

A Strategic Epistemic Logic for Bounded Memory Agents

Sophia Knight

CNRS, Université de Lorraine, Nancy, France

Workshop on Resource Bounded Agents
Barcelona
12 August, 2015

Contents

- 1 Introduction
- 2 Background
- 3 Memory
- 4 The Logic
- 5 Final Remarks

Combining Strategic and Epistemic Reasoning

- The goal is to develop a strategic epistemic logic for agents with *arbitrary* memory abilities.
- Much work has already been done on epistemic ATL, but it is mostly either perfect recall or memoryless situations.
- We allow each agent to have an arbitrary equivalence relation on histories.
- Agents' strategies and knowledge are based on these equivalence relations, so we consider their abilities to act, their knowledge, and the relationship between them.

Combining Strategic and Epistemic Reasoning

- The goal is to develop a strategic epistemic logic for agents with *arbitrary* memory abilities.
- Much work has already been done on epistemic ATL, but it is mostly either perfect recall or memoryless situations.
- We allow each agent to have an arbitrary equivalence relation on histories.
- Agents' strategies and knowledge are based on these equivalence relations, so we consider their abilities to act, their knowledge, and the relationship between them.

Combining Strategic and Epistemic Reasoning

- The goal is to develop a strategic epistemic logic for agents with *arbitrary* memory abilities.
- Much work has already been done on epistemic ATL, but it is mostly either perfect recall or memoryless situations.
- We allow each agent to have an arbitrary equivalence relation on histories.
- Agents' strategies and knowledge are based on these equivalence relations, so we consider their abilities to act, their knowledge, and the relationship between them.

Combining Strategic and Epistemic Reasoning

- The goal is to develop a strategic epistemic logic for agents with *arbitrary* memory abilities.
- Much work has already been done on epistemic ATL, but it is mostly either perfect recall or memoryless situations.
- We allow each agent to have an arbitrary equivalence relation on histories.
- Agents' strategies and knowledge are based on these equivalence relations, so we consider their abilities to act, their knowledge, and the relationship between them.

Our contributions

We develop a strategic, epistemic logic based on uniform strategies.

- We allow agents to have arbitrary equivalence relations on histories,
- Our logic allows different agents to have different memory abilities,
- We present a new version of “perfect recall” for agents, allowing agents to reason about their own actions.

Contents

- 1 Introduction
- 2 Background**
- 3 Memory
- 4 The Logic
- 5 Final Remarks

Epistemic Concurrent Game Structures

$\langle Q, \Pi, \Sigma, \mathcal{B}, \sim, \pi, Av, \delta \rangle$

- Q : states,
- Π : propositions,
- Σ : finite set of agents, $\{a_1, \dots, a_n\}$,
- \mathcal{B} : finite set of actions,
- $\sim : \Sigma \rightarrow \mathcal{P}(Q \times Q)$: equivalence relation on states for each agent,
- $\pi : Q \rightarrow \Pi$: valuation,
- $Av : Q \times \Sigma \rightarrow \mathcal{P}(\mathcal{B})$: the available actions for an agent at a state,
- $\delta : Q \times \Sigma \times \mathcal{B} \rightarrow \mathcal{P}(Q)$: transition function.

Epistemic Concurrent Game Structures

$\langle Q, \Pi, \Sigma, \mathcal{B}, \sim, \pi, Av, \delta \rangle$

Requirements:

- **Indistinguishable states.** If $q_1 \sim_i q_2$, then $Av(q_1, a_i) = Av(q_2, a_i)$: if two states are indistinguishable for a_i , then the same actions are available to a_i .
- **Action availability.** $Av(q, a_i) \neq \emptyset$: every agent has at least one action available at every state.
- **Determinacy:** when every agent chooses an available action, this leads to exactly one next state.

Histories and Strategies

A **history** is a sequence $q_0.b_1^*.q_1.b_2^*.q_2\dots q_{k-1}.b_k^*.q_k$ where each $q_j \in Q$, $b_j \in \mathcal{B}^n$, such that q_j is the b_j^* successor of q_{j-1} (technically, $\{q_j\} = \cap_{i=1}^n \delta(q_{j-1}, a_i, b_i)$)

Given an arbitrary equivalence relation \approx_i on histories, a **uniform strategy** for a_i is a function $f : Hist \rightarrow \mathcal{B}$ satisfying the following requirements:

- for all $h \in Hist$, $f_i(h)$ is an available action for a_i in h , and
- if $h_1 \approx_i h_2$ then $f(h_1) = f(h_2)$.

Contents

- 1 Introduction
- 2 Background
- 3 Memory**
- 4 The Logic
- 5 Final Remarks

Equivalence relations on histories

- Most work on epistemic ATL considers two possibilities for memory: either **perfect recall** or **memoryless**.
- Furthermore, all agents in a system are assumed to have the **same memory capabilities**.
- We generalize the notion of memory in several ways.
 - Allow **arbitrary** equivalence relations on histories.
 - Allow different agents to have different memory capabilities.
 - New notion of perfect recall, taking **actions** into account.

Arbitrary equivalence relations

In other work, agents' equivalence relations on histories are derivable from their equivalence relations on states— usually, memoryless or remembering all past states.

We allow a more general definition: each agent can have **any** equivalence relation on histories.

This means we consider more general systems. Examples:

- Agent remembers all past states except a certain state that is “invisible” to him,
- An agent who remembers half the states the system has been in,
- Agent remembers entire history until the system enters s_0 , which wipes out his memory.

This generalization is similar to allowing agents in traditional, static Kripke models to have arbitrary equivalence relations on the set of states.

Different agents with different memory abilities

In other work, all agents are assumed to have the same memory capabilities- e.g. all memoryless or all with perfect recall.

Since we allow arbitrary equivalence relations, each agent can have a different type of memory.

Practical examples:

- A system where some simple agents with limited memory interact with sophisticated agents who remember everything, or
- A system with friendly agents of known memory ability and adversarial agents of unknown ability. Taking adversaries as perfect recall agents models a worst-case scenario, e.g. to check security properties.

Example I: Agents with different memory abilities

Two agents: a_1 controls a lightswitch, is memoryless and blind.
 a_2 can turn over card: red on one side, green on other. Perfect recall.

Propositions:

r : red

g : green

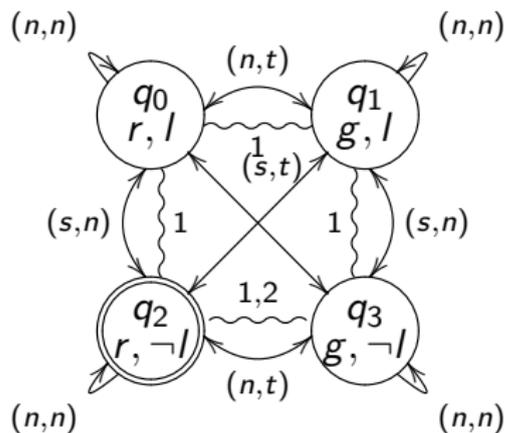
l : light on

Actions:

s : flip lightswitch

t : turn card over

n : do nothing



Example I: Agents with different memory abilities

Formulas:

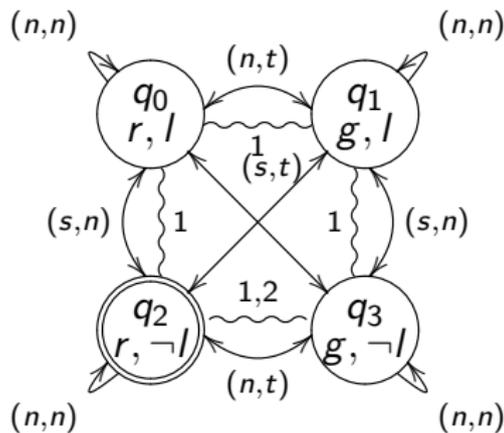
$$q_2 \stackrel{?}{\models} \langle\langle \{a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_2\} \rangle\rangle \circ \langle\langle \{a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_1, a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_1, a_2\} \rangle\rangle \circ \langle\langle \{a_1, a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_1, a_2\} \rangle\rangle \diamond g$$



Example I: Agents with different memory abilities

Formulas:

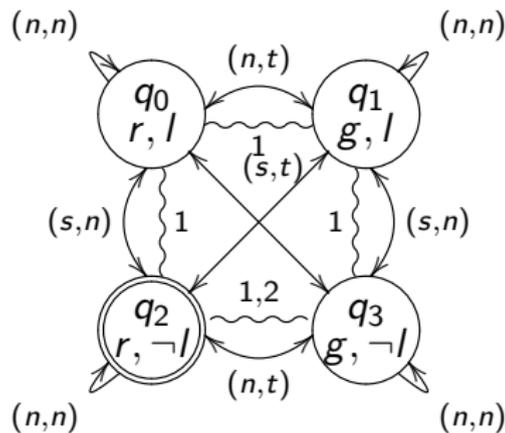
$$q_2 \not\models \langle\langle \{a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_2\} \rangle\rangle \circ \langle\langle \{a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_1, a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_1, a_2\} \rangle\rangle \circ \langle\langle \{a_1, a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_1, a_2\} \rangle\rangle \diamond g$$



Example I: Agents with different memory abilities

Formulas:

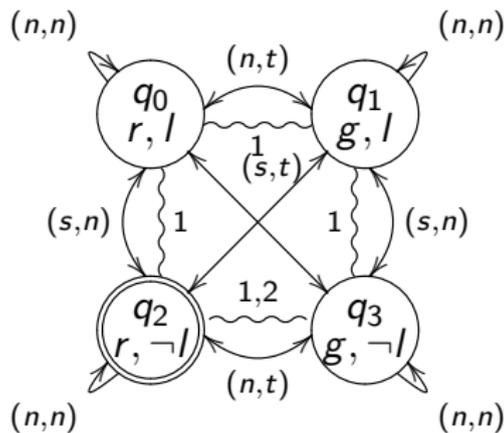
$$q_2 \not\models \langle\langle \{a_2\} \rangle\rangle \circ g$$

$$q_2 \not\models \langle\langle \{a_2\} \rangle\rangle \circ \langle\langle \{a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_1, a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_1, a_2\} \rangle\rangle \circ \langle\langle \{a_1, a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_1, a_2\} \rangle\rangle \diamond g$$



Example I: Agents with different memory abilities

Formulas:

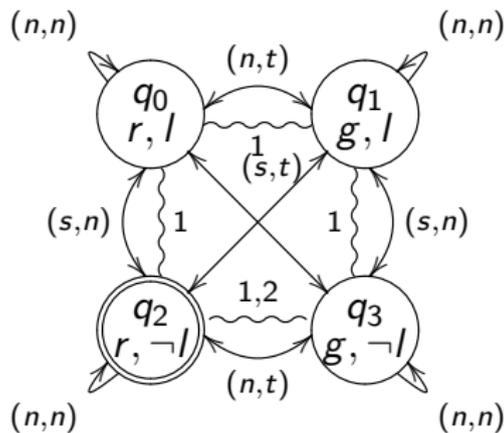
$$q_2 \not\models \langle\langle \{a_2\} \rangle\rangle \circ g$$

$$q_2 \not\models \langle\langle \{a_2\} \rangle\rangle \circ \langle\langle \{a_2\} \rangle\rangle \circ g$$

$$q_2 \not\models \langle\langle \{a_1, a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_1, a_2\} \rangle\rangle \circ \langle\langle \{a_1, a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_1, a_2\} \rangle\rangle \diamond g$$



Example I: Agents with different memory abilities

Formulas:

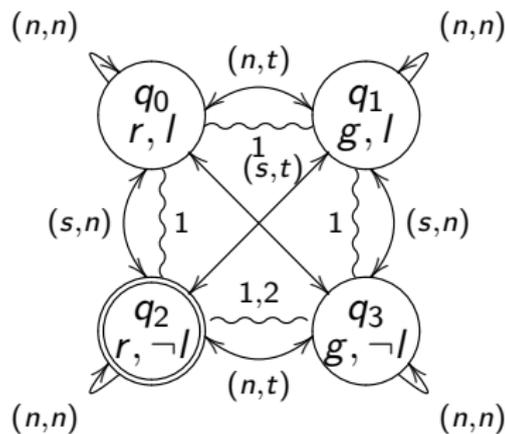
$$q_2 \not\models \langle\langle \{a_2\} \rangle\rangle \circ g$$

$$q_2 \not\models \langle\langle \{a_2\} \rangle\rangle \circ \langle\langle \{a_2\} \rangle\rangle \circ g$$

$$q_2 \not\models \langle\langle \{a_1, a_2\} \rangle\rangle \circ g$$

$$q_2 \models \langle\langle \{a_1, a_2\} \rangle\rangle \circ \langle\langle \{a_1, a_2\} \rangle\rangle \circ g$$

$$q_2 \stackrel{?}{\models} \langle\langle \{a_1, a_2\} \rangle\rangle \diamond g$$



Example I: Agents with different memory abilities

Formulas:

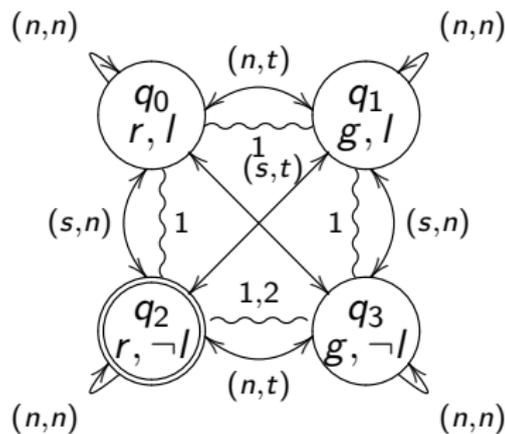
$$q_2 \not\models \langle\langle \{a_2\} \rangle\rangle \circ g$$

$$q_2 \not\models \langle\langle \{a_2\} \rangle\rangle \circ \langle\langle \{a_2\} \rangle\rangle \circ g$$

$$q_2 \not\models \langle\langle \{a_1, a_2\} \rangle\rangle \circ g$$

$$q_2 \models \langle\langle \{a_1, a_2\} \rangle\rangle \circ \langle\langle \{a_1, a_2\} \rangle\rangle \circ g$$

$$q_2 \models \langle\langle \{a_1, a_2\} \rangle\rangle \diamond g$$



New definition of perfect recall

Traditional definition: $q_0.a_1^*.q_1.a_2^*.q_2\dots q_k \approx_i r_0.b_1^*.r_1.b_2^*.r_3\dots r_k$ iff $q_j \sim_i r_j$ for $j \in \{0, \dots, k\}$.

Only looks at **states**.

New definition: two histories are equivalent for a_i if all past states are equivalent **and** a_i took the same action in both histories at every step in the past.

Definition (Perfect recall equivalence)

$h_1 \approx_i h_2$, iff either $h_1 = q_1$ and $h_2 = q_2$ and $q_1 \sim_i q_2$, or $h_1 = q_0.b_1^*.q_1\dots q_{j-1}.b_j^*.q_j$ and $h_2 = r_0.c_1^*.r_1\dots r_{j-1}.c_j^*.r_j$ and:

- 1 $q_0.b_1^*.q_1\dots q_{j-1} \approx_i r_0.c_1^*.r_1\dots r_{j-1}$, and
- 2 $q_j \sim_i r_j$, and
- 3 $b_i = c_i$ where $b_j^* = \langle b_1, \dots, b_n \rangle$ and $c_j^* = \langle c_1, \dots, c_n \rangle$.

New definition of perfect recall

Traditional definition: $q_0.a_1^*.q_1.a_2^*.q_2\dots q_k \approx_i r_0.b_1^*.r_1.b_2^*.r_3\dots r_k$ iff $q_j \sim_i r_j$ for $j \in \{0, \dots, k\}$.

Only looks at **states**.

New definition: two histories are equivalent for a_i if all past states are equivalent **and** a_i took the same action in both histories at every step in the past.

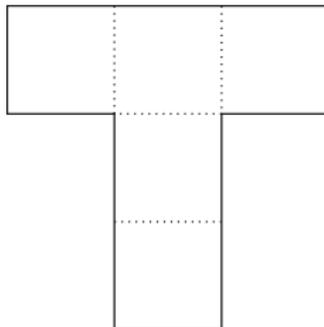
Definition (Perfect recall equivalence)

$h_1 \approx_i h_2$, iff either $h_1 = q_1$ and $h_2 = q_2$ and $q_1 \sim_i q_2$, or $h_1 = q_0.b_1^*.q_1\dots q_{j-1}.b_j^*.q_j$ and $h_2 = r_0.c_1^*.r_1\dots r_{j-1}.c_j^*.r_j$ and:

- 1 $q_0.b_1^*.q_1\dots q_{j-1} \approx_i r_0.c_1^*.r_1\dots r_{j-1}$, and
- 2 $q_j \sim_i r_j$, and
- 3 $b_i = c_i$ where $b_j^* = \langle b_1, \dots, b_n \rangle$ and $c_j^* = \langle c_1, \dots, c_n \rangle$.

Example II: Recalling actions

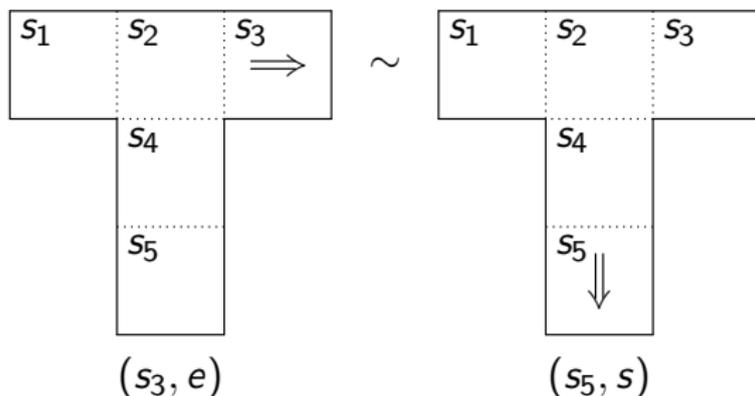
A robot is in a simple maze shaped like this:



Robot has a **position** and an **orientation**.
Robot perceives only the walls around him.

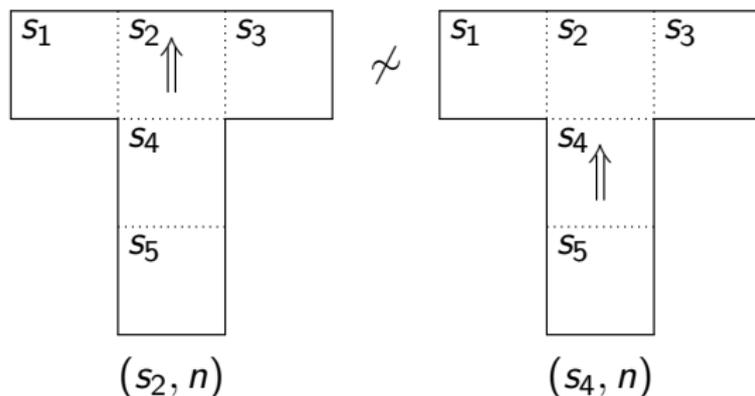
Example II: equivalence relation

So, these two states are indistinguishable for the robot:



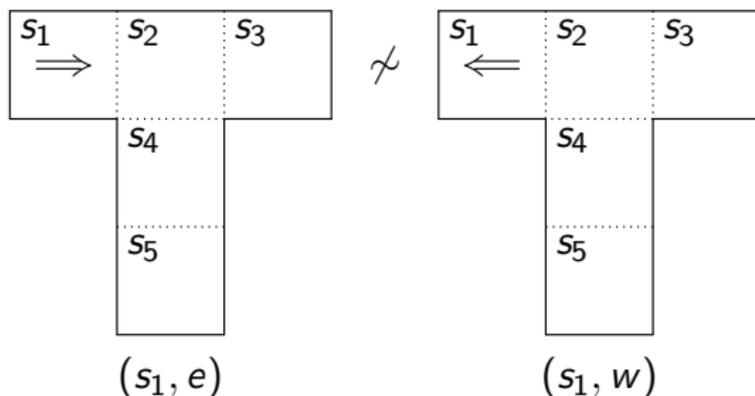
Example II: equivalence relation

But these two states are distinguishable:



Example II: equivalence relation

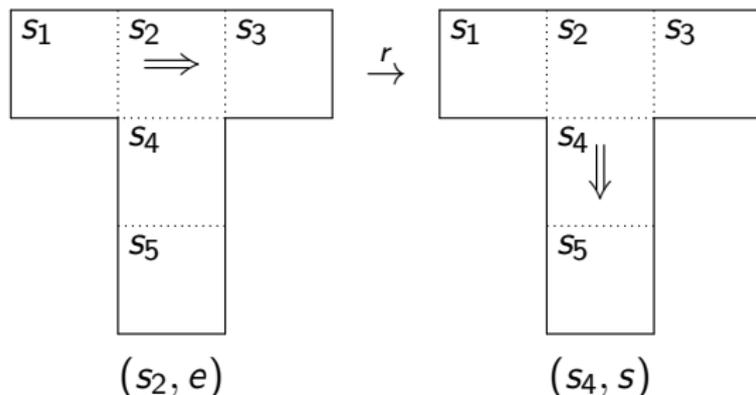
And these two states are distinguishable:



Example II: actions

The robot's actions are go left, right, forward, and back, l, r, f, b .
 The actions l, r , and b change the orientation as you would expect, but f does not change the orientation.

Every action changes the position if the space in that direction is free; otherwise, the position stays the same. E.g. right action:

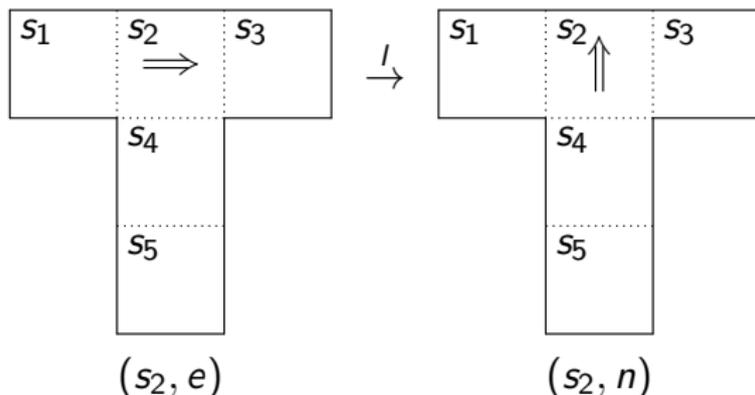


Example II: actions

The robot's actions are go left, right, forward, and back, l, r, f, b .

The actions l, r , and b change the orientation as you would expect, but f does not change the orientation.

Every action changes the position if the space in that direction is free; otherwise, the position stays the same. E.g. left action:



Contents

- 1 Introduction
- 2 Background
- 3 Memory
- 4 The Logic**
- 5 Final Remarks

Syntax

The syntax of uATEL:

$$\phi ::= p \mid \neg\phi \mid \phi \vee \phi \mid K_i\phi \mid C_A\phi \mid \langle\langle A \rangle\rangle \bigcirc \phi \mid \langle\langle A \rangle\rangle \square \phi \mid \langle\langle A \rangle\rangle \phi \mathcal{U} \phi$$

where $p \in \Pi$, $i \in \Sigma$, and $A \subseteq \Sigma$.

K_i : knowledge

C_A : common knowledge

Semantics

In order to express knowledge, we must define our semantics over (finite) *histories*.

$L, h \models p$ iff $p \in \pi(\text{last}(h))$

Semantics for $\neg\phi$, $\phi_1 \vee \phi_2$, $K_i\phi$, and $C_A\phi$ are standard.

$L, h \models \langle\langle A \rangle\rangle \bigcirc \phi$ iff \exists group strategy F_A for A such that $\forall h' \approx_A^* h$,
 $\forall \lambda \in \text{out}(h', F_A)$, $L, \lambda[0, |h'| + 1] \models \phi$,

$L, h \models \langle\langle A \rangle\rangle \square \phi$ iff \exists group strategy F_A for A such that $\forall h' \approx_A^* h$,
 $\forall \lambda \in \text{out}(h', F_A)$, $L, \lambda[0, |h'| + n] \models \phi$ for all $n \geq 0$

$L, h \models \langle\langle A \rangle\rangle \phi_1 \mathcal{U} \phi_2$ iff \exists group strategy F_A for A s.t. $\forall h' \approx_A^* h$,
 $\forall \lambda \in \text{out}(h', F_A)$, $\exists m \in \mathbb{N}$ s.t. $L, \lambda[0, |h'| + m] \models \phi_2$ and
 $\forall n \in \{0, \dots, m - 1\}$, $L, \lambda[0, |h'| + n] \models \phi_1$

Next operator

$L, h \models \langle\langle A \rangle\rangle \bigcirc \phi$ iff \exists group strategy F_A for A such that $\forall h' \approx_A^* h,$
 $\forall \lambda \in \text{out}(h', F_A), L, \lambda[0, |h'| + 1] \models \phi,$

This means that there is a successful group strategy (a set of strategies, one for each agent in the group), and it is **common knowledge** for the group that the strategy will succeed from the current history.

This corresponds to the intuitive notion of a group of agents being able to achieve a goal.

Validities

$L, h \models \langle\langle A \rangle\rangle \bigcirc \phi$ iff \exists group strategy F_A for A such that $\forall h' \approx_A^* h,$
 $\forall \lambda \in \text{out}(h', F_A), L, \lambda[0, |h'| + 1] \models \phi,$

This means that there is a successful group strategy (a set of strategies, one for each agent in the group), and it is **common knowledge** for the group that the strategy will succeed from the current history.

This corresponds to the intuitive notion of a group of agents being able to achieve a goal.

Validities

For perfect recall agents:

$$\langle\langle\Gamma\rangle\rangle\Box\phi \leftrightarrow (C_{\Gamma}\phi \wedge \langle\langle\Gamma\rangle\rangle \bigcirc \langle\langle\Gamma\rangle\rangle\Box\phi).$$

$$\langle\langle\Gamma\rangle\rangle \bigcirc \phi \leftrightarrow \langle\langle\Gamma\rangle\rangle \bigcirc C_{\Gamma}\phi \leftrightarrow C_{\Gamma}\langle\langle\Gamma\rangle\rangle \bigcirc \phi,$$

$$\langle\langle\Gamma\rangle\rangle\Box\phi \leftrightarrow \langle\langle\Gamma\rangle\rangle\Box C_{\Gamma}\phi \leftrightarrow C_{\Gamma}\langle\langle\Gamma\rangle\rangle\Box\phi,$$

and

$$\langle\langle\Gamma\rangle\rangle\phi_1\mathcal{U}\phi_2 \leftrightarrow C_{\Gamma}\langle\langle\Gamma\rangle\rangle\phi_1\mathcal{U}\phi_2.$$

but NOT

$$\langle\langle\Gamma\rangle\rangle_a\phi_1\mathcal{U}\phi_2 \rightarrow \langle\langle\Gamma\rangle\rangle_a C_{\Gamma}\phi_1\mathcal{U}C_{\Gamma}\phi_2.$$

Validities

For perfect recall agents:

$$\langle\langle \Gamma \rangle\rangle \Box \phi \leftrightarrow (C_{\Gamma} \phi \wedge \langle\langle \Gamma \rangle\rangle \bigcirc \langle\langle \Gamma \rangle\rangle \Box \phi).$$

$$\langle\langle \Gamma \rangle\rangle \bigcirc \phi \leftrightarrow \langle\langle \Gamma \rangle\rangle \bigcirc C_{\Gamma} \phi \leftrightarrow C_{\Gamma} \langle\langle \Gamma \rangle\rangle \bigcirc \phi,$$

$$\langle\langle \Gamma \rangle\rangle \Box \phi \leftrightarrow \langle\langle \Gamma \rangle\rangle \Box C_{\Gamma} \phi \leftrightarrow C_{\Gamma} \langle\langle \Gamma \rangle\rangle \Box \phi,$$

and

$$\langle\langle \Gamma \rangle\rangle \phi_1 \mathcal{U} \phi_2 \leftrightarrow C_{\Gamma} \langle\langle \Gamma \rangle\rangle \phi_1 \mathcal{U} \phi_2.$$

but NOT

$$\langle\langle \Gamma \rangle\rangle_a \phi_1 \mathcal{U} \phi_2 \rightarrow \langle\langle \Gamma \rangle\rangle_a C_{\Gamma} \phi_1 \mathcal{U} C_{\Gamma} \phi_2.$$

Validities

For perfect recall agents:

$$\langle\langle \Gamma \rangle\rangle \Box \phi \leftrightarrow (C_{\Gamma} \phi \wedge \langle\langle \Gamma \rangle\rangle \bigcirc \langle\langle \Gamma \rangle\rangle \Box \phi).$$

$$\langle\langle \Gamma \rangle\rangle \bigcirc \phi \leftrightarrow \langle\langle \Gamma \rangle\rangle \bigcirc C_{\Gamma} \phi \leftrightarrow C_{\Gamma} \langle\langle \Gamma \rangle\rangle \bigcirc \phi,$$

$$\langle\langle \Gamma \rangle\rangle \Box \phi \leftrightarrow \langle\langle \Gamma \rangle\rangle \Box C_{\Gamma} \phi \leftrightarrow C_{\Gamma} \langle\langle \Gamma \rangle\rangle \Box \phi,$$

and

$$\langle\langle \Gamma \rangle\rangle \phi_1 \mathcal{U} \phi_2 \leftrightarrow C_{\Gamma} \langle\langle \Gamma \rangle\rangle \phi_1 \mathcal{U} \phi_2.$$

but NOT

$$\langle\langle \Gamma \rangle\rangle_a \phi_1 \mathcal{U} \phi_2 \rightarrow \langle\langle \Gamma \rangle\rangle_a C_{\Gamma} \phi_1 \mathcal{U} C_{\Gamma} \phi_2.$$

Contents

- 1 Introduction
- 2 Background
- 3 Memory
- 4 The Logic
- 5 Final Remarks

Future Work

- Study the decidability of model checking for the logic,
- Complexity if decidable,
- Axiomatization?
- Abilities of subgroups of agents,
- Include memory abilities in the logical language: be able to express properties like “ a_1 is memoryless”
- Use strategy logic to discuss strategies explicitly within formulas.

Thank you

Thank you.

Related Work

van der Hoek & Wooldridge 03 (ATEL); Jamroga & van der Hoek 04 (ATOL); Schobbens 04; Jamroga & Ågotnes 07 (CSL); Ågotnes & Alechina 12 (ECL), Bulling & Jamroga 14, others. Main differences:

- Uniform vs. non-uniform strategies. Non-uniform: ATEL.
- *De re* or *de dicto* strategies. *De dicto*: ECL. *De re*: Sch04, ATOL, us. Both: CSL and BJ14.
- Different coalitional operators. ATOL can express ours and more for memoryless systems.
- Memory abilities: we allow arbitrary equivalence on histories, different memory abilities for different agents and introduce action-based definition of perfect recall. These are new.