A Logic of Knowledge and Strategies with Imperfect Information

Francesco Belardinelli

ongoing work with Wojtek Jamroga

Laboratoire IBISC, Université d'Evry

Journées de l'Intelligence Artificielle Fondamentale – 29 June 2015

Overview

1. Background:
   - from temporal logic to strategy logic
2. The Problem:
   - imperfect information in reasoning about strategies
   - weaker semantical properties (w.r.t. perfect information)
   - failure of relevant fixed-point characterisations of ATL operators
3. The Proposed Solution:
   - Methodology: an agent knows the stragy she's using (at least)
   - E-ATL: an epistemic extension of ATL
4. The Contribution:
   - (partial) characterisations of ATL modalities $\langle\langle\Sigma\rangle\rangle F$, $\langle\langle\Sigma\rangle\rangle G$, $\langle\langle\Sigma\rangle\rangle U$ in contexts of imperfect information
5. Conclusions and Future Work
   - applications to the model checking and satisfiability problems

Background

An essential history of temporal logics in CS

'70: Linear-time Temporal Logic (LTL [Pnu77])

Xp:

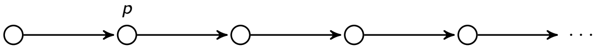

qUp:

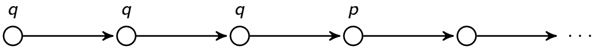

'80: Computation-tree Temporal Logic (CTL [EC82])

$AX\varphi$:

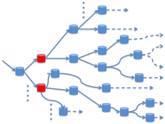

$A\varphi\mathcal{U}\psi$:

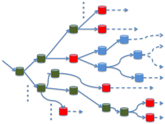

$E\varphi\mathcal{U}\psi$:

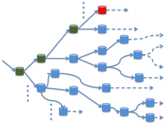

'90: Alternating-time Temporal Logic (ATL [AHK02])

Background

Alternating-time Temporal Logic

ATL: a logic of strategic abilities

- strategy modality $\langle\langle\Sigma\rangle\rangle$ expressing that 'the agents in coalition Σ have a strategy to enforce ...'
- LTL modalities *next* X and *until* U
- interpreted on *Concurrent Game Structures* ...
- ...with a variety of semantical options:
  - perfect v. imperfect information
  - perfect v. imperfect memory
  - objective v. subjective strategies
- Perfect information: fixed-point characterisations of ATL operators

$$\langle\langle\Sigma\rangle\rangle G\phi \quad\leftrightarrow\quad \phi \wedge \langle\langle\Sigma\rangle\rangle X \langle\langle\Sigma\rangle\rangle G\phi \tag{1}$$

$$\langle\langle\Sigma\rangle\rangle F\phi \quad\leftrightarrow\quad \phi \vee \langle\langle\Sigma\rangle\rangle X \langle\langle\Sigma\rangle\rangle F\phi \tag{2}$$

$$\langle\langle\Sigma\rangle\rangle(\phi U\phi') \quad\leftrightarrow\quad \phi' \vee (\phi \wedge \langle\langle\Sigma\rangle\rangle X \langle\langle\Sigma\rangle\rangle(\psi U\phi')) \tag{3}$$

- useful validities: techniques for satisfiability [GS09] and model checking [AHK02, BDJ10]
- The Problem: (1)-(3) do not hold in the imperfect information semantics!

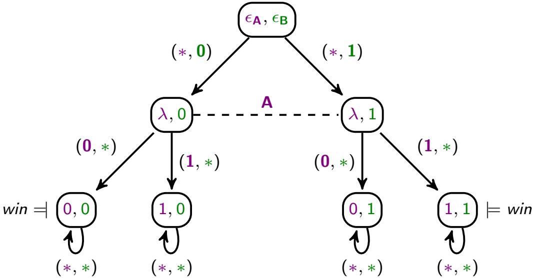

- **B**ob chooses secretly between 0 and 1
- at the next step **A**nne also chooses between 0 and 1
- **A**nne wins the game iff the values provided by the two players coincide
- the dotted line indicates epistemic indistinguishability

The Problem

ATL with Imperfect Information

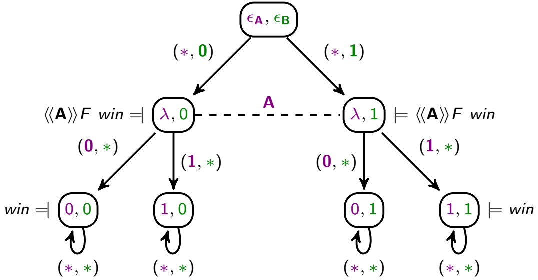

- **B**ob chooses secretly between 0 and 1
- at the next step **A**nne also chooses between 0 and 1
- **A**nne wins the game iff the values provided by the two players coincide
- the dotted line indicates epistemic indistinguishability

The Problem

ATL with Imperfect Information

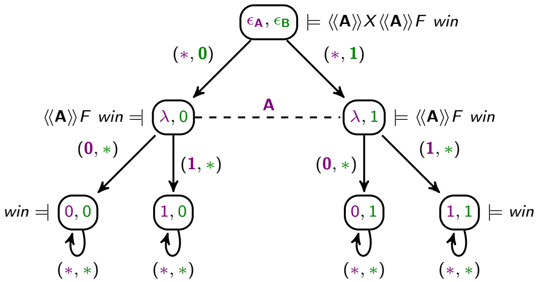

- **B**ob chooses secretly between 0 and 1
- at the next step **A**nne also chooses between 0 and 1
- **A**nne wins the game iff the values provided by the two players coincide
- the dotted line indicates epistemic indistinguishability

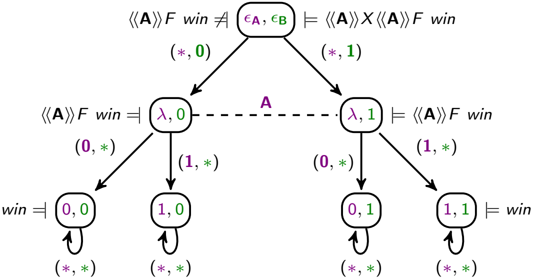

- **B**ob chooses secretly between 0 and 1
- at the next step **A**nne also chooses between 0 and 1
- **A**nne wins the game iff the values provided by the two players coincide
- the dotted line indicates epistemic indistinguishability

The Problem

ATL with Imperfect Information

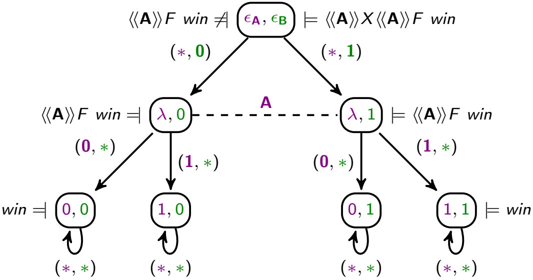

- **B**ob chooses secretly between 0 and 1
- at the next step **A**nne also chooses between 0 and 1
- **A**nne wins the game iff the values provided by the two players coincide
- the dotted line indicates epistemic indistinguishability
- **A**nne knows that there exists a strategy to win the game ...

  ... however, she is not able to point this strategy out

  $\Leftarrow$ **A**nne has *imperfect information* of the game

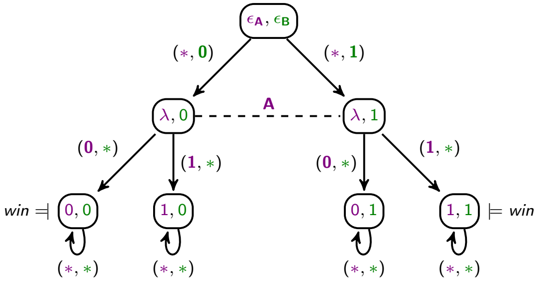

It looks like it's a question of knowledge

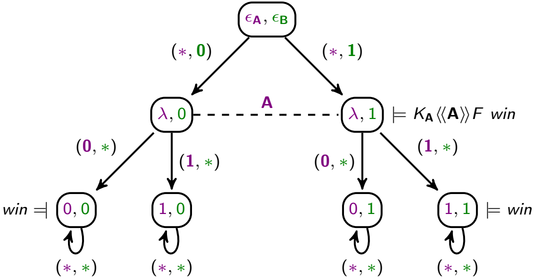

It looks like it's a question of knowledge

- **A**nne knows that there is some strategy to win (knowledge *de dicto*)

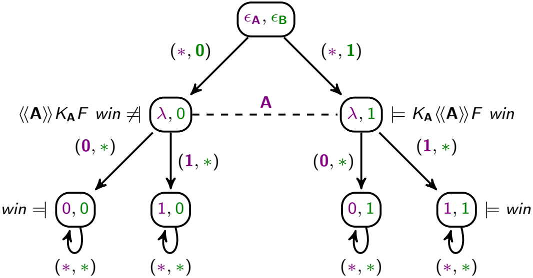

It looks like it's a question of knowledge

- **A**nne knows that there is some strategy to win (knowledge *de dicto*)
- but there is no strategy known to her to guarantee a win (knowledge *de re*)

The Problem

ATL with Imperfect Information

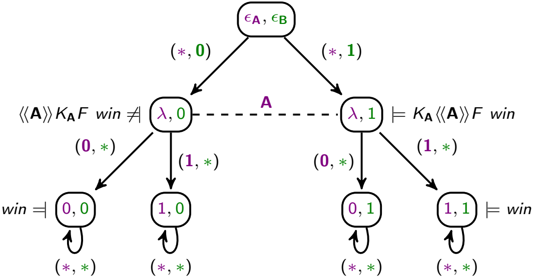

It looks like it's a question of knowledge

- **A**nne knows that there is some strategy to win (knowledge *de dicto*)
- but there is no strategy known to her to guarantee a win (knowledge *de re*)

. . . Let's try and express this distinction explicitly in our language!

- Extensions of logics for reactive systems with epistemic operators to reason about the knowledge agents have of the system's evolution:
  - combinations of CTL and LTL with multi-modal epistemic logic $S5_n$ [HV86, HV89, FHMV95]
  - successfully applied to MAS specification and verification [GvdM04, KNN$^+$08, LQR09]

- Along these lines, [vdHW03] introduced ATEL.
  - spawned a wealth of contributions:
    - imperfect information/uniform strategies [Sch04, JvdH04]
    - constructive knowledge [JÅ07]
    - irrevocable/feasible strategies [AGJ07, Jon03]

- E-ATL: a logic of knowledge and strategies (under imperfect information)
  - not the first attempt to distinguish knowledge *de re*/*de dicto* (ATOL [JvdH04])
  - but here knowledge is not masked by strategy operators

We adopt an agent-oriented perspective.

Definition (Agent)

An *agent* i is

- situated in some *local state* $l_i \in L_i$ and ...
- performs the *actions* in Act_i
- ...according to her *protocol function* $Pr_i : L_i \mapsto 2^{Act_i}$

The setting is reminiscent of the *Interpreted Systems* semantics for MAS [FHMV95].

Example

Anne $= \langle L_{\mathbf{A}}, Act_{\mathbf{A}}, Pr_{\mathbf{A}} \rangle$ is defined as

- $L_{\mathbf{A}} = \{\epsilon_{\mathbf{A}}, \lambda, 0, 1\}$
- $Act_{\mathbf{A}} = \{\mathbf{0}, \mathbf{1}, *\}$, where $*$ is the *skip* action
- $Pr_{\mathbf{A}}(\epsilon_{\mathbf{A}}) = Pr_{\mathbf{A}}(0) = Pr_{\mathbf{A}}(1) = \{*\}$, $Pr_{\mathbf{A}}(\lambda) = \{\mathbf{0}, \mathbf{1}\}$

Epistemic Concurrent Game Models

ECGM

Interactions amongst agents generate ECGM.

- related to CGS [AHK02, MMPV14] and AETS [vdHW03]
- *global states* are not primitive: $s = \langle l_0, \dots, l_\ell \rangle \in G = \Pi_{i \in Ag} L_i$
- *joint actions* are tuples $\sigma = \langle \sigma_0, \dots, \sigma_\ell \rangle \in Act = \Pi_{i \in Ag} Act_i$

Definition (ECGM)

Given

- a set $Ag = \{i_0, \dots, i_\ell\}$ of agents
- a set AP of atomic propositions

an *ECGM* $\mathcal{P}$ includes

- a finite set $I \subseteq G$ of *initial global states*
- a *transition function* $\tau : G \times Act \to G$
- an *interpretation* $\pi : AP \to 2^G$ of atomic propositions

- we denote with S the set of reachable global states
- the *epistemic indistinguishability relation* is not primitive: $s \sim_i s'$ iff $l_i = l'_i$

Epistemic Alternating-time Temporal Logic

E-ATL

E-ATL extends ATL with epistemic operators K_i for individual knowledge.

Definition (E-ATL)

E-ATL state formulas ϕ and path formulas ψ are defined in BNF as follows:

$$\begin{aligned} \phi &::= p \mid \neg\phi \mid \phi \rightarrow \phi \mid \langle\langle \Sigma \rangle\rangle \psi \mid K_i \phi \\ \psi &::= X\phi \mid \phi U \phi \mid K_i \psi \end{aligned}$$

where $p \in AP$, $i \in Ag$ and $\Sigma \subseteq Ag$.

- Syntatically,
  - ATEL $\subset$ E-ATL
  - E-ATL and ATEL* are uncomparable
- $K_{\mathbf{A}} \langle\langle \mathbf{A} \rangle\rangle F\ win$: **A**nne knows that there is some strategy to win the game
- $\neg\langle\langle \mathbf{A} \rangle\rangle K_{\mathbf{A}} F\ win$: but there is no strategy known to her to guarantee a win

Definition (Strategy)

An *i-strategy* $f_i : G^+ \mapsto Act_i$ maps finite sequences of states to *enabled* i-actions (i.e., $f_i(s) \in Pr_i(l_i)$).

- for a group $\Sigma = \{i_0, \ldots, i_\ell\}$ of agents, a group strategy f_Σ is a tuple $\langle f_0, \ldots, f_\ell \rangle$
- a run λ is a sequence $s^0 \to s^1 \to \ldots$ of states s.t. $s^{i+1} = \tau(s^i, \sigma)$ for some joint action $\sigma \in Act$
- a run λ belongs to outcome $out(s, f_\Sigma)$ iff $\lambda(i+1) \in \tau(\lambda(i), (f_\Sigma, f_{\overline{\Sigma}})(\lambda(i)))$ for some $\overline{\Sigma}$-strategy $f_{\overline{\Sigma}}$.

Under imperfect information, strategies depend on the local state of agents only.

Definition (Uniform Strategy [JvdH04])

An i-strategy is *uniform* iff for all states $s, s' \in S$, $s \sim_i s'$ implies $f_i(s) = f_i(s')$.

- A uniform *i-strategy* $f_i : L_i \mapsto Act_i$ maps local states to *enabled* i-actions (i.e., $f_i(l_i) \in Pr_i(l_i)$).

Semantics of E-ATL

Formal definition

Definition (Satisfaction)

An ECGM $\mathcal{P}$ *satisfies* a formula φ in a state s (possibly w.r.t. a strategy profile f_{Ag}) as follows:

$(\mathcal{P}, s) \models p$	iff	$s \in \pi(p)$
$(\mathcal{P}, s) \models \langle\langle \Sigma \rangle\rangle \psi$	iff	for some Σ-strategy f_Σ, for all $\overline{\Sigma}$-strategies $f_{\overline{\Sigma}}$, $(\mathcal{P}, s, (f_\Sigma, f_{\overline{\Sigma}})) \models \psi$
$(\mathcal{P}, s) \models K_i \phi$	iff	for every $s' \in S$, $s \sim_i s'$ implies $(\mathcal{P}, s') \models \phi$
$(\mathcal{P}, s, f_{Ag}) \models X\phi$	iff	for $\lambda = out(s, f_{Ag})$, $(\mathcal{P}, \lambda(1)) \models \phi$
$(\mathcal{P}, s, f_{Ag}) \models \phi U \phi'$	iff	for $\lambda = out(s, f_{Ag})$, for some $k \geq 0$, $(\mathcal{P}, \lambda(k)) \models \phi'$ and $0 \leq j < k$ implies $(\mathcal{P}, \lambda(j)) \models \phi$
$(\mathcal{P}, s, f_{Ag}) \models K_i \psi$	iff	for every $s' \in S$, $s \sim_i s'$ implies $(\mathcal{P}, s', f_{Ag}) \models \psi$

The Working Hypothesis

Fixed-point Characterisations

$$\langle\langle i\rangle\rangle G\phi \quad\leftrightarrow\quad \phi \wedge \langle\langle i\rangle\rangle X(\langle\langle i\rangle\rangle G\phi \wedge (K_i\langle\langle i\rangle\rangle G\phi \rightarrow \langle\langle i\rangle\rangle K_i G\phi)) \tag{4}$$

$$\langle\langle i\rangle\rangle F\phi \quad\leftrightarrow\quad \phi \vee \langle\langle i\rangle\rangle X(\langle\langle i\rangle\rangle F\phi \wedge (K_i\langle\langle i\rangle\rangle F\phi \rightarrow \langle\langle i\rangle\rangle K_i F\phi)) \tag{5}$$

$$\langle\langle i\rangle\rangle(\psi U\phi) \quad\leftrightarrow\quad \phi \vee (\psi \wedge \langle\langle i\rangle\rangle X(\langle\langle i\rangle\rangle(\psi U\phi) \wedge (K_i\langle\langle i\rangle\rangle(\psi U\phi) \rightarrow \langle\langle i\rangle\rangle K_i(\psi U\phi)))) \tag{6}$$

- Single agent case only.
- Also, negations appear in (4)-(6),
  - hence, the corresponding functions are not monotonous.

$\Rightarrow$ Least and greatest fixed points might not exist.

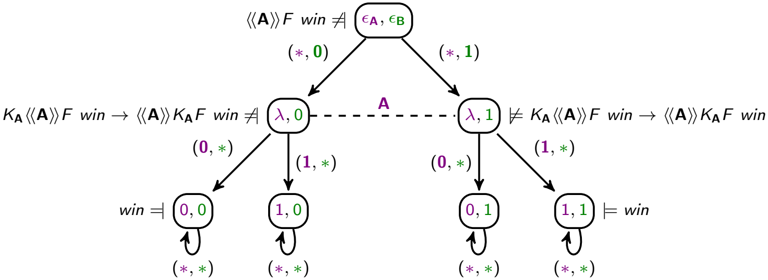

- $(\lambda, 0) \models K_{\mathbf{A}} \langle\langle \mathbf{A} \rangle\rangle F\ win$: **A**nne knows that there is some strategy to win the game
- $(\lambda, 0) \not\models \langle\langle \mathbf{A} \rangle\rangle K_{\mathbf{A}} F\ win$: but there is no strategy known to her to guarantee a win

More Problems ...

... and a first solution

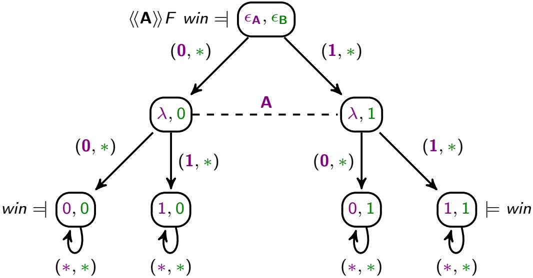

More Problems . . .

. . . and a first solution

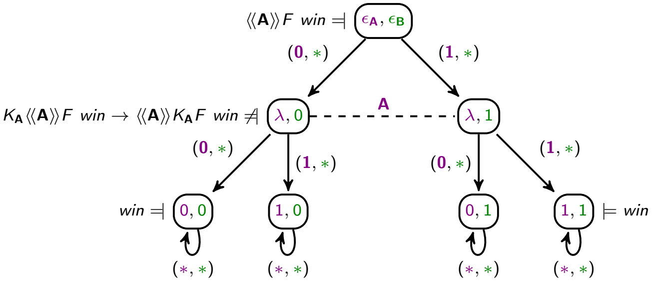

More Problems ...

... and a first solution

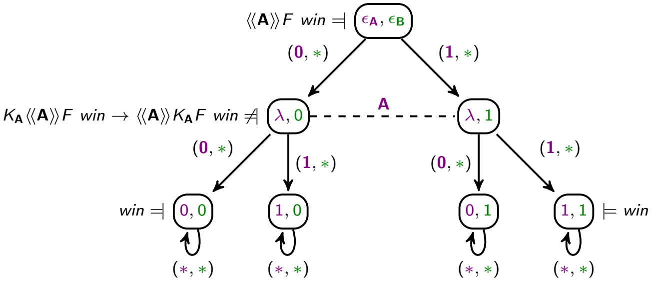

- Methodology: agents know the strategy they are using (*context*)
- ECGM $\mathcal{P}$ satisfies formula φ in state s w.r.t. strategy profile f_{Ag} and context $V_{Ag} = \langle V_0, \ldots, V_\ell \rangle$ iff

$$
\begin{array}{lll}
(\mathcal{P}, s, V_{Ag}) \models K_i \phi & \text{iff} & \text{for every } s' \in V_i,\ s \sim_i s' \text{ implies } (\mathcal{P}, s', V_{Ag}) \models \phi \\
(\mathcal{P}, s, V_{Ag}, f_{Ag}) \models K_i \psi & \text{iff} & \text{for every } s' \in V_i,\ s \sim_i s' \text{ implies } (\mathcal{P}, s', V_{Ag}, f_{Ag}) \models \psi \\
(\mathcal{P}, s, V_{Ag}) \models \langle\langle \Sigma \rangle\rangle \psi & \text{iff} & \text{for some } \Sigma\text{-strategies } f_\Sigma, \text{ for all } \overline{\Sigma}\text{-strategies } f_{\overline{\Sigma}}, \\
 & & (\mathcal{P}, s, out(V_1, f_1), \ldots, out(V_\ell, f_\ell), (f_\Sigma, f_{\overline{\Sigma}})) \models \psi
\end{array}
$$

- A formula ϕ is satisfied at s iff it is satisfied in context $\langle \{s\}, \ldots, \{s\} \rangle$.

Yet More Problems . . .

. . . and a second attempt

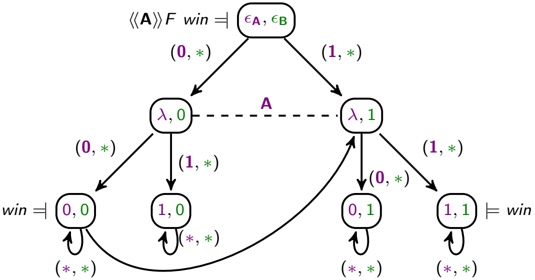

Yet More Problems . . .

. . . and a second attempt

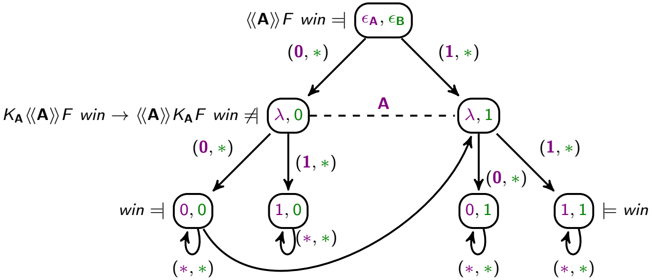

Yet More Problems . . .

. . . and a second attempt

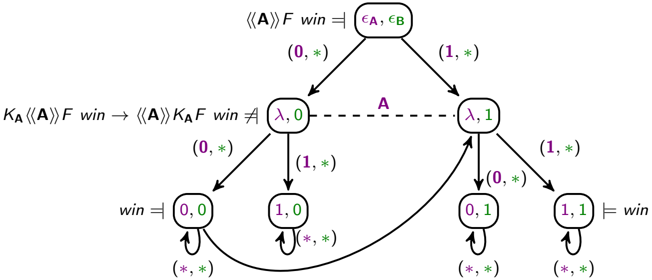

- Let's consider a perfect memory semantics
- ECGM $\mathcal{P}$ satisfies formula φ at *history* h w.r.t. strategy profile f_{Ag} and context $V_{Ag} = \langle V_0, \ldots, V_\ell \rangle$ iff

$(\mathcal{P}, h, V_{Ag}) \models K_i\phi$ iff for every $h' \in V_i$, $h \sim_i h'$ implies $(\mathcal{P}, h', V_{Ag}) \models \phi$

$(\mathcal{P}, h, V_{Ag}, f_{Ag}) \models K_i\psi$ iff for every $h' \in V_i$, $h \sim_i h'$ implies $(\mathcal{P}, h', V_{Ag}, f_{Ag}) \models \psi$

$(\mathcal{P}, h, V_{Ag}) \models \langle\langle \Sigma \rangle\rangle \psi$ iff for some Σ-strategies f_Σ, for all $\overline{\Sigma}$-strategies $f_{\overline{\Sigma}}$, $(\mathcal{P}, h, out(V_1, f_1), \ldots, out(V_\ell, f_\ell), (f_\Sigma, f_{\overline{\Sigma}})) \models \psi$

A (Fixed-point) Characterisation

- by considering a semantics with imperfect information but perfect memory, formulas (4)-(6) are valid.
- actually, they can be reduced to the following equivalences:

$$\begin{aligned}
\langle\langle\Sigma\rangle\rangle G\phi &\quad\leftrightarrow\quad \phi \wedge \langle\langle\Sigma\rangle\rangle X(\langle\langle\Sigma\rangle\rangle E_\Sigma G\phi) \\
\langle\langle\Sigma\rangle\rangle F\phi &\quad\leftrightarrow\quad \phi \vee \langle\langle\Sigma\rangle\rangle X(\langle\langle\Sigma\rangle\rangle E_\Sigma F\phi) \\
\langle\langle\Sigma\rangle\rangle(\phi U\phi') &\quad\leftrightarrow\quad \phi' \vee (\phi \wedge \langle\langle\Sigma\rangle\rangle X(\langle\langle\Sigma\rangle\rangle E_\Sigma(\phi U\phi')))
\end{aligned}$$

Limitations:

- ϕ must be purely temporal!
- no unfolding!

Conclusions

Results:

- E-ATL: a logic for reasoning about knowledge and strategies in a multi-agent setting
- Methodology: agents know the strategy they are using, that is, their *context*
- under perfect memory E-ATL allows us to (partially) recover the characterisation of ATL operators

and Future Work:

- Extension to arbitrary formulas (arbitrary contexts)
- Application (algorithms?) to satisfiability

	Ir	*IR*	*ir*		*iR*	
			sub	*obj*	*sub*	*obj*
SAT	*EXPTIME*	*EXPTIME*	no result	no result	no result	no result

  - perfect (I) and imperfect (i) information
  - perfect (R) and imperfect (r) memory
  - *sub*jective and *obj*ective strategies

Questions?

References

Thomas Agotnes, Valentin Goranko, and Wojciech Jamroga.
Alternating-time temporal logics with irrevocable strategies.
In *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge*, TARK '07, pages 15–24, New York, NY, USA, 2007. ACM.

Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman.
Alternating-time temporal logic.
J. ACM, 49(5):672–713, 2002.

Nils Bulling, Jurgen Dix, and Wojciech Jamroga.
Model checking logics of strategic ability: Complexity*.
In Mehdi Dastani, Koen V. Hindriks, and John-Jules Charles Meyer, editors, *Specification and Verification of Multi-agent Systems*, pages 125–159. Springer US, 2010.

E. A. Emerson and E. M. Clarke.
Using branching-time temporal logic to synthesize synchronization skeletons.
Science of Computer Programming, 2(3):241–266, 1982.

R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi.
Reasoning About Knowledge.
The MIT Press, 1995.

V. Goranko and D. Shkatov.
Tableau-based decision procedures for logics of strategic ability in multiagent systems.
ACM Trans. Comput. Log., 11(1), 2009.

Peter Gammie and Ron van der Meyden.
Mck: Model checking the logic of knowledge.
In Rajeev Alur and Doron Peled, editors, *CAV*, volume 3114 of *Lecture Notes in Computer Science*, pages 479–483. Springer, 2004.

Joseph Y. Halpern and Moshe Y. Vardi.
The complexity of reasoning about knowledge and time: Extended abstract.
In Juris Hartmanis, editor, *STOC*, pages 304–315. ACM, 1986.

Joseph Y. Halpern and Moshe Y. Vardi.
The complexity of reasoning about knowledge and time. i. lower bounds.
J. Comput. Syst. Sci., 38(1):195–237, 1989.