Some Epistemic Issues in Strategy Logic

Sophia Knight, Bastien Maubert, Aniello Murano, Sasha Rubin, Francesco Belardinelli, and Alessio Lomuscio

University of Minnesota Duluth

Online Logic Seminar
24 February 2021

Outline

- Background
- ATL and ATL*, imperfect information ATL, and uniform strategies
- Strategy Logic
- Imperfect information strategy logic
- Strategy logic with uniform strategies
- What do agents know about each other's strategies?

Example

Two robots pushing a cart around a track. Robot 1 pushes clockwise and Robot 2 pushes counterclockwise. From W. Jamroga.

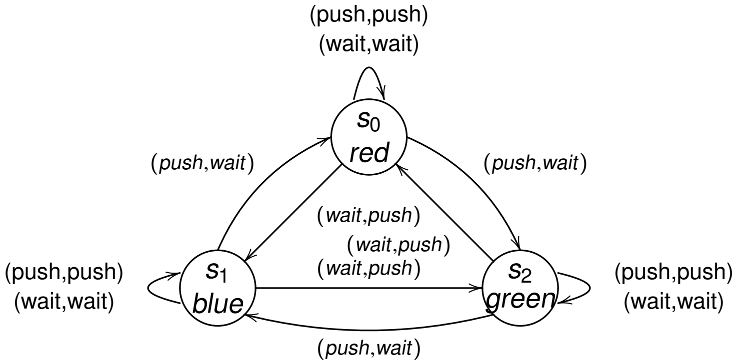

Example

What can each agent accomplish in this setting, starting from s_1?

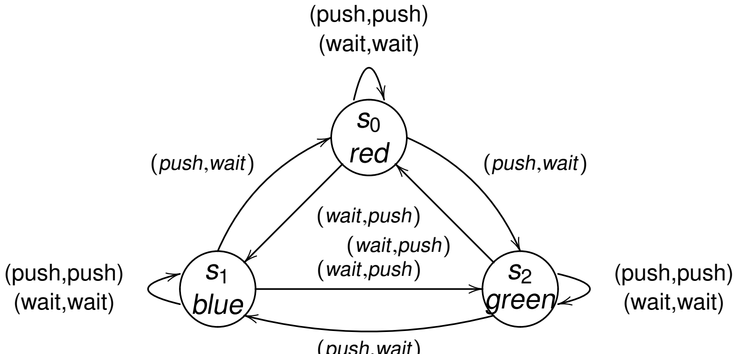

Example

What can each agent accomplish in this setting, starting from s_1?

What can they accomplish individually? What can they accomplish if they cooperate?

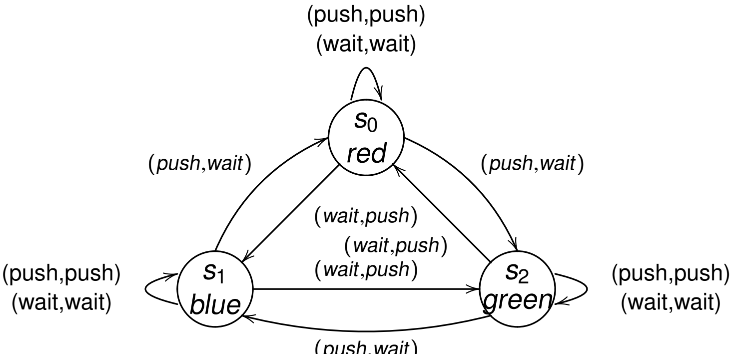

Example

What can each agent accomplish in this setting, starting from s_1?

What can they accomplish individually? What can they accomplish if they cooperate?

Can Robot 1 ensure that the system eventually reaches the red state?

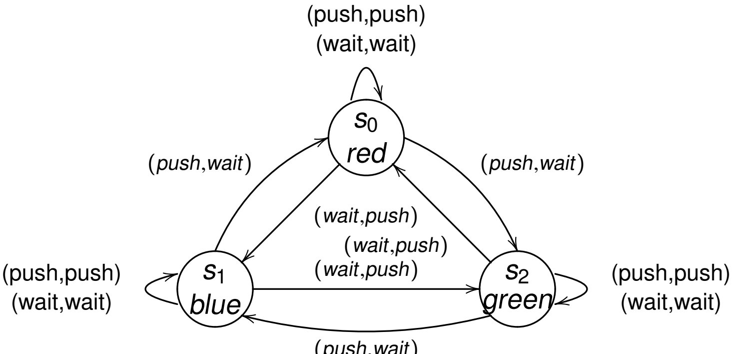

Example

What can each agent accomplish in this setting, starting from s_1?

What can they accomplish individually? What can they accomplish if they cooperate?

Can Robot 1 ensure that the system eventually reaches the red state?

Can Robot 1 ensure that the system *never* reaches the red state?

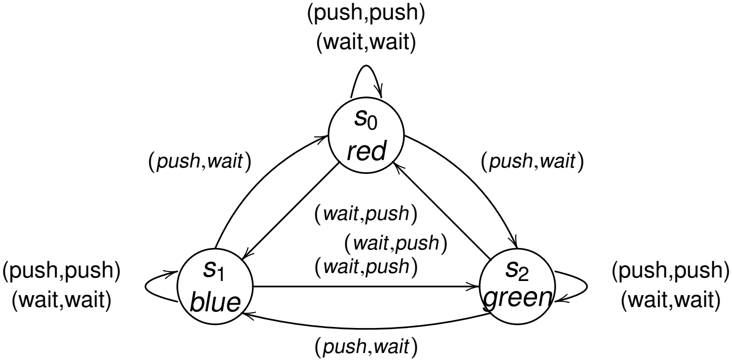

Example

What can each agent accomplish in this setting, starting from s_1?

What can they accomplish individually? What can they accomplish if they cooperate?

Can Robot 1 ensure that the system eventually reaches the red state?

Can Robot 1 ensure that the system *never* reaches the red state?

Can the two robots together ensure that the system eventually reaches the red state?

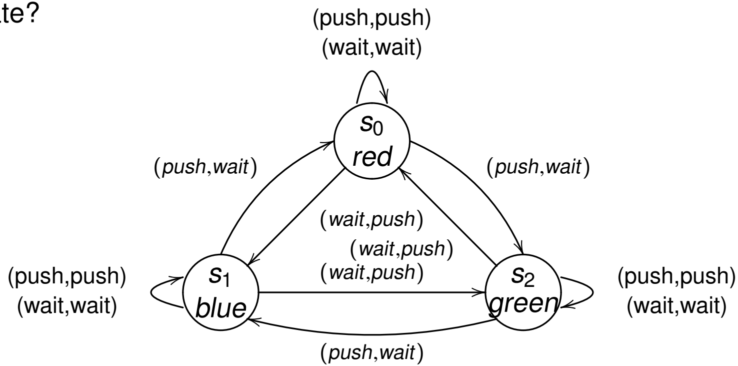

What can each agent accomplish in this setting, starting from s_1?

What can they accomplish individually? What can they accomplish if they cooperate?

What if the robots have *imperfect information* about the state or the history?

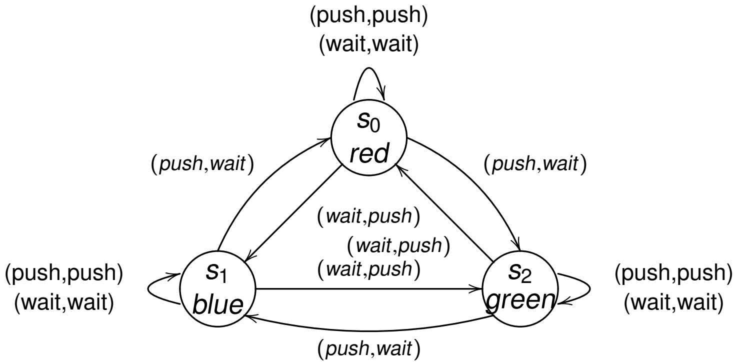

Introduction

- Goal: develop a logic to analyze strategic abilities in multi-agent systems with imperfect information.

Introduction

- Goal: develop a logic to analyze strategic abilities in multi-agent systems with imperfect information.
- Models for the logic concurrent game structures: multi-state, multi-agent systems; each agent acts to determine the next state.

Introduction

- Goal: develop a logic to analyze strategic abilities in multi-agent systems with imperfect information.
- Models for the logic concurrent game structures: multi-state, multi-agent systems; each agent acts to determine the next state.
- We are interested in imperfect information CGS's: each agent has an equivalence relation on the set of states, or histories.

Introduction

- Goal: develop a logic to analyze strategic abilities in multi-agent systems with imperfect information.
- Models for the logic concurrent game structures: multi-state, multi-agent systems; each agent acts to determine the next state.
- We are interested in imperfect information CGS's: each agent has an equivalence relation on the set of states, or histories.
- Before we discuss the logics, there are a few key concepts to clarify.

Background: Histories and Strategies

A strategy for an agent is a function that tells them what move to make depending on the history of the system.

Background: Histories and Strategies

A strategy for an agent is a function that tells them what move to make depending on the history of the system.

E.g. “always push” or ”always wait in the red state, otherwise push” or “alternate between pushing and waiting”

Background: Histories and Strategies

A strategy for an agent is a function that tells them what move to make depending on the history of the system.

E.g. “always push” or ”always wait in the red state, otherwise push” or “alternate between pushing and waiting”

Given a starting state and a strategy for each agent, the outcome of the game is determined- a single sequence of states and actions, called a history.

Background: Histories and Strategies

A strategy for an agent is a function that tells them what move to make depending on the history of the system.

E.g. “always push” or ”always wait in the red state, otherwise push” or “alternate between pushing and waiting”

Given a starting state and a strategy for each agent, the outcome of the game is determined- a single sequence of states and actions, called a history.

For example: $s_0.(push, push).s_0.(push, wait).s_2$.

Background: Histories and Strategies

A strategy for an agent is a function that tells them what move to make depending on the history of the system.

E.g. “always push” or ”always wait in the red state, otherwise push” or “alternate between pushing and waiting”

Given a starting state and a strategy for each agent, the outcome of the game is determined- a single sequence of states and actions, called a history.

For example: $s_0.(push, push).s_0.(push, wait).s_2$.

On the other hand, a strategy for a subset of the agents determines a set of histories: more than one possible outcome.

Concurrent Game Structures with Imperfect Information

Each agent has an equivalence relation $\sim_i$ on the set of states.

If $s \sim_i t$ then agent i cannot distinguish state s from state t.

Concurrent Game Structures with Imperfect Information

Each agent has an equivalence relation $\sim_i$ on the set of states.

If $s \sim_i t$ then agent i cannot distinguish state s from state t.

Example: Robot 1 is colorblind and cannot distinguish green from red. Robot 2 has perfect information.

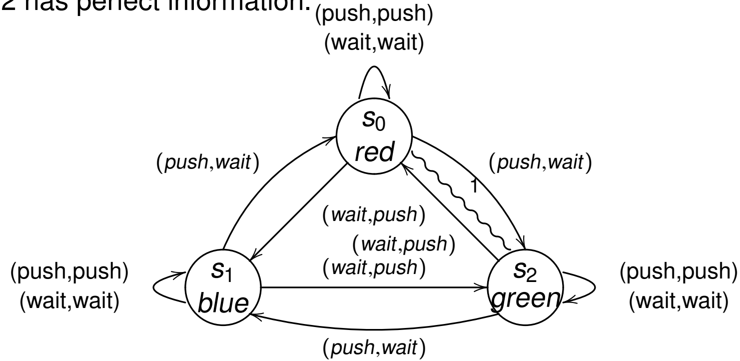

Uniform Strategies

A strategy σ is *uniform* for agent i if whenever $h_1 \sim_i h_2$, $\sigma(h_1) = \sigma(h_2)$.

Uniform Strategies

A strategy σ is *uniform* for agent *i* if whenever $h_1 \sim_i h_2$, $\sigma(h_1) = \sigma(h_2)$.

Agent's action is based on their knowledge.

Note: everything I will talk about today is orthogonal to issues of memory, so for convenience I usually assume agents have perfect recall.

Uniform strategy: example

E.g. for robot 1: $\sigma(s_0) = \sigma(s_2) = push$.

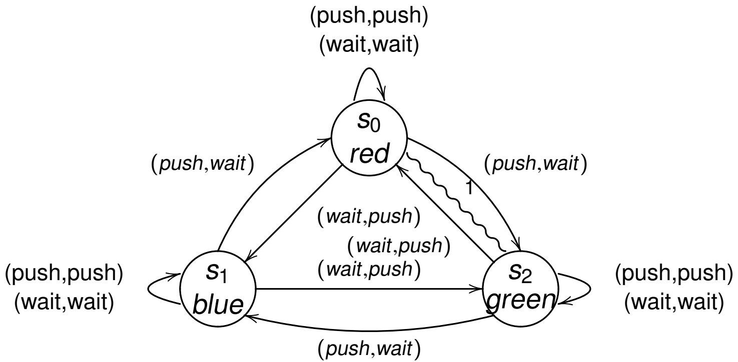

Alternating-time Temporal Logic: Alur, Henzinger, and Kupferman, 1997.

A logic of strategic abilities in concurrent game structures.

ATL Syntax:

$$\phi \quad ::== \quad p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \langle\langle A \rangle\rangle X\phi \mid \langle\langle A \rangle\rangle \Box\phi \mid \langle\langle A \rangle\rangle \phi_1 U \phi_2$$

A is a group of agents.

$\langle\langle A \rangle\rangle X\phi$ says "A has a joint strategy to make ϕ true at the neXt step."
$\langle\langle A \rangle\rangle \Box\phi$ says "A has a strategy to make ϕ always true,"
$\langle\langle A \rangle\rangle \phi_1 U \phi_2$ says "A has a strategy so that ϕ_1 will remain true until ϕ_2 eventually becomes true."

ATL: Example

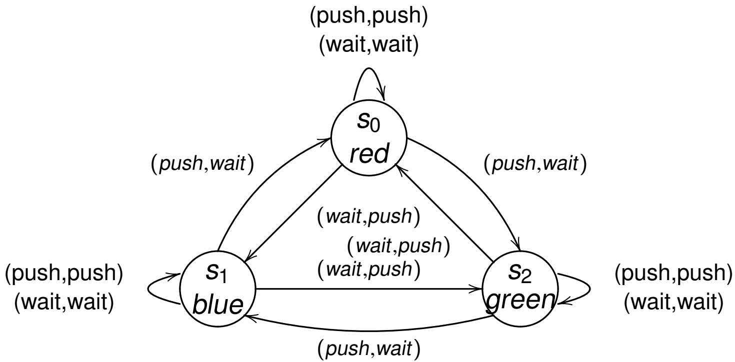

$s_1 \models \langle\!\langle 1 \rangle\!\rangle X(\neg green)$

$s_1 \models \neg\langle\!\langle 1 \rangle\!\rangle X(red)$

$s_1 \models \langle\!\langle 1, 2 \rangle\!\rangle \Box blue$

$s_1 \models \langle\!\langle 1 \rangle\!\rangle [\Box(blue \lor green)]$

$s_1 \models \neg[(\langle\!\langle 1 \rangle\!\rangle \Box blue) \lor (\langle\!\langle 1 \rangle\!\rangle \Box green)]$

ATL: Semantics

The semantics for perfect information ATL are as you would expect.

- $q \models p$, for propositions $p \in \Pi$, iff $p \in \pi(q)$.
- $q \models \neg\varphi$ iff $q \not\models \varphi$.
- $q \models \varphi_1 \vee \varphi_2$ iff $q \models \varphi_1$ or $q \models \varphi_2$.
- $q \models \langle\langle A\rangle\rangle \bigcirc \varphi$ iff there exists a set F_A of strategies, one for each player in A, such that for all computations $\lambda \in out(q, F_A)$, we have $\lambda[1] \models \varphi$.
- $q \models \langle\langle A\rangle\rangle \Box \varphi$ iff there exists a set F_A of strategies, one for each player in A, such that for all computations $\lambda \in out(q, F_A)$ and all positions $i \geq 0$, we have $\lambda[i] \models \varphi$.
- $q \models \langle\langle A\rangle\rangle \varphi_1 \mathcal{U} \varphi_2$ iff there exists a set F_A of strategies, one for each player in A, such that for all computations $\lambda \in out(q, F_A)$, there exists a position $i \geq 0$ such that $\lambda[i] \models \varphi_2$ and for all positions $0 \leq j < i$, we have $\lambda[j] \models \varphi_1$.

From "Alternating-time temporal logic," Alur, Henzinger and Kupferman, JACM 2002.

ATL* decouples $\langle\langle A\rangle\rangle$ operator from X, $\Box$, and $\mathcal{U}$ operators.

Imperfect Information Variants of ATL

How to incorporate imperfect information into ATL?

Imperfect Information Variants of ATL

How to incorporate imperfect information into ATL?
Many different proposals, based on

- uniform vs. non-uniform strategies
- de re vs. de dicto
- agents' cooperation and knowledge-sharing
- memory

Imperfect Information Variants of ATL

How to incorporate imperfect information into ATL?
Many different proposals, based on

- uniform vs. non-uniform strategies
- de re vs. de dicto
- agents' cooperation and knowledge-sharing
- memory

Generally, semantics look like this:
$L, q \models \langle\langle\Gamma\rangle\rangle X\phi$ iff there exists a group strategy F_Γ such that
$\forall\lambda \in out(q, F_\Gamma), L, \lambda[1] \models \phi$ (de dicto)
or
$L, q \models \langle\langle\Gamma\rangle\rangle X\phi$ iff there exists a group strategy F_Γ such that $\forall q' \sim^{?} q$,
$\forall\lambda \in out(q', F_\Gamma), L, \lambda[1] \models \phi$ (de re)

Background: Strategy Logic

- Introduced in Chatterjee, Henzinger and Piterman 2007; more general version Mogavero, Murano, and Vardi 2010.

Background: Strategy Logic

- Introduced in Chatterjee, Henzinger and Piterman 2007; more general version Mogavero, Murano, and Vardi 2010.
- Reasons explicitly about strategies, unlike ATL and ATL*.

Background: Strategy Logic

- Introduced in Chatterjee, Henzinger and Piterman 2007; more general version Mogavero, Murano, and Vardi 2010.
- Reasons explicitly about strategies, unlike ATL and ATL*.
- New formulas: x is a strategy,
  $\langle\langle x \rangle\rangle \phi$: there exists a strategy x such that ϕ holds,
  and $(a, x)\phi$: if agent a follows strategy x, ϕ holds.

Strategy Logic: Syntax and Semantics

Syntax:

$$\phi ::= p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid X\phi \mid \phi_1 U \phi_2 \mid \langle\langle x \rangle\rangle \phi \mid (a, x)\phi$$

Where $p \in Prop$, $x \in Var$ and $a \in Ag$.

Semantics:

G a CGS, χ an assignment including free variables, s a state.

$p, \neg\phi, \phi_1 \wedge \phi_2$	standard
$X\phi$ and $\phi_1 U \phi_2$	as LTL (χ gives free agents strategies - single path)
$(G, \chi, s) \models \langle\langle x \rangle\rangle \phi$	iff $\exists f \in Str$ s.t. $(G, \chi[x \mapsto f], s) \models \phi$
$(G, \chi, s) \models (a, x)\phi$	iff $(G, \chi[a \mapsto \chi(x)], s) \models \phi$

Strategy Logic: formulas

$\langle\langle x \rangle\rangle (a, x) \phi$ means "there is some strategy x that a can follow to make ϕ true."

Strategy Logic: formulas

$\langle\langle x \rangle\rangle (a, x) \phi$ means "there is some strategy x that a can follow to make ϕ true."

Like $\langle\langle a \rangle\rangle \phi$ in ATL*.

Strategy Logic: formulas

$\langle\langle x \rangle\rangle (a, x)\phi$ means “there is some strategy x that a can follow to make ϕ true.”

Like $\langle\langle a \rangle\rangle \phi$ in ATL*.

$\langle\langle x \rangle\rangle (a, x)(b, x)\phi$ means “there is some strategy x that a and b can *both* follow simultaneously to make ϕ true.”

Strategy Logic: formulas

$\langle\langle x \rangle\rangle (a, x)\phi$ means “there is some strategy x that a can follow to make ϕ true.”

Like $\langle\langle a \rangle\rangle \phi$ in ATL*.

$\langle\langle x \rangle\rangle (a, x)(b, x)\phi$ means “there is some strategy x that a and b can *both* follow simultaneously to make ϕ true.”

Not expressible in ATL*.

Strategy Logic: formulas

$\langle\langle x\rangle\rangle(a,x)\phi$ means "there is some strategy x that a can follow to make ϕ true."

Like $\langle\langle a\rangle\rangle\phi$ in ATL*.

$\langle\langle x\rangle\rangle(a,x)(b,x)\phi$ means "there is some strategy x that a and b can *both* follow simultaneously to make ϕ true."

Not expressible in ATL*.

$\langle\langle x\rangle\rangle[\![y]\!]\langle\langle z\rangle\rangle(a,x)(b,y)(c,z)\phi$

Strategy Logic: formulas

$\langle\langle x \rangle\rangle (a, x)\phi$ means "there is some strategy x that a can follow to make ϕ true."

Like $\langle\langle a \rangle\rangle \phi$ in ATL*.

$\langle\langle x \rangle\rangle (a, x)(b, x)\phi$ means "there is some strategy x that a and b can *both* follow simultaneously to make ϕ true."

Not expressible in ATL*.

$\langle\langle x \rangle\rangle [\![y]\!] \langle\langle z \rangle\rangle (a, x)(b, y)(c, z)\phi$

There is a strategy for a so that no matter what b does there is a strategy for c that makes ϕ true.

Strategy Logic vs. ATL*

- ATL* can be translated into SL, specifically into "one-goal" fragment of SL (SL[1G]).
- SL can express the existence of Nash equilibria. Not expressible in ATL*.
- SL is more expressive than ATL*, but satisfiability in SL is undecidable.
- However, satisfiability for SL[1G] is decidable in 2EXPTIME, like ATL* (Mogavero, Murano, Perelli and Vardi, CONCUR '12).
- Thus, translation of ATL* into SL is no worse computationally than ATL*, and SL[1G] is strictly more expressive (i.e. sharing strategies).

Imperfect Information SL

Goal: define a version of SL where the models are CGS's with imperfect information.

Imperfect Information SL

Goal: define a version of SL where the models are CGS's with imperfect information.

Difficulties:

- How to ensure that strategies are uniform when they can be assigned to anyone?

Imperfect Information SL

Goal: define a version of SL where the models are CGS's with imperfect information.

Difficulties:

- How to ensure that strategies are uniform when they can be assigned to anyone?
- What should agents know about one another's strategies?

Uniform strategies

In earlier work on iiSL, there are three approaches to this problem:

Uniform strategies

In earlier work on iiSL, there are three approaches to this problem:

1. Merging quantification and binding: $\langle\langle x \rangle\rangle_a$ means “there exists a strategy x for agent a.”

Uniform strategies

In earlier work on iiSL, there are three approaches to this problem:

1. Merging quantification and binding: $\langle\langle x \rangle\rangle_a$ means “there exists a strategy x for agent a.”
2. Explicit observations on strategies: $\langle\langle x \rangle\rangle^o$ means “strategy x is uniform w.r.t. observations o.”

Uniform strategies

In earlier work on iiSL, there are three approaches to this problem:

1. Merging quantification and binding: $\langle\langle x \rangle\rangle_a$ means "there exists a strategy x for agent a."
2. Explicit observations on strategies: $\langle\langle x \rangle\rangle^o$ means "strategy x is uniform w.r.t. observations o."
3. Syntactic uniformity: in $\langle\langle x \rangle\rangle\phi$, ϕ is required to be uniform for every agent that could play x in ϕ.

Problem

The first two solutions reduce expressivity. The third solution looks good at first, but consider $\phi = \langle\langle s_1 \rangle\rangle [\![s_2]\!] ((a, s_1)(b, s_2)XXp \vee (a, s_2)(b, s_1)XXq)$

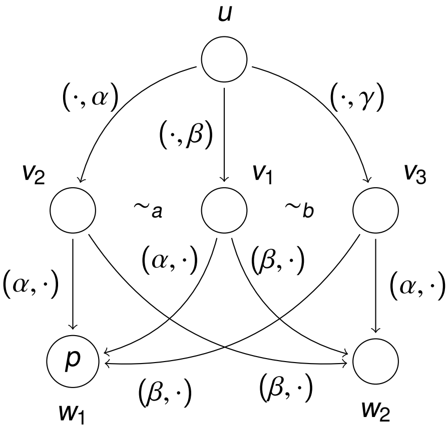

Solution

The solution turns out to be simple and elegant: check uniformity at the time the strategy is assigned.

$$(G,\chi,s) \models (a,x)\phi \qquad \text{iff } (G,\chi[a \mapsto \chi(x)],s) \models \phi \text{ and } \chi(x) \text{ is uniform for } a$$

Agents' knowledge about each other's strategies

Earlier work on SL either assumes that agents all know each other's strategies (informed semantics) or no one knows anyone's strategy (uninformed semantics).

Agents' knowledge about each other's strategies

Earlier work on SL either assumes that agents all know each other's strategies (informed semantics) or no one knows anyone's strategy (uninformed semantics).

Some problems, e.g. distributed synthesis, are decidable with the uninformed semantics but undecidable with the informed semantics.

Our goal is to make this information explicit and more fine-grained.

Example: Coin game

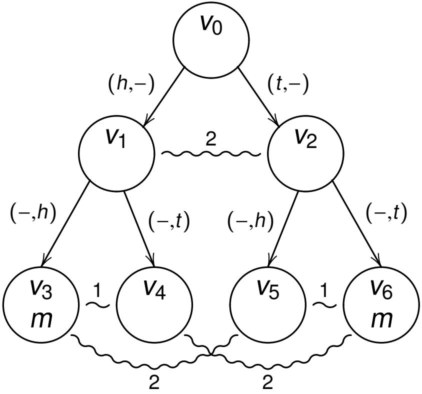

If 2 knows 1's strategy, 2 can ensure that m becomes true, or not (and vice versa).

Knowledge about strategies

How do we incorporate agents' knowledge about one another's strategies into the logic?

Knowledge about strategies

How do we incorporate agents' knowledge about one another's strategies into the logic?

1. Make knowledge explicit: For each a, $\mathcal{A}_a$ is the group of agents whose strategies a knows.

Knowledge about strategies

How do we incorporate agents' knowledge about one another's strategies into the logic?

1. Make knowledge explicit: For each a, $\mathcal{A}_a$ is the group of agents whose strategies a knows.
2. A new equivalence relation: For strategy assignment χ, and histories ρ, ρ', we say $\rho \sim_a^{\chi} \rho'$ iff
    1. $\rho \sim_a \rho'$
    2. ρ and ρ' are both consistent with $\chi \upharpoonright \mathcal{A}_a$.

Knowledge about strategies

How do we incorporate agents' knowledge about one another's strategies into the logic?

1. Make knowledge explicit: For each a, $\mathcal{A}_a$ is the group of agents whose strategies a knows.
2. A new equivalence relation: For strategy assignment χ, and histories ρ, ρ', we say $\rho \sim_a^{\chi} \rho'$ iff
   1. $\rho \sim_a \rho'$
   2. ρ and ρ' are both consistent with $\chi \upharpoonright \mathcal{A}_a$.
3. New definition of uniformity: require strategies to be uniform w.r.t. $\sim_a^{\chi}$.

New definition of uniformity

Now, if $1 \in \mathcal{A}_2$, the strategy $v_1 \mapsto h$ and $v_2 \mapsto t$ is a uniform strategy, because $v_0.v_1$ and $v_0.v_2$ cannot both be consistent with 1's strategy, so they are distinguishable.

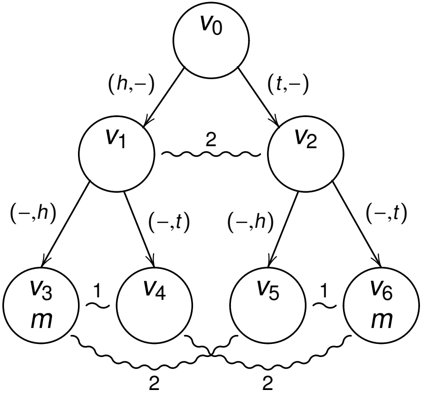

New definition of uniformity

Now the following formula is true if $1 \in \mathcal{A}_2$.

$$\gamma = \langle\langle s_2 \rangle\rangle [\![s_1]\!] (1, s_1)(2, s_2) XXm$$

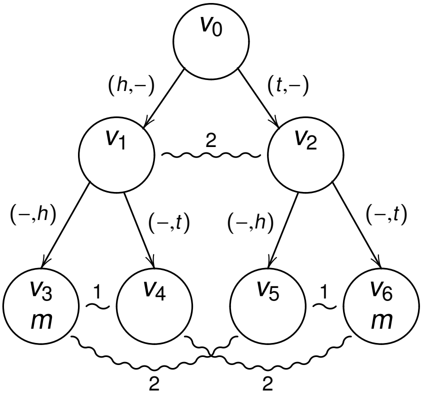

New semantics of Knowledge

Our knowledge operator also must take into account knowledge about strategies"

$$\mathcal{G} \models K_a \phi \quad \text{iff} \quad \begin{array}{l} \forall \rho' \in \mathit{Hist} \text{ s.t. } \mathit{rho} \sim_a^{\chi} \rho' \text{ and} \\ \forall \chi' \text{ s.t. } \chi' \restriction \mathcal{A}_a = \chi \restriction \mathcal{A}_a \\ \mathcal{G}, \chi', \rho' \models \phi \end{array}$$

Formulas

This definition of the semantics makes the following formula true:

$$\alpha = [\![s_1, s_2]\!](1, s_1)(2, s_2)XX(K_1 m \vee K_1 \neg m)$$

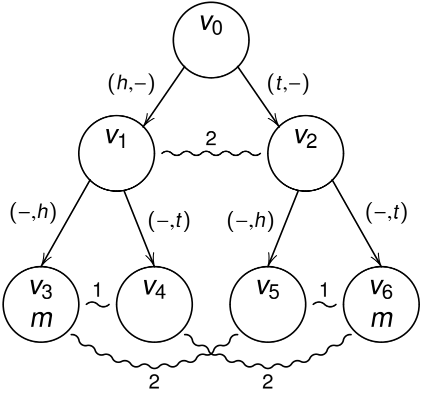

The whole semantics

Definition 3. Let $\mathcal{G} = (\text{Ac}, V, E, \ell, V_\iota, \{\sim\}_{a \in \text{Ag}}, \{\mathcal{A}_a\}_{a \in \text{Ag}})$ be a CGS, χ an assignment, ρ a history and φ an $\mathsf{SLK}_{\text{inf}}[\text{BG}]$-formula. The semantics $\mathcal{G}, \chi, \rho \models \varphi$ is defined as follows:

$\mathcal{G}, \chi, \rho \models p$	if	$p \in \ell(\text{last}(\rho))$
$\mathcal{G}, \chi, \rho \models \neg\varphi$	if	$\mathcal{G}, \chi, \rho \not\models \varphi$
$\mathcal{G}, \chi, \rho \models \varphi \vee \varphi'$	if	$\mathcal{G}, \chi, \rho \models \varphi$ or $\mathcal{G}, \chi, \rho \models \varphi'$
$\mathcal{G}, \chi, \rho \models \exists s.\varphi$	if	$\exists \sigma \in \mathit{Str}$ s.t. $\mathcal{G}, \chi[s \mapsto \sigma], \rho \models \varphi$
$\mathcal{G}, \chi, \rho \models (\text{Ag}, \bar{s})\varphi$	if	for all a, $\chi'(a)$ is $\sim_a^{\chi'}$ -uniform and $\mathcal{G}, \chi', \rho \models \varphi$, where $\chi' = \chi[\text{Ag} \mapsto \chi(\bar{s})]$
$\mathcal{G}, \chi, \rho \models K_a\varphi$	if	for all $\rho' \in \text{Hist}$ s.t. $\rho \sim_a^{\chi} \rho'$ and for all χ' s.t. $\chi' \sim_a \chi$, $\mathcal{G}, \chi', \rho' \models \varphi$

and, writing $\pi = \text{Out}(\chi, \rho)$:

$\mathcal{G}, \chi, \rho \models X\varphi$	if	$\mathcal{G}, \chi, \pi_{\leq \lvert\rho\rvert+1} \models \varphi$
$\mathcal{G}, \chi, \rho \models \varphi U \varphi'$	if	$\exists i \geq 0$ s.t. $\mathcal{G}, \chi, \pi_{\leq \lvert\rho\rvert+i} \models \varphi'$, and $\forall j$ s.t. $0 \leq j < i$, $\mathcal{G}, \chi, \pi_{\leq \lvert\rho\rvert+j} \models \varphi$

Model checking results

Theorem Model checking SLK is undecidable, even with hierarchical knowledge.

Theorem Model checking SLK with public actions is decidable.

Future Work

- Agents who can change their strategies.
- Non-deterministic strategies
- Better understanding of higher-order knowledge
- Model interesting examples, like Hanabi

The End

Thank you!

Questions?