A Hennessy-Milner Theorem for ATL with Imperfect Information

Francesco Belardinelli, Cătălin Dima, Vadim Malvone, Ferucio Țiplea

Imperial College London, LACL – Université Paris-Est Créteil, IBISC – Université d'Evry, University of Iași

LICS 2020 & Highlights 2020

ATL with common knowledge semantics

$1 \models_{subj} \langle\langle alice, bob \rangle\rangle \mathbb{X}\ win$

missing arrows are losing

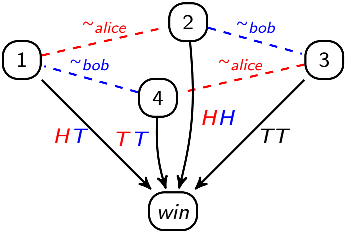

Common knowledge semantics for $s \models \langle\langle A \rangle\rangle \varphi$ requires:

- The existence of a joint strategy profile built over the whole common knowledge neigbourhood $C_A(s)$.
- Which, when applied in each state of $C_A(s)$, produces the objective φ.

ATL with common knowledge semantics

Common knowledge = invariant under $\sim_A^C = \left(\bigcup_{a \in A} \sim_a \right)^*$.

$1 \models_{subj} \langle\!\langle alice, bob \rangle\!\rangle \mathbb{X}\, win$

$1 \not\models_{ck} \langle\!\langle alice, bob \rangle\!\rangle \mathbb{X}\, win$

missing arrows are losing

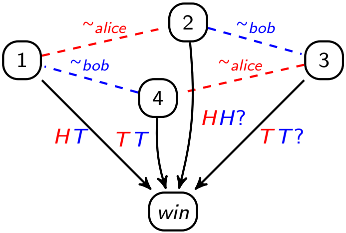

Common knowledge semantics for $s \models \langle\!\langle A \rangle\!\rangle \varphi$ requires:

- The existence of a joint strategy profile built over the whole common knowledge neigbourhood $C_A(s)$.
- Which, when applied in each state of $C_A(s)$, produces the objective φ.

Alternating bisimulation – the idea

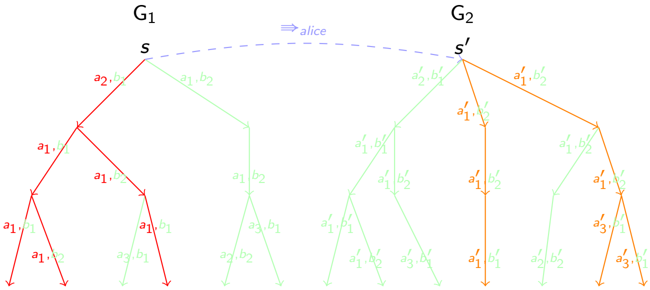

$$\forall \sigma_{alice} \exists \sigma'_{alice} \forall \rho' \in out_{G_2}(\sigma', s') \exists \rho \in out_{G_1}(\sigma, s) \text{ with } \rho'\big|_{AP} = \rho\big|_{AP}$$

Guarantees that, for any ϕ,

$$G_1, s \models \langle\langle alice \rangle\rangle \phi \quad \Rightarrow \quad G_2, s' \models \langle\langle alice \rangle\rangle \phi$$

Alternating (bi)simulation with imperfect information

Definition

$\Rrightarrow_A \subseteq Hist(\mathrm{G}) \times \in Hist(G')$ for which, whenever $h \Rrightarrow_A h'$, then

1. $\pi(h) = \pi'(h')$.
2. [Compatibility with indistinguishability] For each $a \in A$,

$$\forall k' \sim_a h' \; \exists k \sim_a h \text{ with } h' \Rrightarrow_A k'$$

3. [Compatibility with uniform strategies] – $\forall \sigma \in PStr(C_A(h)) \; \exists \sigma \in PStr(C_A(h'))$ with... (see next slide).

History-based version of [Belardinelli, Condurache, D., Jamroga, Jones, Knapik, 2017, 2020]

Alternating bisimulation with imperfect information ($A = \{alice, bob\}$)

$\exists ST : PStr_A(C_A(s_1)) \to PStr_A(C'_A(s'_1))$

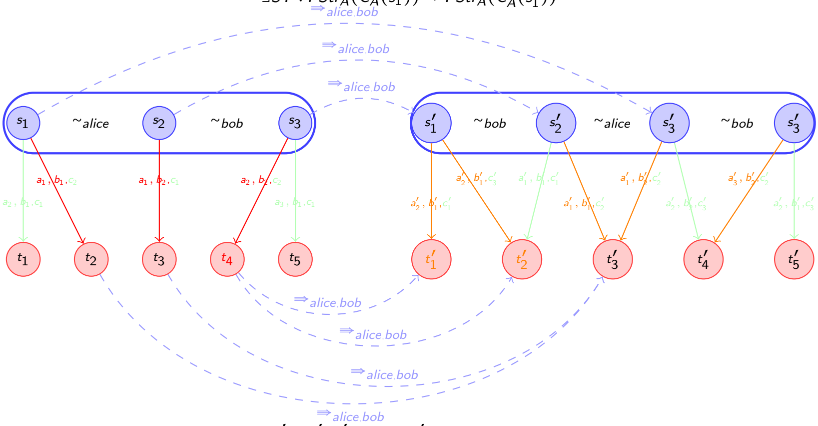

$\forall r \in C_A(s_1), \forall r' \in C'_A(s'_1), r \Rrightarrow_A r'$ implies $\forall \sigma_A \in PStr_A(C_A(s_1))$,

$\forall r' \xrightarrow{ST(\sigma_A)(r')} s', \exists s \in C_A(s_1), r \xrightarrow{\sigma_A(r)} s$ and $s \Rrightarrow_A s'$

The Hennessy-Milner Theorem

Theorem

Assume $\Longleftrightarrow_A$ *is a history-based A-bisimulation between two game structures* G *and* G'*. Let* $h \in Hist(\mathsf{G})$ *and* $h' \in Hist(\mathsf{G}')$ *with* $h \Longleftrightarrow_A h'$*. Then, for every A-formula* ϕ *and* $x \in \{subj, obj, ck\}$*,*

$$(\mathsf{G}, h) \vDash_x \phi \text{ iff } (\mathsf{G}', h') \vDash_x \phi$$

Theorem

$\mathsf{G}(h_0)$ *and* $\mathsf{G}'(h'_0)$ *are A-equivalent for the common knowledge semantics* $\vDash_{ck}$ *if and only if they are A-bisimilar.*

2nd theorem fails for $\vDash_{subj}$.

The Hennessy-Milner Theorem

Theorem

Assume $\Longleftrightarrow_A$ is a history-based A-bisimulation between two game structures G *and* G′*. Let $h \in Hist(\mathsf{G})$ and $h' \in Hist(\mathsf{G}')$ with $h \Longleftrightarrow_A h'$. Then, for every A-formula ϕ and $x \in \{subj, obj, ck\}$,*

$$(\mathsf{G}, h) \models_x \phi \text{ iff } (\mathsf{G}', h') \models_x \phi$$

Theorem

$\mathsf{G}(h_0)$ *and* $\mathsf{G}'(h'_0)$ *are A-equivalent for the common knowledge semantics* $\models_{ck}$ *if and only if they are A-bisimilar.*

2nd theorem fails for $\models_{subj}$.

The bisimulation game

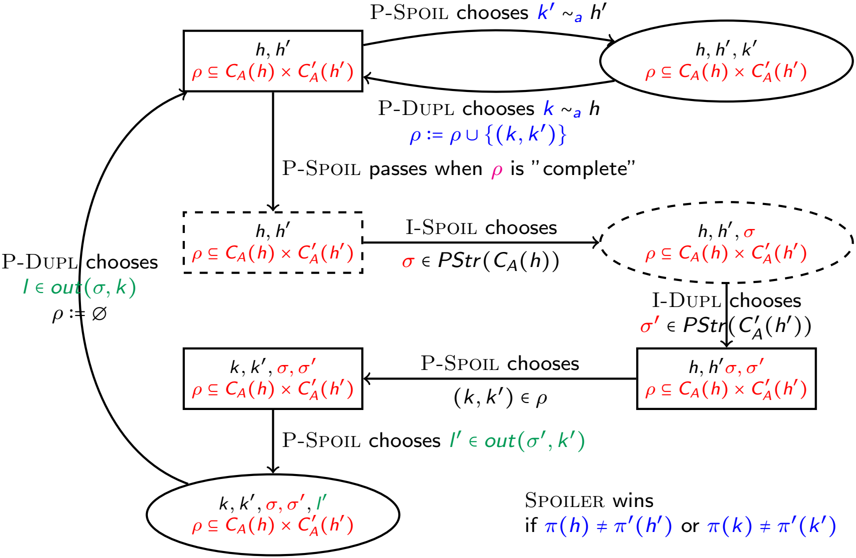

SPOILER wins
if $\pi(h) \neq \pi'(h')$ or $\pi(k) \neq \pi'(k')$

The bisimulation game

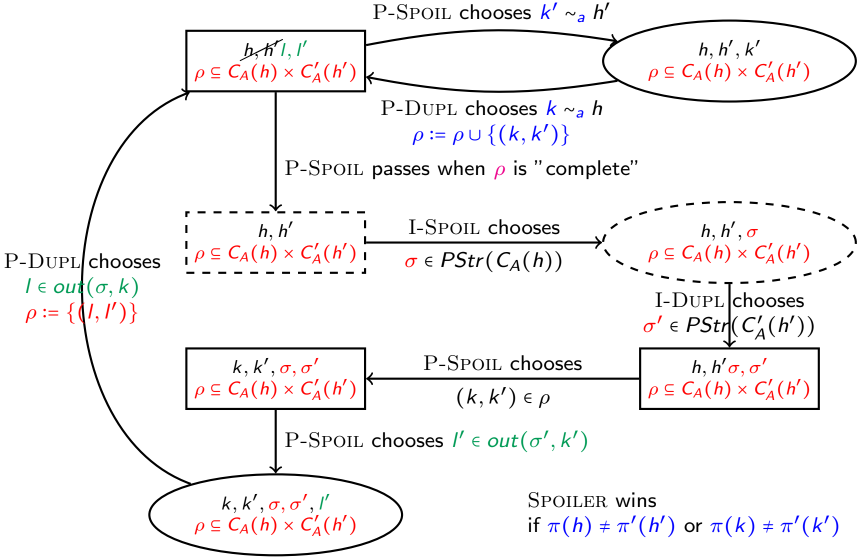

Spoiler wins
if $\pi(h) \neq \pi'(h')$ or $\pi(k) \neq \pi'(k')$

Determinacy for bisimulation games

Gale-Stewart games between 4 players:

- Reachability objective for SPOILERS, safety objective for DUPLICATORS.
- Defensive strategy: does not put the game into a winning state for the opponent coalition.
- Defensive strategies against reachability objectives can be transformed into winning strategies (for the safety objective).
- The game is determined, since when SPOILERS do not win, a defensive strategy for DUPLICATORS exists.
- From a winning strategy for SPOILERS one may build an ATL formula containing the Yesterday modality which is distinguishes the two CGS.

Concluding remarks

- Bisimulations can be adapted to zero-sum game structures with imperfect information.
- Simple combinations of (perfect information) alternating bisimulations and epistemic bisimulations don't work.
- History-based alternating bisimulation is undecidable (see paper).

Further work

- Strategy logic with imperfect information?
- Determinacy for other types of zero-sum multy-player games?