How to Manage a Budget with ATL^{+}

Stéphane Demri
CNRS, LMF

Raine Rönnholm
University of Tampere

KR'23, Rhodes, Sept. 2023

ATL-Like Logics with Resources

- ATL: to reason about the strategic abilities of agents in multi-agent systems. [Alur & Henzinger & Kupferman, JACM 2002]

  (Coalition *A* has a strategy to enforce pro Φ)

Concurrent game structure (CGS):

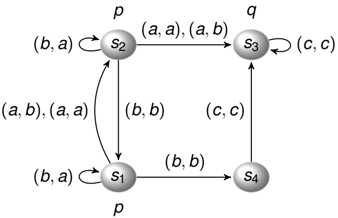

- In resource-aware logics, actions have costs/rewards, e.g.
  - RBTL* [Bulling & Farwer, CLIMA X '09]
  - RB±ATL [Alechina et al., ECAI'14]
  - QATL* [Bulling & Goranko, AAMAS 2022]
  - pe-ATL [Della Monica & Murano, AAMAS 2018]

- Undecidability often happens with resource-aware logics.

Game-Theoretic Semantics for ATL-Like Logics

- Alternative semantics to compositional semantics.
  - for ATL, [Goranko & Kuusisto & Rönnholm, ACM ToCL 2018]
  - for ATL$^+$. [Goranko & Kuusisto & Rönnholm, IC 2021]

- Players are Eloise (**E**) and Abelard (**A**) and **E** tries to verify that a state satisfies a formula.

- Evaluation games come with a transition (sub)game to handle strategy quantifiers.

- Alternative semantics, but also new complexity results:
  - PSPACE upper bound for ATL$^+$ model-checking,
  - PTIME fragments of ATL$^+$.

[Goranko & Kuusisto & Rönnholm,IC 2021]

Our Motivations

- To design game-theoretical semantics for resource-aware ATL-like logics.

- We investigate ATL^+ with one resource (written $\mathsf{ATL}^+(1)$).

  – Often informative to start with the single-resource case and complexity characterisation of $\mathrm{MC}(\mathsf{ATL}^+(1))$ is open.

  – Solving $\mathrm{MC}(\mathsf{ATL}^+)$ is already challenging.

  [Laroussinie & Markey & Oreiby, LMCS 2008]
  [Bulling & Jamroga, AIC 2010]
  [Goranko & Kuusisto & Rönnholm, IC 2021]
  [Cerrito, arXiv 2021]

  – ATL^+ exponentially more succinct than ATL with release.

Concurrent Game Structures with One Resource

- Partial "weight" function $wf : S \times Agt \times Act \to \mathbb{Z}$.

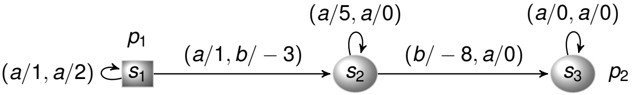

(two agent $wf(s_1, 2, b) = -3$)

- $wf(s, f)$: weight/cost to trigger a transition from s with the (total) joint action $f : Agt \to Act$:

$$wf(s, f) \stackrel{\text{def}}{=} \sum_{a \in Agt} wf(s, a, f(a))$$

$\left(wf(s_1, [1 \mapsto a, 2 \mapsto b]) = -2\right)$

Basic Concepts About Strategies

- Strategy $\sigma : (s_0 \xrightarrow{f_0} s_1 \xrightarrow{f_1} s_2 \cdots \xrightarrow{f_{n-1}} s_n) \mapsto (g : A \rightarrow Act)$.
  $comp(s, \sigma)$: set of inf. computations respecting σ from s.

- Given $b \in \mathbb{N}$ and $\lambda = s_0 \xrightarrow{f_0} s_1 \xrightarrow{f_1} s_2 \ldots$ in $comp(s, \sigma)$, resource availability values $(v_i)_{i \in \mathbb{N}}$:

$$v_0 \stackrel{\text{def}}{=} b; \qquad v_{i+1} \stackrel{\text{def}}{=} v_i + \sum_{a \in Agt} wf(s_i, a, f_i(a)).$$

(al $\quad (s_0, v_0) \xrightarrow{f_0} (s_1, v_1) \xrightarrow{f_1} (s_2, v_2) \ldots)$

- $\lambda \in comp(s, \sigma)$ is b-consistent $\stackrel{\text{def}}{\Leftrightarrow}$ for all i, $0 \leq v_i$.

- σ is b-consistent w.r.t. s $\stackrel{\text{def}}{\Leftrightarrow}$ all the computations from s respecting σ are b-consistent.

Introduction to ATL$^+$(1)

- Design follows principles from [Alechina et al., JCSS 2017].

- ATL$^+$(1) objectives: Boolean combinations of φ, Xφ, φ_1Uφ_2, Gφ, e.g. $\neg(\top \mathsf{U} \neg p_1) \vee \top \mathsf{U} p_2$. (as in CTL$^+$/ATL$^+$)

- State and path formulae: ($b \in \mathbb{N}$, $p \in \mathrm{PROP}$, $A \subseteq Agt$)

  $\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle\langle A^b \rangle\rangle \Phi \qquad \Phi ::= \varphi \mid \neg\Phi \mid \Phi \vee \Phi \mid \mathsf{X}\,\varphi \mid \varphi \mathsf{U} \varphi$

- $M, s \models \langle\langle A^b \rangle\rangle \Phi \overset{\text{def}}{\Leftrightarrow}$ for some b-consistent σ w.r.t. s for A, for all $\lambda \in comp(s, \sigma)$, we have $M, \lambda \models \Phi$.

  (quantitative and qualitative ob

- $M, s_1 \models \langle\langle \{1\}^3 \rangle\rangle (\mathsf{G} p_1 \vee \mathsf{F} p_2)$

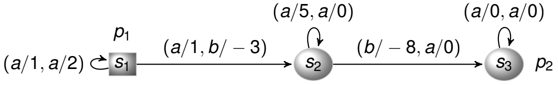

Model-Checking for ATL$^+$ and Variants

- MC(ATL$^+$(1)): model-checking problem for ATL$^+$(1).

  Input: a finite CGS M, a state s and φ in ATL$^+$(1),

  Question: $M, s \models \varphi$?

- MC(ATL$^+$) is PSPACE-hard:
  - proof for LTL games, [Alur & La Torre, LiCS'01]
  - direct proof for ATL$^+$. [Bulling & Jamroga, AIC 2010]

- MC(ATL$^+$) in PSPACE. [Goranko et al., IC 2021]

  See also [Bulling & Jamroga, AIC 2010; Wang et al., ATPLS 2015]

- MC(ATL(1)) is in PSPACE. [Alechina et al., JCSS 2017]

- For all $n \geq 1$, MC(ATL$^+$(n)) is in EXPTIME.

  [Colcombet et al., LiCS'17; Belardinelli & Demri, AI 2021]

Resource-Capped Game $\mathcal{G}_{\textsf{RC}}[\beta](M, s_{in}, \varphi_{in})$

- Resource cap $\beta \in \mathbb{N}$: upper bound on resource values.
- The game begins from $(\mathbf{E}, s_{in}, \varphi_{in})$ and is played according to the following rules.

1. In $(\mathbf{V}, s, p)$ the evaluation game ends and $\mathbf{V}$ wins iff $s \in L(p)$ (else the falsifier $\overline{\mathbf{V}}$ wins).
2. In $(\mathbf{V}, s, \varphi_1 \vee \varphi_2)$, $\mathbf{V}$ selects whether the next position is $(\mathbf{V}, s, \varphi_1)$ or $(\mathbf{V}, s, \varphi_2)$.
3. In $(\mathbf{V}, s, \neg\varphi)$, the game proceeds to the position $(\overline{\mathbf{V}}, s, \varphi)$ (swapping the roles).
4. In $(\mathbf{V}, s, \langle\langle A^b \rangle\rangle \Phi)$, the game enters the transition game $\mathcal{TG}_{\textsf{RC}}[\beta](\mathbf{V}, s, A, b, \Phi)$.

Transition Game $\mathcal{TG}_{RC}[\beta](\mathbf{V}, s, A, b, \Phi)$

- Alternation *verification phase* $\rightleftarrows$ *transition phase*.
  Initial game position (s, b, T_{open}).

- Truth function $T : atoms(\Phi) \to \{\top, \bot, open\}$ with Φ equal to a Boolean combination of "atoms" from $atoms(\Phi)$.

$$atoms((\langle\!\langle A^b \rangle\!\rangle \Phi' \vee p) \wedge \mathsf{X}\varphi_1 \wedge (\varphi_3 \mathsf{U} \langle\!\langle A^b \rangle\!\rangle \Phi'')) =$$

$$\{\langle\!\langle A^b \rangle\!\rangle \Phi' \vee p, \mathsf{X}\varphi_1, \varphi_3 \mathsf{U} \langle\!\langle A^b \rangle\!\rangle \Phi''\}$$

($\approx$ finite memory dedicated to the satisfaction of Φ)

- Definition of $T \models \Psi$ with $atoms(\Psi) \subseteq atoms(\Phi)$.
  - $T \models \Psi \overset{\text{def}}{\Leftrightarrow} T(\Psi) = \top$. $\quad (\Psi \in atoms(\Phi))$
  - $T \models \neg\Psi \overset{\text{def}}{\Leftrightarrow} T \not\models \Psi$; $T \models \Psi_1 \vee \Psi_2 \overset{\text{def}}{\Leftrightarrow} T \models \Psi_1$ or $T \models \Psi_2$.

- $[\top \mathsf{U} \neg p_1 \mapsto open, \top \mathsf{U} p_2 \mapsto \bot] \models \mathsf{G} p_1 \vee \mathsf{F} p_2$.
  (since $\mathsf{G} p_1 = \neg(\top \mathsf{U} \neg p_1)$ and $\not\models \top \mathsf{U} \neg p_1$)

Verification Phase (or how to update truth functions)

- Current game position (s', b', T').
- Any player can pick a temporal atom $\Psi \in \mathit{atoms}(\Phi)$ such that $T'(\Psi) = \mathit{open}$ and attempt to verify or falsify it.
- E.g., either player $\mathbf{P} \in \{\mathbf{E}, \mathbf{A}\}$ may attempt to verify $\varphi_1 \mathrm{U} \varphi_2 \in \mathit{atoms}(\Phi)$, by claiming that φ_2 is true at s'.
- If the opponent $\overline{\mathbf{P}}$ does not object this claim, then T' is updated by mapping $\varphi_1 \mathrm{U} \varphi_2$ to $\top$.
- If $\overline{\mathbf{P}}$ objects, then claim evaluated by exiting the transition game and resuming to evaluation game from $(\mathbf{P}, s', \varphi_2)$.
- Other cases are handled similarly.

Transition Phase (or how to update states and resource values)

- Current game position (s', b', T').
- **V** chooses a joint action $g : A \to Act$ (depending on s'). Then, $\overline{\mathbf{V}}$ chooses $\overline{g} : \overline{A} \to Act$ (depending on s').
- Update of the game position (<u>truncation</u> of values above β):

$$s' := \delta(s', g \oplus \overline{g}) \quad b' := \min(\beta, b' + wf(s', g \oplus \overline{g}))$$

(δ = transition function of M)

- If $b' < 0$, **V** loses the game. Otherwise, the play continues with the new (s', b', T') in the verification phase.
- For infinite plays, **V** wins iff the final T_f satisfies $T_f \models \Phi$.

Relationships with Compositional Semantics

Equivalence between (I) and (II):

(I) $M, s \models \varphi$,
(II) **E** has a winning strategy for $\mathcal{G}_{\mathrm{RC}}[F(M,\varphi)](M, s, \varphi)$ with

- $F(M,\varphi)$ polynomial in card(S), $\|M\|_{\infty}^{-}$ and $tw(\varphi)$,
- $\|M\|_{\infty}^{-} \stackrel{\text{def}}{=} \max\{|wf(s,f)| : s \in S, \text{ joint action } f \text{ at } s, wf(s,f) \leq 0\} \cup \{0\}$,
- $tw(\varphi)$ maximal value card$(atoms(\Phi))$ with Φ occurring in φ.

(similar equivalence for loo__________ evaluation game)

Hints About Small Strategy Skeletons

- Strategy skeletons are restrictions of strategy trees in which pruning occurs when no progress is possible.
- Small strategy skeletons crucial for the equivalence proofs.
- Strategy σ for the agent 1: action a on s_2 only if no state is already visited twice, otherwise b.

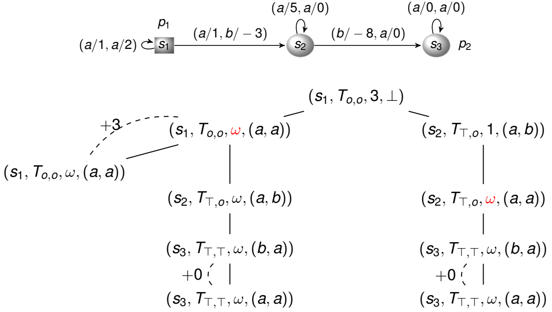

$T_{o,o}$:$\top \mathsf{U} \neg p_1, \top \mathsf{U} p_2 \mapsto$ *open* $T_{\top,o}$:$\top \mathsf{U} \neg p_1 \mapsto \top, \top \mathsf{U} p_2 \mapsto$ *open* $T_{\top,\top}$:$\top \mathsf{U} \neg p_1, \top \mathsf{U} p_2 \mapsto \top$

$T_{o,o}$ and $T_{\top,\top}$ satisfy $\mathsf{G} p_1 \vee \mathsf{F} p_2$

Complexity of $\mathrm{MC}(\mathsf{ATL}^+(1))$

- $\mathrm{MC}(\mathsf{ATL}^+(1))$ is PSPACE-complete.

  (small strategy skeletons computed in APTime)

- Given $K \in \mathbb{N}$ and a polynomial $P(\cdot)$, the subproblem of $\mathrm{MC}(\mathsf{ATL}^+(1))$ with

$$||M||_\infty^- \leq P(\mathsf{size}(M) + \mathsf{size}(\varphi)) \quad \& \quad tw(\varphi) \leq K$$

is in PTIME. (solving Büchi game

ATL$^+$(1) **and Beyond**

- Design of two evaluation games for ATL$^+$(1).

- Complexity analysis: MC(ATL$^+$(1)) is PSPACE-complete and in PTIME if we bound $tw(\varphi)$ and $||M||^-_\infty$.

- By-products for
  - budget synthesis,
  - ATL$^+$ and other variants,
  - energy game with depth-one objectives.

(more detail

- What about ATL$^+$(2)?