Game-Theoretic Semantics for the Alternating-Time Temporal Logic ATL

Valentin Goranko
Stockholm University

Joint work with **Antti Kuusisto** and **Raine Rönnholm**
(first presented at AAMAS'2016)

Highlights of Logic, Games and Automata 2016
Brussels, September 8, 2016

10 sec trailer

Two main approaches for relating logic and games:

10 sec trailer

Two main approaches for relating logic and games:

1. **Analyzing games with logic**

10 sec trailer

Two main approaches for relating logic and games:

1. **Analyzing games with logic**
2. **Analyzing logic with games**

10 sec trailer

Two main approaches for relating logic and games:

1. **Analyzing games with logic**
2. **Analyzing logic with games**

We bring these two approaches together.

Introduction: the story in a nutshell

Introduction: the story in a nutshell

Alternating-time temporal logic (ATL):
a logical formalism for reasoning about strategic abilities of players (agents) and coalitions in multi-player games.

Introduction: the story in a nutshell

Alternating-time temporal logic (ATL):
a logical formalism for reasoning about strategic abilities of players (agents) and coalitions in multi-player games.

Given a temporal objective φ, the ATL formula $\langle\langle A \rangle\rangle \varphi$ says that

"**the agents in the coalition A have a joint strategy to ensure that φ holds, regardless of the behaviour of all other agents.**"

Introduction: the story in a nutshell

Alternating-time temporal logic (ATL):
a logical formalism for reasoning about strategic abilities of players (agents) and coalitions in multi-player games.

Given a temporal objective φ, the ATL formula $\langle\langle A \rangle\rangle \varphi$ says that

"**the agents in the coalition A have a joint strategy to ensure that φ holds, regardless of the behaviour of all other agents.**"

In *game-theoretic semantics* (GTS), truth of a formula φ is determined in a formal dispute, called *evaluation game*, between two players: Eloise, who is trying to *verify* φ, and Abelard, who is trying to *falsify* it.

Introduction: the story in a nutshell

Alternating-time temporal logic (ATL):
a logical formalism for reasoning about strategic abilities of players (agents) and coalitions in multi-player games.

Given a temporal objective φ, the ATL formula $\langle\langle A \rangle\rangle \varphi$ says that

"**the agents in the coalition A have a joint strategy to ensure that φ holds, regardless of the behaviour of all other agents.**"

In *game-theoretic semantics* (GTS), truth of a formula φ is determined in a formal dispute, called *evaluation game*, between two players: Eloise, who is trying to *verify* φ, and Abelard, who is trying to *falsify* it.

GTS defines **truth of φ as existence of a winning strategy for Eloise** in the evaluation game for φ.

Introduction: the story in a nutshell

Alternating-time temporal logic (ATL):
a logical formalism for reasoning about strategic abilities of players (agents) and coalitions in multi-player games.

Given a temporal objective φ, the ATL formula $\langle\langle A \rangle\rangle \varphi$ says that

"**the agents in the coalition A have a joint strategy to ensure that φ holds, regardless of the behaviour of all other agents.**"

In *game-theoretic semantics* (GTS), truth of a formula φ is determined in a formal dispute, called *evaluation game*, between two players: Eloise, who is trying to *verify* φ, and Abelard, who is trying to *falsify* it.

GTS defines **truth of φ as existence of a winning strategy for Eloise** in the evaluation game for φ.

Here we develop game-theoretic semantics for ATL.

Introduction: the story in a nutshell

Alternating-time temporal logic (ATL):
a logical formalism for reasoning about strategic abilities of players (agents) and coalitions in multi-player games.

Given a temporal objective φ, the ATL formula $\langle\langle A \rangle\rangle \varphi$ says that

"**the agents in the coalition A have a joint strategy to ensure that φ holds, regardless of the behaviour of all other agents.**"

In *game-theoretic semantics* (GTS), truth of a formula φ is determined in a formal dispute, called *evaluation game*, between two players: Eloise, who is trying to *verify* φ, and Abelard, who is trying to *falsify* it.

GTS defines **truth of φ as existence of a winning strategy for Eloise** in the evaluation game for φ.

Here we develop game-theoretic semantics for ATL. Two-fold motivation:

– provide a two-way interaction between logic and games.

– eventually, make the ATL formula evaluation a finitary process.

Preliminaries: models and language of ATL

Preliminaries: models and language of ATL

A *concurrent game model* (CGM) for a set of *agents* $\mathbb{A}\mathrm{gt}$ consists of a *state space* St together with 3 mappings v, act, and o, defined on St:

Preliminaries: models and language of ATL

A *concurrent game model* (CGM) for a set of *agents* $\mathbb{A}\mathrm{gt}$ consists of a *state space* St together with 3 mappings v, act, and o, defined on St:

- $v(\mathbf{s})$: the set of atomic propositions that are true at the state $\mathbf{s}$,

Preliminaries: models and language of ATL

A *concurrent game model* (CGM) for a set of *agents* $\mathbb{A}\mathrm{gt}$ consists of a *state space* St together with 3 mappings v, act, and o, defined on St:

- $v(\mathbf{s})$: the set of atomic propositions that are true at the state $\mathbf{s}$,
- for each $a \in \mathbb{A}\mathrm{gt}$, $act(\mathbf{s}, a)$ is the set of *actions* available to a at $\mathbf{s}$,

Preliminaries: models and language of ATL

A *concurrent game model* (CGM) for a set of *agents* $\mathbb{A}\mathrm{gt}$ consists of a *state space* St together with 3 mappings v, act, and o, defined on St:

- $v(\mathbf{s})$: the set of atomic propositions that are true at the state $\mathbf{s}$,
- for each $a \in \mathbb{A}\mathrm{gt}$, $act(\mathbf{s}, a)$ is the set of *actions* available to a at $\mathbf{s}$,
- for every admissible tuple α of actions chosen by the agents in $\mathbb{A}\mathrm{gt}$, $o(\mathbf{s}, \alpha)$ is the unique *outcome / successor state* of $\mathbf{s}$.

Preliminaries: models and language of ATL

A *concurrent game model* (CGM) for a set of *agents* $\mathbb{A}\mathrm{gt}$ consists of a *state space* St together with 3 mappings *v*, *act*, and *o*, defined on St:

- $v(\mathbf{s})$: the set of atomic propositions that are true at the state $\mathbf{s}$,
- for each $a \in \mathbb{A}\mathrm{gt}$, $act(\mathbf{s}, a)$ is the set of *actions* available to a at $\mathbf{s}$,
- for every admissible tuple α of actions chosen by the agents in $\mathbb{A}\mathrm{gt}$, $o(\mathbf{s}, \alpha)$ is the unique *outcome* / *successor state* of $\mathbf{s}$.

Play in the CGM: an infinite sequence of consecutive successor states.

Preliminaries: models and language of ATL

A *concurrent game model* (CGM) for a set of *agents* $\mathbb{A}\mathrm{gt}$ consists of a *state space* St together with 3 mappings v, act, and o, defined on St:

- $v(\mathbf{s})$: the set of atomic propositions that are true at the state $\mathbf{s}$,
- for each $a \in \mathbb{A}\mathrm{gt}$, $act(\mathbf{s}, a)$ is the set of *actions* available to a at $\mathbf{s}$,
- for every admissible tuple α of actions chosen by the agents in $\mathbb{A}\mathrm{gt}$, $o(\mathbf{s}, \alpha)$ is the unique *outcome* / *successor state* of $\mathbf{s}$.

Play in the CGM: an infinite sequence of consecutive successor states.

Formulae of ATL:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \vee \varphi) \mid \langle\langle A \rangle\rangle \mathrm{X}\, \varphi \mid \langle\langle A \rangle\rangle \mathrm{F}\, \varphi \mid \langle\langle A \rangle\rangle \mathrm{G}\, \varphi$$

(Until and Release omitted for simplicity.)

Preliminaries: Compositional semantics of ATL

Preliminaries: Compositional semantics of ATL

Defined via the following clauses:

$\mathcal{M}, \mathbf{s} \models \langle\langle A \rangle\rangle \mathsf{X}\, \varphi$ / $\mathcal{M}, \mathbf{s} \models \langle\langle A \rangle\rangle \mathsf{F}\, \varphi$ / $\mathcal{M}, \mathbf{s} \models \langle\langle A \rangle\rangle \mathsf{G}\, \varphi$

iff the coalition A has a collective strategy which ensures that every possible outcome play starting at $\mathbf{s}$ satisfies φ:

at the next (successor) state / eventually / always.

Example: extended 'Rock-paper-scissors'

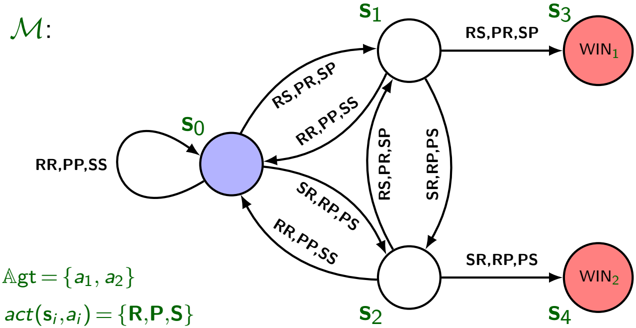

$\mathbb{A}\mathrm{gt} = \{a_1, a_2\}$

$act(\mathbf{s}_i, a_i) = \{\mathbf{R}, \mathbf{P}, \mathbf{S}\}$

The unbounded evaluation game for ATL

Consider a CGM $\mathcal{M}$, a state $\mathbf{s}_0$ in $\mathcal{M}$ and an ATL-formula φ.

Brief description of the *unbounded evaluation game* $G = \mathcal{G}(\mathcal{M}, \mathbf{s}_0, \varphi)$:

The unbounded evaluation game for ATL

Consider a CGM $\mathcal{M}$, a state $\mathbf{s}_0$ in $\mathcal{M}$ and an ATL-formula φ.

Brief description of the *unbounded evaluation game* $G = \mathcal{G}(\mathcal{M}, \mathbf{s}_0, \varphi)$:

A *position* in G is a tuple $(\mathbf{P}, \mathbf{s}, \psi)$,
where $\mathbf{P} \in \{\text{Eloise}, \text{Abelard}\}$, $\mathbf{s}$ is a state in $\mathcal{M}$ and ψ a subformula of φ.

The unbounded evaluation game for ATL

Consider a CGM $\mathcal{M}$, a state $\mathbf{s}_0$ in $\mathcal{M}$ and an ATL-formula φ.

Brief description of the *unbounded evaluation game* $G = \mathcal{G}(\mathcal{M}, \mathbf{s}_0, \varphi)$:

A *position* in G is a tuple $(\mathbf{P}, \mathbf{s}, \psi)$,
where $\mathbf{P} \in \{\text{Eloise}, \text{Abelard}\}$, $\mathbf{s}$ is a state in $\mathcal{M}$ and ψ a subformula of φ.

The game G begins from the *initial position* $(\text{Eloise}, \mathbf{s}_0, \varphi)$
and proceeds according to specific rules for each logical connective.

Rules for the unbounded evaluation game, informally

Rules for the unbounded evaluation game, informally

– For the propositional connectives, the rules are as in classical logic.

Rules for the unbounded evaluation game, informally

– For the propositional connectives, the rules are as in classical logic.

– For $\langle\langle A \rangle\rangle \mathsf{X}$ the rule invokes playing a one-step subgame $\mathbf{g}_{\mathsf{X}}$, where **P** chooses actions for the agents in A, then the opponent $\overline{\mathbf{P}}$ for the rest.

Rules for the unbounded evaluation game, informally

– For the propositional connectives, the rules are as in classical logic.

– For $\langle\langle A \rangle\rangle \mathsf{X}$ the rule invokes playing a one-step subgame $\mathbf{g}_{\mathsf{X}}$, where $\mathbf{P}$ chooses actions for the agents in A, then the opponent $\overline{\mathbf{P}}$ for the rest.

– For $\langle\langle A \rangle\rangle \mathsf{F}$ the rule invokes playing an embedded subgame $\mathbf{g}_{\mathsf{F}}$, in which $\mathbf{g}_{\mathsf{X}}$ is played repeatedly, until the controlling player $\mathbf{P}$ stops it.

Rules for the unbounded evaluation game, informally

– For the propositional connectives, the rules are as in classical logic.

– For $\langle\langle A \rangle\rangle \mathsf{X}$ the rule invokes playing a one-step subgame $\mathbf{g}_{\mathsf{X}}$, where **P** chooses actions for the agents in A, then the opponent $\overline{\mathbf{P}}$ for the rest.

– For $\langle\langle A \rangle\rangle \mathsf{F}$ the rule invokes playing an embedded subgame $\mathbf{g}_{\mathsf{F}}$, in which $\mathbf{g}_{\mathsf{X}}$ is played repeatedly, until the controlling player **P** stops it. If the embedded game goes on forever, **P** loses.

Rules for the unbounded evaluation game, informally

– For the propositional connectives, the rules are as in classical logic.

– For $\langle\langle A \rangle\rangle \mathsf{X}$ the rule invokes playing a one-step subgame $\mathbf{g}_{\mathsf{X}}$, where $\mathbf{P}$ chooses actions for the agents in A, then the opponent $\overline{\mathbf{P}}$ for the rest.

– For $\langle\langle A \rangle\rangle \mathsf{F}$ the rule invokes playing an embedded subgame $\mathbf{g}_{\mathsf{F}}$, in which $\mathbf{g}_{\mathsf{X}}$ is played repeatedly, until the controlling player $\mathbf{P}$ stops it. If the embedded game goes on forever, $\mathbf{P}$ loses.

– For $\langle\langle A \rangle\rangle \mathsf{G}$: likewise, but in the embedded subgame $\mathbf{g}_{\mathsf{G}}$ the controlling player is the opponent $\overline{\mathbf{P}}$.

Addendum: rules of the unbounded evaluation game for ATL

- In a position $(\mathbf{P}, \mathbf{s}, p)$ for an atomic proposition p, the game ends. $\mathbf{P}$ wins the game if $p \in v(\mathbf{s})$. Otherwise $\overline{\mathbf{P}}$ wins the game.
- In a position $(\mathbf{P}, \mathbf{s}, \psi \vee \theta)$, the player $\mathbf{P}$ chooses a disjunct $\alpha \in \{\psi, \theta\}$, and then the game continues from $(\mathbf{P}, \mathbf{s}, \alpha)$.
- In a position $(\mathbf{P}, \mathbf{s}, \neg\psi)$, the game moves to the position $(\overline{\mathbf{P}}, \mathbf{s}, \psi)$.
- In a position $(\mathbf{P}, \mathbf{s}, \langle\langle A \rangle\rangle \mathsf{X}\, \psi)$, the following *one-step game* $\mathbf{g}_{\mathsf{X}}$ is played with starting position $(\mathbf{P}, A, \mathbf{s})$, as follows:
  1. $\mathbf{P}$ chooses actions for the agents in A.
  2. $\overline{\mathbf{P}}$ chooses actions for the agents in $\mathbb{A}\mathrm{gt} \setminus A$.
     Thus, an action profile $\vec{\alpha}$ is selected.

  The one-step game $\mathbf{g}_{\mathsf{X}}$ ends in position $(\mathbf{P}, \mathbf{s}', \psi)$, where $\mathbf{s}' = o(\mathbf{s}, \vec{\alpha})$ is the resulting state.

Addendum: rules of the unbounded evaluation game (cont.)

- In a position $(\mathbf{P}, \mathbf{s}, \langle\langle A \rangle\rangle \mathsf{F}\, \psi)$, the game enters an *embedded subgame* $\mathbf{g}_{\mathsf{F}}$. It begins from the state $\mathbf{s}$ and proceeds by playing repeatedly the one-step game $\mathbf{g}_{\mathsf{X}}$, starting from $(\mathbf{P}, A, \mathbf{s})$.

  The player $\mathbf{P}$ is the *controlling player* in $\mathbf{g}_{\mathsf{F}}$, who may decide to end $\mathbf{g}_{\mathsf{F}}$ at any state $\mathbf{s}'$ that is reached.
  When (if) $\mathbf{g}_{\mathsf{F}}$ ends, in a position $(\mathbf{P}, \mathbf{s}', \psi)$, the evaluation game then continues from that position.

  If $\mathbf{g}_{\mathsf{F}}$ goes on forever, the controlling player $\mathbf{P}$ loses the entire evaluation game.

- In a position $(\mathbf{P}, \mathbf{s}, \langle\langle A \rangle\rangle \mathsf{G}\, \psi)$, the players enter a dual *embedded subgame* $\mathbf{g}_{\mathsf{G}}$, just like $\mathbf{g}_{\mathsf{F}}$, but now the controlling player is $\overline{\mathbf{P}}$.

Unbounded game-theoretic semantics for ATL

Unbounded game-theoretic semantics for ATL

Unbounded game-theoretic semantics for ATL :

$\mathcal{M}, \mathbf{s} \models_{GTS} \varphi$ *iff Eloise has a winning strategy in* $\mathcal{G}(\mathcal{M}, \mathbf{s}, \varphi)$.

Unbounded game-theoretic semantics for ATL

Unbounded game-theoretic semantics for ATL :

$\mathcal{M}, \mathbf{s} \models_{GTS} \varphi$ *iff Eloise has a winning strategy in* $\mathcal{G}(\mathcal{M}, \mathbf{s}, \varphi)$.

Theorem (1)

The unbounded GTS *for* ATL *is equivalent to the standard (compositional) semantics of* ATL.

Unbounded game-theoretic semantics for ATL

Unbounded game-theoretic semantics for ATL :

$\mathcal{M}, \mathbf{s} \models_{GTS} \varphi$ *iff Eloise has a winning strategy in* $\mathcal{G}(\mathcal{M}, \mathbf{s}, \varphi)$.

Theorem (1)

The unbounded GTS *for* ATL *is equivalent to the standard (compositional) semantics of* ATL.

The unbounded evaluation games are determined, but possibly infinite.

Unbounded game-theoretic semantics for ATL

Unbounded game-theoretic semantics for ATL :

$\mathcal{M}, \mathbf{s} \models_{GTS} \varphi$ *iff Eloise has a winning strategy in* $\mathcal{G}(\mathcal{M}, \mathbf{s}, \varphi)$.

Theorem (1)

The unbounded GTS *for* ATL *is equivalent to the standard (compositional) semantics of* ATL.

The unbounded evaluation games are determined, but possibly infinite. Can we make them finite?

Finitely bounded evaluation games for ATL

We modify the evaluation game by adding for every embedded subgame ($\mathbf{g}_{\mathsf{F}}$ or $\mathbf{g}_{\mathsf{G}}$) a (finite) *time limit* $n \in \mathbb{N}$.

Finitely bounded evaluation games for ATL

We modify the evaluation game by adding for every embedded subgame ($\mathbf{g}_{\mathsf{F}}$ or $\mathbf{g}_{\mathsf{G}}$) a (finite) *time limit* $n \in \mathbb{N}$.

The time limit is selected and announced by the controlling player at the beginning of the embedded subgame.

Finitely bounded evaluation games for ATL

We modify the evaluation game by adding for every embedded subgame ($\mathbf{g}_\mathsf{F}$ or $\mathbf{g}_\mathsf{G}$) a (finite) *time limit* $n \in \mathbb{N}$.

The time limit is selected and announced by the controlling player at the beginning of the embedded subgame.

After every round of playing $\mathbf{g}_\mathsf{X}$ the time limit decreases by 1.

Finitely bounded evaluation games for ATL

We modify the evaluation game by adding for every embedded subgame ($\mathbf{g}_{\mathsf{F}}$ or $\mathbf{g}_{\mathsf{G}}$) a (finite) *time limit* $n \in \mathbb{N}$.

The time limit is selected and announced by the controlling player at the beginning of the embedded subgame.

After every round of playing $\mathbf{g}_{\mathsf{X}}$ the time limit decreases by 1.

The controlling player must end the subgame before the time runs out.

Finitely bounded evaluation games for ATL

We modify the evaluation game by adding for every embedded subgame ($\mathbf{g}_{\mathsf{F}}$ or $\mathbf{g}_{\mathsf{G}}$) a (finite) *time limit* $n \in \mathbb{N}$.

The time limit is selected and announced by the controlling player at the beginning of the embedded subgame.

After every round of playing $\mathbf{g}_{\mathsf{X}}$ the time limit decreases by 1.

The controlling player must end the subgame before the time runs out.

Evaluation games with time limits provide *finitely bounded* GTS for ATL.

Finitely bounded evaluation games for ATL

We modify the evaluation game by adding for every embedded subgame ($\mathbf{g}_\mathsf{F}$ or $\mathbf{g}_\mathsf{G}$) a (finite) *time limit* $n \in \mathbb{N}$.

The time limit is selected and announced by the controlling player at the beginning of the embedded subgame.

After every round of playing $\mathbf{g}_\mathsf{X}$ the time limit decreases by 1.

The controlling player must end the subgame before the time runs out.

Evaluation games with time limits provide *finitely bounded* GTS for ATL.

Theorem (2)

The finitely bounded and unbounded GTS *are generally non-equivalent.*

Finitely bounded evaluation games for ATL

We modify the evaluation game by adding for every embedded subgame ($\mathbf{g}_{\mathsf{F}}$ or $\mathbf{g}_{\mathsf{G}}$) a (finite) *time limit* $n \in \mathbb{N}$.

The time limit is selected and announced by the controlling player at the beginning of the embedded subgame.

After every round of playing $\mathbf{g}_{\mathsf{X}}$ the time limit decreases by 1.

The controlling player must end the subgame before the time runs out.

Evaluation games with time limits provide *finitely bounded* GTS for ATL.

Theorem (2)

The finitely bounded and unbounded GTS *are generally non-equivalent.*

However, they are equivalent on finitely branching models.

Example

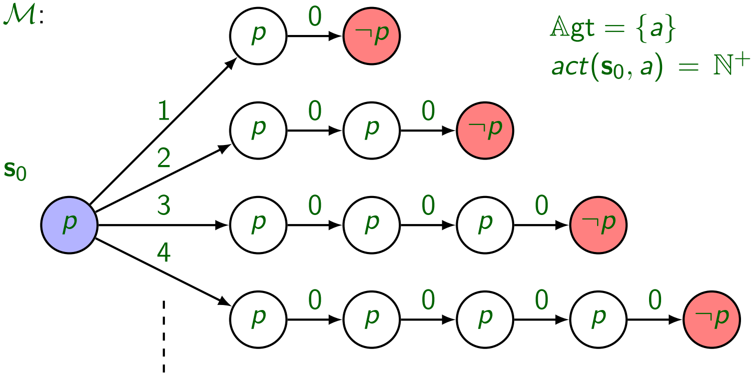

$\mathbb{A}\text{gt} = \{a\}$

$act(\mathbf{s}_0, a) = \mathbb{N}^+$

Example

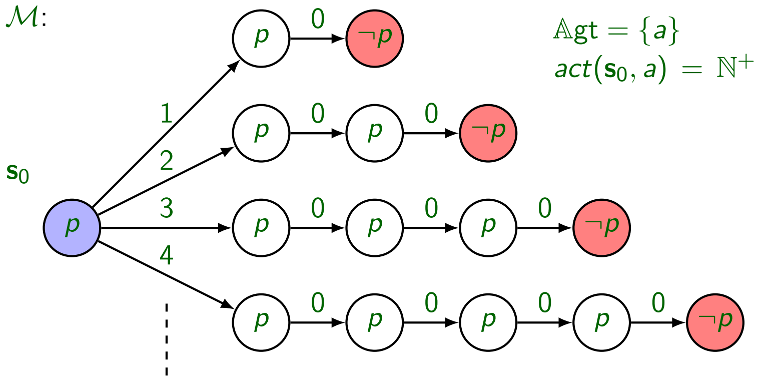

Here Eloise can win $\mathcal{G}(\mathcal{M}, \mathbf{s}_0, \langle\langle a \rangle\rangle \, \mathrm{G}\, p)$ with finite time limits.

Example

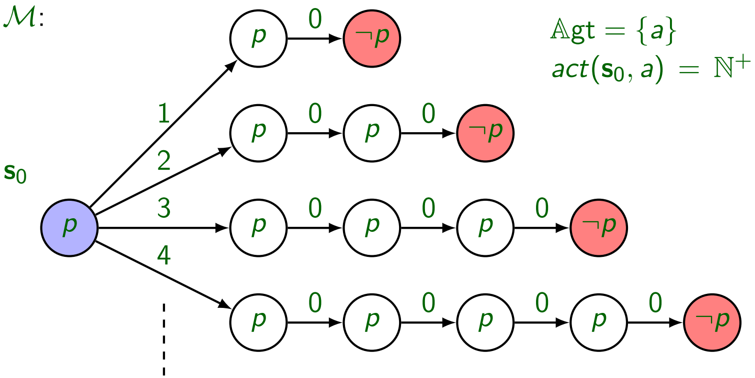

Here Eloise can win $\mathcal{G}(\mathcal{M}, \mathbf{s}_0, \langle\langle a \rangle\rangle \, \mathsf{G} \, p)$ with finite time limits.
But she not win in the unbounded GTS.

V Goranko

Example

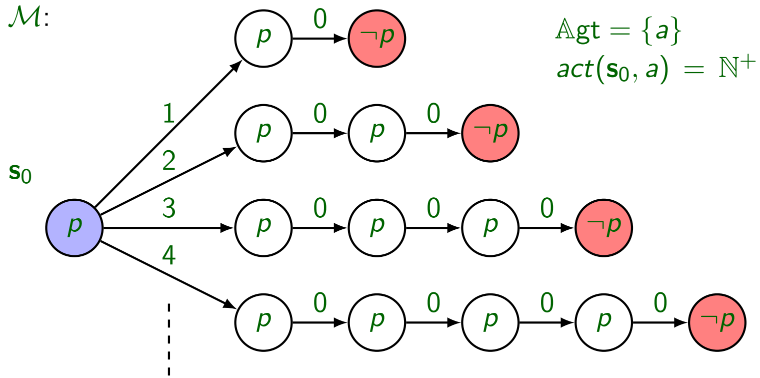

Here Eloise can win $\mathcal{G}(\mathcal{M}, \mathbf{s}_0, \langle\langle a\rangle\rangle \,\mathsf{G}\, p)$ with finite time limits.
But she not win in the unbounded GTS.

Therefore, $\mathcal{M}, \mathbf{s}_0 \models^{fin}_{GTS} \varphi$ but $\mathcal{M}, \mathbf{s}_0 \not\models_{GTS} \varphi$.

Ordinal bounded evaluation games for ATL

We now generalize, by allowing time limits for the embedded subgames to be any ordinals.

Ordinal bounded evaluation games for ATL

We now generalize, by allowing time limits for the embedded subgames to be any ordinals.

The initial time limit of the subgame is chosen by the controlling player.

Ordinal bounded evaluation games for ATL

We now generalize, by allowing time limits for the embedded subgames to be any ordinals.

The initial time limit of the subgame is chosen by the controlling player.

After every round of $\mathbf{g}_X$, the time limit decreases:
– if a successor ordinal, by 1;
– if a limit ordinal, the controlling player must choose a strictly smaller one.

Ordinal bounded evaluation games for ATL

We now generalize, by allowing time limits for the embedded subgames to be any ordinals.

The initial time limit of the subgame is chosen by the controlling player.

After every round of $\mathbf{g}_X$, the time limit decreases:
– if a successor ordinal, by 1;
– if a limit ordinal, the controlling player must choose a strictly smaller one.

Ordinal bounded GTS: uses evaluation games with ordinal time limits.

Ordinal bounded evaluation games for ATL

We now generalize, by allowing time limits for the embedded subgames to be any ordinals.

The initial time limit of the subgame is chosen by the controlling player.

After every round of $\mathbf{g}_X$, the time limit decreases:
– if a successor ordinal, by 1;
– if a limit ordinal, the controlling player must choose a strictly smaller one.

Ordinal bounded GTS: uses evaluation games with ordinal time limits.

NB. Since ordinals are well-founded, the ordinal bounded semantics guarantees that evaluation games end in finite number of rounds.

Ordinal bounded evaluation games for ATL

We now generalize, by allowing time limits for the embedded subgames to be any ordinals.

The initial time limit of the subgame is chosen by the controlling player.

After every round of $\mathbf{g}_{\mathsf{X}}$, the time limit decreases:
– if a successor ordinal, by 1;
– if a limit ordinal, the controlling player must choose a strictly smaller one.

Ordinal bounded GTS: uses evaluation games with ordinal time limits.

NB. Since ordinals are well-founded, the ordinal bounded semantics guarantees that evaluation games end in finite number of rounds.

Theorem (3)

The ordinal bounded and the unbounded GTS *are equivalent.*

Example

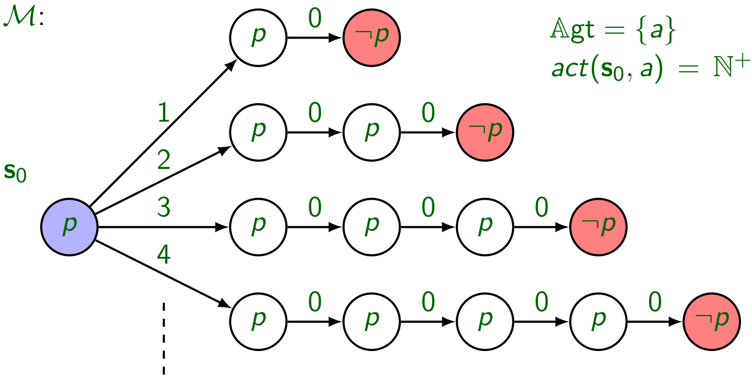

$\mathbb{A}\mathrm{gt} = \{a\}$

$act(\mathbf{s}_0, a) = \mathbb{N}^+$

Example

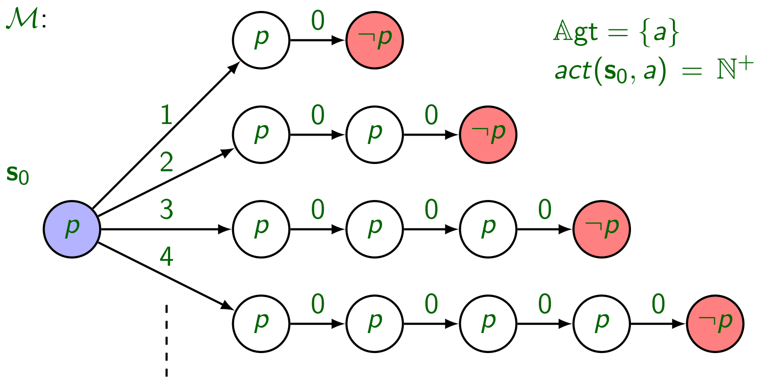

Here Abelard can win $\mathcal{G}(\mathcal{M}, \mathbf{s}_0, \langle\langle a \rangle\rangle \, \mathrm{G} \, p)$ with ordinal bounded GTS by choosing the ordinal ω as the initial time limit.

Example

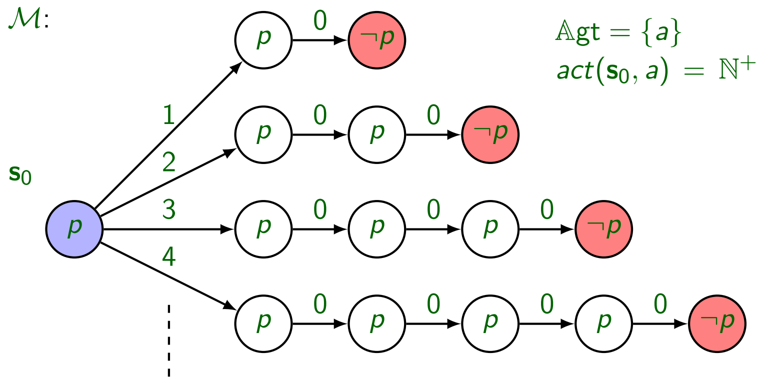

Here Abelard can win $\mathcal{G}(\mathcal{M}, \mathbf{s}_0, \langle\langle a \rangle\rangle \, \mathsf{G} \, p)$ with ordinal bounded GTS by choosing the ordinal ω as the initial time limit.

After the 1st round, where Eloise chooses an action,
Abelard can always reduce the time limit appropriately to win.

Concluding remarks

Concluding remarks

- The game-theoretic semantics for ATL is a natural framework for reasoning about games by using games.

Concluding remarks

- The game-theoretic semantics for ATL is a natural framework for reasoning about games by using games.
- Limitations on time resources in evaluation games lead to interesting semantic variants.

Concluding remarks

- The game-theoretic semantics for ATL is a natural framework for reasoning about games by using games.
- Limitations on time resources in evaluation games lead to interesting semantic variants.

Ongoing work:

Concluding remarks

- The game-theoretic semantics for ATL is a natural framework for reasoning about games by using games.
- Limitations on time resources in evaluation games lead to interesting semantic variants.

Ongoing work:

- develop game-theoretic semantics for extensions of ATL, incl. ATL* and the alternating time mu-calculus.

Concluding remarks

- The game-theoretic semantics for ATL is a natural framework for reasoning about games by using games.
- Limitations on time resources in evaluation games lead to interesting semantic variants.

Ongoing work:

- develop game-theoretic semantics for extensions of ATL, incl. ATL* and the alternating time mu-calculus.
- consider evaluation games with other limited resources, e.g. memory.

Concluding remarks

- The game-theoretic semantics for ATL is a natural framework for reasoning about games by using games.
- Limitations on time resources in evaluation games lead to interesting semantic variants.

Ongoing work:

- develop game-theoretic semantics for extensions of ATL, incl. ATL* and the alternating time mu-calculus.
- consider evaluation games with other limited resources, e.g. memory.

Thank you for your attention!