q, A, \overline{m})} \mathbf{p}_{q'}) \Leftrightarrow \mathbf{p}_{\text{out}})) \right) \right] \end{aligned}$$

where \overline{m} is a move $(m_j)_{a \in A} \in \text{Mov}(q, A)$ for A and $\mathbf{p}_{\overline{m}}$ is the propositional formula $\bigwedge_{a \in A} \mathbf{m}_j^a$ characterizing \overline{m} . Formula $\Phi_{\text{strat}}(A)$ ensures that the labeling of propositions $\mathbf{m}_i^{a_j}$ describes a feasible strategy for A , while formula $\Phi_{\text{out}}(A)$ ensures that the outcomes of the strategy for A combined with the strategy context for C are labeled by the atomic proposition \mathbf{p}_{out} . Note that $\Phi_{\text{out}}(A)$ is based on the transition table Edge of \mathcal{C} (via $\text{Next}(q, A, m)$) and its size in

$O(|Q|^2 \cdot |\mathcal{M}|^{|A|})$ (i.e., in $O(|Q| \cdot |\text{Edge}|)$). For $\overline{(\bar{A})} \varphi^C$ and $\overline{\langle \bar{A} \rangle} \varphi_p^C$, we simply replace A with $\text{Agt} \setminus A$ and apply the same definitions as above.

Each strategy quantifier in the original ATL_{sc}^* formula induces a strategy quantifier in the QCTL^* formula. There are two exceptions to this rule, in which the translation can be adapted to not involve quantification:

- the first one is for formulas of the form $\overline{\langle \emptyset \rangle} \varphi_p^\emptyset$. For such formulas, the translation can be changed to $\mathbf{A} \overline{\varphi_p^\emptyset}$;
- the second case is for formulas of the form $\overline{\langle \emptyset \rangle} \varphi_p^C$ for which the labelling with \mathbf{p}_{out} does not have to be recomputed. This is the case if this formula is in the scope of another strategy quantifier with no $\langle _ \rangle$ operator inbetween. Then $\overline{\langle \emptyset \rangle} \varphi_p^C$ can be defined as $\mathbf{A}(\mathbf{G} \mathbf{p}_{\text{out}} \Rightarrow \overline{\varphi_p^C})$.

Such occurrences of modality $\langle \emptyset \rangle$ are said to be *trivial*.

Let $\mathcal{T}_C = \langle T, \ell \rangle$ be the execution tree of the Kripke structure \mathcal{S}_C , ρ be a path in \mathcal{S}_C , n be a position along ρ , and f_A be a strategy for some coalition A . Let ℓ' be a labelling extending ℓ with propositions $(\mathbf{m}_i^a)_{a \in A, 1 \leq i \leq k}$ and \mathbf{p}_{out} . We say that ℓ' is an f_A -labelling after $\rho_{\leq n}$ if, for every finite path π containing $\rho_{\leq n}$ as a prefix, it holds $\mathbf{m}_i^a \in \ell'(\pi)$ if, and only if, $f_A(a)(\pi) = m_i$, and $\mathbf{p}_{\text{out}} \in \ell'(\pi)$ if, and only if, π is compatible with f_A after $\rho_{\leq n}$.

For such an f_A -labelling ℓ' , we clearly have $\langle T, \ell' \rangle, \rho, n \models \Phi_{\text{strat}}(A) \wedge \Phi_{\text{out}}(A)$. The converse is also true: if $\langle T, \ell' \rangle, \rho, n \models \Phi_{\text{strat}}(A) \wedge \Phi_{\text{out}}(A)$, then ℓ' is an f_A -labelling after $\rho_{\leq n}$. Indeed, consider a path π from $\rho_{\leq n}$ and a position $i \geq n$. Then formula $\Phi_{\text{strat}}(A)$ enforces that node $\pi_{\leq i}$ is labeled with some $\bar{m} \in \text{Mov}(q, A)$. Now, by induction, one easily shows that the node corresponding to finite paths that are compatible with f_A after $\rho_{\leq n}$ are labeled with \mathbf{p}_{out} , while the other nodes are not labeled with \mathbf{p}_{out} .

The following result is a direct consequence of the above:

Theorem 11. *Let ρ be an infinite path in a CGS \mathcal{C} , and n be a position along ρ . Let Φ be an ATL_{sc}^* formula, and f_C be a strategy context for some coalition C . Let $\mathcal{T}_C(\rho(0)) = \langle T, \ell \rangle$ be the execution tree \mathcal{S}_C from $\rho(0)$, and ℓ' be an f_C -labelling extending ℓ . Then $\mathcal{C}, \rho, n \models_{f_C} \Phi$ if, and only if, $\langle T, \ell_{f_C} \rangle, \rho, n \models \overline{\Phi}^C$.*

Combined with the (non-elementary) decision procedure for QCTL^* model checking, we get a model-checking algorithm for model checking ATL_{sc}^* .

We now consider complexity issues more precisely. We have

$$\begin{aligned} |\overline{\Phi}^\emptyset| &= O(|\Phi| \cdot |Q|(|\text{Agt}| \cdot |\mathcal{M}|^2 + |Q| \cdot |\mathcal{M}|^{|\text{Agt}|})) \\ &= O(|\Phi| \cdot |Q|(|\text{Agt}| \cdot |\mathcal{M}|^2 + |\text{Edge}|)). \end{aligned}$$

which is polynomial in $|\Phi|$ and $|\mathcal{C}|$. Moreover, $\overline{\Phi}^\emptyset$ belongs to Q^kCTL^* , where k is the *depth of $\langle \cdot \rangle$* in Φ (which we define as the maximal number of nested *non-trivial* $\langle \cdot \rangle$ quantifiers). Given an ATL_{sc}^* formula Φ of depth k and a CGS \mathcal{C} , the reduction yields a model checking algorithm running in $(k + 1)$ -EXPTIME [14].

Finally note that when starting from an ATL_{sc} formula, the QCTL^* formula we obtain can easily be translated into QCTL : it contains the CTL^+ formula $\mathbf{A}(\mathbf{G} \mathbf{p}_{\text{out}} \Rightarrow \varphi_p)$, which can be succinctly written in CTL (for instance, when $\varphi_p = \varphi_s \mathbf{U} \psi_s$, $\mathbf{A}(\mathbf{G} \mathbf{p}_{\text{out}} \Rightarrow \varphi_p)$ is equivalent to $\mathbf{A}(\varphi_s \wedge \mathbf{p}_{\text{out}}) \mathbf{U} (\psi_s \vee \neg \mathbf{p}_{\text{out}})$). Thus it provides a Q^kCTL formula whose size is in $O(|\Phi| \cdot |Q|(|\text{Agt}| \cdot |\mathcal{M}|^2 + |\text{Edge}|))$ if Φ is an ATL_{sc} formula of $\langle \cdot \rangle$ -depth k . This yields a model-checking algorithm in k -EXPTIME.

4.2. From QCTL^* back to ATL_{sc}^*

We now propose a reduction in the converse direction, from an instance of the QCTL^* model-checking problem into an instance of the ATL_{sc}^* model-checking problem. Intuitively, strategies in the resulting game will correspond to labeling with atomic propositions in the original Kripke structure.

Let Φ be a QCTL formula and $\mathcal{S} = \langle Q, R, \ell \rangle$ be a Kripke structure. W.l.o.g., we assume that every atomic proposition in Φ is quantified at most once. We write $\text{AP}_f(\Phi)$ for the set of free atomic propositions in Φ (which are intended to already label the Kripke structure), and $\text{AP}_Q(\Phi) = \{P_1, \dots, P_k\}$ for the set of atomic propositions that are quantified in Φ . We build a TBGS $\mathcal{C}_\mathcal{S}$ and an ATL_{sc} formula $\tilde{\Phi}$ such that $\mathcal{S}, q \models_t \Phi$ if, and only if, $\mathcal{C}_\mathcal{S}, q \models \tilde{\Phi}$.

The game $\mathcal{C}_\mathcal{S} = \langle Q', R', \ell', \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge} \rangle$ is defined as follows. The set of agents is $\text{Agt} = \{a_0, \dots, a_k\}$: Player a_0 is in charge of selecting the transitions in \mathcal{S} , while Player a_i (for $i \geq 1$) has to decide the truth value of P_i . The set of states Q' is $Q \cup \{c_{q,i} \mid i = 1, \dots, k\} \cup \{p_i \mid i = 0, \dots, k\}$: states $q \in Q$ are controlled by a_0 , while states $c_{q,i}$ are controlled by a_i . States p_i only carry a self-loop, and are not explicitly controlled. The transition set R' contains every transition $(q, q') \in R$, and also transitions $(q, c_{q,i})$, $(c_{q,i}, p_i)$, and $(c_{q,i}, p_0)$ for $i = 1, \dots, k$ and $q \in Q$. The labeling ℓ' is as follows: $\ell'(q) = \ell(q) \cup \{P_Q\}$ if

$q \in Q$ (P_Q is a fresh atomic proposition), $\ell'(c_{q,i}) = \ell'(p_0) = \emptyset$, and $\ell'(p_i) = P_i$ for $i \geq 1$.

In a state $q \in Q$, a_0 can choose either a successor state q' (i.e., an \mathcal{S} -transition (q, q')) or some $c_{q,i}$, the latter choice being used to check whether P_i holds true in q . Indeed in $c_{q,i}$, a_i has two available moves: move \mathbf{m}_1 goes to P_i , and while mode \mathbf{m}_0 goes to P_0 . Thus as soon as a_i has selected a strategy, $c_{q,i}$ has a unique successor; this encodes the labeling for atomic proposition $P_i \in \text{AP}_{\mathcal{Q}}(\Phi)$. Note also that for any path in $\mathcal{C}_{\mathcal{S}}$ of the form $\rho \cdot c_{q,i}$, ρ is a path in \mathcal{S} ending in q . Finally note that as $\mathcal{C}_{\mathcal{S}}$ is a TBGS, its size is in $O(|Q| \cdot |\Phi| + |R|)$, i.e. in $O(|\mathcal{S}| \cdot |\Phi|)$.

Following the ideas above, we define $\widetilde{\Phi}$ inductively:

$$\begin{aligned} \widetilde{\neg\psi} &= \neg\widetilde{\varphi} & \widetilde{\varphi \wedge \psi} &= \widetilde{\varphi} \wedge \widetilde{\psi} & \widetilde{\exists P_i. \varphi} &= \langle a_i \rangle \widetilde{\varphi} \\ \widetilde{\mathbf{X}\varphi} &= \mathbf{X}\widetilde{\varphi} & \widetilde{\varphi \mathbf{U} \psi} &= \widetilde{\varphi} \mathbf{U} \widetilde{\psi} & \widetilde{\mathbf{E}\varphi} &= \langle a_0 \rangle (\mathbf{G} P_Q \wedge \widetilde{\varphi}) \\ \widetilde{P_i} &= \begin{cases} \langle a_0 \rangle \mathbf{X} \langle a_0 \rangle \mathbf{X} P & \text{if } P \in \text{AP}_{\mathcal{Q}}(\Phi) \\ P & \text{otherwise} \end{cases} \end{aligned}$$

The size of $\widetilde{\Phi}$ is in $O(|\Phi|)$. We state the correctness of the reduction as follows:

Proposition 12. *Let Φ be a QCTL* formula with $\text{AP}_{\mathcal{Q}}(\Phi) = \{P_1, \dots, P_k\}$ and ψ be a Φ -subformula. Let \mathcal{I} be the indices of propositions in $\text{AP}_f(\psi) \cap \text{AP}_{\mathcal{Q}}(\Phi)$. Let $\mathcal{S} = \langle Q, R, \ell \rangle$ be a KS, ρ a path in \mathcal{S} and n a position along ρ . Let $\mathcal{T}_{\mathcal{S}}(\rho(0)) = \langle T, \ell_{\mathcal{I}} \rangle$ be the execution tree from $\rho(0)$ with a labeling function $\ell_{\mathcal{I}}$ that extends ℓ for $\{P_i \mid i \in \mathcal{I}\}$. Let f be the strategy context for $\{a_i\}_{i \in \mathcal{I}}$ such that: $f(a_i)(\rho' \cdot c_{q,i}) = \mathbf{m}_1$ iff $\ell_{\mathcal{T}}(\rho') \ni P_i$ for every \mathcal{T} -node ρ' . Then we have:*

$$\langle T, \ell_{\mathcal{I}} \rangle, \rho, n \models_s \psi \quad \text{iff} \quad \mathcal{C}_{\mathcal{S}}, \rho, n \models_f \widetilde{\psi}$$

Proof. The proof is based on the fact that any strategy for agent a_i in $\mathcal{C}_{\mathcal{S}}$ corresponds to a (unique) P_i labeling of T . Indeed such a strategy is a mapping from paths of the form $\rho' \cdot c_{q,i}$ with $\rho' \in Q^+$ and as noticed above, we have: for any such a $\mathcal{C}_{\mathcal{S}}$ path $\rho' \cdot c_{q,i}$, ρ' is a path in \mathcal{S} ending in q which implies that ρ is a state of T . Using this observation, the proof is straightforward. \square

A complexity lower bound for ATL_{sc} model checking. Now consider a model checking instance $\mathcal{C} \models \Phi$ with $\Phi \in \text{EQ}^k\text{CTL}$. This problem is k -EXPTIME-complete [14]. Now we can adapt the previous reduction from QCTL to ATL_{sc} in order to obtain a formula with $\langle _ \rangle$ -depth equals to k :

- First we can replace all occurrences of the quantifier $\langle a_0 \rangle$ in $\widetilde{\Phi}$ with $\neg \langle \emptyset \rangle \neg$, because every such subformula is interpreted in a context where every agent in $\{a_1, \dots, a_k\}$ has a fixed strategy, so that quantifying on the ability of the last agent (a_0) to select a path in the structure is equivalent to looking for a path in the outcomes of the context (and $\mathbf{E}_c \equiv \neg \langle \emptyset \rangle \neg$). Notice that such $\langle \emptyset \rangle$ correspond to *trivial* quantification, because no $(_)$ is used in the formula.
- Second, formula $\widetilde{\mathbf{E}}\varphi$ is in ATL_{sc}^* , as it involves a conjunction of path formulas in the scope of a strategy quantifier. However, this is easily overcome in pretty much the same way as we did at the end of Section 4.1.

From this reduction, we get a k -EXPTIME-hardness lower bound for ATL_{sc} model checking (notice that this already holds for TBGSs). The same approach can be followed for EQ^kCTL^* and ATL_{sc}^* , yielding a $(k + 1)$ -EXPTIME-hardness bound. As a result:

Theorem 13. *Model checking the fragment of ATL_{sc} (resp. ATL_{sc}^*) with at most k non-trivial nested strategy quantifiers is k -EXPTIME-hard (resp. $(k + 1)$ -EXPTIME-hard), even for TBGS.*

To sum up our results about model checking, we have:

Corollary 14. *Model checking ATL_{sc} and ATL_{sc}^* is Tower-complete. More precisely, model checking the fragment of ATL_{sc} (resp. ATL_{sc}^*) with k non-trivial nested strategy quantifiers is k -EXPTIME-complete (resp. $(k + 1)$ -EXPTIME-complete).*

5. Satisfiability

The above translation to QCTL^* works for model checking, but does not extend to satisfiability: the QCTL^* formula we build depends both on the formula and on the structure. Actually, as was proved recently in [24], satisfiability is undecidable for ATL_{sc} , both when looking for infinite or finite CGS. For the sake of completeness, and because it has interesting consequences, we sketch a (modified) proof of this result in this section.

We then establish decidability of satisfiability in two different settings: first when restricting to turn-based games, and then in the case where the set of actions allowed to the players is fixed. A consequence of our decidability

proofs (based on automata constructions) is that in both settings, ATL_{sc} does have the finite-model property (thanks to Rabin's regularity theorem).

Before we proceed to the algorithms for satisfiability, we prove a generic result about the number of players needed in a game in order to satisfy a formula involving a given set of agents. This result has already been proved for ATL (*e.g.* in [26]). Given a formula $\Phi \in \text{ATL}_{sc}^*$, we use $\text{Agt}(\Phi)$ to denote the set of all agents involved in the strategy quantifiers of Φ .

Proposition 15. *An ATL_{sc}^* formula Φ is satisfiable if, and only if, it is satisfiable in a CGS with $|\text{Agt}(\Phi)| + 1$ agents.*

Proof. Assume that Φ is satisfied in a CGS $\mathcal{C} = \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge} \rangle$. If $|\text{Agt}| \leq |\text{Agt}(\Phi)|$, it suffices to add extra idle players in \mathcal{C} . Otherwise, if $|\text{Agt}| > |\text{Agt}(\Phi)| + 1$, we can replace the agents $\{b_1, \dots, b_k\}$ in Agt that do not belong to $\text{Agt}(\Phi)$ by a unique agent a mimicking the actions of the removed players. Notice that this requires extending the set of moves for Player a to \mathcal{M}^k . \square

5.1. General case

In [24], Troquard and Walther show that satisfiability of ATL_{sc} is undecidable. The proof consists in reducing the satisfiability problem for the modal logic $S5^n$ to the satisfiability problem for ATL_{sc} . The construction is elegant and induces several important result, which is the reason why we include it here (with a few changes).

The multi-modal logic $S5^n$. The logic $S5^n$ is a multidimensional modal logic [13], whose formulas are built from boolean operators, atomic propositions $P \in \text{AP}$ and modalities \diamond_i . These formulas are interpreted over models $\mathcal{M} = \langle \mathcal{F}, \mathcal{V} \rangle$ where \mathcal{F} is a product frame $W_1 \times \dots \times W_n$, and \mathcal{V} is a valuation for atomic propositions over \mathcal{F} . The (implicit) transition relation over W_i is universal: for any world $w = (w_1, \dots, w_n)$ and any $w'_i \in W_i$, there is an i -transition to $(w_1, \dots, w_{i-1}, w'_i, w_{i+1}, \dots, w_n)$. This provides the semantics of $\diamond_i \varphi$: $\mathcal{M}, w \models \diamond_i \varphi$ if, and only if, there exists $w'_i \in W_i$ such that $\mathcal{M}, w[w_i \rightarrow w'_i] \models \varphi$. Satisfiability (both over finite and infinite models) is undecidable for $S5^n$ [25]. Additionally, $S5^n$ does not have the finite-model property [12].

Let Φ be an $S5^n$ formula. From Φ , we build an ATL_{sc} formula $\widehat{\Phi}$ inductively as follows:

$$\widehat{\varphi \wedge \psi} = \widehat{\varphi} \wedge \widehat{\psi} \quad \widehat{\neg \psi} = \neg \widehat{\psi} \quad \widehat{P} = \langle \emptyset \rangle \mathbf{X} P \quad \widehat{\diamond_i \psi} = \langle a_i \rangle \widehat{\psi}$$

The following result connects both satisfiability problems:

Proposition 16 ([24]). *Let Φ be an $S5^n$ formula and $\widehat{\Phi}$ be the resulting ATL_{sc} formula, obtained as above. Then Φ is satisfiable in a finite (resp. infinite) $S5^n$ model if, and only if, $\langle \emptyset \rangle \widehat{\Phi}$ is satisfiable in a finite (resp. infinite) CGS.*

Proof. First assume that there exists a model $\mathcal{M} = (\mathcal{F}, \mathcal{V})$ for Φ , with $\mathcal{F} = W_1 \times \dots \times W_n$. Take w such that $\mathcal{M}, w \models \Phi$. For every i , the states in W_i are denoted w_1^i, w_2^i, \dots . We define a CGS $\mathcal{C}_{\mathcal{M}} = \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge} \rangle$ with \mathcal{M} as its underlying transition system: $Q = \mathcal{F}$, $R = Q \times Q$, and $\ell(w) = \mathcal{V}(w)$. We let $\text{Agt} = \{a_1, \dots, a_n\}$. The action alphabet \mathcal{M} is $\{1, \dots, \max_{1 \leq i \leq n} |W_i|\}$ if \mathcal{M} is finite, and $\mathbb{N}_{>0}$ otherwise. In every world $w \in Q$, Player a_i can choose the next position in W_i : in other terms, $\text{Mov}(w, a_i) = \{1, \dots, |W_i|\}$ if W_i is finite, or $\mathbb{N}_{>0}$ otherwise, and $\text{Edge}(w, \bar{m}) = w_{\bar{m}}$ with $w_{\bar{m}} = \langle w_{m_1}^1, w_{m_2}^2, \dots, w_{m_n}^n \rangle$ when \bar{m} is $\langle m_1, \dots, m_n \rangle$. Note that as \mathcal{M} is universal, the transition table does not depend on the current state w . As a world w in \mathcal{M} is also a state in $\mathcal{C}_{\mathcal{M}}$ and also corresponds to a move in the game structure, we might abusively write $\text{Edge}(w, \bar{m}) = \bar{m}$ or $\text{Edge}(w, w') = w'$ in the sequel.

In our reduction, formula $\widehat{\Phi}$ only involves non-nested occurrences of the \mathbf{X} modality. Hence we are only interested in the first move proposed by strategies. Given a world w in \mathcal{F} , we write F_w for the class of all strategies for Agt such that $F_w(\emptyset) = w$. In other terms, any strategy in F_w enforces the first transition to go to w . Now we can easily see that, for every $S5^n$ formula φ and for any two worlds w and w' :

$$\mathcal{M}, w \models \varphi \quad \Leftrightarrow \quad \mathcal{C}_{\mathcal{M}}, w' \models_{F_w} \widehat{\varphi}$$

The proof is by structural induction over φ . The cases of boolean operators are direct. We only consider \diamond_i and propositions P :

- when $\varphi = P$: $\mathcal{M}, w \models P$ is equivalent to $P \in \mathcal{V}(w)$, which in turn is equivalent to $\mathcal{C}_{\mathcal{M}}, w' \models_{F_w} \langle \emptyset \rangle \mathbf{X}P$ because the strategy F_w ensures a first transition to w , where P holds true.
- when $\varphi = \diamond_i \psi$: if $\mathcal{M}, w \models \diamond_i \psi$, then there exists w'' such that $w''_j = w_j$ for $j \neq i$ and $\mathcal{M}, w'' \models \psi$. By induction hypothesis, we get $\mathcal{C}_{\mathcal{M}}, w' \models_{F_{w''}} \widehat{\psi}$, which implies $\mathcal{C}_{\mathcal{M}}, w' \models_{F_w} \langle a_i \rangle \widehat{\psi}$ because $F_{w''}$ is clearly equivalent (when considering only the first transition) to F_w modified by a new move for Player a_i .

Now consider any state w' in Q . Since $\mathcal{M}, w \models \Phi$, we have $\mathcal{C}_{\mathcal{M}}, w' \models \langle \overline{\emptyset} \rangle \widehat{\Phi}$, since the strategy quantifier $\langle \overline{\emptyset} \rangle$ (i.e. $\langle \text{Agt} \rangle$) allows us to select a strategy in F_w to ensure $\mathcal{C}_{\mathcal{M}}, w' \models_{F_w} \widehat{\Phi}$. Thus $\langle \overline{\emptyset} \rangle \widehat{\Phi}$ is satisfiable.

Now assume that there exists $\mathcal{C} = \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge} \rangle$ and $q \in Q$ such that $\mathcal{C}, q \models \langle \overline{\emptyset} \rangle \widehat{\Phi}$. We first show that there is such a structure involving n agents. We then deduce that the $S5^n$ formula Φ is satisfiable.

Assume that \mathcal{C} involves $n + 1$ players $\{a_0, a_1, \dots, a_n\}$ (following Proposition 15). As $\mathcal{C}, q \models \langle \overline{\emptyset} \rangle \widehat{\Phi}$, there exists a strategy F for $\{a_0, a_1, \dots, a_n\}$ such that $\mathcal{C}, q \models_F \widehat{\Phi}$. Pick such an F . In $\widehat{\Phi}$, the strategy quantifiers only deal with Players a_1 to a_n . The strategy (and in particular the first move) for a_0 is fixed by F , and is not updated by $\widehat{\Phi}$. Now, consider the structure $\mathcal{C}' = \langle Q, R', \ell, \{a_1, \dots, a_n\}, \mathcal{M}, \text{Mov}', \text{Edge}' \rangle$, in which $\text{Edge}'(q, \langle m_1, \dots, m_n \rangle)$ is defined as $\text{Edge}(q, \langle m_0, m_1, \dots, m_n \rangle)$, where m_0 is the first move proposed by F_{a_0} . Then $\mathcal{C}', q \models \langle \overline{\emptyset} \rangle \widehat{\Phi}$.

We now pick an n -player CGS \mathcal{C} such that $\mathcal{C}, q \models \langle \overline{\emptyset} \rangle \widehat{\Phi}$. We define $\mathcal{M}_{\mathcal{C}} = (\mathcal{F}, \mathcal{V})$ as follows: $\mathcal{F} = W_1 \times \dots \times W_n$ with $W_i = \text{Mov}(q, a_i)$, and $\mathcal{V}(\overline{m}) = \ell(\text{Edge}(q, \overline{m}))$. In other terms, the states of $\mathcal{M}_{\mathcal{C}}$ are the move vectors of \mathcal{C} , and they are labeled with the atomic propositions of the state they lead to. Associated with a universal relation, this defined an $S5^n$ model. We now show that it is a model for Φ .

Given a strategy F in \mathcal{C} for every player, $F(q)$ defines a unique move vector $\overline{m}_{F(q)}$ from q ; it also correspond to some world $w_{F(q)}$ in $\mathcal{M}_{\mathcal{C}}$. We clearly have:

$$\mathcal{C}, q \models_F \widehat{\varphi} \quad \Leftrightarrow \quad \mathcal{M}_{\mathcal{C}}, w_{F(q)} \models \varphi$$

The proof works exactly as in the previous case. Finally we have $\mathcal{C}, q \models \langle \overline{\emptyset} \rangle \widehat{\Phi}$, and then there exists a complete strategy F such that $\mathcal{C}, q \models_F \widehat{\Phi}$. From the previous result, we get $\mathcal{M}_{\mathcal{C}}, w_{F(q)} \models \Phi$, and then Φ is satisfiable.

Finally note that \mathcal{M} is infinite if, and only if, $\mathcal{C}_{\mathcal{M}}$ has an infinite action alphabet. Conversely, \mathcal{C} has an infinite action alphabet if, and only if, $\mathcal{M}_{\mathcal{C}}$ is infinite. Therefore Φ is finitely (resp. infinitely) satisfiable if, and only if, $\langle \overline{\emptyset} \rangle \widehat{\Phi}$ is satisfiable in a finite (resp. infinite) CGS¹ \square

Proposition 16 and its proof entail the following results:

¹Notice that finiteness refers here to the total size of the CGS, not only its number of states.

Corollary 17. • *Satisfiability of ATL_{sc} is undecidable for finite or infinite CGS [24].*

- ATL_{sc} does not have the finite-model property.
- ATL_{sc} does not have the finite-branching property.

5.2. Turn-based games

In this section, we consider the restriction of satisfiability to turn-based games: given an ATL_{sc} formula, we look for a *turn-based* game structure satisfying Φ . As we now explain, this problem turn out to be decidable.

Let Φ be an ATL_{sc} formula. Write $\text{Agt}(\Phi) = \{a_1, \dots, a_n\}$ for the set of players involved in Φ . Following Prop. 15, let Agt be the set $\text{Agt}(\Phi) \cup \{a_0\}$, where a_0 is an additional player. Pick a TBGS \mathcal{C} , and consider its execution tree $\mathcal{T}_{\mathcal{C}} = \langle T, \ell \rangle$. Since \mathcal{C} is turn-based, we may assume that ℓ labels each node of the tree with the owner of the corresponding state. Formally, for each node $\pi \in T$, we have $\ell(\pi) \ni \text{turn}_j$ if, and only if, $a_j = \text{Own}(\text{last}(\pi))$. A strategy for an agent a_j is then encoded using an atomic proposition mov_j : indeed, a strategy for a_j precisely corresponds to a selection of one *successor* of every turn_j -state (notice that this is a crucial difference with CGSs). The outcomes of a strategy of Player a_j are the runs in which every turn_j -state is followed by a mov_j -state; The decidability proof now consists in using this encoding of strategies together with the translation from ATL_{sc}^* into QCTL^* .

We again consider the execution tree $\mathcal{T}_{\mathcal{C}} = \langle T, \ell \rangle$ of \mathcal{C} . Given a coalition $B \subseteq \text{Agt}$ and a strategy f_B for coalition B , a labeling function ℓ' extending ℓ is an f_B -*tb-labeling* if ℓ' labels T with propositions mov_j for all $a_j \in B$ as dictated by f_B , and with proposition \mathbf{p}_{out} in order to mark outcomes. More precisely, (1) for any node π labeled with $\text{turn}_j \in B$, we have $\ell'(\pi \cdot q) = \text{mov}_j$ if, and only if, $f_B(j)(\pi) = q$, and (2) when a state is labeled with \mathbf{p}_{out} and turn_j for some $a_j \in B$, then only its mov_j -successor is labeled with \mathbf{p}_{out} ; states labeled with \mathbf{p}_{out} and turn_k with $a_k \notin B$, then all its successors are labeled with \mathbf{p}_{out} ; finally, all the successors of nodes not labeled with \mathbf{p}_{out} are not labeled with \mathbf{p}_{out} .

Given a coalition C (which we intend to represent the agents that have a strategy in the current context), we translate an ATL_{sc}^* formula Φ into a

QCTL* formula $\widehat{\Phi}^C$ inductively as follows:

$$\begin{array}{lll} \widehat{\langle A \rangle \varphi}^C = \widehat{\varphi}^{C \setminus A} & \widehat{\varphi \wedge \psi}^C = \widehat{\varphi}^C \wedge \widehat{\psi}^C & \widehat{\neg \psi}^C = \neg \widehat{\varphi}^C \\ \widehat{\varphi \mathbf{U} \psi}^C = \widehat{\varphi}^C \mathbf{U} \widehat{\psi}^C & \widehat{\mathbf{X} \varphi}^C = \mathbf{X} \widehat{\varphi}^C & \widehat{P}^C = P \end{array}$$

For formulas of the form $\langle A \rangle \varphi$ with $A = \{a_{j_1}, \dots, a_{j_l}\}$, we let:

$$\begin{aligned} \widehat{\langle A \rangle \varphi}^C &= \exists \text{mov}_{j_1} \dots \text{mov}_{j_l} \cdot \mathbf{p}_{\text{out}} \cdot \\ &\quad \left(\Phi_{\text{strat}}^{\text{tb}}(A) \wedge \Phi_{\text{out}}^{\text{tb}}(A \cup C) \wedge \mathbf{A} \left(\mathbf{G} \mathbf{p}_{\text{out}} \Rightarrow \widehat{\varphi}_p^{C \cup A} \right) \right) \end{aligned}$$

with:

$$\begin{aligned} \Phi_{\text{strat}}^{\text{tb}}(A) &= \mathbf{AG} \bigwedge_{a_j \in A} (\text{turn}_j \Rightarrow \mathbf{EX}_1 \text{mov}_j) \\ \Phi_{\text{out}}^{\text{tb}}(A) &= \mathbf{p}_{\text{out}} \wedge \mathbf{AG} \left[\neg \mathbf{p}_{\text{out}} \Rightarrow \mathbf{AX} \neg \mathbf{p}_{\text{out}} \right] \wedge \\ &\quad \mathbf{AG} \left[\mathbf{p}_{\text{out}} \Rightarrow \left(\bigwedge_{a_j \in A} (\text{turn}_j \Rightarrow \mathbf{AX} (\text{mov}_j \Leftrightarrow \mathbf{p}_{\text{out}})) \right) \right] \end{aligned}$$

where $\mathbf{EX}_1 \alpha$ is a shorthand for $\mathbf{EX} \alpha \wedge \forall p. (\mathbf{EX} (\alpha \wedge p) \Rightarrow \mathbf{AX} (\alpha \Rightarrow p))$, specifying the existence of a unique successor satisfying α .

Now we have the following proposition, whose proof is done by structural induction over the formula:

Proposition 18. *Let $\Phi \in \text{ATL}_{sc}^*$, and $\text{Agt} = \text{Agt}(\Phi) \cup \{a_0\}$ as above. Let \mathcal{C} be a TBGS, ρ be a path of \mathcal{C} , n be a position along ρ , and f_B be a strategy context whose domain is $B \subseteq \text{Agt}$. Let $\mathcal{T}_{\mathcal{C}}(\rho(0)) = \langle T, \ell_{f_B} \rangle$ be the execution tree of the Kripke structure underlying \mathcal{C} from $\rho(0)$, labeled with an f_B -tb-labeling ℓ_{f_B} . Then we have:*

$$\mathcal{C}, \rho, n \models_{f_B} \Phi \quad \text{iff} \quad \langle T, \ell_{f_B} \rangle, \rho, n \models \widehat{\Phi}^B$$

Proof. The proof is by structural induction over Φ . The cases of atomic propositions, boolean operators and temporal modalities are straightforward.

- If $\Phi = \langle A \rangle \varphi_p$: assume $\mathcal{C}, \rho, n \models_{f_B} \Phi$. Then there exists $g_A \in \text{Strat}(A)$ such that for any $\rho' \in \text{Out}(\rho_{\leq n}, g_A \circ f_B)$, we have $\mathcal{C}, \rho', n \models_{(g_A \circ f_B)} \varphi$.

Let $\ell_{g_A \circ f_B}$ be some $(g_A \circ f_B)$ -tb-labeling built from ℓ_{f_B} by updating the labeling for propositions $(\mathbf{mov}_j)_{a_j \in A}$ and \mathbf{p}_{out} . By induction hypothesis, for any $\rho' \in \mathbf{Out}(\rho_{\leq n}, g_A \circ f_B)$, we have $\langle T, \ell_{(g_A \circ f_B)} \rangle, \rho', n \models \widehat{\varphi}_p^{A \cup B}$. Then, by definition of $\ell_{g_A \circ f_B}$, the outcomes generated by $g_A \circ f_B$ are exactly the run satisfying $\mathbf{G} \mathbf{p}_{\text{out}}$, and then we clearly have $\langle T, \ell_{(g_A \circ f_B)} \rangle, \rho, n \models \mathbf{A}(\mathbf{G} \mathbf{p}_{\text{out}} \Rightarrow \widehat{\varphi}_p^{A \cup B})$. Moreover $\Phi_{\text{strat}}^{\text{tb}}(A)$ and $\Phi_{\text{out}}^{\text{tb}}(A \cup B)$ also hold true for $\langle T, \ell_{(g_A \circ f_B)} \rangle, \rho, n$. Therefore we have $\langle T, \ell_{f_B} \rangle, \rho, n \models \widehat{\langle A \rangle}_{\varphi_p}^B$, since it suffices to extend ℓ_{f_B} in order to encode the strategy g_A and update the truth value of \mathbf{p}_{out} accordingly.

Conversely, assume $\langle T, \ell_{f_B} \rangle, \rho, n \models \widehat{\langle A \rangle}_{\varphi_p}^B$. Then there exists a labeling ℓ' for $(\mathbf{mov}_j)_{a_j \in A}$ and for \mathbf{p}_{out} so that (1) $\Phi_{\text{strat}}^{\text{tb}}(A)$ holds true, which ensures that the labeling with $(\mathbf{mov}_j)_{a_j \in A}$ corresponds to a strategy g_A for A from $\rho_{\leq n}$, and (2) $\Phi_{\text{out}}^{\text{tb}}(A \cup B)$ also holds true, which ensures that \mathbf{p}_{out} marks the outcomes from $\rho_{\leq n}$ induced by $g_A \circ f_B$. This implies that ℓ' is a $(g_A \circ f_B)$ -tb-labeling. Finally we also know that $\mathbf{A}(\mathbf{G} \mathbf{p}_{\text{out}} \Rightarrow \widehat{\varphi}_p^{A \cup B})$ holds for $\langle T, \ell_{f_B} \rangle, \rho, n$, which entails that every outcome of $g_A \circ f_B$ satisfies $\widehat{\varphi}_p^{A \cup B}$. The induction hypothesis entails the expected result.

- If $\Phi = \langle A \rangle \varphi_s$: assume $\mathcal{C}, \rho, n \models_{f_B} \Phi$. Then $\mathcal{C}, \rho, n \models_{f_B \setminus A} \varphi_s$. Applying the induction hypothesis, we get $\langle T, \ell_{f_B \setminus A} \rangle, \rho, n \models \widehat{\varphi}_s^{B \setminus A}$. It follows that $\langle T, \ell_{f_B} \rangle, \rho, n \models \widehat{\varphi}_s^{B \setminus A}$, because the labeling of strategies for coalition A in f_B is not used for evaluating $\widehat{\varphi}_s^{B \setminus A}$, and the labeling with proposition \mathbf{p}_{out} will be updated at the next occurrence of a $\langle _ \rangle$ quantifier.

Conversely, assume that $\langle T, \ell_{f_B} \rangle, \rho, n \models \widehat{\varphi}_s^{B \setminus A}$. For the same reason as above, we have $\langle T, \ell_{f_B \setminus A} \rangle, \rho, n \models \widehat{\varphi}_s^{B \setminus A}$. Applying the induction hypothesis, we get $\mathcal{C}, \rho, n \models_{f_B \setminus A} \varphi_s$, and then $\mathcal{C}, \rho, n \models_{f_B} \Phi$. \square

Finally, it remains to enforce that the Kripke structure satisfying $\widehat{\Phi}^\emptyset$ corresponds to a turn-based game structure. This is achieved by also requiring

$$\Phi_{\text{tb}} = \mathbf{AG} \left[\bigvee_{a_j \in \text{Agt}} \left(\text{turn}_j \wedge \bigwedge_{a_l \neq a_j} \neg \text{turn}_l \right) \right].$$

Finally, we let $\widehat{\Phi}$ be the formula $\Phi_{\text{tb}} \wedge \widehat{\Phi}^\emptyset$.

Proposition 19. *Let Φ be an ATL_{sc}^* formula and $\widehat{\Phi}$ be the QCTL* formula defined as above. Then Φ is satisfiable in a TBGS if, and only if, $\widehat{\Phi}$ is satisfiable (in the tree semantics).*

Proof. If Φ is satisfiable in a TBGS, then there exists such a structure \mathcal{C} with $|\text{Agt}(\Phi)|+1$ agents. Pick such a structure \mathcal{C} , and a path ρ such that $\mathcal{C}, \rho, 0 \models \Phi$. Now consider the execution tree $\mathcal{T}_{\mathcal{C}}(\rho(0)) = \langle T, \ell \rangle$. From Proposition 18, we have $\langle T, \ell \rangle, \rho, 0 \models \widehat{\Phi}^\emptyset$. Thus clearly $\langle T, \ell \rangle, \rho, 0 \models \widehat{\Phi}$.

Conversely assume $\mathcal{T} \models \widehat{\Phi}$. Thus $\mathcal{T} \models \Phi_{\text{tb}} \wedge \widehat{\Phi}^\emptyset$: the first part of the formula ensures that every state of the Kripke structure can be assigned to a unique agent, hence defining a TBGS. The second part ensures that Φ holds in the corresponding game, thanks to Proposition 18. \square

The above translation from ATL_{sc}^* into QCTL* transforms a formula with k strategy quantifiers into a formula with at most $k+1$ nested blocks of quantifiers. By slightly modifying the definition of $\widehat{\langle A \rangle \varphi_p}^{\mathcal{C}}$, we can obtain a translation from ATL_{sc} into QCTL with the same property. Satisfiability of a QCTL* (resp. QCTL) formula with $k+1$ blocks of quantifiers is in $(k+3)$ -EXPTIME (resp. $(k+2)$ -EXPTIME) [14]. Hence the algorithm is in Tower.

We now prove that this high complexity cannot be avoided:

Proposition 20. *Satisfiability of ATL_{sc} and ATL_{sc}^* formulas over TBGSs is Tower-hard (i.e., it is k -EXPTIME-hard, for all k).*

Proof (sketch). Model checking ATL_{sc} over turn-based games is Tower-hard (Theorem 13), and it can easily be encoded as a satisfiability problem. Indeed, let $\mathcal{C} = \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge} \rangle$ be a TBGS, and Φ be an ATL_{sc} formula. Let $(\mathbf{p}_q)_{q \in Q}$ be fresh atomic propositions. We define an ATL_{sc} formula $\Psi_{\mathcal{C}}$ to describe the game \mathcal{C} as follows:

$$\Psi_{\mathcal{C}} = \mathbf{AG} \left(\bigvee_{q \in Q} (\mathbf{p}_q \wedge \bigwedge_{q' \neq q} \neg \mathbf{p}_{q'} \wedge \bigwedge_{P \in \ell(q)} P \wedge \bigwedge_{P' \notin \ell(q)} \neg P') \right) \wedge \mathbf{AG} \left[\bigwedge_{q \in Q} \left(\mathbf{p}_q \Rightarrow \left(\bigwedge_{q \rightarrow q'} \langle \langle \text{Own}(q) \rangle \rangle \mathbf{X} \mathbf{p}_{q'} \wedge \bigwedge_{q'. q \not\rightarrow q'} \neg \langle \langle \text{Own}(q) \rangle \rangle \mathbf{X} \mathbf{p}_{q'} \right) \right) \right]$$

where $q \rightarrow q'$ denotes the existence of a transition from q to q' in \mathcal{C} . Any TBGS \mathcal{C}' satisfying $\Psi_{\mathcal{C}}$ corresponds to a kind of unfolding of \mathcal{C} with possibly duplications of transitions (and of the corresponding moves). First note that

duplicating transitions does not change the truth value of ATL_{sc} formula: in a turn-based structure, duplicating a transition consists in adding a new move for the owner of the source state and this move is completely equivalent to the previous move. Thus we can assume that \mathcal{C}' corresponds to some unfolding of \mathcal{C} (with extra labeling for propositions \mathbf{p}_q), thus it yields the same execution tree as \mathcal{C} ; this ensures that both structures satisfy the same ATL_{sc} formulas when the extra propositions are not used. In particular, \mathcal{C}' satisfies Φ if, and only if, \mathcal{C} does.

Finally we clearly have that $\mathcal{C}, q \models \Phi$ if, and only if, $\Psi_{\mathcal{C}} \wedge \mathbf{p}_q \wedge \Phi$ is satisfiable in a turn-based structure. \square

Theorem 21. *Satisfiability for ATL_{sc} and ATL_{sc}^* over TBGSs is Tower-complete.*

We conclude this section with proving that in TBGS, ATL_{sc}^* has the finite-model property:

Proposition 22. *If an ATL_{sc}^* formula Φ is satisfiable in a TBGS, then there exists a finite TBGS satisfying Φ .*

Proof. Assume Φ is satisfiable, then the QCTL^* formula $\widehat{\Phi}$ is satisfiable and there exists a tree satisfying $\widehat{\Phi}$. Such a tree \mathcal{T} can be chosen to be regular (QCTL^* models can be characterized by alternating parity tree automata [14]); we can consider the underlying finite Kripke structure $K_{\mathcal{T}}$, and apply the same construction as we did for proving Proposition 19, obtaining a finite TBGS satisfying Φ . \square

5.3. Games with a bounded action alphabet

We now consider another setting where the reduction to QCTL^* can be used to solve the satisfiability for ATL_{sc}^* : we look for CGSs involving a given set of moves, and a given set of players². Formally, the problem is defined as follows:

- Problem:** $(\text{Agt}, \mathcal{M})$ -satisfiability
- Input:** a finite set of moves \mathcal{M} , a set of agents Agt , and an ATL_{sc}^* formula Φ involving the agents in Agt ;
- Question:** does there exist a CGS $\mathcal{C} = \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge} \rangle$ and a state $q \in Q$ such that $\mathcal{C}, q \models \Phi$.

²Notice that when the set of moves is fixed, we cannot apply Prop. 15.

We fix $\mathcal{M} = \{m_1, \dots, m_k\}$ and $\mathbf{Agt} = \{a_1, \dots, a_n\}$. Moreover we assume w.l.o.g. that every agent may choose all k moves from every state (*i.e.*, $\mathbf{Mov}(q, a) = \mathcal{M}$ for any $q \in Q$ and $a \in \mathbf{Agt}$). Therefore we know that we are looking for a CGS whose execution tree is a $[k^n]$ -tree. We encode the transition table of the CGS in its execution tree as follows: for every agent a_i and move m_j in \mathcal{M} , we label a node with an atomic proposition $\mathbf{after}_{m_j}^{a_i}$ to specify that agent a_i has played move m_j in the *parent* node. The resulting execution tree satisfies formula

$$\Phi_{\text{Edge}} = \mathbf{AG} \left(\bigwedge_{\bar{m} \in \mathcal{M}^n} \mathbf{EX}_1 \mathbf{after}_{\bar{m}} \right)$$

where $\mathbf{after}_{\bar{m}}$ stands for $\bigwedge_{a_i \in \mathbf{Agt}} \mathbf{after}_{m_j}^{a_i}$.

From an \mathbf{ATL}_{sc}^* formula Φ , we now define a \mathbf{QCTL}^* formula $\widehat{\Phi}^C$ in a similar way as we defined formula $\overline{\Phi}^C$ in Section 4 (when reducing the \mathbf{ATL}_{sc}^* model-checking problem to the \mathbf{QCTL} model-checking problem). Here, the subformulas Φ_{strat} and Φ_{out} are defined using propositions \mathbf{after} instead of the transition table (which is not known here). We let

$$\begin{aligned} \widehat{\langle A \rangle \varphi_p}^C &= \exists m_1^{a_{j_1}} \dots m_k^{a_{j_1}} \dots m_1^{a_{j_l}} \dots m_k^{a_{j_l}} \cdot \mathbf{p}_{\text{out}} \cdot \\ &\quad \left(\Phi_{\text{strat}}^{\text{ba}}(A) \wedge \Phi_{\text{out}}^{\text{ba}}(A \cup C) \wedge \mathbf{A} \left(\mathbf{G} \mathbf{p}_{\text{out}} \Rightarrow \widehat{\varphi_p}^{C \cup A} \right) \right) \end{aligned}$$

with:

$$\begin{aligned} \Phi_{\text{strat}}^{\text{ba}}(A) &= \bigwedge_{a \in A} \mathbf{AG} \left(\bigvee_{m_i \in \mathcal{M}} (m_i^a \wedge \bigwedge_{j \neq i} \neg m_j^a) \right) \\ \Phi_{\text{out}}^{\text{ba}}(A) &= \mathbf{p}_{\text{out}} \wedge \mathbf{AG} \left(\neg \mathbf{p}_{\text{out}} \Rightarrow \mathbf{AX} \neg \mathbf{p}_{\text{out}} \right) \wedge \\ &\quad \mathbf{AG} \left[\mathbf{p}_{\text{out}} \Rightarrow \bigvee_{\bar{m} \in \mathcal{M}^A} \left(\mathbf{p}_{\bar{m}} \wedge \mathbf{AX} (\mathbf{after}_{\bar{m}} \Leftrightarrow \mathbf{p}_{\text{out}}) \right) \right] \end{aligned}$$

where, given $\bar{m} = (m_j^{a_i})_{a_i \in A}$, $\mathbf{p}_{\bar{m}}$ stands for $\bigwedge_{a_i \in A} m_j^{a_i}$. This yields a formula whose size is in $O(|\Phi| \cdot |\mathbf{Agt}| \cdot |\mathcal{M}|^{|\mathbf{Agt}|})$.

In the following, we consider execution trees whose labeling functions include propositions \mathbf{after}_m^a . As for the turn-based case, given a strategy f_B , we say that a labeling function ℓ is an f_B -ba-labeling when propositions $(m_i^{a_j})_{1 \leq i \leq k, a_j \in B}$ describe the moves proposed by f_B , and when \mathbf{p}_{out} labels the outcomes of f_B . Our translation has the expected property (we omit the proof as it follows exactly the same lines as for Proposition 18):

Proposition 23. Let $\Phi \in \text{ATL}_{sc}^*$. Fix a finite set $\mathcal{M} = \{m_1, \dots, m_k\}$ of actions, and a finite set $\text{Agt} = \{a_1, \dots, a_n\}$ of agents (containing $\text{Agt}(\Phi)$). Let \mathcal{C} be a CGS based on \mathcal{M} and Agt . Let ρ be a path of \mathcal{C} , n be a position along ρ , and f_B be a strategy context whose domain is $B \subseteq \text{Agt}$. Let $\mathcal{T}_{\mathcal{C}}(\rho(0)) = \langle T, \ell_{f_B} \rangle$ be an execution tree of the Kripke structure underlying \mathcal{C} from $\rho(0)$, labeled with an f_B -ba-labeling ℓ_{f_B} . Then we have:

$$\mathcal{C}, \rho, n \models_{f_B} \Phi \quad \text{iff} \quad \langle T, \ell_{f_B} \rangle, \rho, n \models \widehat{\Phi}^B$$

We can then relate the truth values of Φ and of $\widehat{\Phi} = \Phi_{\text{Edge}} \wedge \widehat{\Phi}^\emptyset$:

Proposition 24. Let Φ be an ATL_{sc}^* formula, $\text{Agt} = \{a_1, \dots, a_n\}$ be a finite set of agents, $\mathcal{M} = \{m_1, \dots, m_k\}$ be a finite set of moves, and $\widehat{\Phi}$ be the formula defined above. Then Φ is $(\text{Agt}, \mathcal{M})$ -satisfiable in a CGS if, and only if, the QCTL^* formula $\widehat{\Phi}$ is satisfiable (in the tree semantics).

Sketch of proof. If Φ is $(\text{Agt}, \mathcal{M})$ -satisfiable, then we can derive a tree satisfying $\widehat{\Phi}$ (thanks to Proposition 23). Conversely if $\widehat{\Phi}$ is satisfied in some tree \mathcal{T} , then this tree corresponds to some CGS based on Agt and \mathcal{M} : indeed, formula Φ_{Edge} ensures that for every node, the labeling of the successor nodes with $(\text{after}_{m_j}^{a_i})_{m_j \in \mathcal{M}, a_i \in \text{Agt}}$ defines a transition table. The end of the proof is similar to proof for the turn-based case. \square

Notice that if Φ has $k \geq 1$ nested quantifiers, then so does $\widehat{\Phi}$. However, the size of $\widehat{\Phi}$ is exponential in $|\text{Agt}|$. Satisfiability of Q^kCTL^* (resp. Q^kCTL) formula being $(k + 2)$ -EXPTIME-complete (resp. $(k + 1)$ -EXPTIME-complete), we end up with an algorithm in $(k + 3)$ -EXPTIME (resp. $(k + 2)$ -EXPTIME) for ATL_{sc}^* (resp. ATL_{sc}) formulas involving at most $k \geq 1$ nested quantifiers.

Proposition 25. The $(\text{Agt}, \mathcal{M})$ -satisfiability problem for ATL_{sc} and ATL_{sc}^* is Tower-hard.

Proof. The proof uses similar ideas as for the case of TBGSs. Since the reduction is from the model-checking problem, we already know the set of agents and actions. The main difficulty is to require that the satisfying CGS be turn-based. This can be achieved using the following formula:

$$\mathbf{AG} \left[\bigwedge_{q \in Q} \bigvee_{a \in \text{Agt}} \left(p_q \Rightarrow \left[\left(\bigwedge_{q \rightarrow q'} \langle\langle a \rangle\rangle \mathbf{X} p_{q'} \right) \wedge \left(\llbracket a \rrbracket \bigvee_{q \rightarrow q'} \langle\langle \emptyset \rangle\rangle \mathbf{X} p_{q'} \right) \right] \right) \right]$$

The rest of the proof is similar to the proof of Proposition 20. \square

Theorem 26. $(\text{Agt}, \mathcal{M})$ -satisfiability for ATL_{sc} and ATL_{sc}^* is Tower-complete.

6. Extensions of ATL_{sc}

In this section, we explain how our technique of using QCTL^* applies in two other settings: first, for the variant of ATL_{sc}^* where strategy quantifiers are restricted to range over *memoryless* strategies; second, for strategy logic (SL), a different formalism for expressing properties of multi-agent systems.

6.1. ATL_{sc} with memoryless strategy quantifiers

In this section, we consider the logic obtained from ATL_{sc} by restricting strategy quantifiers to range over *memoryless* strategies. Notice that the memoryless requirement only applies to explicitly quantified strategies: for instance, $\langle A \rangle_0 \varphi$ states that coalition A has a memoryless strategy to enforce φ , whatever the other players do, even if they have memory:

$$\mathcal{C}, \pi, n \models_f \langle A \rangle_0 \varphi \quad \Leftrightarrow \quad \exists f_A \in \text{Strat}_0(A). \forall \pi' \in \text{Out}(\pi_{\leq n}, f_A \circ f). \mathcal{C}, \pi', n \models_{f_A \circ f} \varphi.$$

Enforcing memoryless strategies for the opponent coalition can be achieved by expliciting the strategy quantification: we would then write $\langle A \rangle_0 [B]_0 \langle \emptyset \rangle_0 \varphi$ to force players in B to play memoryless. Notice that because of this difference between implicit and explicit strategy quantification, our translation from ATL_{sc}^* to ATL_{sc} does not apply for memoryless strategies.

It is well-known that the strategies witnessing an ATL property can be chosen memoryless: in other terms, $\langle\langle A \rangle\rangle \varphi$ and $\langle\langle A \rangle\rangle_0 \varphi$ are equivalent (when $\langle\langle A \rangle\rangle \varphi$ is interpreted with the classical semantics of ATL). Moreover, $\langle a_1 \rangle_0 (\langle a_2 \rangle_0 \mathbf{X} a \wedge \langle a_2 \rangle_0 \mathbf{X} b)$ is true on only one of the two CGSs of Fig. 1. It follows that $\text{ATL}_{sc,0}$ is strictly more expressive than ATL. Actually, $\text{ATL}_{sc,0}$ can even distinguish CGSs that ATL_{sc} cannot: consider the two one-player CGSs \mathcal{S} and \mathcal{S}' of Figure 2; they involve only one player, and can be seen as Kripke structure; as Kripke structures, they are bisimilar, so that they satisfy the same CTL^* formulas, and consequently also the same ATL_{sc} formulas (any ATL_{sc} formula is easily translated into an equivalent CTL^* formula for these one-player models). But the $\text{ATL}_{sc,0}$ formula $\langle a \rangle_0 (\mathbf{X} \neg a \wedge \mathbf{X} \mathbf{X} a)$ holds true in state ℓ'_0 of \mathcal{S}' , while it fails to hold in ℓ_0 in \mathcal{S} .

We now turn to model checking and satisfiability. The number of memoryless strategies for one player being bounded (with $|\mathcal{M}|^{|\mathcal{Q}|}$), we can easily enumerate all of them, and store each strategy within polynomial space. Hence:




Fig. 2: Two CGS that can be distinguished by $\text{ATL}_{sc,0}$, but not by ATL_{sc}

Theorem 27. *Model checking $\text{ATL}_{sc,0}$ and $\text{ATL}_{sc,0}^*$ is PSPACE-complete.*

Proof. We again rely on our translation to QCTL^* , but this time in the *structure semantics* [14]: instead of ranging over labellings of the execution tree, propositional quantification then ranges over labellings of the Kripke structure. Indeed, a memoryless strategy is simply a function mapping each state of the game to an available move.

However, we cannot directly reuse the translation of Section 4.1: indeed, in this translation, we quantify over atomic proposition \mathbf{p}_{out} to mark the outcomes of the selected strategies. This is not correct in the structure semantics of QCTL^* , since this would only range over ultimately-periodic outcomes.

To overcome this problem, we propose a slightly different translation: instead of quantifying over \mathbf{p}_{out} , we use a CTL^* formula to characterize the outcomes: the translation of $\langle A \rangle_0 \varphi$ in a context with domain C now reads as follows (reusing the notations of Section 4.1):

$$\overline{\langle A \rangle_0 \varphi}_p^C = \exists \mathbf{m}_1^{a_{j_1}} \dots \mathbf{m}_k^{a_{j_1}} \dots \mathbf{m}_1^{a_{j_l}} \dots \mathbf{m}_k^{a_{j_l}}. \left(\Phi_{\text{strat}}(A) \wedge \mathbf{A} \left[(\Phi'_{\text{out}}(A \cup C)) \Rightarrow \overline{\varphi}_p^{CUA} \right] \right)$$

where

$$\Phi'_{\text{out}}(C) = \mathbf{G} \left(\bigwedge_{q \in Q} \bigwedge_{\bar{m} \in \text{Mov}(q, C)} \left[(\mathbf{p}_q \wedge \mathbf{p}_{\bar{m}}) \Rightarrow \bigvee_{q' \in \text{Next}(q, C, \bar{m})} \mathbf{X} \mathbf{p}_{q'} \right] \right).$$

Formula $\Phi'_{\text{out}}(C)$ characterizes the outcomes of the strategies in use for coalition C . In the end, the QCTL^* formula $\overline{\Phi}^\emptyset$ has size $O(|\Phi| \cdot |Q| \cdot (|\text{Agt}| \cdot |\mathcal{M}|^2 + |Q| \cdot |\text{Edge}|))$. Using the PSPACE algorithm for model checking QCTL^* in

the structure semantics, we obtain a PSPACE algorithm for model checking $\text{ATL}_{sc,0}^*$.

Finally, hardness in PSPACE is obtained for $\text{ATL}_{sc,0}$ by a straightforward encoding of QBF. \square

While restricting to memoryless strategies makes model checking easier, it actually makes satisfiability harder:

Theorem 28. *Satisfiability of $\text{ATL}_{sc,0}$ (with memoryless-strategy quantification) is undecidable, even when restricting to turn-based games or when the set of agents and actions is fixed.*

The proof of this result being long and technical, we postponed it to Appendix Appendix B. While this result may look surprising given our previous results, it is the natural counterpart of the fact that QCTL satisfiability is undecidable over finite graphs. The proof of Theorem 28 uses the same ideas as for the undecidability of QCTL satisfiability over graphs [14].

6.2. Strategy Logic

Strategy Logic (SL) [7, 19] extends LTL with explicit quantification and use of strategies. SL allows first-order quantification over strategies, and those strategies are then assigned to players.

Formula $\langle\langle x \rangle\rangle \varphi$ expresses the existence of a strategy enforcing φ ; the strategy is stored in variable x for later use in φ : the *agent binding* operator (a, x) can be used to assign strategy x to agent a . An assignment χ is a partial function from $\text{Agt} \cup \text{Var}$ to Strat . An SL formula φ is interpreted over pairs (χ, q) where q is a state of some CGS and χ is an assignment such that any free³ strategy variable or agent occurring in φ belongs to $\text{dom}(\chi)$. Note that we must have $\text{Agt} \subseteq \text{dom}(\chi)$ when temporal modalities \mathbf{X} and \mathbf{U} are interpreted: this implies that the set of outcomes is restricted to a single execution generated by all the strategies assigned to players in Agt , and the temporal modalities are therefore interpreted along this execution. We refer to [19] for a complete definition of SL.

In the following we assume w.l.o.g. that every quantifier $\langle\langle x \rangle\rangle$ introduces a fresh strategy variable x : this allows us to permanently use variable x to

³We use the standard notion of freeness for strategy variables, with the hypothesis that $\langle\langle x \rangle\rangle$ binds x , and for the agents with the hypothesis that (a, x) binds a and that every agent in Agt is free in temporal subformula (*i.e.*, with \mathbf{U} or \mathbf{X} as root).

denote the selected strategy for a . Moreover, we require that every player may play any move in any state ($\text{Mov}(q, a) = \mathcal{M}$): this rules out the problem whether a selected strategy can be assigned to a player when evaluating a formula. We omit the formal proofs of the results stated in this part, as they closely follow the same arguments as for ATL_{sc}^* .

Model checking. Following the ideas developed for ATL_{sc}^* model checking, we reduce the model-checking problem for **SL** to the model-checking problem for **QCTL***. Consider a CGS $\mathcal{C} = \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge} \rangle$ with $\mathcal{M} = \{m_1, \dots, m_k\}$ and $\text{Agt} = \{a_1, \dots, a_n\}$, and where **Mov** constantly returns \mathcal{M} . Let $\Phi \in \text{SL}$ and V a partial function $V: \text{Agt} \rightarrow \text{Var}$ assigning strategy variables to some of the agents. We build a **QCTL*** formula $\overline{\Phi}_V$ as follows:

$$\begin{aligned} \overline{\varphi \wedge \psi}^V &= \overline{\varphi}^V \wedge \overline{\psi}^V & \overline{\neg \varphi}^V &= \neg \overline{\varphi}^V & \overline{\varphi \mathbf{U} \psi}^V &= \mathbf{A} \left(\Phi_{\text{out}}(V) \Rightarrow \overline{\varphi}_p^V \mathbf{U} \overline{\psi}^V \right) \\ \overline{(a, x)\varphi}^V &= \overline{\varphi}^{V[a \rightarrow x]} & \overline{p}^V &= p & \overline{\mathbf{X}\varphi_p}^V &= \mathbf{A} \left(\Phi_{\text{out}}(V) \Rightarrow \mathbf{X}\overline{\varphi}_p^V \right) \end{aligned}$$

Strategy quantification is handled as follows:

$$\overline{\langle\langle x \rangle\rangle \varphi}^V = \exists m_1^x \dots m_k^x. \left(\Phi_{\text{strat}}(x) \wedge \overline{\varphi} \right)$$

with:

$$\begin{aligned} \Phi_{\text{strat}}(x) &= \mathbf{AG} \left(\bigvee_{m_i \in \mathcal{M}} (m_i^x \wedge \bigwedge_{j \neq i} \neg m_j^x) \right) \\ \Phi_{\text{out}}(V) &= \mathbf{G} \left[\bigvee_{q \in Q} \left(\mathbf{p}_q \wedge \bigvee_{\overline{m} \in \mathcal{M}^{\text{Agt}}} (\mathbf{p}_{\overline{m}}^V \wedge \mathbf{X} \mathbf{p}_{\text{Edge}(q, \overline{m})}) \right) \right] \end{aligned}$$

where $\mathbf{p}_{\overline{m}}^V$ stands for $\bigwedge_{a \in \text{Agt}} m_{\alpha_i}^{V(a)}$ when $\overline{m} = (m_{\alpha_1}, \dots, m_{\alpha_n})$.

We then have $\mathcal{C}, q \models \Phi$ if, and only if, $\langle T, \ell \rangle, q \models \overline{\Phi}^{V_\emptyset}$, where $\langle T, \ell \rangle$ is the execution tree of \mathcal{C} with each state q being labeled with a corresponding atomic proposition \mathbf{p}_q .

Theorem 29. *Let \mathcal{C} be a CGS $= \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Edge} \rangle$ and let $\langle T, \ell \rangle$ be the execution tree of \mathcal{C} such that ℓ includes a labeling for propositions \mathbf{p}_q . Let Φ be an **SL** formula and $\overline{\Phi}^{V_\emptyset}$ be the **QCTL*** formula defined as above. Then $\mathcal{C}, q \models \Phi$ if, and only if, $\langle T, \ell \rangle q \models \overline{\Phi}^{V_\emptyset}$*

Satisfiability for turn-based case. One easily sees that ATL_{sc}^* can be expressed in SL. It follows that satisfiability is undecidable for SL. We thus restrict to our two decidable cases (turn-based games and bounded set of moves and players), and prove decidability of satisfiability for SL in both cases.

Given an SL formula Φ and a partial function $V: \text{Agt} \rightarrow \text{Var}$, we define a QCTL* formula $\widehat{\Phi}^V$ inductively as follows (boolean cases omitted):

$$\widehat{\langle\langle x \rangle\rangle} \varphi^V = \exists \text{mov}_x. \left[\mathbf{AG} \left(\mathbf{EX}_1 \text{mov}_x \right) \wedge \widehat{\varphi}^V \right] \quad \widehat{(a, x)} \varphi^V = \widehat{\varphi}^{V[a \rightarrow x]}$$

Note that in this case we require that *every* reachable state has a (unique) successor labeled with mov_x : indeed when one quantifies over a strategy x , the agent(s) who will use this strategy are not known a priori. However, in the turn-based case, a given strategy should be dedicated to a single agent: there is no natural way to share a strategy between two different agents (or the other way around, any two strategies for two different agents can be seen as a single strategy), as they are not playing in the same states. When the strategy x is assigned to some agent a , only the choices made in the a -states are considered.

The temporal modalities are treated as follows:

$$\widehat{\mathbf{U}} \varphi^V = \mathbf{A} \left[\mathbf{G} \left(\bigwedge_{a_j \in \text{Agt}} (\text{turn}_j \Rightarrow \mathbf{X} \text{mov}_{V(a_j)}) \right) \Rightarrow \widehat{\varphi}^V \mathbf{U} \widehat{\psi}^V \right]$$

$$\widehat{\mathbf{X}} \varphi^V = \mathbf{A} \left[\mathbf{G} \left(\bigwedge_{a_j \in \text{Agt}} (\text{turn}_j \Rightarrow \mathbf{X} \text{mov}_{V(a_j)}) \right) \Rightarrow \mathbf{X} \widehat{\varphi}^V \right]$$

Finally, we let $\widehat{\Phi}$ be the formula $\Phi_{\text{tb}} \wedge \widehat{\Phi}^{V_\emptyset}$. Then we have the following theorem:

Theorem 30. *Let Φ be an SL formula and $\widehat{\Phi}$ be the QCTL* formula defined as above. Then Φ is satisfiable in a TBGS if, and only if, $\widehat{\Phi}$ is satisfiable (in the tree semantics).*

Satisfiability for bounded action alphabet Let \mathcal{M} be $\{m_1, \dots, m_k\}$ and $\text{Agt} = \{a_1, \dots, a_n\}$. The reduction carried out for ATL_{sc} can also be adapted for SL in this case. Given an SL formula Φ and a partial function $V: \text{Agt} \rightarrow \text{Var}$, we define the QCTL* formula $\widehat{\Phi}^V$ inductively as follows:

$$\widehat{\langle\langle x \rangle\rangle} \varphi^V = \exists m_1^x \dots \exists m_k^x. \Phi_{\text{strat}}(x) \wedge \widehat{\varphi}^V \quad \widehat{(a, x)} \varphi^V = \widehat{\varphi}^{V[a \rightarrow x]}$$

The temporal modalities are handled as follows:

$$\begin{aligned}\widehat{\varphi} \mathbf{U} \widehat{\psi}^V &= \mathbf{A} \left[\left(\mathbf{G} \bigwedge_{\overline{m} \in \mathcal{M}^{\text{Agt}}} (\mathbf{p}_{\overline{m}}^V \Rightarrow \mathbf{X} \text{ after}_{\overline{m}}) \right) \Rightarrow (\widehat{\varphi}^V \mathbf{U} \widehat{\psi}^V) \right] \\ \widehat{\mathbf{X}} \widehat{\varphi}^C &= \mathbf{A} \left[\left(\mathbf{G} \bigwedge_{\overline{m} \in \mathcal{M}^{\text{Agt}}} (\mathbf{p}_{\overline{m}}^V \Rightarrow \mathbf{X} \text{ after}_{\overline{m}}) \right) \Rightarrow (\mathbf{X} \widehat{\varphi}^V) \right]\end{aligned}$$

where $\mathbf{m}_{\overline{m}}^V$ stands for $\bigwedge_{a \in \text{Agt}} \mathbf{m}_{\alpha_i}^{V(a)}$ when $\overline{m} = (m_{\alpha_1}, \dots, m_{\alpha_n})$.

Finally, let $\widehat{\Phi}$ be the formula $\Phi_{\text{Edge}} \wedge \widehat{\Phi}_{\emptyset}^V$. We have:

Theorem 31. *Let Φ be an SL formula based on the set $\text{Agt} = \{a_1, \dots, a_n\}$, let $\mathcal{M} = \{m_1, \dots, m_k\}$ be a finite set of moves, and $\widehat{\Phi}$ be the QCTL* formula defined as above. Then Φ is $(\text{Agt}, \mathcal{M})$ -satisfiable if, and only if, $\widehat{\Phi}$ is satisfiable (in the tree semantics).*

7. Conclusion

We developed a tight link between the extension of ATL with strategy contexts and the logic QCTL. We believe that our logic ATL_{sc} (and similar formalisms such as SL) is very well-suited for reasoning about complex, multi-agent systems: it can express useful properties in non-zero-sum games, and provide much better granularity than Nash equilibria and similar solution concepts. But the technical formalism of games blurs the setting, and we believe that QCTL is the formalism of choice for fully understanding ATL_{sc} and related logics.

Our translation to QCTL provides us with a uniform presentation of verification algorithms for ATL_{sc} —when such algorithms exist. In view of this, we will keep on developing our knowledge and understanding of QCTL, for instance in terms of the behavioral equivalence it characterizes.

References

- [1] T. Ågotnes, V. Goranko, and W. Jamroga. Alternating-time temporal logics with irrevocable strategies. In D. Samet, editor, *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK'07)*, p. 15–24, June 2007.
- [2] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, Sept. 2002.

- [3] R. Alur, T. A. Henzinger, O. Kupferman, and M. Y. Vardi. Alternating refinement relations. In D. Sangiorgi and R. de Simone, editors, *Proceedings of the 9th International Conference on Concurrency Theory (CONCUR'98)*, Lecture Notes in Computer Science 1466, p. 163–178. Springer-Verlag, Sept. 1998.
- [4] Ch. Baier, T. Brázdil, M. Größer, and A. Kučera. Stochastic game logic. In *Proceedings of the 4th International Conference on Quantitative Evaluation of Systems (QEST'07)*, p. 227–236. IEEE Comp. Soc. Press, Sept. 2007.
- [5] Th. Brihaye, A. Da Costa, F. Laroussinie, and N. Markey. ATL with strategy contexts and bounded memory. In S. N. Artemov and A. Nerode, editors, *Proceedings of the International Symposium Logical Foundations of Computer Science (LFCS'09)*, Lecture Notes in Computer Science 5407, p. 92–106. Springer-Verlag, Jan. 2009.
- [6] K. Chatterjee, T. A. Henzinger, and M. Jurdziński. Games with secure equilibria. *Theoretical Computer Science*, 365(1-2):67–82, Nov. 2006.
- [7] K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. In L. Caires and V. T. Vasconcelos, editors, *Proceedings of the 18th International Conference on Concurrency Theory (CONCUR'07)*, Lecture Notes in Computer Science 4703, p. 59–73. Springer-Verlag, Sept. 2007.
- [8] A. Da Costa, F. Laroussinie, and N. Markey. ATL with strategy contexts: Expressiveness and model checking. In K. Lodaya and M. Mahajan, editors, *Proceedings of the 30th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'10)*, Leibniz International Proceedings in Informatics 8, p. 120–132. Leibniz-Zentrum für Informatik, Dec. 2010.
- [9] A. Da Costa, F. Laroussinie, and N. Markey. Quantified CTL: Expressiveness and model checking. In M. Koutny and I. Ulidowski, editors, *Proceedings of the 23rd International Conference on Concurrency Theory (CONCUR'12)*, Lecture Notes in Computer Science 7454, p. 177–192. Springer-Verlag, Sept. 2012.
- [10] T. French. Decidability of quantified propositional branching time logics. In M. Stumptner, D. Corbett, and M. Brooks, editors, *Proceedings of the 14th Australian Joint Conference on Artificial Intelligence (AJCAI'01)*, Lecture Notes in Computer Science 2256, p. 165–176. Springer-Verlag, Dec. 2001.
- [11] O. Kupferman. Augmenting branching temporal logics with existential quantification over atomic propositions. In P. Wolper, editor, *Proceedings of the 7th*

- International Conference on Computer Aided Verification (CAV'95)*, Lecture Notes in Computer Science 939, p. 325–338. Springer-Verlag, July 1995.
- [12] A. Kurucz. $S5 \times S5 \times S5$ lacks the finite model property. In M. de Rijke, H. Wansing, F. Wolter, and M. Zakharyashev, editors, *Proceedings of the 3rd Workshop on Advances in Modal Logic (AIML'00)*, p. 321–327. World Scientific, 2002.
 - [13] A. Kurucz, F. Wolter, M. Zakharyashev, and D. M. Gabbay. *Many-Dimensional Modal Logics: Theory and Applications*. Elsevier Science, 2003.
 - [14] F. Laroussinie and N. Markey. Quantified CTL: expressiveness and complexity. Research Report LSV-13-07, Lab. Spécification & Vérification, ENS Cachan, France, Apr. 2013.
 - [15] F. Laroussinie and N. Markey. Satisfiability of atl with strategy contexts. In G. Puppis and T. Villa, editors, *Proceedings of the 4th International Symposium on Games, Automata, Logics and Formal Verification (GandALF'13)*, Electronic Proceedings in Theoretical Computer Science 119, p. 208–223, Aug. 2013.
 - [16] F. Mogavero, A. Murano, G. Perelli, and M. Y. Vardi. Reasoning about strategies: On the model-checking problem. Research Report cs.LO/1112.6275, arXiv, Nov. 2012.
 - [17] F. Mogavero, A. Murano, G. Perelli, and M. Y. Vardi. What makes ATL* decidable? a decidable fragment of strategy logic. In M. Koutny and I. Ulidowski, editors, *Proceedings of the 23rd International Conference on Concurrency Theory (CONCUR'12)*, Lecture Notes in Computer Science 7454, p. 193–208. Springer-Verlag, Sept. 2012.
 - [18] F. Mogavero, A. Murano, and L. Sauro. On the boundary of behavioral strategies. In *Proceedings of the 28th Annual Symposium on Logic in Computer Science (LICS'13)*, p. 263–272. IEEE Comp. Soc. Press, June 2013.
 - [19] F. Mogavero, A. Murano, and M. Y. Vardi. Reasoning about strategies. In K. Lodaya and M. Mahajan, editors, *Proceedings of the 30th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'10)*, Leibniz International Proceedings in Informatics 8, p. 133–144. Leibniz-Zentrum für Informatik, Dec. 2010.
 - [20] J. F. Nash, Jr. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, Jan. 1950.

- [21] S. Pinchinat. A generic constructive solution for concurrent games with expressive constraints on strategies. In K. Namjoshi, T. Yoneda, T. Higashino, and Y. Okamura, editors, *Proceedings of the 5th International Symposium on Automated Technology for Verification and Analysis (ATVA '07)*, Lecture Notes in Computer Science 4762, p. 253–267. Springer-Verlag, Oct. 2007.
- [22] S. Schmitz. Complexity hierarchies beyond elementary. Research Report cs.CC/1312.5686, arXiv, Dec. 2013.
- [23] A. P. Sistla. *Theoretical Issues in the Design and Verification of Distributed Systems*. PhD thesis, Harvard University, Cambridge, Massachusetts, USA, 1983.
- [24] N. Troquard and D. Walther. On satisfiability in ATL with strategy contexts. In L. Fariñas del Cerro, A. Herzig, and J. Mengin, editors, *Proceedings of the 13th European Conference in Logics in Artificial Intelligence (JELIA '12)*, Lecture Notes in Computer Science 7519, p. 398–410. Springer-Verlag, Sept. 2012.
- [25] Y. Venema. Rectangular games. *Journal of Symbolic Logic*, 63(4):1549–1564, Dec. 1998.
- [26] D. Walther, C. Lutz, F. Wolter, and M. Wooldridge. ATL satisfiability is indeed EXPTIME-complete. *Journal of Logic and Computation*, 16(6):765–787, Dec. 2006.
- [27] F. Wang, C.-H. Huang, and F. Yu. A temporal logic for the interaction of strategies. In J.-P. Katoen and B. König, editors, *Proceedings of the 22nd International Conference on Concurrency Theory (CONCUR '11)*, Lecture Notes in Computer Science 6901, p. 466–481. Springer-Verlag, Sept. 2011.

Appendix A. Proof of Theorem 9

Theorem 9. *Given a formula $\varphi \in \text{ATL}_{sc}^*$ and a coalition B' , there exists an ATL_{sc} formula $\widehat{\varphi}^{[\emptyset, B']}$, involving only players in $\text{Agt}_\varphi \cup B'$, such that for any strategy context f with $\text{dom}(f) = B'$, φ and $\widehat{\varphi}^{[\emptyset, B']}$ are equivalent under context f .*

This result is proved through the following two lemmas.

Lemma 32. *Let \mathcal{C} be a CGS with set of agents $\text{Agt}_\mathcal{C}$, ρ be a path and n be a position along ρ . For any state-formula $\varphi \in \text{ATL}_{sc}^*$, for any strategy contexts f, g and g' such that $\text{dom}(g) \subseteq \text{dom}(g')$, $g'_{|\text{dom}(g)} = g$ (in other terms, g' extends g), and $\text{dom}(f) \cap \text{dom}(g') = \emptyset$, for any coalitions B and B' s.t. $\text{dom}(f) = (\text{Agt}_\mathcal{C} \setminus B) \cap B'$, and for any outcome $\pi \in \text{Out}(\rho_{\leq n}, g' \circ f)$, we have:*

$$\mathcal{C}, \pi, n \models_{g \circ f} \widehat{\varphi}^{[B, B']} \Leftrightarrow \mathcal{C}, \pi, n \models_{g' \circ f} \widehat{\varphi}^{[B, B']}.$$

Proof. We prove this result by structural induction, omitting the easy cases of atomic propositions and boolean operators.

For the case where $\varphi = \langle A \rangle \psi$, assume $\mathcal{C}, \pi, n \models_{g \circ f} \widehat{\langle A \rangle \psi}^{[B, B']}$. Then by definition, $\mathcal{C}, \pi, n \models_{g \circ f} \langle A \rangle [B \setminus A] [\overline{B' \cup A}] \widehat{\psi}^{[B \setminus A, B' \cup A]}$. Thus there exists $f_A \in \text{Strat}(A)$ such that for any $f' \in \text{Strat}((B \setminus A) \cup (\overline{B' \cup A}))$, it holds $\mathcal{C}, \pi', n \models_{f' \circ f_A \circ g \circ f} \widehat{\psi}^{[B \setminus A, B' \cup A]}$, where π' is the unique path in $\text{Out}(\rho_{\leq n}, f' \circ f_A \circ g \circ f)$. Equivalently, with the above f_A and for any $f' \in \text{Strat}((B \setminus A) \cup (\overline{B' \cup A}))$, $\mathcal{C}, \pi', n \models_{f' \circ f_A \circ g' \circ f} \widehat{\psi}^{[B \setminus A, B' \cup A]}$ because $\text{dom}(g) \subseteq \text{dom}(g') \subseteq \text{dom}(f) \subseteq \text{dom}(f') \cup A$. In the end, we get $\mathcal{C}, \pi, n \models_{g' \circ f} \widehat{\varphi}^{[B, B']}$, as expected. The converse implication is proven similarly, as well as the case where $\varphi = \langle \overline{A} \rangle \psi$.

Now assume $\varphi = \langle A \rangle \psi$. If $\mathcal{C}, \pi, n \models_{g \circ f} \widehat{\varphi}^{[B, B']}$, then $\mathcal{C}, \pi, n \models_{g \circ f} \widehat{\psi}^{[B \cup A, B']}$, so that $\mathcal{C}, \pi, n \models_{g' \circ f} \widehat{\psi}^{[B \cup A, B']}$ by induction hypothesis (as ψ is a state-formula). Thus $\mathcal{C}, \pi, n \models_{g' \circ f} \widehat{\varphi}^{[B, B']}$. The converse direction follows the same lines. The proof for $\varphi = \langle \overline{A} \rangle \psi$ is similar. \square

Lemma 33. *Let \mathcal{C} be a CGS with set of agents $\text{Agt}_\mathcal{C}$, ρ be a path and n be a position along ρ , and f be a strategy context with $\text{dom}(f) \subseteq \text{Agt}_\Phi$. Then for any ATL_{sc}^* formula φ , for any strategy context g s.t. $\text{dom}(g) = \text{Agt}_\mathcal{C} \setminus \text{dom}(f)$, for any outcome $\pi \in \text{Out}(\rho_{\leq n}, g \circ f)$, and for any coalitions B and B' s.t. $\text{dom}(f) = (\text{Agt}_\mathcal{C} \setminus B) \cap B'$, it holds: $\mathcal{C}, \pi, n \models_f \varphi \Leftrightarrow \mathcal{C}, \pi, n \models_{g \circ f} \widehat{\varphi}^{[B, B']}$. Moreover, if φ is a state-formula, this result extends to any strategy context g s.t. $\text{dom}(g) \cap \text{dom}(f) = \emptyset$.*

Proof. We prove the result by induction on the structure of φ . The cases of atomic propositions and boolean connectives are straightforward.

- If $\varphi = \mathbf{X} \psi$, then $\mathcal{C}, \pi, n \models_f \varphi$ is equivalent to $\mathcal{C}, \pi, n+1 \models_f \psi$. Applying the induction hypothesis, this is equivalent to $\mathcal{C}, \pi, n+1 \models_{g \circ f} \widehat{\psi}^{[B, B']}$. This means that $\mathcal{C}, \pi, n \models_{g \circ f} \mathbf{X} \widehat{\psi}^{[B, B']}$, which in turn is equivalent to $\mathcal{C}, \pi, n \models_{g \circ f} \langle \cdot \rangle \mathbf{X} \widehat{\psi}^{[B, B']}$ because $\text{Out}^\infty(\ell, G \circ F) = \{\pi\}$.
- If $\varphi = \psi_1 \mathbf{U} \psi_2$: this case can be handled in a similar way as for the previous case, and we omit it.
- If $\varphi = \langle A \rangle \psi$: as this is a state formula, we prove the second, more general statement. Let g be a strategy context with $\text{dom}(g) \cap \text{dom}(f) = \emptyset$, and π be an outcome of $g \circ f$ from $\rho_{\leq n}$. Finally, fix B and B' such that $\text{dom}(f) = (\text{Agt}_{\mathcal{C}} \setminus B) \cap B'$.

$$\begin{aligned}
& \mathcal{C}, \pi, n \models_f \langle A \rangle \psi \\
& \Leftrightarrow \exists f_A \in \text{Strat}(A). \forall \pi' \in \text{Out}(\pi_{\leq n}, f_A \circ f). \mathcal{C}, \pi', n \models_{f_A \circ f} \psi \\
& \Leftrightarrow \exists f_A \in \text{Strat}(A). \forall f' \in \text{Strat}(\text{Agt}_{\mathcal{C}} \setminus \text{dom}(f_A \circ f)). \\
& \quad \forall \pi' \in \text{Out}(\pi_{\leq n}, f' \circ f_A \circ f). \mathcal{C}, \pi', n \models_{f_A \circ f} \psi \\
& \Leftrightarrow \exists f_A \in \text{Strat}(A). \forall f' \in \text{Strat}(\text{Agt}_{\mathcal{C}} \setminus \text{dom}(f_A \circ f)). \\
& \quad \forall \pi' \in \text{Out}(\pi_{\leq n}, f' \circ f_A \circ f). \mathcal{C}, \pi', n \models_{f' \circ f_A \circ f} \widehat{\psi}^{[B \setminus A, B' \cup A]} \\
& \quad (\text{by i.h., because } \text{dom}(f_A \circ f) = (\text{Agt}_{\mathcal{C}} \setminus (B \setminus A)) \cap (B' \cup A)) \\
& \Leftrightarrow \exists f_A \in \text{Strat}(A). \forall f' \in \text{Strat}(\text{Agt}_{\mathcal{C}} \setminus \text{dom}(f_A \circ f)). \\
& \quad \forall \pi' \in \text{Out}(\pi_{\leq n}, f' \circ f_A \circ f). \mathcal{C}, \pi', n \models_{f' \circ f_A \circ g \circ f} \widehat{\psi}^{[B \setminus A, B' \cup A]} \\
& \quad (\text{because } f' \circ f_A \circ g \circ f = f' \circ f_A \circ f) \\
& \Leftrightarrow \exists f_A \in \text{Strat}(A). \forall f' \in \text{Strat}(\text{Agt}_{\mathcal{C}} \setminus \text{dom}(f_A \circ f)). \\
& \quad \exists \pi' \in \text{Out}(\pi_{\leq n}, f' \circ f_A \circ f). \mathcal{C}, \pi', n \models_{f' \circ f_A \circ g \circ f} \widehat{\psi}^{[B \setminus A, B' \cup A]} \\
& \quad (\text{because } |\text{Out}(\pi_{\leq n}, f' \circ f_A \circ g \circ f)| = 1) \\
& \Leftrightarrow \exists f_A \in \text{Strat}(A). \mathcal{C}, \pi, n \models_{f_A \circ g \circ f} [B \setminus A] [\overline{B' \cup A}] \widehat{\psi}^{[B \setminus A, B' \cup A]} \\
& \quad (\text{because } \text{dom}(f') = (B \setminus A) \cup (\text{Agt}_{\mathcal{C}} \setminus (B' \cup A))) \\
& \Leftrightarrow \mathcal{C}, \pi, n \models_{g \circ f} \langle A \rangle [B \setminus A] [\overline{B' \cup A}] \widehat{\psi}^{[B \setminus A, B' \cup A]} \\
& \Leftrightarrow \mathcal{C}, \pi, n \models_{g \circ f} \widehat{\varphi}^{[B, B']}
\end{aligned}$$

- If $\varphi = \langle \bar{A} \rangle \psi$: a similar sequence of equivalences applies, but now $\text{dom}(f_{\bar{A}} \circ f) = \text{Agt}_{\mathcal{C}} \setminus [(B \cap A) \cup (A \setminus B')]$.
- If $\varphi = \langle A \rangle \psi$: again, we prove the second statement. Let g be a strategy context with $\text{dom}(g) \cap \text{dom}(f) = \emptyset$, and π be an outcome of $g \circ f$ after $\rho_{\leq n}$. Let B and B' such that $\text{dom}(f) = (\text{Agt}_{\mathcal{C}} \setminus B) \cap B'$. We have:

$$\begin{aligned}
\mathcal{C}, \pi, n \models_f \langle A \rangle \psi &\Leftrightarrow \mathcal{C}, \pi, n \models_{f \setminus A} \psi \\
&\Leftrightarrow \mathcal{C}, \pi, n \models_{g \circ f \setminus A} \widehat{\psi}^{[B \cup A, B']} \\
&\text{(by i.h., because } \text{dom}(f \setminus A) = (\text{Agt}_{\mathcal{C}} \setminus (B \cup A)) \cap B') \\
&\Leftrightarrow \mathcal{C}, \pi, n \models_{g \circ f|_A \circ f \setminus A} \widehat{\psi}^{[B \cup A, B']} \quad \text{(by Lemma 32)} \\
&\Leftrightarrow \mathcal{C}, \pi, n \models_{g \circ f} \widehat{\psi}^{[B \cup A, B']}
\end{aligned}$$

- If $\varphi = \langle \bar{A} \rangle \psi$, a similar sequence of equivalences applies. □

Appendix B. Proof of Theorem 28

Theorem 28. *Satisfiability of $\text{ATL}_{sc,0}$ (with memoryless-strategy quantification) is undecidable, even when restricting to turn-based games or when the set of agents and actions is fixed.*

Proof. We prove the result for infinite-state turn-based games, by adapting the corresponding proof for QCTL under the structure semantics [10], which consists in encoding the problem of tiling a quadrant. The result for finite-state turn-based games can be obtained using similar (but more involved) ideas, by encoding the problem of tiling all finite grids (see [14] for the corresponding proof for QCTL).

We consider a finite set T of tiles, and two binary relations H and V indicating which tile(s) may appear on the right and above (respectively) a given tile. Our proof consists in writing a formula that is satisfiable only on a grid-shaped (turn-based) game structure representing a tiling of the quadrant (*i.e.*, of $\mathbb{N} \times \mathbb{N}$). The reduction involves two players: Player 1 controls square states (which are labeled with \square), while Player 2 controls circle states (labeled with \circ). Each state of the grid is intended to represent one cell of the quadrant to be tiled. For technical reasons, the reduction is not that simple, and our game structure will have three kinds of states (see Fig. B.3):

- the “main” states (controlled by Player 2), which form the grid. Each state in this main part has a *right* neighbour and a *top* neighbour, which we assume we can identify: more precisely, we make use of two atomic propositions v_1 and v_2 which alternate along the horizontal lines of the grid. The *right* successor of a v_1 -state is labeled with v_2 , while its *top* successor is labeled with v_1 ;
- the “tile” states, labeled with one item of T (seen as atomic propositions). Each tile state only has outgoing transition(s) to a tile state labeled with the same tile;
- the “choice” states, which appear between “main” states and “tile” states: there is one choice state associated with each main state, and each choice state has a transition to each tile state. Choice states are controlled by Player 1.




Fig. B.3: The turn-based game encoding the tiling problem

Assuming that we have such a structure, a tiling of the grid corresponds to a *memoryless* strategy of Player 1 (who only plays in the “choice” states). Once such a memoryless strategy for Player 1 has been selected, that it corresponds to a valid tiling can be expressed easily: for instance, in any cell of the grid (assumed to be labeled with v_1), there must exist a pair of tiles $(t_1, t_2) \in H$ such that $v_1 \wedge \langle 2 \rangle_0 \mathbf{X} \mathbf{X} t_1 \wedge \langle 2 \rangle_0 \mathbf{X} (v_2 \wedge \mathbf{X} \mathbf{X} t_2)$. This would be written as follows:

$$\langle 1 \rangle_0 \mathbf{G} \left[\begin{array}{l} v_1 \Rightarrow \bigvee_{(t_1, t_2) \in H} \langle 2 \rangle_0 \mathbf{X} \mathbf{X} t_1 \wedge \langle 2 \rangle_0 \mathbf{X} (v_2 \wedge \mathbf{X} \mathbf{X} t_2) \\ \wedge \\ v_2 \Rightarrow \bigvee_{(t_1, t_2) \in H} \langle 2 \rangle_0 \mathbf{X} \mathbf{X} t_1 \wedge \langle 2 \rangle_0 \mathbf{X} (v_1 \wedge \mathbf{X} \mathbf{X} t_2) \end{array} \right].$$

The same can be imposed for vertical constraints, and for imposing a fairness constraint on the base line (under the same memoryless strategy for Player 1).




Fig. B.4: The cell gadget




Fig. B.5: Several cells forming (part of) a grid

It remains to build a formula characterising an infinite grid. This requires a slight departure from the above description of the grid: each main state will in fact be a gadget composed of four states, as depicted on Fig. B.4. The first state of each gadget will give the opportunity to Player 1 to *color* the state with either α or β . This will be used to enforce “confluence” of several transitions to the same state (which we need to express that the two successors of any cell of the grid share a common successor).

We now start writing our formula, which we present as a conjunction of several subformulas. We require that the main states be labeled with m , the choice states be labeled with c , and the tile states be labeled with the names of the tiles. We let $\mathbf{AP}' = \{m, c\} \cup T$ and $\mathbf{AP} = \mathbf{AP}' \cup \{v_1, v_2, \alpha, \beta, \square, \circ\}$. The first part of the formula reads as follows (where universal path quantification can be encoded, as long as the context is empty, using $\langle \cdot \rangle_0$):

$$\begin{aligned}
 & \mathbf{A}(m \mathbf{W} c) \wedge \mathbf{AG} \left[\bigvee_{p \in \mathbf{AP}'} p \wedge \bigwedge_{p' \in \mathbf{AP}' \setminus \{p\}} \neg p' \right] \wedge \mathbf{AG} (\square \Leftrightarrow \neg \circ) \wedge \\
 & \mathbf{AG} \left[\begin{array}{l} c \Rightarrow \left(\square \wedge \bigwedge_{t \in T} \langle 1 \rangle_0 \mathbf{X} t \wedge \mathbf{AX} \left(\bigvee_{t \in T} \mathbf{AG} t \right) \right) \\ \square \Rightarrow \left(\bigwedge_{p \in \mathbf{AP}} (\mathbf{EX} p \Leftrightarrow \langle 1 \rangle_0 \mathbf{X} p) \right) \\ \circ \Rightarrow \left(\bigwedge_{p \in \mathbf{AP}} (\mathbf{EX} p \Leftrightarrow \langle 2 \rangle_0 \mathbf{X} p) \right) \end{array} \right] \quad (\text{B.1})
 \end{aligned}$$

This formula enforces that each state is labeled with exactly one proposition from \mathbf{AP}' . It also enforces that any path will wander through the main part until it possibly goes to a choice state (this is expressed as $\mathbf{A}(m \mathbf{W} c)$, where $m \mathbf{W} c$ means $\mathbf{G} m \vee m \mathbf{U} c$, and can be expressed a negated-until formula). Finally, the second part of the formula enforces the witnessing structures to be turn-based.

Now we have to impose that the m -part has the shape of a grid: intuitively, each cell has three successors: one “to the right” and one “to the top” in the main part of the grid, and one c -state which we will use for associating a tile with this cell. For technical reasons, the situation is not that simple, and each cell is actually represented by the gadget depicted on Fig. B.4. Each state of the gadget is labeled with m . We constrain the form of the cells as follows:

$$\mathbf{AG} \left[\begin{array}{l} m \Rightarrow ((\square \wedge \neg \alpha \wedge \neg \beta) \vee (\circ \wedge \neg(\alpha \wedge \beta))) \\ \wedge \\ ((m \wedge \square) \Rightarrow (v_1 \Leftrightarrow \neg v_2)) \wedge ((v_1 \vee v_2) \Rightarrow (m \wedge \square)) \\ \wedge \\ (m \wedge \square) \Rightarrow [\mathbf{AX}(m \wedge \circ \wedge (\alpha \vee \beta)) \wedge \mathbf{AX}(m \wedge \circ \wedge \neg \alpha \wedge \neg \beta)] \wedge \\ \langle \cdot 1 \cdot \rangle_0 \mathbf{X} \alpha \wedge \langle \cdot 1 \cdot \rangle_0 \mathbf{X} \beta \end{array} \right] \quad (\text{B.2})$$

This says that there are four types of states in each cell, and specifies the possible transitions within such cells. We now express constraints on the transitions leaving a cell:

$$\begin{aligned} \mathbf{AG} \left[(\mathbf{EX} c \vee \mathbf{EX} v_1 \vee \mathbf{EX} v_2) \Rightarrow (m \wedge \circ \wedge \neg \alpha \wedge \neg \beta) \right] \wedge \\ \mathbf{AG} \left[(m \wedge \circ \wedge \neg \alpha \wedge \neg \beta) \Rightarrow (\mathbf{EX} c \wedge \mathbf{EX} v_1 \wedge \mathbf{EX} v_2 \wedge \mathbf{AX}(c \vee v_1 \vee v_2)) \right] \end{aligned} \quad (\text{B.3})$$

It remains to enforce that the successor of the α and β states are the same. This is obtained by the following formula:

$$\mathbf{AG} \left[(m \wedge \square) \Rightarrow [2]_0 \left(\langle \cdot \emptyset \cdot \rangle_0 \mathbf{X}^3(c \vee v_1) \vee \langle \cdot \emptyset \cdot \rangle_0 \mathbf{X}^3(c \vee v_2) \right) \right] \quad (\text{B.4})$$

Indeed, assume that some cell has two different “final” states; then there would exist a strategy for Player 2 (consisting in playing differently in those two final states) that would violate Formula (B.4). Hence each cell as a single final state.

We now impose that each cell in the main part has exactly two m -successors, and these two m -successors have an m -successor in common. For

the former property, Formula (B.3) already imposes that each cell has at least two m -successors (one labeled with v_1 and one with v_2). We enforce that there cannot be more than two:

$$\mathbf{AG} \left[(m \wedge \square) \Rightarrow [1]_0 \left[\left(\langle 2 \rangle_0 \mathbf{X}^3 (v_1 \wedge \mathbf{X} \alpha) \wedge \langle 2 \rangle_0 \mathbf{X}^3 (v_2 \wedge \mathbf{X} \alpha) \right) \Rightarrow \right. \right. \\ \left. \left. [2]_0 \langle \emptyset \rangle_0 \mathbf{X}^3 \mathbf{X} \alpha \right] \right]. \quad (\text{B.5})$$

Notice that $[2]_0 \langle \emptyset \rangle_0 \varphi$ means that φ has to hold along any outcome of any *memoryless* strategy of Player 2. Assume that a cell has three (or more) successor cells. Then at least one is labeled with v_1 and at least one is labeled with v_2 . There is a strategy for Player 1 to color one v_1 -successor cell and one v_2 -successor cell with α , and a third successor cell with β , thus violating Formula (B.5) (as Player 2 has a strategy to reach a successor cell colored with β)

For the latter property (the two successors have a common successor), we add the following formula (as well as its v_2 -counterpart):

$$[1]_0 \langle \emptyset \rangle_0 \mathbf{G} \left[(m \wedge \square \wedge v_1) \Rightarrow \left(\left[\langle 2 \rangle_0 \mathbf{X}^3 (v_1 \wedge [2]_0 \mathbf{X}^3 \mathbf{X} \alpha) \right] \Rightarrow \right. \right. \\ \left. \left. \left[\langle 2 \rangle_0 \mathbf{X}^3 (\neg v_1 \wedge \mathbf{X}^3 (\neg v_1 \wedge \mathbf{X} \alpha)) \right] \right) \right] \quad (\text{B.6})$$

In this formula, the initial (universal) quantification over strategies of Player 1 fixes a color for each cell. The formula claims that whatever this choice, if we are in some v_1 -cell and can move to another v_1 -cell whose two successors have color α , then also we can move to a v_2 -cell having one α successor (which we require to be a v_2 -cell). As this must hold for any coloring, both successors of the original v_1 -cell share a common successor. Notice that this does not prevent the grid to be collapsed: this would just indicate that there is a *regular* infinite tiling.

We conclude by requiring that the initial state be in a square state of a cell in the main part. \square