Reasoning about Knowledge and Strategies: Epistemic Strategy Logic

Francesco Belardinelli

Laboratoire IBISC, Université d'Evry

Strategic Reasoning – 5 April 2014

Overview

1. Motivation and Background
   - logics for reasoning about strategies and knowledge
2. Epistemic Strategy Logic
   - semantics
   - syntax
3. Main Contribution
   - model checking ESL is no harder than SL
4. Imperfect Information
   - benefits of combining epistemic and strategy modalities
5. Conclusions and Future Work

Motivation and Background

Logics of strategic abilities

- Logics for strategic reasoning are a thriving area of research in AI and MAS.

Motivation and Background

- Logics for strategic reasoning are a thriving area of research in AI and MAS.
- Two lines of research:
  1. multi-modal logics to formalise strategic abilities and behaviours of individual agents and groups:
     - **A**lternating-time **T**emporal **L**ogic [AHK02]
     - **C**oalition **L**ogic [Pau02]
     - **S**trategy **L**ogic [CHP10, MMV10]
  2. 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]

Motivation and Background

Logics of strategic abilities

- Logics for strategic reasoning are a thriving area of research in AI and MAS.
- Two lines of research:
  1. multi-modal logics to formalise strategic abilities and behaviours of individual agents and groups:
     - **A**lternating-time **T**emporal **L**ogic [AHK02]
     - **C**oalition **L**ogic [Pau02]
     - **S**trategy **L**ogic [CHP10, MMV10]
  2. 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]

Motivation and Background

Logics of strategic abilities

- Logics for strategic reasoning are a thriving area of research in AI and MAS.
- Two lines of research:
  1. multi-modal logics to formalise strategic abilities and behaviours of individual agents and groups:
     - **A**lternating-time **T**emporal **L**ogic [AHK02]
     - **C**oalition **L**ogic [Pau02]
     - **S**trategy **L**ogic [CHP10, MMV10]
  2. 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]

Epistemic **S**trategy **L**ogic = strategies + knowledge

  - topic of interest [HvdM14b, HvdM14a]

- **A**nne and **B**ob can either **C**ooperate or **D**efect
- payoff ordering: $a > b > c > d$

		Bob	
		Cooperate	**D**efect
Anne	**C**ooperate	b, b	d, a
	Defect	a, d	c, c

- can **A**nne achieve payoff a?
- does **A**nne know whether she can achieve payoff a?
- does **B**ob know whether **A**nne has a strategy to achieve payoff a?
- do **A**nne and **B**ob know (*de dicto*) whether they can reach a Nash equilibrium?
- are there strategies such that **A**nne and **B**ob know (*de re*) that they can reach a Nash equilibrium?

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 (Prisoner's Dilemma)

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

- $L_{\mathbf{A}} = \{\epsilon_{\mathbf{A}}, a, b, c, d\}$
- $Act_{\mathbf{A}} = \{\mathbf{C}, \mathbf{D}, *\}$, where $*$ is the *skip* action
- $Pr_{\mathbf{A}}(\epsilon_{\mathbf{A}}) = \{\mathbf{C}, \mathbf{D}\}$ and $Pr_{\mathbf{A}}(a) = Pr_{\mathbf{A}}(b) = Pr_{\mathbf{A}}(c) = Pr_{\mathbf{A}}(d) = \{*\}$

The definition of agent **B**ob is symmetric.

Epistemic Concurrent Game Models

ECGM

The interactions amongst agents generate ECGM.

- related to CGS [AHK02, MMV10] and AETS [vdHW03]
- *global states* are not primitive: $s = \langle l_0, \ldots, l_\ell \rangle \in G = \Pi_{i \in Ag} L_i$
- *joint actions* are tuples $\sigma = \langle \sigma_0, \ldots, \sigma_\ell \rangle \in Act = \Pi_{i \in Ag} Act_i$

Definition (ECGM)

Given

- a set $Ag = \{i_0, \ldots, i_\ell\}$ 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

- the *epistemic indistinguishability relation* is not primitive: $s \sim_i s'$ iff $l_i = l'_i$

The Prisoner's Dilemma as an ECGM

Let $AP = \{a_i, b_i, c_i, d_i\}$ for $i \in \{\mathbf{A}, \mathbf{B}\}$.

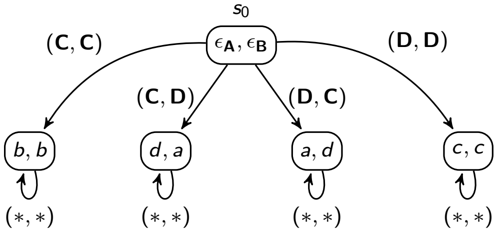

Example (ECGM $\mathcal{P}_{pd}$)

For the set $Ag = \{\mathbf{A}, \mathbf{B}\}$ of agents, the prisoner's dilemma ECGM $\mathcal{P}_{pd}$ includes

- the set $I = \{s_0\}$ of initial states, with $s_0 = (\epsilon_\mathbf{A}, \epsilon_\mathbf{B})$
- the transition function τ, given as
  - $\tau(s_0, (\mathbf{C}, \mathbf{C})) = (b, b)$
  - $\tau(s_0, (\mathbf{C}, \mathbf{D})) = (d, a)$
  - $\tau(s_0, (\mathbf{D}, \mathbf{C})) = (a, d)$
  - $\tau(s_0, (\mathbf{D}, \mathbf{D})) = (c, c)$
  - $\tau(s, (*, *)) = s$, for every state s different from s_0
- the interpretation π s.t. a state $(l_\mathbf{A}, l_\mathbf{B})$ belongs to $\pi(p_i)$ iff $l_i = p$.

Epistemic Strategy Logic

ESL

ESL extends SL with epistemic operators K_i for individual knowledge.

- we introduce a set Var_i of strategy variables for each agent $i \in Ag$

Definition (ESL)

ESL formulas are defined in BNF as follows:

$$\phi \quad ::= \quad p \mid \neg\phi \mid \phi \rightarrow \phi \mid X\phi \mid \phi U \phi \mid \exists x_i \phi \mid K_i \phi$$

- we consider a multi-agent setting ($\neq$ [CHP10])
- the language does not include the binding operator (α, x) ($\neq$ [MMV10])

The questions above can be recast as model checking problems:

$$\mathcal{P}_{pd} \overset{?}{\models} \exists x_{\mathbf{A}} F a_{\mathbf{A}}$$

$$\mathcal{P}_{pd} \overset{?}{\models} K_{\mathbf{A}}(\exists x_{\mathbf{A}} F a_{\mathbf{A}} \vee \neg \exists x_{\mathbf{A}} F a_{\mathbf{A}})$$

$$\mathcal{P}_{pd} \overset{?}{\models} K_{\mathbf{B}}(\exists x_{\mathbf{A}} F a_{\mathbf{A}} \vee \neg \exists x_{\mathbf{A}} F a_{\mathbf{A}})$$

Definition (Strategy)

An *A-strategy* is a mapping $f_A : G^+ \mapsto Act_A$ from finite sequences of states to *enabled* A-actions.

- a run λ is a sequence $s^0 \to s^1 \to \ldots$ of global states
- a run λ belongs to outcome $out(s, f_A)$ iff $\lambda(i+1) \in \hat{\tau}(\lambda(i), f_A(\lambda[\ldots, i]))$

  $\Rightarrow$ a group strategy is really the composition of its members' strategies
- an assignment χ maps each agent $i \in Ag$ to an i-strategy f_i
  - f^χ is the Ag-strategy $\chi(i_0) \times \ldots \times \chi(i_\ell)$

Definition (Satisfaction)

An ECGM $\mathcal{P}$ *satisfies* an ESL formula φ in a state s for an assignment χ, iff

$(\mathcal{P}, s, \chi) \models p$	iff	$s \in \pi(p)$
$(\mathcal{P}, s, \chi) \models X\psi$	iff	for $\lambda = out(s, f^\chi)$, $(\mathcal{P}, \lambda(1), \chi) \models \psi$
$(\mathcal{P}, s, \chi) \models \psi U \psi'$	iff	for $\lambda = out(s, f^\chi)$ there is $k \geq 0$ s.t. $(\mathcal{P}, \lambda(k), \chi) \models \psi'$ and $0 \leq j < k$ implies $(\mathcal{P}, \lambda(j), \chi) \models \psi$
$(\mathcal{P}, s, \chi) \models \exists x_i \psi$	iff	there exists an i-strategy f_i s.t. $(\mathcal{P}, s, \chi^i_{f_i}) \models \psi$
$(\mathcal{P}, s, \chi) \models K_i \psi$	iff	for every $s \in S$, $s \sim_i s'$ implies $(\mathcal{P}, s', \chi) \models \psi$

Expressiveness

Knowledge of Nash Equilibria

- given an n-player game in normal form with payoff ordering $a_1 > \ldots > a_k$, define

$$\psi_{NE} \quad ::= \quad \bigwedge_{i=1}^{n} \bigwedge_{i=1}^{k} \left(\left(\bigwedge_{j=1}^{i-1} \neg \exists y_j X a_j \right) \wedge \exists y_i X a_i \rightarrow X a_i \right)$$

Proposition

$$(\mathcal{P}_{pd}, s_0, \chi) \models \psi_{NE} \quad \textit{iff} \quad (\chi(1)(s_0), \ldots, \chi(n)(s_0)) \textit{ is a Nash equilibrium}$$

- for the prisoner's dilemma,

$$\begin{aligned} (\mathcal{P}_{pd}, s_0, \chi) \models \psi_{NE} \quad & \text{iff} \quad (\chi(1)(s_0), \chi(2)(s_0)) \text{ is a Nash equilibrium} \\ & \text{iff} \quad \chi(1)(s_0) = \chi(2)(s_0) = \mathbf{D} \end{aligned}$$

The questions above can be recast as model checking problems:

$$\begin{aligned} \mathcal{P}_{pd} \quad & \overset{?}{\models} \quad K_{\mathbf{A}} \exists x_{\mathbf{A}}, x_{\mathbf{B}} \psi_{NE} \wedge K_{\mathbf{B}} \exists x_{\mathbf{A}}, x_{\mathbf{B}} \psi_{NE} \\ \mathcal{P}_{pd} \quad & \overset{?}{\models} \quad \exists x_{\mathbf{A}}, x_{\mathbf{B}} (K_{\mathbf{A}} \psi_{NE} \wedge K_{\mathbf{B}} \psi_{NE}) \end{aligned}$$

- knowledge *de re* $\Rightarrow$ knowledge *de dicto*:

  $$\models \exists x_i K_j \phi \rightarrow K_j \exists x_i \phi$$

  also, knowledge *de dicto* $\Rightarrow$ knowledge *de re*:

  $$\models K_j \exists x_i \phi \rightarrow \exists x_i K_j \phi$$

  indeed, agents have *perfect information* of the game

- individual strategies depend on global states [JvdH04]

Model Checking ESL

Theorem (Hardness)

The model checking problem for ESL is NON-ELEMENTARYSPACE*-hard.*

- reduction to satisfiability for quantified propositional temporal logic (QPTL)
- differently from [MMV10] the syntax does not include the binding operator

Theorem (Completeness)

The model checking problem for ESL is PTIME*-complete w.r.t. the size of the model and* NON-ELEMENTARY *w.r.t. the size of the formula.*

- reduction to non-emptyness for alternating tree automata [MMV10]

$\Rightarrow$ The model checking problem is no harder for ESL than for SL.

Imperfect Information

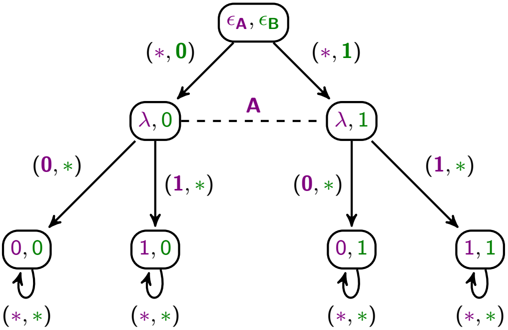

- **B**ob chooses secretly between 0 and 1
- at the next step **A**nne also chooses between 0 and 1
- **A**nne wins the game iff the values provided by the two players coincide
- the dotted line indicates epistemic indistinguishability

- **A**nne knows that there exists a strategy to win the game ...
  ... however, she is not able to point this strategy out

  ⇒ **A**nne has *imperfect information* of the game

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

Definition (Uniform Strategies [JvdH04])

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

- **A**nne knows that there exists a strategy to win the game ...

$$(\mathcal{Q}, s_{\lambda 0}) \quad \models_{ii} \quad K_{\mathbf{A}}\, \exists x_{\mathbf{A}}\; X\; win$$

... however, there is no strategy that she knows to be winning:

$$(\mathcal{Q}, s_{\lambda 0}) \quad \not\models_{ii} \quad \exists x_{\mathbf{A}}\; K_{\mathbf{A}}\; X\; win$$

Epistemic modalities allow us to recover some fixed-point characterisations of ATL operators.

- for $A = \{i_0, \dots, i_\ell\}$, $\vec{x}_A = x_{i_0}, \dots, x_{i_\ell}$ and $\bar{A} = Ag \setminus A$, define

$$\langle\langle A\rangle\rangle \phi \quad ::= \quad \exists \vec{x}_A \forall \vec{x}_{\bar{A}} \phi$$

- under imperfect information we have that

$$\begin{aligned}
\langle\langle A\rangle\rangle G\phi &\quad \not\Leftrightarrow \quad \phi \wedge \langle\langle A\rangle\rangle X \langle\langle A\rangle\rangle G\phi \\
\langle\langle A\rangle\rangle F\phi &\quad \not\Leftrightarrow \quad \phi \vee \langle\langle A\rangle\rangle X \langle\langle A\rangle\rangle F\phi \\
\langle\langle A\rangle\rangle \phi U\phi' &\quad \not\Leftrightarrow \quad \phi' \vee (\phi \wedge \langle\langle A\rangle\rangle X \langle\langle A\rangle\rangle \phi U\phi')
\end{aligned}$$

- by using epistemic modalities we can recover fixed-point characterisations:

$$\begin{aligned}
\langle\langle i\rangle\rangle G\phi &\quad \Leftrightarrow \quad \phi \wedge \langle\langle i\rangle\rangle X \langle\langle i\rangle\rangle (G\phi \wedge K_i(\langle\langle i\rangle\rangle G\phi \rightarrow G\phi)) \\
\langle\langle i\rangle\rangle F\phi &\quad \Leftrightarrow \quad \phi \vee \langle\langle i\rangle\rangle X \langle\langle i\rangle\rangle (F\phi \wedge K_i(\langle\langle i\rangle\rangle F\phi \rightarrow F\phi)) \\
\langle\langle i\rangle\rangle \phi U\phi' &\quad \Leftrightarrow \quad \phi' \vee (\phi \wedge \langle\langle i\rangle\rangle X \langle\langle i\rangle\rangle (\phi U\phi' \wedge K_i(\langle\langle i\rangle\rangle \phi U\phi' \rightarrow \phi U\phi')))
\end{aligned}$$

Conclusions

Results:

- ESL: a logic for reasoning about knowledge and strategies in a multi-agent setting
- the model checking problem is no harder than for SL
- under imperfect information ESL allows us to recover the fixed-point characterisation of ATL operators

and Future Work:

- fragments of ESL: better computational complexity?
- epistemic operators for group knowledge (distributed, common knowledge, etc.)
- imperfect information

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.

Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman.
Strategy logic.
Inf. Comput., 208(6):677–693, 2010.

Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi.
Reasoning About Knowledge.
The MIT Press, 1995.

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.

X. Huang and R. van der Meyden.
An epistemic strategy logic (extended abstract).
In *Proceedings of 2nd International Workshop on Strategic Reasoning*, 2014.

X. Huang and R. van der Meyden.
A temporal logic of strategic knowledge.
In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning*, 2014.

Wojciech Jamroga and Thomas Ågotnes.