Game-Theoretic Semantics for ATL^+ with Applications to Model Checking

Valentin Goranko
Stockholm University, Sweden
and Univ. of Johannesburg,
South Africa (visiting professor)
valentin.goranko@philosophy.su.se

Antti Kuusisto
University of Bremen
Germany
kuusisto@uni-bremen.de

Raine Rönnholm
University of Tampere
Finland
raine.ronnholm@uta.fi

ABSTRACT

We develop a game-theoretic semantics (GTS) for the fragment ATL^+ of the Alternating-time Temporal Logic ATL^*, essentially extending a recently introduced GTS for ATL. We show that the new game-theoretic semantics is equivalent to the standard compositional semantics of ATL^+ (with perfect-recall strategies). Based on the new semantics, we provide an analysis of the memory and time resources needed for model checking ATL^+ and show that strategies of the verifier that use only a very limited amount of memory suffice. Furthermore, using the GTS we provide a new algorithm for model checking ATL^+ and identify a natural hierarchy of tractable fragments of ATL^+ that extend ATL.

Keywords

Logic and game theory; Logics for agents and multi-agent systems; Argumentation-based dialogue and protocols

1. INTRODUCTION

The full Alternating-time Temporal Logic ATL^* [3] is one of the main logical systems used for formalising and verifying strategic reasoning about agents in multi-agent systems. It is very expressive, and that expressiveness comes at the high (2-EXPTIME) price of computational complexity of model checking. Its basic fragment ATL has, on the other hand, tractable model checking but its expressiveness is rather limited. In particular, ATL only allows expressing strategic objectives of the type $\langle\langle A\rangle\rangle\,\Phi$ where Φ is a simple temporal goal involving a single temporal operator. The intermediate fragment ATL^+ naturally emerges as a good alternative, extending ATL so that it is possible to directly express strategic objectives which are Boolean combinations of simple temporal goals. The price for this is the reasonably higher computational complexity of model checking ATL^+, viz. PSPACE-completeness [5]. Still, the PSPACE-completeness result alone gives a rather crude estimate of the amount of memory needed for model checking ATL^+.

In this paper we take an alternative approach to semantic analysis and model checking of ATL^+, based not on the standard compositional semantics but on game-theoretic semantics GTS. The main aims and contributions of this paper are three-fold:

Appears in: *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

1. We introduce an adequate GTS for ATL^+ equivalent to the standard (perfect-recall) compositional semantics.
2. We propose new model checking algorithms for ATL^+ and some of its fragments, using the GTS developed here, rather than the standard semantics. We also analyse the use of memory resources in ATL^+ via GTS.
3. We apply the GTS-based approach to model checking in order to identify new tractable fragments of ATL^+.

The main part of the paper consists of a detailed presentation and analysis of the new GTS for ATL^+. We obtain similar results as in our earlier work [10] where we defined GTS for ATL. We show, in particular, that it is always sufficient to construct *finite paths only* when formulae are evaluated by GTS (even on infinite models). However, for ATL^+, a range of new technical ideas and mechanisms are needed for the correct evaluation of multiple temporal goals pursued simultaneously by the proponent coalition.

The approach via GTS enables us, inter alia, to perform a more precise analysis on the memory resources needed in evaluating ATL^+-formulae than the algorithm from [5] which employs a mix of a path construction procedure for checking strategic formulae $\langle\langle A\rangle\rangle\,\Phi$ on one hand, and the standard labelling algorithm on the other hand. Our model checking algorithm for ATL^+ follows uniformly a procedure directly based on GTS and in fact enables us, inter alia, to identify and correct a flaw in the model checking procedure of [5] and some of the claims on which it is based (see Section 5). However, the PSPACE upper bound result of [5] is easily confirmed by our algorithm, and we provide a new simple proof of that result. In addition to new methods, we use some ideas from [5]. As a new complexity result obtained via GTS, we identify in Section 5 a natural hierarchy of fragments of ATL^+ that extend ATL and have a *tractable* model checking. The hierarchy is based on bounding the *Boolean strategic width* (cf. Section 5) of formulae.

We note that a GTS for ATL^+ alternative to ours could be obtained via a GTS for coalgebraic fixed point logic [16, 8, 9], but such a semantics (being designed for more powerful logics) would not directly lead to our GTS that is custom-made for ATL^+ and thereby enables the complexity analysis we require. Also, the alternative approach would not give a semantics where the construction of only finite paths suffices.

The current paper expands the results in [10] in various non-trivial ways. Several new ideas and technical notions, such as the *role of a seeker* and the use of a *truth function*, will be introduced in order to enable the transition from ATL to ATL^+ in the GTS setting. Also, a connection of our GTS with Büchi games will be established; the connection applies trivially also to the games of [10]. Most importantly, we can directly use the new upgraded semantics in a model checking procedure for ATL^+ and the fragments ATL^k. This would not be possible with with the semantics in [10]. We mention here a few other relevant works: [1, 2, 6, 7, 13, 15, 18]. An extended version of this paper, with technical details and full proofs, is provided in the technical report [11].

2. PRELIMINARIES

Definition 2.1. *A **concurrent game model** (CGM) is a tuple $\mathcal{M} := (\mathbb{A}\mathrm{gt}, \mathrm{St}, \Pi, \mathrm{Act}, d, o, v)$ which consists of:*
*– The following non-empty sets: **agents** $\mathbb{A}\mathrm{gt} = \{a_1, \dots, a_k\}$, **states** St, **proposition symbols** Π, **actions** Act;*
*– The following functions: an **action function** d (such that $d : \mathbb{A}\mathrm{gt} \times \mathrm{St} \to \mathcal{P}(\mathrm{Act}) \setminus \{\emptyset\}$) which assigns a non-empty set of actions available to each agent at each state; a **transition function** o which assigns an **outcome state** $o(q, \vec{\alpha})$ to each state $q \in \mathrm{St}$ and action profile (a tuple of actions $\vec{\alpha} = (\alpha_1, \dots, \alpha_k)$ such that $\alpha_i \in d(a_i, q)$ for each $a_i \in \mathbb{A}\mathrm{gt}$); and finally, a **valuation function** $v : \Pi \to \mathcal{P}(\mathrm{St})$.*

*We use symbols $p, p_0, p_1, \dots$ to denote proposition symbols and $q, q_0, q_1, \dots$ to denote states. Sets of agents are called **coalitions**. The complement $\overline{A} = \mathbb{A}\mathrm{gt} \setminus A$ of a coalition A is the **opposing coalition** of A. The set $\mathsf{action}(A, q)$ of action tuples* available *to coalition A at state $q \in \mathrm{St}$ is defined as $\mathsf{action}(A, q) := \{(\alpha_i)_{a_i \in A} \mid \alpha_i \in d(a_i, q)$ for each $a_i \in A\}$.*

Example 2.2. *Let $\mathcal{M}^* = (\mathbb{A}\mathrm{gt}, \mathrm{St}, \Pi, \mathrm{Act}, d, o, v)$, where: $\mathbb{A}\mathrm{gt} = \{a_1, a_2\}$, $\mathrm{St} = \{q_0, q_1, q_2, q_3, q_4\}$, $\Pi = \{p_1, p_2, p_3\}$, $\mathrm{Act} = \{\alpha, \beta\}$, and d, o and v defined as shown below:*

$\mathcal{M}^*$:

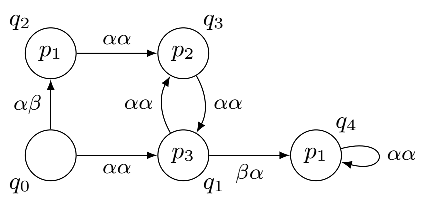

Definition 2.3. *Let $\mathcal{M} = (\mathbb{A}\mathrm{gt}, \mathrm{St}, \Pi, \mathrm{Act}, d, o, v)$ be a* CGM. *A **path** in $\mathcal{M}$ is a sequence $\Lambda : \mathbb{N} \to \mathrm{St}$ of states such that for each $n \in \mathbb{N}$, we have $\Lambda[n+1] = o(\Lambda[n], \vec{\alpha})$ for some admissible action profile $\vec{\alpha}$ in $\Lambda[n]$. A **finite path** (aka **history**) is a finite prefix sequence of a path in $\mathcal{M}$. We let $\mathsf{paths}(\mathcal{M})$ denote the set of all paths in $\mathcal{M}$ and $\mathsf{paths}_{fin}(\mathcal{M})$ the set of all finite paths in $\mathcal{M}$.*

*A (perfect-recall) **strategy** of agent $a \in \mathbb{A}\mathrm{gt}$ is a function $s_a : \mathsf{paths}_{fin}(\mathcal{M}) \to \mathrm{Act}$ such that $s_a(\lambda) \in d(a, \lambda[k])$ for each $\lambda \in \mathsf{paths}_{fin}(\mathcal{M})$ where $\lambda[k]$ is the last state in λ. A **collective strategy** S_A for $A \subseteq \mathbb{A}\mathrm{gt}$ is a tuple of individual strategies, one for each agent in A. We let $\mathsf{paths}(q, S_A)$ denote the set of all paths that can be formed when the agents in A play according to the strategy S_A, beginning from q.*

The syntax of ATL^+ is given by the following grammar.

State formulae: $\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle\!\langle A \rangle\!\rangle\, \Phi \quad (p \in \Pi)$
Path formulae: $\Phi ::= \varphi \mid \neg\Phi \mid \Phi \vee \Phi \mid \mathsf{X}\, \varphi \mid \varphi\, \mathsf{U}\, \varphi$

Other Boolean connectives are defined as usual, and furthermore, $\mathsf{F}\,\varphi$, $\mathsf{G}\,\varphi$ and $\varphi\,\mathsf{R}\,\psi$ are abbreviations for $\top\,\mathsf{U}\,\varphi$, $\neg(\top\,\mathsf{U}\,\neg\varphi)$, and $\neg(\neg\varphi\,\mathsf{U}\,\neg\psi)$ respectively. Φ and Ψ denote path formulae only; φ, ψ, χ denote any formulae.

Definition 2.4. *Let $\mathcal{M}$ be a* CGM. ***Truth** of state and path formulae of ATL^+ is defined, respectively, with respect to states $q \in St$ and paths $\Lambda \in \mathsf{paths}(\mathcal{M})$, as follows:*

- *$\mathcal{M}, q \models p$ iff $q \in v(p)$ (for $p \in \Pi$).*
- *$\mathcal{M}, q \models \neg\varphi$ iff $\mathcal{M}, q \not\models \varphi$.*
- *$\mathcal{M}, q \models \varphi \vee \psi$ iff $\mathcal{M}, q \models \varphi$ or $\mathcal{M}, q \models \psi$.*
- *$\mathcal{M}, q \models \langle\!\langle A \rangle\!\rangle\, \Phi$ iff there exists a (perfect-recall) strategy S_A such that $\mathcal{M}, \Lambda \models \Phi$ for each $\Lambda \in \mathsf{paths}(q, S_A)$.*
- *$\mathcal{M}, \Lambda \models \varphi$ iff $\mathcal{M}, \Lambda[0] \models \varphi$ (where φ is a state formula).*
- *$\mathcal{M}, \Lambda \models \mathsf{X}\,\varphi$ iff $\mathcal{M}, \Lambda[1] \models \varphi$.*
- *$\mathcal{M}, \Lambda \models \neg\Phi$ iff $\mathcal{M}, \Lambda \not\models \Phi$.*
- *$\mathcal{M}, \Lambda \models \Phi \vee \Psi$ iff $\mathcal{M}, \Lambda \models \Phi$ or $\mathcal{M}, \Lambda \models \Psi$.*
- *$\mathcal{M}, \Lambda \models \varphi\,\mathsf{U}\,\psi$ iff there exists $i \in \mathbb{N}$ such that $\mathcal{M}, \Lambda[i] \models \psi$ and $\mathcal{M}, \Lambda[j] \models \varphi$ for all $j < i$.*

The **set of subformulae**, $\mathrm{SUB}(\varphi)$, of a formula φ is defined as usual. Subformulae with a temporal operator as the main connective will be called **temporal subformulae**, while subformulae with $\langle\!\langle\rangle\!\rangle$ as the main connective are **strategic subformulae**. The subformula Ψ of a formula $\varphi = \langle\!\langle A \rangle\!\rangle\, \Psi$ is called the **temporal objective of** φ. We also define the set $At(\Phi)$ of **relative atoms** of Φ as follows:

- $At(\chi \vee \chi') = At(\chi) \cup At(\chi')$ and $At(\neg\chi) = At(\chi)$.
- $At(\langle\!\langle A \rangle\!\rangle\, \chi) = \{\langle\!\langle A \rangle\!\rangle\, \chi\}$ and $At(p) = \{p\}$ for $p \in \Pi$.
- $At(\chi\,\mathsf{U}\,\chi') = \{\chi\,\mathsf{U}\,\chi'\}$ and $At(\mathsf{X}\,\chi) = \{\mathsf{X}\,\chi\}$.

We say that $\chi \in At(\Phi)$ occurs **positively** (resp. **negatively**) **in** Φ if χ has an occurrence in the scope of an even (resp. odd) number of negations in Φ. We denote by $\mathrm{SUB}_{At}(\Phi)$ the subset of $\mathrm{SUB}(\Phi)$ that contains all the relative atoms of Φ and also all the Boolean combinations χ of these relative atoms such that $\chi \in \mathrm{SUB}(\Phi)$.

Example 2.5. *Let $\varphi^* := \langle\!\langle a_1 \rangle\!\rangle\, \Psi$, where*

$$\Psi := (\neg\mathsf{X}\, p_3 \wedge \langle\!\langle a_2 \rangle\!\rangle\, \mathsf{X}\, p_1) \vee (\mathsf{F}\, p_1 \wedge (\neg p_1)\,\mathsf{U}\, p_2).$$

Written without using abbreviations, Ψ becomes

$$\neg(\neg\neg\mathsf{X}\, p_3 \vee \neg\langle\!\langle a_2 \rangle\!\rangle\, \mathsf{X}\, p_1) \vee \neg(\neg(\top\,\mathsf{U}\, p_1) \vee \neg((\neg p_1)\,\mathsf{U}\, p_2)).$$

Here $At(\Psi) = \{\mathsf{X}\, p_3, \langle\!\langle a_2 \rangle\!\rangle\, \mathsf{X}\, p_1, \top\,\mathsf{U}\, p_1, (\neg p_1)\,\mathsf{U}\, p_2\}$, where $\langle\!\langle a_2 \rangle\!\rangle\, \mathsf{X}\, p_1$ is a state formula and the rest are path formulae. The formula $\mathsf{X}\, p_3$ occurs negatively in Ψ and the rest of the formulae in $At(\Psi)$ occur positively in Ψ.

3. GAME-THEORETIC SEMANTICS

In this section we define *bounded*, *finitely bounded* and *unbounded evaluation games* for ATL^+. These games give rise to three different semantic systems, namely, the *bounded*, *finitely bounded* and *unbounded* GTS for ATL^+. We use some terminology and notational conventions introduced in [10].

3.1 Evaluation games: informal description

Given a CGM $\mathcal{M}$, a state q_{in} and a formula φ, the **evaluation game** $\mathcal{G}(\mathcal{M}, q_{in}, \varphi)$ is, intuitively, a formal debate between two opponents, **Eloise** (**E**) and **Abelard** (**A**), about whether the formula φ is true at the state q_{in} in the model $\mathcal{M}$. Eloise claims that φ is true, so she (initially) adopts the role of a **verifier** in the game, and Abelard tries to prove the formula false, so he is (initially) the **falsifier**. These roles (verifier, falsifier) can swap in the course of the game when negations are encountered in the formula. If $\mathbf{P} \in \{\mathbf{E}, \mathbf{A}\}$, then $\overline{\mathbf{P}}$ denotes the **opponent** of $\mathbf{P}$, i.e., $\overline{\mathbf{P}} \in \{\mathbf{E}, \mathbf{A}\} \setminus \{\mathbf{P}\}$.

We now provide an intuitive account of the *bounded* evaluation game and thus the *bounded* GTS for ATL^+. The intuitions underlying the finitely bounded and unbounded GTS are similar. A reader unfamiliar with the concept of

GTS may find it useful to consult, for example, [12] for GTS in general and [10] for ATL-specific GTS. The particular GTS for ATL^+ presented here follows the general principles of GTS, the main original feature here being the treatment of strategic formulae $\langle\langle A\rangle\rangle\,\Phi$. We first give an *informal* account of the way such formulae are treated in our evaluation games. Formal definitions and some concrete examples will be given further, beginning from Section 3.2.

The evaluation of ATL^+ formulae of the type $\langle\langle A\rangle\rangle\,\Phi$ in a given model is based on constructing finite paths in that model. The following two main ideas are central.

Firstly, the path formula Φ in $\langle\langle A\rangle\rangle\,\Phi$ can be *divided* into *goals for the verifier* ($\mathbf{V}$), these being the relative atoms $\psi \in At(\Phi)$ that occur *positively* in Φ, and *goals for the falsifier* ($\overline{\mathbf{V}}$), these being the relative atoms $\psi \in At(\Phi)$ that occur *negatively* in Φ. (Some formulae may be goals for both players.) For simplicity, let us assume for now that Φ is in negation normal form and all the atoms in $At(\Phi)$ are temporal formulae of the type $\mathsf{F}\,p$. Then the verifier's goals are eventuality statements $\mathsf{F}\,p$, while the falsifier's goals are statements $\mathsf{F}\,p'$ that occur negated, and thus correspond to safety statements $\mathsf{G}\,\neg p'$. The verifier wishes to *verify* her/his goals. The falsifier, likewise, wants to *verify* her/his goals, i.e., (s)he wishes to *falsify* the related safety statements.

Secondly, every temporal goal has a unique "finite determination point" on any given path, meaning the following. If a goal $\mathsf{F}\,p$ is true on an infinite path π, then there exists an earliest point q on that path where the fact that $\mathsf{F}\,p$ holds on π becomes *verified* simply because p is true at q. Once $\mathsf{F}\,p$ has been verified, it will remain *true on* π, no matter what happens on the path after q. Similarly, if a statement $\mathsf{G}\,\neg p'$ is false on an infinite path, there is a unique point where $\mathsf{G}\,\neg p'$ first becomes *falsified*. Furthermore, $\mathsf{G}\,\neg p'$ will remain false on the path no matter what happens later on. (Note that there is no analogous finite determination point for ATL^*-formulae such as $\langle\langle A\rangle\rangle\,\mathsf{G}\,\mathsf{F}\,p$ on a given infinite path.)

Now, the game-theoretic evaluation procedure of an ATL^+-formula $\langle\langle A\rangle\rangle\,\Phi$ proceeds roughly as follows. The verifier is controlling the agents in the coalition A and the falsifier the agents in the **opposing coalition** $\overline{A} = \mathbb{A}\mathrm{gt} \setminus A$. The players start constructing a path. (Each transition from one state to another is carried out according to the process "Step phase" defined formally in Section 3.2.2.) The verifier is first given a change to verify some of her/his goals in Φ. The falsifier tries to prevent this and to possibly verify some her/his own goals instead. During this path construction/verification process, the verifier is said to have the role of the **seeker**. A player is allowed to stay as the seeker for only a finite number of rounds. This is ensured by requiring the seeker to announce an ordinal, called **timer**, before the path construction process begins, and then lower the ordinal each time a new state is reached. The process ends when the ordinal becomes zero or when the seeker is satisfied, having verified some of her/his goals. Since ordinals are well-founded, the process must terminate.

After the verifier has ended her/his seeker turn, the falsifier may either end the game or take the role of the seeker. If (s)he decides to become the seeker, then (s)he sets a new timer and the path construction process continues for some finite number of rounds. When the falsifier is satisfied, having verified some of her/his goals, the verifier may again take the seeker's role, and so on. Thus the verifier and falsifier take turns being the seeker, trying to reach (verify) their goals. The number of these alternations is bounded by a **seeker turn counter** which is a finite number that equals the total number of goals in Φ. (The formal description of seeker turn alternation is given in "Deciding whether to continue and adjusting the timer" in Section 3.2.2.)

Each time a goal in Φ becomes verified, this is recorded in a **truth function** T. (The recording of verified goals is described formally in the process "Adjusting the truth function" defined in Section 3.2.2.) The truth function carries the following information at any stage of the game:

- The verifier's goals that have been verified.
- The falsifier's goals that have been verified.
- Other goals remain **open**.

When neither of the players wants to become the seeker, or when the seeker turn counter becomes zero, the path construction process *ends* and the players play a standard Boolean evaluation game on Φ by using the values given by T; the open goals are given truth values as follows:

- The verifier's open goals are *(so far) not verified* and thus considered *false*.
- Likewise, the falsifier's open goals are *(so far) not verified* and thus considered *false*. (Recall that the falsifiers goals are negated.)

Next we consider the conditions when a player is "satisfied" with the current status of the truth function T—and thus wants to end the game—and when (s)he is "unsatisfied" and wants to continue the game as the seeker. Note that when the path construction ends, then every goal is given a Boolean truth value based on the truth function T, as described above. With these values, the formula Φ is either true or false. If Φ is true with the current values based on T, then the verifier can win the Boolean game for Φ; dually, if Φ is not true with the values based on T, then the falsifier can win the Boolean game for Φ. Hence the players want to take the role of the seeker in order to modify the truth function T in such a way that the truth of Φ with respect to T changes from false to true (whence $\mathbf{V}$ is satisfied) or from true to false (whence $\overline{\mathbf{V}}$ is satisfied).

The thruth value of Φ with respect to T can keep changing when T is modified, but only a finite number of changes is possible. Indeed, the maximum number of such truth alternations is the total number of goals in Φ.

3.2 Evaluation games: formal description

Now we will present the **bounded evaluation game** which uses the **bounded transition game** as a subgame for evaluating strategic subformulae. Interleaved with the definition we will provide, in *italics*, a running example that uses $\mathcal{M}^*$ and φ^* from Examples 2.2 and 2.5 respectively.

3.2.1 Rules of the bounded evaluation game

Let $\mathcal{M} = (\mathbb{A}\mathrm{gt}, \mathrm{St}, \Pi, \mathrm{Act}, d, o, v)$ be a CGM, $q_{in} \in \mathrm{St}$ a state, φ a state formula and $\Gamma > 0$ an ordinal called a **timer bound**. The **Γ-bounded evaluation game** $\mathcal{G}(\mathcal{M}, q_{in}, \varphi, \Gamma)$ between the players $\mathbf{A}$ and $\mathbf{E}$ is defined as follows.

A **location** of the game is a tuple $(\mathbf{P}, q, \psi, T)$ where $\mathbf{P} \in \{\mathbf{A}, \mathbf{E}\}$, $q \in \mathrm{St}$ is a state, ψ is a subformula of φ and T is a **truth function**, mapping some subset of $\mathrm{SUB}(\varphi)$ into $\{\top, \bot, \mathsf{open}\}$. ($T$ can also be called a *truth history function*.)

The **initial location** of the game is $(\mathbf{E}, q_{in}, \varphi, T_{in})$, where T_{in} is the empty function. In every location $(\mathbf{P}, q, \psi, T)$, the player $\mathbf{P}$ is called the **verifier** and $\overline{\mathbf{P}}$ the **falsifier** for that location. Intuitively, q is the current state of the game and T encodes truth values of formulae on a path that has been constructed earlier in the game.

Each location is associated with exactly one of the rules **1–6** given below. First we provide the rules for locations $(\mathbf{P}, q, \psi, T)$ where ψ is either a proposition symbol or has a Boolean connective as its main operator:

1. A location $(\mathbf{P}, q, p, T)$, where $p \in \Pi$, is an **ending location** of the evaluation game. If $T \neq \emptyset$, then $\mathbf{P}$ wins the game if $T(p) = \top$ and else $\overline{\mathbf{P}}$ wins. Respectively, if $T = \emptyset$, then $\mathbf{P}$ wins if $q \in v(p)$ and else $\overline{\mathbf{P}}$ wins.

2. From a location $(\mathbf{P}, q, \neg\psi, T)$ the game moves to the location $(\overline{\mathbf{P}}, q, \psi, T)$.

3. In a location $(\mathbf{P}, q, \psi \vee \theta, T)$ the player $\mathbf{P}$ chooses one of the locations $(\mathbf{P}, q, \psi, T)$ and $(\mathbf{P}, q, \theta, T)$, which becomes the next location of the game.

We then define the rules of the evaluation game for locations of type $(\mathbf{P}, q, \langle\!\langle A \rangle\!\rangle\, \Phi, T)$ as follows.

4. Suppose a location $(\mathbf{P}, q, \langle\!\langle A \rangle\!\rangle\, \Phi, T)$ is reached.

- If $T \neq \emptyset$, then this location is an ending location where $\mathbf{P}$ wins if $T(\langle\!\langle A \rangle\!\rangle\, \Phi) = \top$ and else $\overline{\mathbf{P}}$ wins.
- If $T = \emptyset$, then the evaluation game enters a **transition game** $\mathbf{g}(\mathbf{P}, q, \langle\!\langle A \rangle\!\rangle\, \Phi, \Gamma)$. The transition game is a subgame to be defined later on. The transition game eventually reaches an **exit location** $(\mathbf{P}', q', \psi, T')$, and the evaluation game continues from that location. Note that an *exit location* only ends the transition game, so exit locations of transition games and *ending locations* of the evaluation game are different concepts.

The rules for temporal formulae are defined using the truth function T (updated in an earlier transition game) as follows.

5. A location $(\mathbf{P}, q, \varphi \,\mathsf{U}\, \psi, T)$ is an ending location of the evaluation game. $\mathbf{P}$ wins if $T(\varphi \,\mathsf{U}\, \psi) = \top$ and else $\overline{\mathbf{P}}$ wins.

6. Likewise, a location $(\mathbf{P}, q, \mathsf{X}\, \varphi, T)$ is an ending location. $\mathbf{P}$ wins if $T(\mathsf{X}\, \varphi) = \top$ and otherwise $\overline{\mathbf{P}}$ wins.

These are the rules of the evaluation game. We note that the timer bound Γ will be used only in transition games. If $\Gamma = \omega$, we say that the evaluation game is **finitely bounded**.

The initial location of the finitely bounded evaluation game $\mathcal{G}(\mathcal{M}^*, q_0, \varphi^*, \omega)$ (see Examples 2.2 and 2.5) is $(\mathbf{E}, q_0, \langle\!\langle a_1 \rangle\!\rangle\, \Psi, \emptyset)$, from where the transition game $\boldsymbol{g}(\mathbf{E}, q_0, \langle\!\langle a_1 \rangle\!\rangle\, \Psi, \omega)$ begins.

3.2.2 Rules of the transition game

Now we give a detailed description of transition games. A **transition game** $\mathbf{g}(\mathbf{V}, q_0, \langle\!\langle A \rangle\!\rangle\, \Phi, \Gamma)$, where $\mathbf{V} \in \{\mathbf{A}, \mathbf{E}\}$, $q_0 \in \mathrm{St}$, $\langle\!\langle A \rangle\!\rangle\, \Phi \in \mathsf{ATL}^+$ and $\Gamma > 0$ is an ordinal, is defined as follows. $\mathbf{V}$ is called **the verifier in the transition game**. The game $\mathbf{g}(\mathbf{V}, q_0, \langle\!\langle A \rangle\!\rangle\, \Phi, \Gamma)$ is based on **configurations**, i.e., tuples $(\mathbf{S}, q, T, n, \gamma, x)$, where the player $\mathbf{S} \in \{\mathbf{E}, \mathbf{A}\}$ is called the **seeker**; q is the **current state**; $T : At(\Phi) \to \{\top, \bot, \mathsf{open}\}$ is a **truth function**; $n \in \mathbb{N}$ is a **seeker turn counter** ($n \leq |At(\Phi)|$); $\gamma < \Gamma$ is an ordinal called **timer**; and $x \in \{\, \mathbf{i}, \mathbf{ii}, \mathbf{iii} \,\}$ is an index showing the current **phase** of the transition game. The game $\mathbf{g}(\mathbf{V}, q_0, \langle\!\langle A \rangle\!\rangle\, \Phi, \Gamma)$ begins at the **initial configuration** $(\mathbf{V}, q_0, T_0, |At(\Phi)|, \Gamma, \mathbf{i})$, with $T_0(\chi) = \mathsf{open}$ for all $\chi \in At(\Phi)$.

The transition game $\boldsymbol{g}(\mathbf{E}, q_0, \langle\!\langle a_1 \rangle\!\rangle\, \Psi, \omega)$ begins from the initial configuration $(\mathbf{E}, q_0, T_0, 4, \omega, \mathbf{i})$, since $|At(\Psi)| = 4$.

The transition game then proceeds by iterating the following phases **i**, **ii** and **iii** which we *first* describe informally; detailed formal definitions are given afterwards.

i. Adjusting the truth function: In this phase the players make claims on the truth of state formulae at the current state q. If $\mathbf{P}$ makes some claim, then the opponent $\overline{\mathbf{P}}$ may either: 1) accept the claim, whence truth function is updated accordingly, or 2) challenge the claim. In the latter case the *transition game ends* and truth of the claim is verified in a continued evaluation game.

ii. Deciding whether to continue and adjusting the timer: Here the current seeker $\mathbf{S}$ may either continue her/his seeker turn and lower the value of the timer, or end her/his seeker turn. If $\mathbf{S}$ chooses the latter option, then the opponent $\overline{\mathbf{S}}$ of the seeker may either 1) take the role of the seeker and announce a new value for the timer or 2) end the transition game, whence the formula Φ is evaluated based on current values of the truth function.

iii. Step phase: Here the verifier $\mathbf{V}$ chooses actions for the agents in the coalition in A at the current state q. Then $\overline{\mathbf{V}}$ chooses actions for the agents in the opposing coalition $\overline{A}$. After the resulting transition to a new state q' has been made, the game continues again with phase **i**.

We now describe the phases **i**, **ii** and **iii** in detail:

i. Adjusting the truth function.

Suppose the current configuration is $(\mathbf{S}, q, T, n, \gamma, \mathbf{i})$. Then the truth function T is updated by considering, one by one, each formula $\chi \in At(\Phi)$ (in some fixed order). If $T(\chi) \neq \mathsf{open}$, then the value χ cannot be updated. Else the value of χ may be modified according to the rules **A** – **C** below.

A. *Updating T on temporal formulae*: Suppose that we have $\varphi \,\mathsf{U}\, \psi \in At(\Phi)$. Now first the verifier $\mathbf{V}$ may *claim that ψ is true* at the current state q. If $\mathbf{V}$ makes this claim, then $\overline{\mathbf{V}}$ chooses either of the following:

- $\overline{\mathbf{V}}$ *accepts* the claim of $\mathbf{V}$, whence the truth function is updated such that $\varphi \,\mathsf{U}\, \psi \mapsto \top$ ($\varphi \,\mathsf{U}\, \psi$ becomes **verified**).
- $\overline{\mathbf{V}}$ *challenges* the claim of $\mathbf{V}$, whence the transition game ends at the **exit location** $(\mathbf{V}, q, \psi, \emptyset)$. (We note that here, and elsewhere, when a transition game ends, the evaluation game will be continued from the related exit location and the evaluation game will *never* return to the same exited transition game any more.)

If $\mathbf{V}$ does not claim that ψ is true at q, then $\overline{\mathbf{V}}$ may make the same claim (that ψ is true at q). If $\overline{\mathbf{V}}$ makes this claim, then the same two steps above concerning *accepting* and *challenging* are followed, but with $\mathbf{V}$ and $\overline{\mathbf{V}}$ swapped everywhere.

Suppose then that neither of the players claims that ψ is true at q. Then first $\mathbf{V}$ can *claim that φ is false* at q. If $\mathbf{V}$ makes this claim, then $\overline{\mathbf{V}}$ chooses either of the following:

- $\overline{\mathbf{V}}$ accepts the claim, whence the truth function is updated such that $\varphi \,\mathsf{U}\, \psi \mapsto \bot$ ($\varphi \,\mathsf{U}\, \psi$ becomes **falsified**).
- $\overline{\mathbf{V}}$ challenges the claim, whence the transition game ends at the exit location $(\overline{\mathbf{V}}, q, \varphi, \emptyset)$.

If $\mathbf{V}$ does not claim that φ is false at q, then $\overline{\mathbf{V}}$ may make the same claim. If (s)he does, then the same steps as those above are followed, but with $\mathbf{V}$ and $\overline{\mathbf{V}}$ swapped.

B. *Updating T on proposition symbols and strategic formulae*: The truth function can be updated on proposition symbols $p \in At(\Phi)$ and strategic formulae $\langle\!\langle A' \rangle\!\rangle\, \Psi \in At(\Phi)$ only when the phase **i** is executed for the first time (whence we have $q = q_0$). In this case, given such a formula χ, first $\mathbf{V}$ can claim that χ is true at q. Now, if $\overline{\mathbf{V}}$ accepts this claim, then the truth function is updated s.t. $\chi \mapsto \top$. If $\overline{\mathbf{V}}$

challenges the claim, then the transition game ends at the exit location $(\mathbf{V}, q, \chi, \emptyset)$. If $\mathbf{V}$ does not claim that χ is true at q, then $\overline{\mathbf{V}}$ may make the same claim. If (s)he does, then the same steps are followed, but with $\mathbf{V}$ and $\overline{\mathbf{V}}$ swapped.

C. *Updating T on formulae with* X: The truth function can be updated on formulae of type $\mathsf{X}\,\psi \in At(\Phi)$ only when phase **i** is executed for the second time in the transition game (whence q is some successor of q_0). First $\mathbf{V}$ can claim that ψ is true at q. If $\overline{\mathbf{V}}$ accepts this claim, then the truth function is updated s.t. $\mathsf{X}\,\psi \mapsto \top$. If $\overline{\mathbf{V}}$ challenges the claim, then the transition game ends at the exit location $(\mathbf{V}, q, \psi, \emptyset)$. If $\mathbf{V}$ does not claim that ψ is true at q, then $\overline{\mathbf{V}}$ can make the same claim. If (s)he does, the same steps are followed, but with $\mathbf{V}$ and $\overline{\mathbf{V}}$ swapped.

If neither player makes any claim which would update the value of a formula $\chi \in At(\Phi)$, then the value of χ is left open. Once the values of the truth function T have been updated (or left as they are) for all formulae in $At(\Phi)$, a new truth function T' is obtained. The transition game then moves to the new configuration $(\mathbf{S}, q, T', n, \gamma, \mathbf{ii})$.

In the configuration $(\mathbf{E}, q_0, T_0, 4, \omega, \mathbf{i})$ the players begin adjusting T_0 for which initially $T_0(\chi) = \mathsf{open}$ for every $\chi \in At(\Psi)$. Since it is the first round of the transition game, the value of $\mathsf{X}\,p_3$ cannot be modified, but the value of $\langle\!\langle a_2 \rangle\!\rangle \mathsf{X}\,p_1$ can be modified. Suppose that Eloise claims that $\langle\!\langle a_2 \rangle\!\rangle \mathsf{X}\,p_1$ is true at the current state q_0. Now Abelard could challenge the claim, whence the transition game ends and the evaluation game continues from location $(\mathbf{E}, q_0, \langle\!\langle a_2 \rangle\!\rangle \mathsf{X}\,p_1, \emptyset)$ (which leads to a new transition game $\boldsymbol{g}(\mathbf{E}, q_0, \langle\!\langle a_2 \rangle\!\rangle \mathsf{X}\,p_1, \omega)$). Suppose Abelard does not challenge the claim, whence $\langle\!\langle a_2 \rangle\!\rangle \mathsf{X}\,p_1$ is mapped to $\top$.

Since $\mathsf{F}\,p_1$ and $(\neg p_1)\,\mathsf{U}\,p_2$ occur positively in Φ, Eloise has interest only to verify them and Abelard has interest only to falsify them. Eloise could verify $\mathsf{F}\,p_1$ by claiming that p_1 is true, or verify $(\neg p_1)\,\mathsf{U}\,p_2$ by claiming that p_2 is true. But if Eloise makes either of these claims, then Abelard wins the whole evaluation game by challenging, since $q_0 \notin v(p_1) \cup v(p_2)$. Suppose that Eloise does not make any claims. Now, Abelard could claim that $\neg p_1$ is not true, in order to falsify $(\neg p_1)\,\mathsf{U}\,p_2$. But if he does this, he loses the evaluation game if Eloise challenges, since $q_0 \notin v(p_1)$. Suppose that Abelard does not make any claims either. Then the transition game proceeds to configuration $(\mathbf{E}, q_0, T, 4, \omega, \mathbf{ii})$, where $T(\langle\!\langle a_2 \rangle\!\rangle \mathsf{X}\,p_1) = \top$ and $T(\chi) = \mathsf{open}$ for the other $\chi \in At(\Psi)$.

ii. Deciding whether to continue and adjusting the timer.
Suppose a configuration $(\mathbf{S}, q, T, n, \gamma, \mathbf{ii})$ has been reached. Assume first that $\gamma \neq 0$. Then the seeker $\mathbf{S}$ can choose whether to continue the transition game as the seeker. If yes, then $\mathbf{S}$ chooses some ordinal $\gamma' < \gamma$ and the transition game continues from $(\mathbf{S}, q, T, n, \gamma', \mathbf{iii})$. If $\mathbf{S}$ does not want to continue, or if $\gamma = 0$, then one of the following applies.

a) Assume that $n \neq 0$. Then the player $\overline{\mathbf{S}}$ chooses whether (s)he wishes to continue the transition game. If yes, then $\overline{\mathbf{S}}$ chooses an ordinal $\gamma' < \Gamma$ (note that $\overline{\mathbf{S}}$ indeed *resets* the timer value) and the transition game continues from $(\overline{\mathbf{S}}, q, T, n-1, \gamma', \mathbf{iii})$. Otherwise the transition game ends at the exit location $(\mathbf{V}, q, \Phi, T)$.
b) Assume $n = 0$. Then the transition game ends at the exit location $(\mathbf{V}, q, \Phi, T)$.

In $(\mathbf{E}, q_0, T, 4, \omega, \mathbf{ii})$ Eloise may decide whether to continue the transition game as the seeker. Suppose that Eloise does not continue, whence Abelard may now become the seeker and continue the transition game, or end it. If Abelard ends the transition game, then the evaluation game is continued from $(\mathbf{E}, q_0, \Psi, T)$. But since $T(\mathsf{X}\,p_3) = \mathsf{open}$ and $T(\langle\!\langle a_2 \rangle\!\rangle \mathsf{X}\,p_1) = \top$, Eloise can then win the evaluation game by choosing the left disjunct of Ψ. Suppose that Abelard decides to become the seeker, whence he chooses some $m < \omega$ and the next configuration is $(\mathbf{A}, q_0, T, 3, m, \mathbf{iii})$.

iii. Step phase.
Suppose that the configuration is $(\mathbf{S}, q, T, n, \gamma, \mathbf{iii})$.

a) First $\mathbf{V}$ chooses an action $\alpha_i \in d(a_i, q)$ for each $a_i \in A$.
b) Then $\overline{\mathbf{V}}$ chooses an action $\alpha_i \in d(a_i, q)$ for each $a_i \in \overline{A}$.

The resulting action profile produces a **successor state** $q' := o(q, \alpha_1, \ldots, \alpha_k)$. The transition game then moves to the configuration $(\mathbf{S}, q', T, n, \gamma, \mathbf{i})$.

In the configuration $(\mathbf{A}, q_0, T, 3, m, \mathbf{iii})$ Eloise (who is the verifier $\mathbf{V}$) first chooses action for agent a_1, then Abelard chooses action for agent a_2, which produces either successor state q_1 or q_2. Then the transition game continues from the configuration $(\mathbf{A}, q_j, T, 3, m, \mathbf{i})$, where $j \in \{1, 2\}$.

This concludes the definition of the rules for the phases **i, ii** and **iii** in the transition game $\mathbf{g}(\mathbf{V}, q_0, \langle\!\langle A \rangle\!\rangle\, \Phi, \Gamma)$.

Suppose first that the transition game is continued from $(\mathbf{A}, q_2, T, 3, m, \mathbf{i})$. Since it is the second round, Abelard could now try to verify $\mathsf{X}\,p_3$ by claiming that p_3 is true at q_2. However, then Eloise would win by challenging. But if Abelard does not try to verify $\mathsf{X}\,p_3$ now, then the value of $\mathsf{X}\,p_3$ will stay open. In that case Eloise will win the evaluation game simply by not making any more claims in the transition game.

Suppose then that the game continues from $(\mathbf{A}, q_1, T, 3, m, \mathbf{i})$. Suppose that Abelard verifies $\mathsf{X}\,p_3$ by claiming that p_3 is true and that Eloise does not challenge. If the transition game now ended at $(\mathbf{E}, q_1, \Psi, T')$ with $T'(\mathsf{X}\,p_3) = \top$, Abelard would win. Thus, suppose that Abelard ends his seeker turn and Eloise chooses some finite timer, say 2. At $(\mathbf{E}, q_1, T', 2, 2, \mathbf{iii})$ Eloise can force the resulting state q_3 by choosing α for a_1. At $(\mathbf{E}, q_3, T', 2, 2, \mathbf{i})$ Eloise can verify $(\neg p_1)\,\mathsf{U}\,p_2$ by claiming that p_2 is true at q_3. Furthermore, Eloise can move via q_1 to q_4 and verify $\mathsf{F}\,p_1$ there, before timer reaches 0. When the evaluation game is eventually continued, Eloise wins by choosing the right disjunct of Ψ.

3.2.3 The unbounded evaluation game

Let $\mathcal{G}(\mathcal{M}, q, \varphi, \Gamma)$ be a Γ-bounded evaluation game. We can define a corresponding **unbounded evaluation game**, $\mathcal{G}(\mathcal{M}, q, \varphi)$, by replacing transition games $\mathbf{g}(\mathbf{V}, q, \langle\!\langle A \rangle\!\rangle\, \Phi, \Gamma)$ with **unbounded transition games**, $\mathbf{g}(\mathbf{V}, q, \langle\!\langle A \rangle\!\rangle\, \Phi)$; these are played with the same rules as $\mathbf{g}(\mathbf{P}, q_0, \langle\!\langle A \rangle\!\rangle\, \Phi, \Gamma)$ except that timers γ are not used in them. Instead, the players can keep the role of a seeker for arbitrarily long and thus the game may last for an infinite number of rounds. In the case of an infinite play, the player who took the last seeker turn loses the entire evaluation game. (Recall that the number of seeker alternations is bounded by the number $|At(\Phi)|$.)

3.3 Defining the game theoretic semantics

REMARK 3.1. *The description of transition games above is based on a simplified notion of configurations. The phases **i–iii** consist of several "subphases" and more information should be encoded into configurations. The full notion of configuration should also include: In phase **i**, a counter indicating the relative atom currently under consideration by the players; flags for each player indicating whether and what claim (s)he has made on the truth of the current relative atom; a 3-bit flag indicating if it is the first, second or some later round in the transition game. For phase **ii**, a flag* *whether the current seeker wants to continue, and for phase* ***iii****, a record of the current choice of actions for the agents in* A *by* $\mathbf{V}$*. For simplicity, we mainly omit these details.*

Hereafter a **position** in an evaluation game will mean either a location of the form $(\mathbf{P}, q, \varphi, T)$ or a configuration in the fully extended form described in the remark above. Note that by this definition, at every position, only one of the players (Abelard or Eloise) has a move to choose. Thus the entire evaluation game—including transition games as subgames—is a turn-based game of perfect information.

The formal definitions of players' memory-based strategies in the evaluation games games are defined as expected, based on histories of positions. As usual, a strategy for a player $\mathbf{P}$ is called **winning** if, following that strategy, $\mathbf{P}$ is guaranteed to win regardless of how $\overline{\mathbf{P}}$ plays. A strategy is **positional** if it depends only on the current position.

Definition 3.2. *Let* $\mathcal{M}$ *be a* CGM*,* $q \in \mathrm{St}$*,* $\varphi \in \mathsf{ATL}^+$ *and* Γ *an ordinal.* ***Truth of*** φ ***according to*** Γ***-bounded*** $(\Vdash_\Gamma)$ *and* ***unbounded*** $(\Vdash)$ GTS *is defined as follows:*

$\mathcal{M}, q \Vdash_\Gamma \varphi$ *(resp.* $\mathcal{M}, q \Vdash \varphi$*) iff Eloise has a positional winning strategy in* $\mathcal{G}(\mathcal{M}, q, \varphi, \Gamma)$ *(resp.* $\mathcal{G}(\mathcal{M}, q, \varphi)$*).*

Example 3.3. *Consider the following* CGM

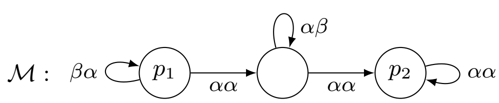

Here we have $\mathcal{M} = (\mathbb{A}\mathrm{gt}, \mathrm{St}, \Pi, \mathrm{Act}, d, o, v)$*, where* $\mathbb{A}\mathrm{gt} = \{a_1, a_2\}$*,* $\Pi = \{p_1, p_2\}$*,* $\mathrm{St} = \{q_0, q_1, q_2\}$*,* $\mathrm{Act} = \{\alpha, \beta\}$*, and the transition, outcome and valuation functions are defined as above.*

Let $\varphi := \langle\!\langle a_2 \rangle\!\rangle (\mathsf{G}\, p_1 \vee \mathsf{F}\, p_2)$ *(here* $\mathsf{G}\, p_1 = \neg \mathsf{F} \neg p_1$*). We describe a winning strategy for Eloise in the unbounded evaluation game* $\mathcal{G}(\mathcal{M}, q_0, \varphi)$*. Eloise immediately ends her seeker's turn and does not make claims while being at* q_0*. If Abelard makes claims at* q_0*, she challenges those claims. If Abelard ends the transition game at* q_0*, Eloise wins the evaluation game by choosing* $\neg \mathsf{F} \neg p_1$*, as now the value of* $\mathsf{F} \neg p_1$ *is* open*. Suppose that Abelard forces a transition to* q_1 *by choosing* α *for* a_1*. If he claims* $\neg p_1$ *is true at* q_1*, Eloise does not challenge. If Abelard ends his seeker turn at* q_1*, Eloise becomes the seeker. At* q_1 *she forces a transition to* q_2*, by choosing* α *for* a_2*. Then she verifies* $\mathsf{F}\, p_2$ *by claiming that* p_2 *is true at* q_2*. If the transition game ends at* q_2*, she wins by choosing* $\mathsf{F}\, p_2$*, whose value is* $\top$*. Note that by following this strategy, Eloise cannot stay as a seeker for infinitely long.*

4. RESULTS ON EVALUATION GAMES

4.1 Positional determinacy

Proposition 4.1. *Bounded evaluation games are determined and the winner has a positional winning strategy.*

Proof. (Sketch) Since ordinals are well-founded and they must decreased during transition games, it is easy to see that the game tree is well-founded. Thus positional determinacy follows essentially by backward induction. □

Proposition 4.2. *Unbounded evaluation games are determined and the winner has a positional winning strategy.*

Proof. (Sketch) This claim can be proved in a similar way as Gale-Stewart theorem. Another way to prove the claim is to show that unbounded evaluation games are essentially Büchi-games (see, e.g., [14] for Büchi-games).

The details of the proof via Büchi-games are in [11], but the principal idea is to set up a Büchi condition such that Eloise wins the Büchi game if the set of positions visited infinitely often is included in the union of configurations of the transition games where Abelard is the seeker and positions of the evaluation game where Eloise has already won. □

We say that Eloise (Abelard) has a *winning strategy in a transition game*, if she (he) can force that game to end at an exit location where she (he) has a winning strategy in the evaluation game that continues from there. By the positional determinacy, we have the following consequence: If Eloise (Abelard) has a perfect recall strategy in a bounded or unbounded evaluation game (or transition game), then she (he) has a positional winning strategy in that game.

4.2 Finding stable timer bounds

We consider a "semi-bounded" variant of the transition game in which one player must use timers when being the seeker and the other is allowed to play without timers. A timer bound Γ is **stable** for an unbounded transition game $\mathbf{g}(\mathbf{V}, q_0, \langle\!\langle A \rangle\!\rangle \Phi)$ if the player with a winning strategy in $\mathbf{g}(\mathbf{V}, q_0, \langle\!\langle A \rangle\!\rangle \Phi)$ can in fact win using timers below Γ.

Proposition 4.3. *Let* $\mathcal{M}$ *be a finite* CGM*,* $q_0 \in \mathrm{St}$ *a state and* $\Phi \in \mathsf{ATL}^+$ *a path formula. Then* $k := |\,\mathrm{St}\,| \cdot |At(\Phi)|$ *is a stable timer bound for* $\boldsymbol{g}(\mathbf{V}, q_0, \langle\!\langle A \rangle\!\rangle \Phi)$*.*

Proof. (Sketch) Let $c = (\mathbf{E}, q, T, n, x)$ be a configuration (for an *unbounded* game, so no timer is listed). Suppose that exit location $(\mathbf{V}, q, \Phi, T)$ is not a winning location for Eloise. Then she wants to stay as the seeker until the truth function is modified to T' so that T' makes Φ true. Since T is updated state-wise, it is not beneficial for Eloise to go in loops such that T is not updated. Hence, if Eloise has a winning strategy from c, then she has a winning strategy in which T is updated at least once every $|\,\mathrm{St}\,|$ rounds. Since T can be updated at most $|At(\Phi)|$ times, we see that a timer greater than $k = |\,\mathrm{St}\,| \cdot |At(\Phi)|$ is not needed. □

Corollary 4.4. *If* $\mathcal{M}$ *is a finite* CGM*, the unbounded* GTS *is equivalent on* $\mathcal{M}$ *to the* $(|\,\mathrm{St}\,| \cdot |\varphi|)$*-bounded* GTS*.*

In order to find stable timer bounds for infinite models, we give the following definition (cf. Def 4.12 in [10]).

Definition 4.5. *Let* $\mathcal{M}$ *be a* CGM *and let* $q \in \mathrm{St}$*. We define the* ***branching degree of*** q*,* $\mathsf{BD}(q)$*, as the cardinality of the set of outcome states from* q*:* $\mathsf{BD}(q) := \mathsf{card}(\{o(q, \vec{\alpha}) \mid \vec{\alpha} \in \mathsf{action}(\mathbb{A}\mathrm{gt}, q)\})$*. We define the* ***regular branching bound of*** $\mathcal{M}$*, or* $\mathsf{RBB}(\mathcal{M})$*, as the smallest infinite regular cardinal* κ *such that* $\kappa > \mathsf{BD}(q)$ *for every* $q \in \mathrm{St}$*. Note that* $\mathsf{RBB}(\mathcal{M}) = \omega$ *if and only if* $\mathcal{M}$ *is* ***image finite****.*

Proposition 4.6. *Let* $\mathcal{M}$ *be a* CGM*,* $q_0 \in \mathrm{St}$ *and* $\Phi \in \mathsf{ATL}^+$ *a path formula. Then* $\mathsf{RBB}(\mathcal{M})$ *is a stable timer bound for* $\boldsymbol{g}(\mathbf{V}, q_0, \langle\!\langle A \rangle\!\rangle \Phi)$*.*

Proof. (Sketch) Let $\mathbf{P}$ be the player with a winning strategy τ in $\mathbf{g}(\mathbf{V}, q_0, \langle\!\langle A \rangle\!\rangle \Phi)$. Let T_τ be the game tree corresponding to τ. We associate each configuration c in T_τ, in which $\mathbf{P}$ is a seeker, with an ordinal γ_c in s.t. c is a winning configuration for $\mathbf{P}$ with the timer γ_c. We do this by attaching 0 to the configurations in which $\mathbf{P}$ ends her/his seeker turn. For other configurations c we take the supremum of the timers attached to the configurations that follow c. Since $\mathsf{RBB}(\mathcal{M})$ is a regular cardinal, we have $\gamma_c \leq \mathsf{RBB}(\mathcal{M})$ for each γ_c. For a complete proof, see [11]. □

Consequently, finite timers suffice in image finite models. However, the finitely bounded GTS ($\Gamma = \omega$) is not generally equivalent to the unbounded GTS (see Example 3.7 in [10]).

COROLLARY 4.7. *Suppose that* $\Gamma \geq \mathsf{RBB}(\mathcal{M})$. *Then the unbounded* GTS *is equivalent on* $\mathcal{M}$ *to the* Γ*-bounded* GTS.

PROOF. Suppose first that $\mathcal{M}, q \Vdash \varphi$. By Proposition 4.6 Eloise can win the evaluation game using timers smaller than Γ when being the seeker. Hence clearly $\mathcal{M}, q \Vdash_\Gamma \varphi$.

Suppose then $\mathcal{M}, q \nVdash \varphi$. By Proposition 4.2, Abelard has a winning strategy in $\mathcal{G}(\mathcal{M}, q, \varphi)$. Thus, by Proposition 4.6, Abelard can win $\mathcal{G}(\mathcal{M}, q, \varphi)$ using timers smaller than Γ when being the seeker. Hence Abelard clearly has a winning strategy in $\mathcal{G}(\mathcal{M}, q, \varphi, \Gamma)$ and thus $\mathcal{M}, q \nVdash_\Gamma \varphi$. □

4.3 GTS vs compositional semantics for ATL$^+$

We now define a so-called **finite path semantics**, to be used later. See [5] for a similar definition. We define the **length** $lgt(\lambda)$ of a finite path λ as the number of transitions in λ (whence the last state of λ is $\lambda[lgt(\lambda)]$). If λ is a prefix sequence of λ', we write $\lambda \preceq \lambda'$.

DEFINITION 4.8. *Let* $\mathcal{M}$ *be a* CGM *and* $\lambda \in \mathsf{paths}_{fin}(\mathcal{M})$. ***Truth*** *of an* ATL^+ *path formula* Φ *on the finite path* λ *is defined as expected, the non-obvious clauses being as follows:*

- $\mathcal{M}, \lambda \models \mathsf{X}\,\varphi$ *iff* $lgt(\lambda) \geq 1$ *and* $\mathcal{M}, \lambda[1] \models \varphi$.
- $\mathcal{M}, \lambda \models \varphi\,\mathsf{U}\,\psi$ *iff there exists some* $i \leq lgt(\lambda)$ *such that* $\mathcal{M}, \lambda[i] \models \psi$ *and* $\mathcal{M}, \lambda[j] \models \varphi$ *for all* $j < i$.

DEFINITION 4.9. *Let* $\mathcal{M}$ *be a* CGM, $\lambda \in \mathsf{paths}(\mathcal{M})$ *and* Φ *a path formula of* ATL^+. *An index* $i \geq 1$ *is a* ***truth swap point*** *of* Φ *on* λ *if we have* $\mathcal{M}, \lambda[i-1] \models \Phi \Leftrightarrow \mathcal{M}, \lambda[i] \not\models \Phi$.

We define the ***truth swap number*** *of* Φ *on* λ *to be* $TSN(\Phi, \lambda) := \mathsf{card}(\{i \mid i \text{ is a truth swap point of } \Phi \text{ on } \lambda\})$.

The claims of the following lemma are easy to prove. Similar observations have been made in [5].

LEMMA 4.10. *Let* $\mathcal{M}$ *be a* CGM, $\lambda \in \mathsf{paths}(\mathcal{M})$ *and* Φ *a path formula of* ATL^+. *Now the following claims hold:*

1. $TSN(\Phi, \lambda) \leq |\{\Psi \in At(\Phi) \mid \Psi \text{ is a temporal subformula}\}|$.
2. $\mathcal{M}, \lambda \models \Phi$ *iff there is some* $k \in \mathbb{N}$ *s.t.* $\mathcal{M}, \lambda_0 \models \Phi$ *for every finite* $\lambda_0 \preceq \lambda$ *for which* $lgt(\lambda_0) \geq k$.

THEOREM 4.11. *The unbounded* GTS *is equivalent to the standard (perfect-recall) compositional semantics of* ATL^+.

PROOF. (Sketch) We prove by induction on ATL^+ state formulae φ that $\mathcal{M}, q \Vdash \varphi$ iff $\mathcal{M}, q \models \varphi$. The cases $\varphi = p$ and $\varphi = \psi \vee \theta$ are easy and $\varphi = \neg\psi$ follows from the inductive hypothesis for ψ and determinacy of the evaluation games.

Consider the case $\varphi = \langle\!\langle A \rangle\!\rangle\, \Phi$. It suffices to show that Eloise has a winning strategy in the (unbounded) transition game $\mathbf{g}(\mathbf{E}, q, \langle\!\langle A \rangle\!\rangle\, \Phi)$ iff the coalition A has a (perfect-recall) strategy S_A s.t. $\mathcal{M}, \lambda \models \Phi$ for every $\lambda \in \mathsf{paths}(q, S_A)$.

Suppose that $\mathbf{E}$ has a (positional) winning strategy τ in $\mathbf{g}(\mathbf{E}, q, \langle\!\langle A \rangle\!\rangle\, \Phi)$. Let $T_\mathbf{g}$ be the game tree that is formed by all of those paths of *states* that can be encountered with τ. We define S_A essentially using the actions according to τ for every finite path in $T_\mathbf{g}$.

Let $\lambda \in \mathsf{paths}(q, S_A)$, whence $\lambda \in T_\mathbf{g}$. Let $k \in \mathbb{N}$ be s.t. Eloise neither does any further claims nor becomes a seeker after the state $\lambda[k]$ (in the infinite play that follows λ). Now we can show that $\mathcal{M}, \lambda_0 \models \Phi$ for every finite $\lambda_0 \preceq \lambda$ s.t. $lgt(\lambda_0) \geq k$. Hence $\mathcal{M}, \lambda \models \Phi$ by Lemma 4.10(2).

Suppose then that there is an S_A s.t. $\mathcal{M}, \lambda \models \Phi$ for every $\lambda \in \mathsf{paths}(q, S_A)$. We define a strategy τ for Eloise as follows (τ will not be positional, but since unbounded transition games are positionally determined, a positional winning strategy τ' for Eloise will exist). Suppose the game is at some configuration c that is reached with a finite path λ_0 such that q_0 is the last state of λ_0.

- When adjusting the truth function, $\mathbf{E}$ makes all the valid claims and all valid challenges; these are made according to the *compositional* truth at the current state q_0.
- If $\mathbf{E}$ is the seeker in c and $\mathcal{M}, \lambda_0 \models \Phi$, then $\mathbf{E}$ decides to end her seeker turn; else, $\mathbf{E}$ continues as a seeker.
- If $\mathbf{A}$ ends seeking at c and $\mathcal{M}, \lambda_0 \not\models \Phi$, then $\mathbf{E}$ decides to become the seeker; else $\mathbf{E}$ ends the transition game at c.
- If $\mathbf{E}$ needs to choose actions for the agents in A at c, she chooses them according to S_A for λ_0.

As Eloise chooses actions for A according to S_A, every path of *states* formed with τ is a prefix sequence of some path $\lambda \in \mathsf{paths}(q, S_A)$. Since $\mathcal{M}, \lambda \models \Phi$ for every $\lambda \in \mathsf{paths}(q, S_A)$, by Lemma 4.10 and the definition of τ, Eloise cannot stay as the seeker forever when playing with τ. If Abelard stays as the seeker forever, then Eloise wins. And since Eloise manages the truth function according to the compositional truth, we can show by the inductive hypothesis that every exit location is a winning location for Eloise. Hence τ is a winning strategy for Eloise. See [11] for more details. □

COROLLARY 4.12. *If* $\Gamma \geq \mathsf{RBB}(\mathcal{M})$, *then the* Γ*-bounded* GTS *is equivalent on* $\mathcal{M}$ *with the standard (perfect recall) compositional semantics of* ATL^+.

5. MODEL CHECKING ATL$^+$ USING GTS

5.1 Revisiting the PSPACE upper bound proof

As mentioned earlier, the PSPACE upper bound proof for the model checking of ATL^+ in [5] contains a flaw. Indeed, the claim of Theorem 4 in [5] is incorrect and a counterexample to it can be extracted from our Example 3.3, where $\mathcal{M}, q_0 \models \varphi$ for $\varphi = \langle\!\langle a_2 \rangle\!\rangle\, (\mathsf{G}\, p_1 \vee \mathsf{F}\, p_2)$. In the notation of [5], since $|St_\mathcal{M}| = 3$ and $\mathcal{APF}(\varphi) = 2$, by the claim there must be a 6-witness strategy for the agent 2 for $(\mathcal{M}, q_0, \mathsf{G}\, p_1 \vee \mathsf{F}\, p_2)$. However, this is not the case, since the player 1 can choose to play at q_0 4 times β, and then α. Then $\mathcal{M}, \lambda \not\models^6 (\mathsf{G}\, p_1 \vee \mathsf{F}\, p_2)$ on any resulting path λ.

The reason for the problem indicated above is that compositional semantics easily ignores the role and power of the falsifier (Abelard) in the formula evaluation process. Still, using the GTS introduced above, we will demonstrate in a simple way that the upper bound result is indeed correct.

The input to the model checking problem of ATL^+ is an ATL^+ formula φ, a finite CGM $\mathcal{M}$ and a state q in $\mathcal{M}$. We assume that $\mathcal{M}$ is encoded in the standard, explicit way (cf. [3, 5]) that provides a full explicit description of the transition function o. We do not assume any bounds on the number of proposition symbols or agents in the input. We only consider here the semantics of ATL^+ based on perfect information and perfect-recall strategies.

THEOREM 5.1 ([5]). *The* ATL^+ *model checking problem is* PSPACE*-complete.*

PROOF. We get the lower bound directly from [5], so we only prove the upper bound here. By Theorem 4.11 and

Proposition 4.3, if $\mathcal{M}$ is a finite CGM, we have $\mathcal{M}, q \models \varphi$ iff Eloise has a positional winning strategy in $\mathcal{G}(\mathcal{M}, q, \varphi, N)$ with $N = |\,\mathrm{St}\,| \cdot |\varphi|$. It is routine to construct an alternating Turing machine TM that simulates $\mathcal{G}(\mathcal{M}, q, \varphi, N)$ such that the positions for Eloise correspond to existential states of TM and Abelard's positions to universal states. Due to the timer bound N, the machine runs in polynomial time. It is clear that if Eloise has a (positional or not) winning strategy in the evaluation game, then TM accepts. Conversely, if TM accepts, we can read a non-positional winning strategy for Eloise from the the computation tree (with only one successful move for existential states recorded everywhere) which demonstrates that TM accepts. By Proposition 4.1, Eloise thus also has a positional winning strategy in the evaluation game. Since APTIME = PSPACE, the claim follows. □

5.2 A hierarchy of tractable fragments of ATL^+

We now identify a natural hierarchy of tractable fragments of ATL^+. Let k be a positive integer. Define ATL^k to be the fragment of ATL^+ where all formulae $\langle\!\langle A \rangle\!\rangle\, \Phi$ have the property that $|At(\Phi)| \leq k$. Note that ATL^1 is essentially the same as ATL (with Release). Note also that the number of non-equivalent formulae of ATL^k is not bounded for any k *even in the special case where the number of propositions and actions is constant*, because nesting of strategic operators $\langle\!\langle A \rangle\!\rangle$ is not limited. Still, we will show that the model checking problem for ATL^k is PTIME-complete for any fixed k. Again CGMs are encoded explicitly and no restrictions on the number of propositions or actions is assumed.

THEOREM 5.2. *For any fixed $k \in \mathbb{N}$, the model checking problem for ATL^k is PTIME-complete.*

PROOF. (Sketch) The claim is well-known for ATL (see [3]), so we have the lower bound for free for any k. One possible proof strategy for the upper bound would involve using alternating LOGSPACE-machines, but here we argue via Büchi-games instead. See the details of the reduction of unbounded evaluation games to Büchi-games in the technical report [11] (the proof for Proposition 4.2).

Consider a triple $(\mathcal{M}, q, \varphi)$, where $\varphi \in \mathsf{ATL}^k$. By the proof of Proposition 4.2 (see also [11]), there exists a Büchi game BG such that Eloise wins the unbounded evaluation game $\mathcal{G}(\mathcal{M}, q, \varphi)$ iff she wins BG from the state of BG that corresponds to the beginning position of the evaluation game. We then observe that since we are considering ATL^k for a *fixed* k, the domain size of each truth function T used in the evaluation game is at most k, and thus the number of positions in $\mathcal{G}(\mathcal{M}, q, \varphi)$ is *polynomial* in the size of the input $(\mathcal{M}, q, \varphi)$. (Check Remark 3.1 for all the information that should be encoded in a position in *bounded* evaluation games; here we only use the simpler unbounded games.) Thus also the size of BG is polynomial in the input size.

We note that In order to avoid blow-ups, it is essential that the maximum domain size k of truth functions T is fixed. We also note—as mentioned already in [3]—that the number of transitions in $\mathcal{M}$ is not bounded by the square of the number of states of $\mathcal{M}$. In fact, already because we impose no limit on the number of actions (other than finiteness) in $\mathcal{M}$, the number of transitions in relation to states is arbitrary. However, this is no problem to us since an explicit encoding of $\mathcal{M}$—which lists all transitions explicitly—is part of the input to the model checking problem. Since Büchi games can be solved in PTIME, the claim follows. □

5.3 Bounded memory semantics for ATL^k

Here we show that to capture the compositional (perfect-recall) semantics for ATL^k, it suffices to consider agents' strategies that use only a limited amount of memory.

Strategies with bounded memory for ATL^* can be naturally defined using *finite state transducers.* (For a transducer-based definition of bounded memory strategies, see e.g. [17], and see [4] for more on this topic.) Using such strategies, an agent's moves are determined both by the current state in the model and by the current state (*memory cell*) of the agent's transducer. Then transitions take place both in the model and in the state of the transducer. In the compositional m-**bounded memory semantics** ($\models^m$) for ATL^+, agents are allowed to use at most m memory cells, i.e., strategies defined by transducers with at most m states.

Observing that the use of the truth function T in our GTS is analogous to the use of memory cells in m-bounded memory semantics, we obtain the following result.

THEOREM 5.3. *For ATL^k, the unbounded* GTS *is equivalent to the m-bounded memory semantics for $m = 3^k - 2^k$.*

PROOF. (Sketch) Let $m := 3^k - 2^k$ and $\varphi \in \mathsf{ATL}^k$. We show that $\mathcal{M}, q \Vdash \varphi$ iff $\mathcal{M}, q \models^m \varphi$. The implication from right to left is immediate by Theorem 4.11. We prove the other direction by induction on φ. The only interesting case is when $\varphi = \langle\!\langle A \rangle\!\rangle\, \Phi$. Suppose that Eloise has a winning strategy in $\mathbf{g}(\mathbf{E}, q, \langle\!\langle A \rangle\!\rangle\, \Phi)$.

We define a memory transducer $\mathcal{T}$ that can be used for the collective strategy of A. We fix the set of states C of $\mathcal{T}$ to be the set of all truth functions T for $At(\Phi)$ s.t. $T(\chi) = \mathsf{open}$ for at least one $\chi \in At(\Phi)$, whence $|C| \leq 3^k - 2^k = m$. The initial state of $\mathcal{T}$ is T_0 where $T_0(\chi) = \mathsf{open}$ for all $\chi \in At(\Phi)$. The transitions in $\mathcal{T}$ are defined according to how Eloise updates T. However, when T becomes fully updated (i.e. $T(\chi) \neq \mathsf{open}$ for every $\chi \in At(\Phi)$), no further transitions are made (then all the relative atoms have been verified/falsified and the truth of Φ on the path is fixed).

The strategy for each $a \in A$ is defined on $C \times \mathrm{St}$: for a pair (T, q), the agent a follows the action prescribed by Eloise's winning strategy for the corresponding step phase in the transition game. It is now easy to show that $\mathcal{M}, \lambda \models^m \Phi$ for any path λ that is consistent with the resulting collective strategy for A. See [11] for more details. □

Consequently, by Theorem 4.11, the compositional perfect-recall semantics and $(3^k - 2^k)$-bounded memory semantics are equivalent for ATL^k. This extends the known fact that positional strategies (using 1 memory cell) suffice for the semantics of ATL (which is essentially the same as ATL^1).

Conclusion

The GTS for ATL^+ developed here has both conceptual and technical significance, as it explains better how the memory-based strategies in the compositional semantics can be generated, and thus also provides better insight on the algorithmic aspects of that semantics. A natural extension of the present work would be to develop GTS for the full ATL^*.

Acknowledgements

The work of Valentin Goranko was supported by a research grant 2015-04388 of the Swedish Research Council. The work of Antti Kuusisto was supported by the ERC grant 647289 "CODA." We thank the referees for valuable remarks.

REFERENCES

[1] T. Ågotnes, V. Goranko, and W. Jamroga. Alternating-time temporal logics with irrevocable strategies. In D. Samet, editor, *Proceedings of the 11th International Conference on Theoretical Aspects of Rationality and Knowledge (TARK XI)*, pages 15–24, Univ. Saint-Louis, Brussels, 2007. Presses Universitaires de Louvain.

[2] T. Ågotnes, W. van der Hoek, and M. Wooldridge. Quantified coalition logic. *Synthese*, 165(2):269–294, 2008.

[3] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002.

[4] T. Brihaye, A. D. C. Lopes, F. Laroussinie, and N. Markey. ATL with strategy contexts and bounded memory. In S. Artëmov and A. Nerode, editors, *Proc. of LFCS 2009*, volume 5407 of *LNCS*, pages 92–106. Springer, 2009.

[5] N. Bulling and W. Jamroga. Verifying agents with memory is harder than it seemed. *AI Commun.*, 23(4):389–403, 2010.

[6] P. Cermák, A. Lomuscio, F. Mogavero, and A. Murano. MCMAS-SLK: A model checker for the verification of strategy logic specifications. In *Proc. of CAV 2014*, volume 8559 of *LNCS*, pages 525–532, 2014.

[7] K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. In *Proceedings of CONCUR*, pages 59–73, 2007.

[8] C. Cîrstea, C. Kupke, and D. Pattinson. EXPTIME tableaux for the coalgebraic μ-calculus. In *CSL 2009*, pages 179–193, 2009.

[9] G. Fontaine, R. A. Leal, and Y. Venema. Automata for coalgebras: An approach using predicate liftings. In *ICALP 2010*, pages 381–392, 2010.

[10] V. Goranko, A. Kuusisto, and R. Rönnholm. Game-theoretic semantics for alternating-time temporal logic. In *Proc. of AAMAS 2016*, pages 671–679, 2016.

[11] V. Goranko, A. Kuusisto, and R. Rönnholm. Game-theoretic semantics for ATL$^+$ with applications to model checking. Technical report, 2017. arXiv:1702.08405 [math.LO].

[12] J. Hintikka and G. Sandu. Game-theoretical semantics. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 361–410. Elsevier, 1997.

[13] A. Lomuscio, H. Qu, and F. Raimondi. MCMAS : A model checker for the verification multi-agent systems. In *Proc. of CAV 2009*, volume 5643 of *LNCS*, pages 682–688, 2009.

[14] R. Mazala. Infinite games. In E. Grädel, W. Thomas, and T. Wilke, editors, *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume 2500 of *LNCS*, pages 23–42. Springer, 2001.

[15] F. Mogavero, A. Murano, G. Perelli, and M. Vardi. Reasoning about strategies: On the model-checking problem. *ACM Transactions on Computational Logic*, 15(4):1–42, 2014.

[16] Y. Venema. Automata and fixed point logic: A coalgebraic perspective. *Inf. Comput.*, 204(4):637–678, 2006.

[17] S. Vester. Alternating-time temporal logic with finite-memory strategies. In *Proc, of GandALF 2013*, volume 119 of *EPTCS*, pages 194–207, 2013.

[18] D. Walther, W. van der Hoek, and M. Wooldridge. Alternating-time temporal logic with explicit strategies. In D. Samet, editor, *Proceedings TARK XI*, pages 269–278. Presses Universitaires de Louvain, 2007.