Coalitional Games

Stéphane Airiau and Wojtek Jamroga

Stéphane: ILLC @ University of Amsterdam
Wojtek: ICR @ University of Luxembourg

European Agent Systems Summer School
Torino, Italy, September 2009

Part 2. Reasoning about Coalitions

Reasoning about Coalitions

Outline

- In the previous chapter, we showed how coalitions can be rationally formed

Outline

- In the previous chapter, we showed how coalitions can be rationally formed
- In this chapter, we show how one can use modal logic to reason about their play and their outcome.

2.1 Modal Logic

Why logic at all?

- framework for thinking about systems,
- makes one realise the implicit assumptions,
- . . . and then we can:
- investigate them, accept or reject them,
- relax some of them and still use a part of the formal and conceptual machinery;

Why logic at all?

- framework for thinking about systems,
- makes one realise the implicit assumptions,
- . . . and then we can:
- investigate them, accept or reject them,
- relax some of them and still use a part of the formal and conceptual machinery;

- reasonably expressive but simpler and more rigorous than the full language of mathematics.

Why logic at all?

- Verification: check specification against implementation
- Executable specifications
- Planning as model checking

Why logic at all?

- Verification: check specification against implementation
- Executable specifications
- Planning as model checking

- Game solving, mechanism design, and reasoning about games have natural interpretation as logical problems

Modal logic is an extension of classical logic by new connectives $\Box$ and $\Diamond$: necessity and possibility.

Modal logic is an extension of classical logic by new connectives $\Box$ and $\Diamond$: necessity and possibility.

- "$\Box p$ is true" means p is necessarily true, i.e. true in every possible scenario,
- "$\Diamond p$ is true" means p is possibly true, i.e. true in at least one possible scenario.

Various modal logics:

- knowledge → epistemic logic,
- beliefs → doxastic logic,
- obligations → deontic logic,
- actions → dynamic logic,
- time → temporal logic,

- and combinations of the above
  Most famous multimodal logic: BDI logic of beliefs, desires, intentions (and time)

Definition 2.1 (Kripke Semantics)

Kripke model (possible world model):

$$M = \langle \mathcal{W}, R, \pi \rangle,$$

- $\mathcal{W}$ is a set of possible worlds
- $R \subseteq \mathcal{W} \times \mathcal{W}$ is an accessibility relation
- $\pi : \mathcal{W} \rightarrow \mathcal{P}(\Pi)$ is a valuation of propositions.

Definition 2.1 (Kripke Semantics)

Kripke model (possible world model):

$$M = \langle \mathcal{W}, R, \pi \rangle,$$

- $\mathcal{W}$ is a set of possible worlds
- $R \subseteq \mathcal{W} \times \mathcal{W}$ is an accessibility relation
- $\pi : \mathcal{W} \to \mathcal{P}(\Pi)$ is a valuation of propositions.

$M, w \models \Box\varphi$ iff for every $w' \in \mathcal{W}$ with wRw' we have that $M, w' \models \varphi$.

An Example

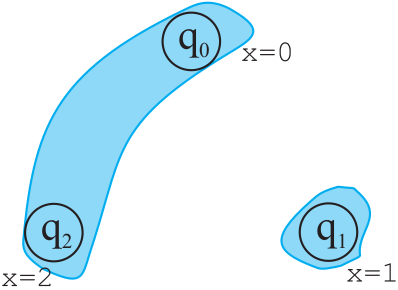

An Example

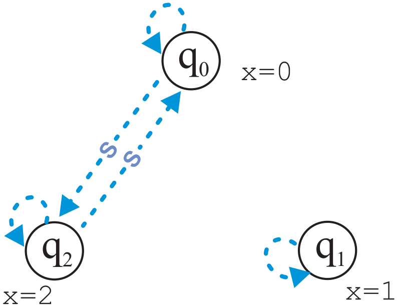

An Example

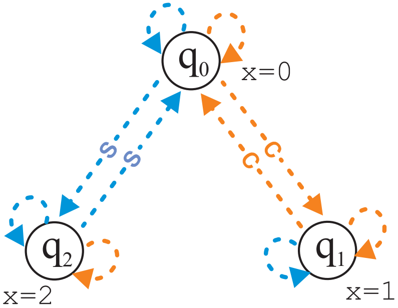

An Example

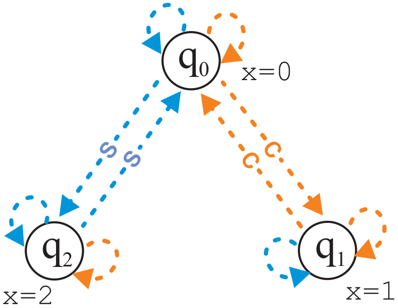

$$x \doteq 1 \rightarrow K_s x \doteq 1$$

2.2 ATL

ATL: What Agents Can Achieve

- ATL: Agent Temporal Logic [Alur et al. 1997]
- Temporal logic meets game theory
- Main idea: cooperation modalities

ATL: What Agents Can Achieve

- ATL: Agent Temporal Logic [Alur et al. 1997]
- Temporal logic meets game theory
- Main idea: cooperation modalities

$\langle\!\langle A \rangle\!\rangle \Phi$: coalition A has a collective strategy to enforce Φ

- $\langle\langle jamesbond \rangle\rangle \Diamond$win:
  "James Bond has an infallible plan to eventually win"

- $\langle\langle jamesbond \rangle\rangle \Diamond$win:
  "James Bond has an infallible plan to eventually win"
- $\langle\langle jamesbond, bondsgirl \rangle\rangle$fun $\mathcal{U}$ shot:
  "James Bond and his girlfriend are able to have fun until someone shoots at them"

- $\langle\langle jamesbond\rangle\rangle\Diamond$win:
  "James Bond has an infallible plan to eventually win"
- $\langle\langle jamesbond, bondsgirl\rangle\rangle$fun $\mathcal{U}$ shot:
  "James Bond and his girlfriend are able to have fun until someone shoots at them"

- "Vanilla" ATL: every temporal operator preceded by exactly one cooperation modality;
- ATL*: no syntactic restrictions;

ATL Models: Concurrent Game Structures

- Agents, actions, transitions, atomic propositions
- Atomic propositions + interpretation
- Actions are abstract

Definition 2.2 (Concurrent Game Structure)

A concurrent game structure is a tuple
$M = \langle \mathbb{A}\mathrm{gt}, St, \pi, Act, d, o \rangle$, where:

Definition 2.2 (Concurrent Game Structure)

A concurrent game structure is a tuple $M = \langle \mathbb{A}\mathrm{gt}, St, \pi, Act, d, o \rangle$, where:

- $\mathbb{A}\mathrm{gt}$: a finite set of all agents

Definition 2.2 (Concurrent Game Structure)

A concurrent game structure is a tuple $M = \langle \mathbb{A}\mathrm{gt}, St, \pi, Act, d, o \rangle$, where:

- $\mathbb{A}\mathrm{gt}$: a finite set of all agents
- St: a set of states

Definition 2.2 (Concurrent Game Structure)

A concurrent game structure is a tuple
$M = \langle \mathbb{A}\mathrm{gt}, St, \pi, Act, d, o \rangle$, where:

- $\mathbb{A}\mathrm{gt}$: a finite set of all agents
- St: a set of states
- π: a valuation of propositions

Definition 2.2 (Concurrent Game Structure)

A concurrent game structure is a tuple
$M = \langle \mathbb{A}\mathrm{gt}, St, \pi, Act, d, o \rangle$, where:

- $\mathbb{A}\mathrm{gt}$: a finite set of all agents
- St: a set of states
- π: a valuation of propositions
- Act: a finite set of (atomic) actions

Definition 2.2 (Concurrent Game Structure)

A concurrent game structure is a tuple $M = \langle \mathbb{A}\mathrm{gt}, St, \pi, Act, d, o \rangle$, where:

- $\mathbb{A}\mathrm{gt}$: a finite set of all agents
- St: a set of states
- π: a valuation of propositions
- Act: a finite set of (atomic) actions
- $d : \mathbb{A}\mathrm{gt} \times St \to \mathcal{P}(Act)$ defines actions available to an agent in a state
- o: a deterministic transition function that assigns outcome states $q' = o(q, \alpha_1, \ldots, \alpha_k)$ to states and tuples of actions

Example: Robots and Carriage

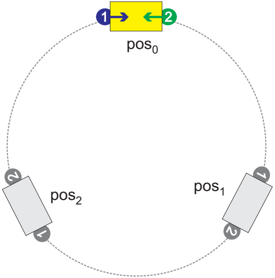

Example: Robots and Carriage

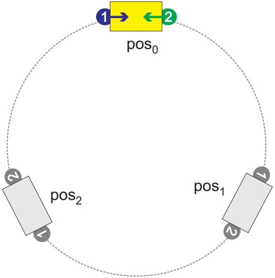

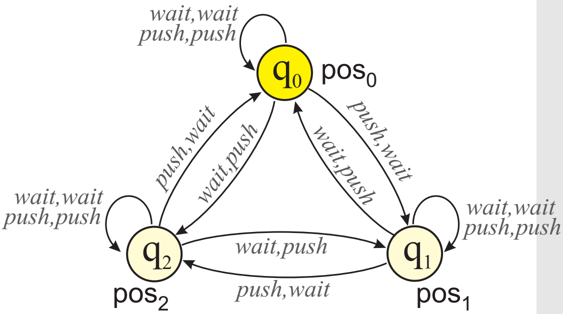

Definition 2.3 (Strategy)

A strategy is a conditional plan.

Definition 2.3 (Strategy)

A strategy is a conditional plan.
We represent strategies by functions $s_a : St \rightarrow Act$.

Definition 2.3 (Strategy)

A strategy is a conditional plan.
We represent strategies by functions $s_a : St \rightarrow Act$.

Function $\mathbf{out(q, S_A)}$ returns the set of all paths that may result from agents A executing strategy S_A from state q onward.

Definition 2.4 (Semantics of ATL)

$M, q \models \langle\langle A \rangle\rangle \Phi$ iff there is a collective strategy S_A such that, for every path $\lambda \in out(q, S_A)$, we have $M, \lambda \models \Phi$.

Definition 2.4 (Semantics of ATL)

$M, q \models \langle\langle A \rangle\rangle \Phi$ iff there is a collective strategy S_A such that, for every path $\lambda \in out(q, S_A)$, we have $M, \lambda \models \Phi$.

$M, \lambda \models \bigcirc \varphi$ iff $M, \lambda[1] \models \varphi$;

Definition 2.4 (Semantics of ATL)

$M, q \models \langle\langle A \rangle\rangle \Phi$	iff there is a collective strategy S_A such that, for every path $\lambda \in out(q, S_A)$, we have $M, \lambda \models \Phi$.
$M, \lambda \models \bigcirc \varphi$	iff $M, \lambda[1] \models \varphi$;
$M, \lambda \models \Diamond \varphi$	iff $M, \lambda[i] \models \varphi$ for some $i \geq 0$;

Definition 2.4 (Semantics of ATL)

$M, q \models \langle\langle A \rangle\rangle \Phi$ iff there is a collective strategy S_A such that, for every path $\lambda \in out(q, S_A)$, we have $M, \lambda \models \Phi$.

$M, \lambda \models \bigcirc\varphi$ iff $M, \lambda[1] \models \varphi$;
$M, \lambda \models \Diamond\varphi$ iff $M, \lambda[i] \models \varphi$ for some $i \geq 0$;
$M, \lambda \models \Box\varphi$ iff $M, \lambda[i] \models \varphi$ for all $i \geq 0$;

Definition 2.4 (Semantics of ATL)

$M, q \models \langle\langle A \rangle\rangle \Phi$ iff there is a collective strategy S_A such that, for every path $\lambda \in out(q, S_A)$, we have $M, \lambda \models \Phi$.

$M, \lambda \models \bigcirc\varphi$ iff $M, \lambda[1] \models \varphi$;

$M, \lambda \models \Diamond\varphi$ iff $M, \lambda[i] \models \varphi$ for some $i \geq 0$;

$M, \lambda \models \Box\varphi$ iff $M, \lambda[i] \models \varphi$ for all $i \geq 0$;

$M, \lambda \models \varphi\,\mathcal{U}\,\psi$ iff $M, \lambda[i] \models \psi$ for some $i \geq 0$, and $M, \lambda[j] \models \varphi$ forall $0 \leq j \leq i$.

Definition 2.4 (Semantics of ATL)

$M, q \models p$	iff p is in $\pi(q)$;
$M, q \models \varphi \wedge \psi$	iff $M, q \models \varphi$ and $M, q \models \psi$;
$M, q \models \langle\langle A \rangle\rangle \Phi$	iff there is a collective strategy S_A such that, for every path $\lambda \in out(q, S_A)$, we have $M, \lambda \models \Phi$.
$M, \lambda \models \bigcirc \varphi$	iff $M, \lambda[1] \models \varphi$;
$M, \lambda \models \Diamond \varphi$	iff $M, \lambda[i] \models \varphi$ for some $i \geq 0$;
$M, \lambda \models \Box \varphi$	iff $M, \lambda[i] \models \varphi$ for all $i \geq 0$;
$M, \lambda \models \varphi \, \mathcal{U} \, \psi$	iff $M, \lambda[i] \models \psi$ for some $i \geq 0$, and $M, \lambda[j] \models \varphi$ forall $0 \leq j \leq i$.

Example: Robots and Carriage

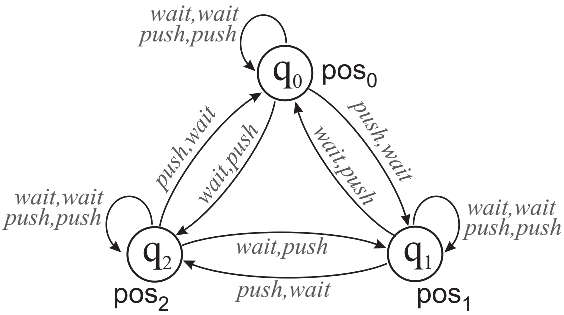

$$pos_0 \rightarrow \langle\langle 1 \rangle\rangle \Box \neg pos_1$$

Example: Robots and Carriage

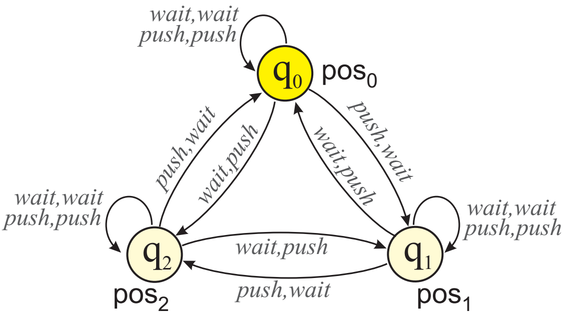

$$pos_0 \longrightarrow \langle\langle 1 \rangle\rangle \Box \neg pos_1$$

Example: Robots and Carriage

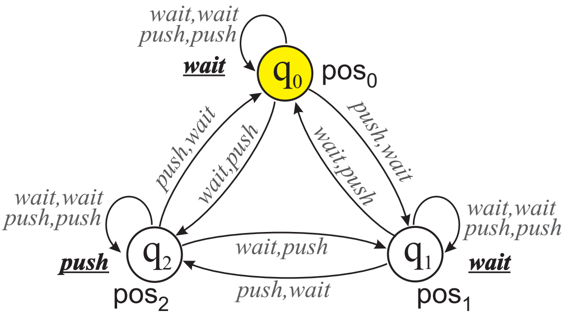

$$pos_0 \longrightarrow \langle\langle 1 \rangle\rangle \Box \neg pos_1$$

Example: Robots and Carriage

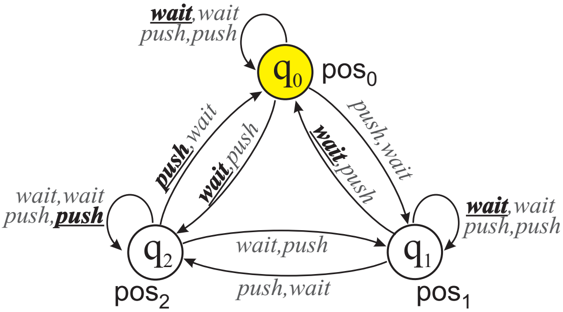

$$pos_0 \longrightarrow \langle\langle 1 \rangle\rangle \Box \neg pos_1$$

Example: Robots and Carriage

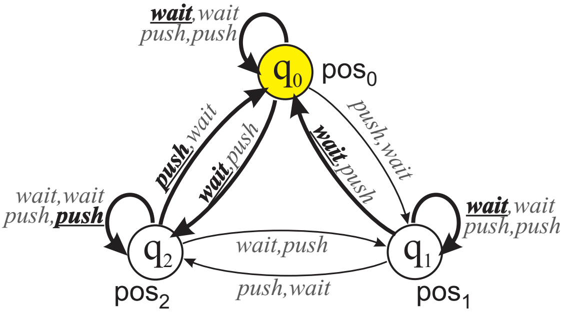

$$pos_0 \rightarrow \langle\langle 1 \rangle\rangle \Box \neg pos_1$$

Example: Robots and Carriage

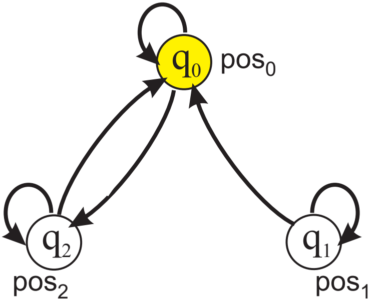

$$pos_0 \rightarrow \langle\langle 1 \rangle\rangle \Box \neg pos_1$$

Example: Robots and Carriage

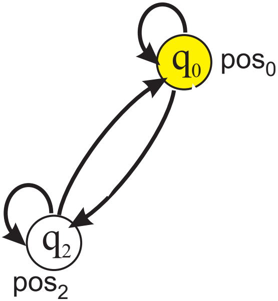

$$pos_0 \rightarrow \langle\langle 1 \rangle\rangle \Box \neg pos_1$$

Temporal operators allow a number of useful concepts to be formally specified

Temporal operators allow a number of useful concepts to be formally specified

- safety properties
- liveness properties
- fairness properties

Safety (maintenance goals):

"something bad will not happen"
"something good will always hold"

Safety (maintenance goals):

"something bad will not happen"
"something good will always hold"

Typical example:

$\Box \neg \mathrm{bankrupt}$

Safety (maintenance goals):

"something bad will not happen"
"something good will always hold"

Typical example:

$\Box\neg$bankrupt

Usually: $\Box\neg\ldots$

Safety (maintenance goals):

"something bad will not happen"
"something good will always hold"

Typical example:

$\Box\neg$bankrupt

Usually: $\Box\neg\ldots$

In ATL:

$\langle\langle os\rangle\rangle\Box\neg$crash

Liveness (achievement goals):

"something good will happen"

Liveness (achievement goals):

"something good will happen"

Typical example:

$\Diamond$rich

Usually: $\Diamond$....

Liveness (achievement goals):

“something good will happen”

Typical example:

$\Diamond$rich

Usually: $\Diamond \ldots$

In ATL:

$\langle\langle alice, bob \rangle\rangle \Diamond$paperAccepted

Fairness (service goals):

"if something is attempted/requested, then it will be successful/allocated"

Fairness (service goals):

"if something is attempted/requested, then it will be successful/allocated"

Typical examples:

$\Box(\text{attempt} \rightarrow \Diamond\text{success})$

$\Box\Diamond\text{attempt} \rightarrow \Box\Diamond\text{success}$

Fairness (service goals):

"if something is attempted/requested, then it will be successful/allocated"

Typical examples:

$\Box(\text{attempt} \rightarrow \Diamond\text{success})$

$\Box\Diamond\text{attempt} \rightarrow \Box\Diamond\text{success}$

In ATL* (!):

$\langle\langle prod, dlr\rangle\rangle\Box(\text{carRequested} \rightarrow \Diamond\text{carDelivered})$

Connection to Games

- Concurrent game structure = generalized extensive game

Connection to Games

- Concurrent game structure = generalized extensive game
- $\langle\langle A \rangle\rangle \gamma$: $\langle\langle A \rangle\rangle$ splits the agents into proponents and opponents
- γ defines the winning condition

Connection to Games

- Concurrent game structure = generalized extensive game

- $\langle\langle A\rangle\rangle\gamma$: $\langle\langle A\rangle\rangle$ splits the agents into proponents and opponents
- γ defines the winning condition
  $\rightsquigarrow$ infinite 2-player, binary, zero-sum game

Connection to Games

- Concurrent game structure = generalized extensive game

- $\langle\langle A \rangle\rangle \gamma$: $\langle\langle A \rangle\rangle$ splits the agents into proponents and opponents
- γ defines the winning condition
  $\rightsquigarrow$ infinite 2-player, binary, zero-sum game

- Flexible and compact specification of winning conditions

- Solving a game $\approx$ checking if $M, q \models \langle\langle A \rangle\rangle \gamma$

- Solving a game $\approx$ checking if $M, q \models \langle\langle A \rangle\rangle \gamma$
- But: do we really want to consider all the possible plays?

2.3 Rational Play (ATLP)

Game-theoretical analysis of games:

- *Solution concepts* define rationality of players

Game-theoretical analysis of games:

- *Solution concepts* define rationality of players
  - maxmin
  - Nash equilibrium
  - subgame-perfect Nash
  - undominated strategies
  - Pareto optimality

Game-theoretical analysis of games:

- *Solution concepts* define rationality of players
  - maxmin
  - Nash equilibrium
  - subgame-perfect Nash
  - undominated strategies
  - Pareto optimality
- Then: we assume that players play rationally
- ...and we ask about the outcome of the game under this assumption

Game-theoretical analysis of games:

- *Solution concepts* define rationality of players
  - maxmin
  - Nash equilibrium
  - subgame-perfect Nash
  - undominated strategies
  - Pareto optimality
- Then: we assume that players play rationally
- ...and we ask about the outcome of the game under this assumption

- Role of rationality criteria: constrain the possible game moves to "sensible" ones

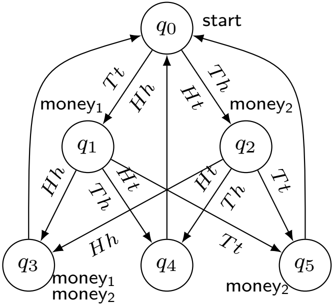

q_0
start
Tt
Hh
Ht
Th
money$_1$
money$_2$
q_1
q_2
Hh
Th
Ht
Ht
Th
Tt
Hh
Tt
q_3
q_4
q_5
money$_1$
money$_2$
money$_2$

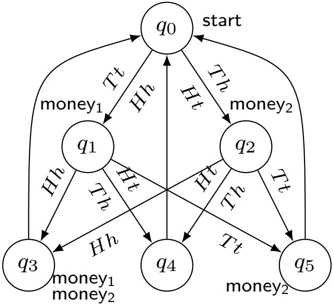

$$\mathsf{start} \rightarrow \neg\langle\langle 1 \rangle\rangle \Diamond \mathsf{money}_1$$

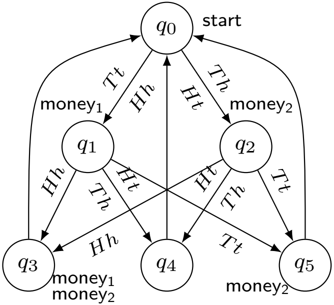

$$\mathsf{start} \rightarrow \neg\langle\langle 1\rangle\rangle\Diamond \mathsf{money}_1$$

$$\mathsf{start} \rightarrow \neg\langle\langle 2\rangle\rangle\Diamond \mathsf{money}_2$$

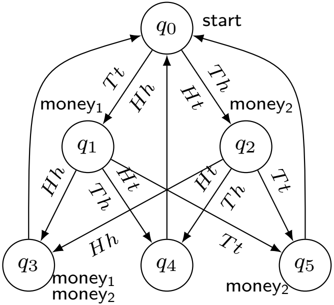

start $\rightarrow \neg\langle\langle 1 \rangle\rangle \Diamond$ money$_1$

start $\rightarrow \neg\langle\langle 2 \rangle\rangle \Diamond$ money$_2$

ATL + Plausibility (ATLP)

ATL: reasoning about *all* possible behaviors.

$\langle\langle A\rangle\rangle\varphi$: agents A have some collective strategy to enforce φ against any response of their opponents.

ATL + Plausibility (ATLP)

ATL: reasoning about *all* possible behaviors.

$\langle\langle A \rangle\rangle \varphi$: agents A have some collective strategy to enforce φ against any response of their opponents.

ATLP: reasoning about *plausible* behaviors.

$\langle\langle A \rangle\rangle \varphi$: agents A have a *plausible* collective strategy to enforce φ against any *plausible* response of their opponents.

ATL + Plausibility (ATLP)

ATL: reasoning about *all* possible behaviors.

$\langle\langle A\rangle\rangle\varphi$: agents A have some collective strategy to enforce φ against any response of their opponents.

ATLP: reasoning about *plausible* behaviors.

$\langle\langle A\rangle\rangle\varphi$: agents A have a *plausible* collective strategy to enforce φ against any *plausible* response of their opponents.

Important

The possible strategies of both A and $\mathbb{A}\mathrm{gt}\setminus A$ are restricted.

New in ATLP:

$(\mathbf{set\text{-}pl}\ \omega)$: the set of plausible profiles is set/reset to the strategies described by ω.
Only plausible strategy profiles are considered!

New in ATLP:

$(\textbf{set-pl}\ \omega)$: the set of plausible profiles is set/reset to the strategies described by ω.
Only plausible strategy profiles are considered!

Example: $(\textbf{set-pl}\ greedy_1)\langle\langle 2\rangle\rangle\Diamond \mathsf{money}_2$

Concurrent game structures with plausibility

$M = (\mathbb{A}\mathrm{gt}, St, \Pi, \pi, Act, d, \delta, \Upsilon, \Omega, \|\cdot\|)$

Concurrent game structures with plausibility

$M = (\mathbb{A}\mathrm{gt}, St, \Pi, \pi, Act, d, \delta, \Upsilon, \Omega, \|\cdot\|)$

- $\Upsilon \subseteq \Sigma$: set of (plausible) strategy profiles

	Deny	Confess
Deny	$-2, -2$	$-5, -1$
Confess	$-1, -5$	$-4, -4$

Υ

Concurrent game structures with plausibility

$M = (\mathbb{A}\mathrm{gt}, St, \Pi, \pi, Act, d, \delta, \Upsilon, \Omega, \|\cdot\|)$

- $\Upsilon \subseteq \Sigma$: set of (plausible) strategy profiles

	Deny	Confess
Deny	$-2, -2$	$-5, -1$
Confess	$-1, -5$	$-4, -4$

Υ

- $\Omega = \{\omega_1, \omega_2, \dots\}$: set of plausibility terms

Example: ω_{NE} may stand for all Nash equilibria

Concurrent game structures with plausibility

$M = (\mathbb{A}\mathrm{gt}, St, \Pi, \pi, Act, d, \delta, \Upsilon, \Omega, \|\cdot\|)$

- $\Upsilon \subseteq \Sigma$: set of (plausible) strategy profiles

	Deny	Confess
Deny	$-2, -2$	$-5, -1$
Confess	$-1, -5$	$-4, -4$

Υ

- $\Omega = \{\omega_1, \omega_2, \dots\}$: set of plausibility terms

  Example: ω_{NE} may stand for all Nash equilibria

- $\|\cdot\| : St \to (\Omega \to \mathcal{P}(()\Sigma))$: plausibility mapping

  Example: $\|\omega_{NE}\|_q = \{(\text{confess}, \text{confess})\}$

Outcome = Paths that may occur when agents A perform s_A

Outcome = Paths that may occur when agents A perform s_A when only plausible strategy profiles from Υ are played

Outcome = Paths that may occur when agents A perform s_A when only plausible strategy profiles from Υ are played

$$out_{\Upsilon}(q, s_A) = \{\lambda \in St^+ \mid \exists t \in \Upsilon(s_A)\ \forall i \in \mathbb{N}\ \big(\lambda[i+1] = \delta(\lambda[i], t(\lambda[i]))\big)\}$$

Outcome = Paths that may occur when agents A perform s_A when only plausible strategy profiles from Υ are played

$$out_\Upsilon(q, s_A) = \{\lambda \in St^+ \mid \exists t \in \Upsilon(s_A)\ \forall i \in \mathbb{N}\ \big(\lambda[i+1] = \delta(\lambda[i], t(\lambda[i]))\big)\}$$

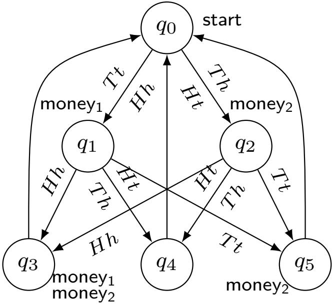

Outcome = Paths that may occur when agents A perform s_A when only plausible strategy profiles from Υ are played

$$out_\Upsilon(q, s_A) = \{\lambda \in St^+ \mid \exists t \in \Upsilon(s_A)\ \forall i \in \mathbb{N}\ \big(\lambda[i+1] = \delta(\lambda[i], t(\lambda[i]))\big)\}$$

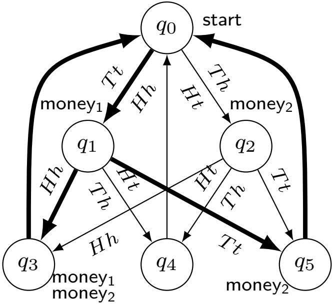

P: the players always show same sides of their coins

Outcome = Paths that may occur when agents A perform s_A when only plausible strategy profiles from Υ are played

$$out_\Upsilon(q, s_A) = \{\lambda \in St^+ \mid \exists t \in \Upsilon(s_A)\ \forall i \in \mathbb{N}\ (\lambda[i+1] = \delta(\lambda[i], t(\lambda[i])))\}$$

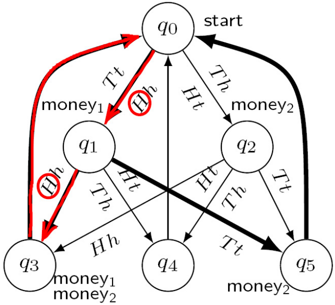

P: the players always show same sides of their coins

s_1: always show "heads"

Semantics of ATLP

$M, q \models \langle\langle A \rangle\rangle \gamma$ iff there is a strategy s_A consistent with Υ such that $M, \lambda \models \gamma$ for all $\lambda \in out_\Upsilon(q, s_A)$

$M, q \models (\textbf{set-pl}\ \omega)\varphi$ iff $M^\omega, q \models \varphi$ where the new model M^ω is equal to M but the new set Υ^ω of plausible strategy profiles is set to $\|\omega\|_q$.

Example: Pennies Game

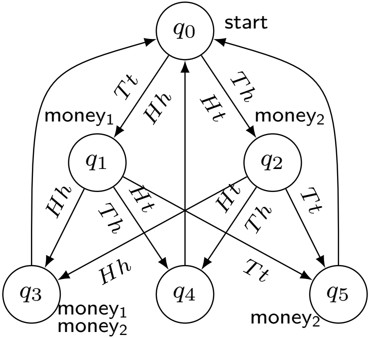

$$M, q_0 \models (\textbf{set-pl}\ \omega_{NE})\langle\langle 2 \rangle\rangle \Diamond \text{money}_2$$

Example: Pennies Game

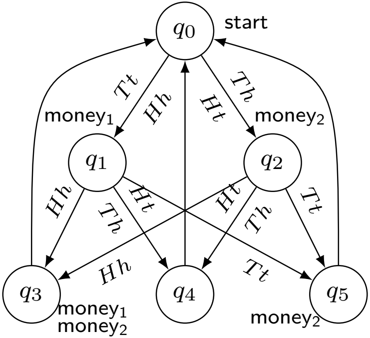

$M, q_0 \models (\textbf{set-pl}\ \omega_{NE})\langle\langle 2\rangle\rangle\Diamond \mathsf{money}_2$

What is a Nash equilibrium in this game?

We need some kind of winning criteria!

Agent 1 "wins", if $\gamma_1 \equiv \Box(\neg\mathsf{start} \to \mathsf{money}_1)$ is satisfied.
Agent 2 "wins", if $\gamma_2 \equiv \Diamond\mathsf{money}_2$ is satisfied.

Agent 1 "wins", if $\gamma_1 \equiv \Box(\neg \mathsf{start} \rightarrow \mathsf{money}_1)$ is satisfied.
Agent 2 "wins", if $\gamma_2 \equiv \Diamond \mathsf{money}_2$ is satisfied.

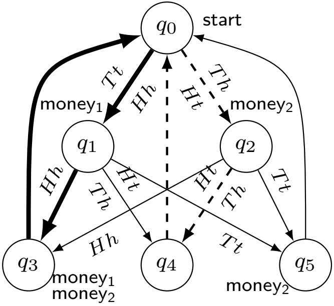

Agent 1 “wins”, if $\gamma_1 \equiv \Box(\neg\mathsf{start} \rightarrow \mathsf{money}_1)$ is satisfied.
Agent 2 “wins”, if $\gamma_2 \equiv \Diamond\mathsf{money}_2$ is satisfied.

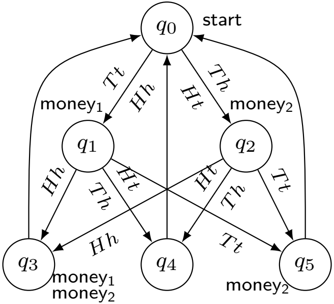

$\gamma_1 \backslash \gamma_2$	hh	ht	th	tt
HH	**1, 1**	0, 0	0, 1	0, 1
HT	0, 0	0, 1	0, 1	0, 1
TH	0, 1	0, 1	**1, 1**	0, 0
TT	0, 1	0, 1	0, 0	0, 1

Agent 1 “wins”, if $\gamma_1 \equiv \Box(\neg\text{start} \rightarrow \text{money}_1)$ is satisfied.
Agent 2 “wins”, if $\gamma_2 \equiv \Diamond\text{money}_2$ is satisfied.

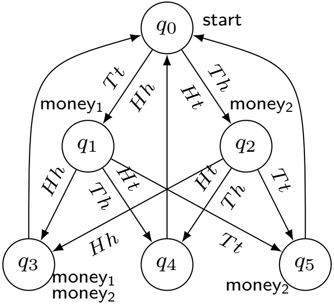

$\gamma_1 \backslash \gamma_2$	hh	ht	th	tt
HH	**1, 1**	0, 0	0, 1	0, 1
HT	0, 0	0, 1	0, 1	0, 1
TH	0, 1	0, 1	**1, 1**	0, 0
TT	0, 1	0, 1	0, 0	0, 1

Now we have a qualitative notion of success.

Agent 1 "wins", if $\gamma_1 \equiv \Box(\neg\text{start} \rightarrow \text{money}_1)$ is satisfied.
Agent 2 "wins", if $\gamma_2 \equiv \Diamond\text{money}_2$ is satisfied.

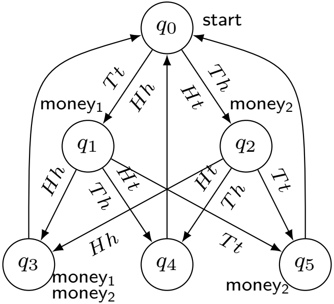

$\gamma_1 \backslash \gamma_2$	hh	ht	th	tt
HH	**1, 1**	0, 0	0, 1	0, 1
HT	0, 0	0, 1	0, 1	0, 1
TH	0, 1	0, 1	**1, 1**	0, 0
TT	0, 1	0, 1	0, 0	0, 1

Now we have a qualitative notion of success.

$$M, q_0 \models (\textbf{set-pl}\ \omega_{NE})\langle\langle 2\rangle\rangle\Box(\neg\text{start} \rightarrow \text{money}_1)$$

where $\|\omega_{NE}\|_{q_0}$ = "all profiles belonging to grey cells".

How to obtain plausibility terms?

How to obtain plausibility terms?

Idea

Formulae that describe plausible strategies!

(**set-pl** $\sigma.\theta)\varphi$: "suppose that θ characterizes rational strategy profiles, then φ holds".

How to obtain plausibility terms?

Idea

Formulae that describe plausible strategies!

$(\textbf{set-pl}\ \sigma.\theta)\varphi$: "suppose that θ characterizes rational strategy profiles, then φ holds".

Sometimes quantifiers are needed...

E.g.: $(\textbf{set-pl}\ \sigma.\ \forall\sigma'\ dominates(\sigma, \sigma'))$

Characterization of Nash Equilibrium

σ_a is a's best response to σ (wrt $\vec{\gamma}$):

$$BR_a^{\vec{\gamma}}(\sigma) \equiv (\textbf{set-pl}\ \sigma[\mathbb{A}\mathrm{gt} \setminus \{a\}]) \big(\langle\langle a \rangle\rangle \gamma_a \rightarrow (\textbf{set-pl}\ \sigma) \langle\langle \emptyset \rangle\rangle \gamma_a \big)$$

Characterization of Nash Equilibrium

σ_a is a's best response to σ (wrt $\vec{\gamma}$):

$$BR_a^{\vec{\gamma}}(\sigma) \equiv (\textbf{set-pl}\ \sigma[\mathbb{A}\mathrm{gt}\setminus\{a\}])\big(\langle\langle a\rangle\rangle\gamma_a \rightarrow (\textbf{set-pl}\ \sigma)\langle\langle\emptyset\rangle\rangle\gamma_a\big)$$

σ is a Nash equilibrium:

$$NE^{\vec{\gamma}}(\sigma) \equiv \bigwedge_{a\in\mathbb{A}\mathrm{gt}} BR_a^{\vec{\gamma}}(\sigma)$$

Example: Pennies Game revisited

$\gamma_1 \equiv \Box(\neg \mathsf{start} \rightarrow \mathsf{money}_1)$; $\gamma_2 \equiv \Diamond \mathsf{money}_2$.

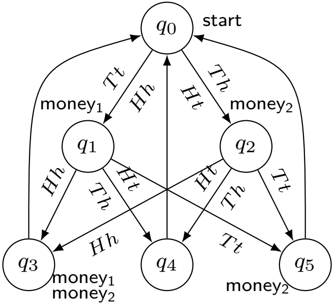

$\gamma_1 \backslash \gamma_2$	hh	ht	th	tt
HH	**1, 1**	0, 0	0, 1	0, 1
HT	0, 0	0, 1	0, 1	0, 1
TH	0, 1	0, 1	**1, 1**	0, 0
TT	0, 1	0, 1	0, 0	0, 1

$$M_1, q_0 \models (\textbf{set-pl}\ \sigma.NE^{\gamma_1,\gamma_2}(\sigma))\langle\langle 2\rangle\rangle\Box(\neg \mathsf{start} \rightarrow \mathsf{money}_1)$$

...where $NE^{\gamma_1,\gamma_2}(\sigma)$ is defined as on the last slide.

Characterizations of Other Solution Concepts

σ is a subgame perfect Nash equilibrium:

$$SPN^{\vec{\gamma}}(\sigma) \equiv \langle\langle\emptyset\rangle\rangle\Box NE^{\vec{\gamma}}(\sigma)$$

σ is Pareto optimal:

$$PO^{\vec{\gamma}}(\sigma) \equiv \forall\sigma'\Big(\bigwedge_{a\in\mathbb{A}\mathrm{gt}} ((\textbf{set-pl } \sigma')\langle\langle\emptyset\rangle\rangle\gamma_a \rightarrow (\textbf{set-pl } \sigma)\langle\langle\emptyset\rangle\rangle\gamma_a) \vee \bigvee_{a\in\mathbb{A}\mathrm{gt}} ((\textbf{set-pl } \sigma)\langle\langle\emptyset\rangle\rangle\gamma_a \wedge \neg(\textbf{set-pl } \sigma')\langle\langle\emptyset\rangle\rangle\gamma_a \Big).$$

σ is undominated:

$$UNDOM^{\vec{\gamma}}(\sigma) \equiv \forall\sigma_1\forall\sigma_2\exists\sigma_3 \Big(\big((\textbf{set-pl}\ \langle\sigma_1^{\{a\}}, \sigma_2^{\mathbb{A}\mathrm{gt}\setminus\{a\}}\rangle)\langle\langle\emptyset\rangle\rangle\gamma_a \rightarrow (\textbf{set-pl}\ \langle\sigma^{\{a\}}, \sigma_2^{\mathbb{A}\mathrm{gt}\setminus\{a\}}\rangle)\langle\langle\emptyset\rangle\rangle\gamma_a \big) \vee \big((\textbf{set-pl}\ \langle\sigma^{\{a\}}, \sigma_3^{\mathbb{A}\mathrm{gt}\setminus\{a\}}\rangle)\langle\langle\emptyset\rangle\rangle\gamma_a \wedge \neg(\textbf{set-pl}\ \langle\sigma_1^{\{a\}}, \sigma_3^{\mathbb{A}\mathrm{gt}\setminus\{a\}}\rangle)\langle\langle\emptyset\rangle\rangle\gamma_a \big) \Big).$$

Theorem 2.5

The characterizations coincide with game-theoretical solution concepts in the class of game trees.

2.4 Imperfect Information

How can we reason about extensive games with imperfect information?

How can we reason about extensive games with imperfect information?

Let's put ATL and epistemic logic in one box.

How can we reason about extensive games with imperfect information?

Let's put ATL and epistemic logic in one box.

$\rightsquigarrow$ Problems!

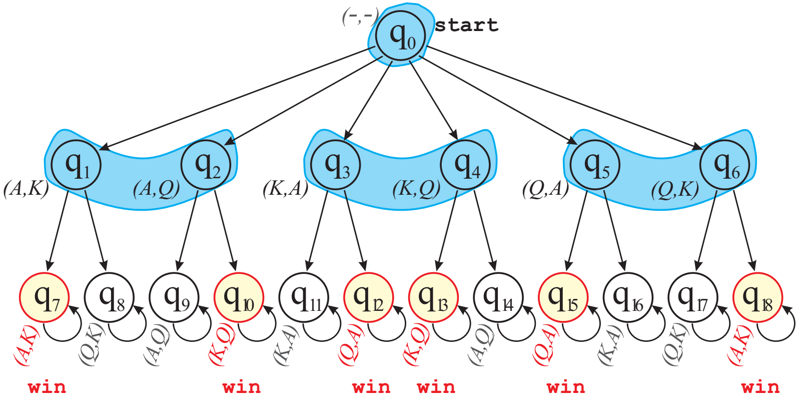
(-,-)
q0
start
q1
q2
q3
q4
q5
q6
(A,K)
(A,Q)
(K,A)
(K,Q)
(Q,A)
(Q,K)
q7
q8
q9
q10
q11
q12
q13
q14
q15
q16
q17
q18
(A,K)
(Q,K)
(A,Q)
(K,Q)
(K,A)
(Q,A)
(K,Q)
(A,Q)
(Q,A)
(K,A)
(Q,K)
(A,K)
win
win
win
win
win
win

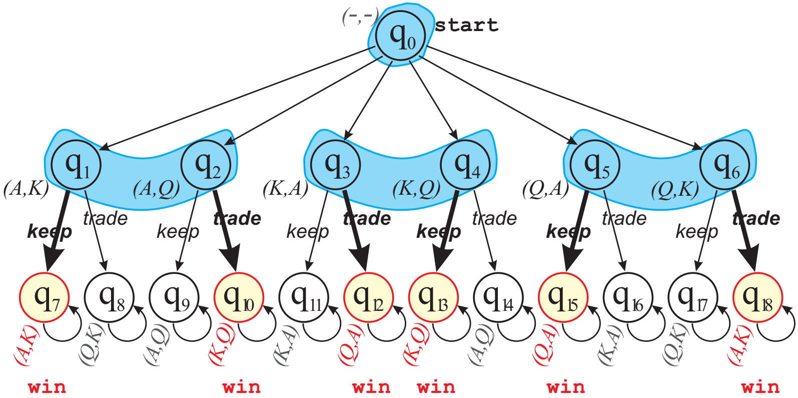
(-,-)
q0
start
q1
q2
q3
q4
q5
q6
(A,K)
(A,Q)
(K,A)
(K,Q)
(Q,A)
(Q,K)
keep
trade
keep
trade
keep
trade
keep
trade
keep
trade
keep
trade
q7
q8
q9
q10
q11
q12
q13
q14
q15
q16
q17
q18
(A,K)
(Q,K)
(A,Q)
(K,Q)
(K,A)
(Q,A)
(K,Q)
(A,Q)
(Q,A)
(K,A)
(Q,K)
(A,K)
win
win
win
win
win
win

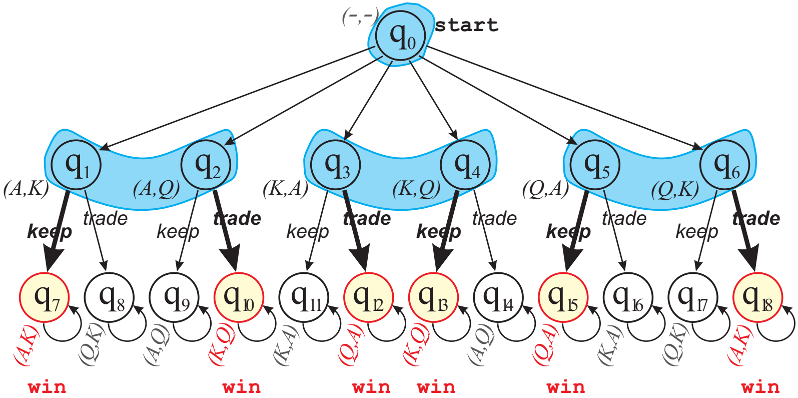

$$start \rightarrow \langle\langle a \rangle\rangle \Diamond win$$

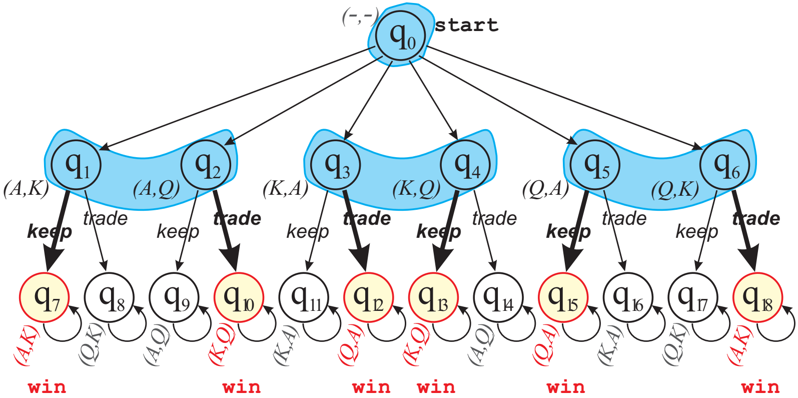

$$start \rightarrow \langle\langle a \rangle\rangle \Diamond win$$

$$start \rightarrow K_a \langle\langle a \rangle\rangle \Diamond win$$

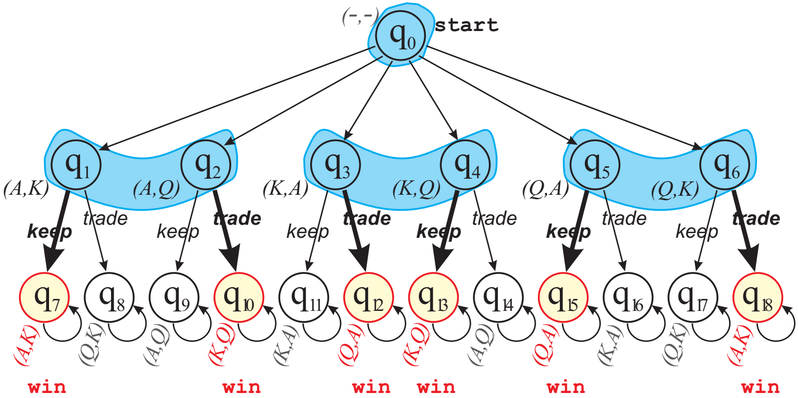

$$start \rightarrow \langle\langle a \rangle\rangle \Diamond win$$

$$start \rightarrow K_a \langle\langle a \rangle\rangle \Diamond win$$

Does it make sense?

Problem:

Strategic and epistemic abilities are *not* independent!

Problem:

Strategic and epistemic abilities are *not* independent!

$\langle\langle A \rangle\rangle \Phi$ = A can enforce Φ

Problem:

Strategic and epistemic abilities are *not* independent!

$\langle\langle A \rangle\rangle \Phi$ = A can enforce Φ

It should at least mean that A are able to identify and execute the right strategy!

Problem:

Strategic and epistemic abilities are *not* independent!

$\langle\langle A \rangle\rangle \Phi$ = A can enforce Φ

It should at least mean that A are able to identify and execute the right strategy!

Executable strategies = uniform strategies

Definition 2.6 (Uniform strategy)

Strategy s_a is uniform iff it specifies the same choices for indistinguishable situations:

- (no recall:) if $q \sim_a q'$ then $s_a(q) = s_a(q')$
- (perfect recall:) if $\lambda \approx_a \lambda'$ then $\Rightarrow s_a(\lambda) = s_a(\lambda)$, where $\lambda \approx_a \lambda'$ iff $\lambda[i] \sim_a \lambda'[i]$ for every i.

Definition 2.6 (Uniform strategy)

Strategy s_a is uniform iff it specifies the same choices for indistinguishable situations:

- (no recall:) if $q \sim_a q'$ then $s_a(q) = s_a(q')$
- (perfect recall:) if $\lambda \approx_a \lambda'$ then $\Rightarrow s_a(\lambda) = s_a(\lambda)$, where $\lambda \approx_a \lambda'$ iff $\lambda[i] \sim_a \lambda'[i]$ for every i.

A collective strategy is uniform iff it consists only of uniform individual strategies.

Note:

Having a successful strategy does not imply knowing that we have it!

Note:

Having a successful strategy does not imply knowing that we have it!

Knowing that a successful strategy exists does not imply knowing the strategy itself!

Levels of Strategic Ability

From now on, we restrict our discussion to uniform memoryless strategies.

Levels of Strategic Ability

From now on, we restrict our discussion to uniform memoryless strategies.

Our cases for $\langle\langle A \rangle\rangle \Phi$ under incomplete information:

2 There is σ such that, for every execution of σ, Φ holds

Levels of Strategic Ability

From now on, we restrict our discussion to uniform memoryless strategies.

Our cases for $\langle\langle A\rangle\rangle\Phi$ under incomplete information:

2. There is σ such that, for every execution of σ, Φ holds
3. A know that there is σ such that, for every execution of σ, Φ holds

Levels of Strategic Ability

From now on, we restrict our discussion to uniform memoryless strategies.

Our cases for $\langle\langle A\rangle\rangle\Phi$ under incomplete information:

2. There is σ such that, for every execution of σ, Φ holds
3. A know that there is σ such that, for every execution of σ, Φ holds
4. There is σ such that A know that, for every execution of σ, Φ holds

Case [4]: knowing how to play

Case [4]: knowing how to play

- Single agent case: we take into account the paths starting from indistinguishable states (i.e., $\bigcup_{q' \in \mathrm{img}(q,\sim_a)} out(q, s_A)$)

Case [4]: knowing how to play

- Single agent case: we take into account the paths starting from indistinguishable states (i.e., $\bigcup_{q' \in \mathrm{img}(q, \sim_a)} out(q, s_A)$)

- What about coalitions?
- Question: in what sense should they know the strategy? Common knowledge (C_A), mutual knowledge (K_A), distributed knowledge (D_A)?

Given strategy σ, agents A can have:

- Common knowledge that σ is a winning strategy. This requires the least amount of additional communication (agents from A may agree upon a total order over their collective strategies at the beginning of the game and that they will always choose the maximal winning strategy with respect to this order)

Given strategy σ, agents A can have:

- Common knowledge that σ is a winning strategy. This requires the least amount of additional communication (agents from A may agree upon a total order over their collective strategies at the beginning of the game and that they will always choose the maximal winning strategy with respect to this order)
- Mutual knowledge that σ is a winning strategy: everybody in A knows that σ is winning

- Distributed knowledge that σ is a winning strategy: if the agents share their knowledge at the current state, they can identify the strategy as winning

- Distributed knowledge that σ is a winning strategy: if the agents share their knowledge at the current state, they can identify the strategy as winning
- "The leader": the strategy can be identified by agent $a \in A$

- Distributed knowledge that σ is a winning strategy: if the agents share their knowledge at the current state, they can identify the strategy as winning
- "The leader": the strategy can be identified by agent $a \in A$
- "Headquarters' committee": the strategy can be identified by subgroup $A' \subseteq A$

- Distributed knowledge that σ is a winning strategy: if the agents share their knowledge at the current state, they can identify the strategy as winning
- "The leader": the strategy can be identified by agent $a \in A$
- "Headquarters' committee": the strategy can be identified by subgroup $A' \subseteq A$
- "Consulting company": the strategy can be identified by some other group B

Many subtle cases...

Many subtle cases...

$\leadsto$ Solution: constructive knowledge operators

Constructive Strategic Logic (CSL)

- $\langle\langle A \rangle\rangle \Phi$: A have a uniform memoryless strategy to enforce Φ

Constructive Strategic Logic (CSL)

- $\langle\langle A \rangle\rangle \Phi$: A have a uniform memoryless strategy to enforce Φ
- $K_a \langle\langle a \rangle\rangle \Phi$: a has a strategy to enforce Φ, and knows that he has one
- For groups of agents: $C_A, E_A, D_A, \ldots$

Constructive Strategic Logic (CSL)

- $\langle\langle A\rangle\rangle\Phi$: A have a uniform memoryless strategy to enforce Φ
- $K_a\langle\langle a\rangle\rangle\Phi$: a has a strategy to enforce Φ, and knows that he has one
- For groups of agents: $C_A, E_A, D_A, \ldots$
- $\mathbb{K}_a\langle\langle a\rangle\rangle\Phi$: a has a strategy to enforce Φ, and knows that this is a winning strategy
- For groups of agents: $\mathbb{C}_A, \mathbb{E}_A, \mathbb{D}_A, \ldots$

Non-standard semantics:

- Formulae are evaluated in sets of states
- $M, Q \models \langle\langle A \rangle\rangle \Phi$: A have a single strategy to enforce Φ from all states in Q

Non-standard semantics:

- Formulae are evaluated in sets of states
- $M, Q \models \langle\langle A \rangle\rangle \Phi$: A have a single strategy to enforce Φ from all states in Q

Additionally:

- $out(Q, S_A) = \bigcup_{q \in Q} out(q, S_A)$
- $\mathrm{img}(Q, \mathcal{R}) = \bigcup_{q \in Q} \mathrm{img}(q, \mathcal{R})$

Non-standard semantics:

- Formulae are evaluated in sets of states
- $M, Q \models \langle\langle A \rangle\rangle \Phi$: A have a single strategy to enforce Φ from all states in Q

Additionally:

- $out(Q, S_A) = \bigcup_{q \in Q} out(q, S_A)$
- $\mathrm{img}(Q, \mathcal{R}) = \bigcup_{q \in Q} \mathrm{img}(q, \mathcal{R})$

- $M, q \models \varphi$ iff $M, \{q\} \models \varphi$

Definition 2.7 (Semantics of CSL)

$M, Q \models p$ iff $p \in \pi(q)$ for every $q \in Q$;

Definition 2.7 (Semantics of CSL)

$M, Q \models p$ iff $p \in \pi(q)$ for every $q \in Q$;

$M, Q \models \neg\varphi$ iff not $M, Q \models \varphi$;

Definition 2.7 (Semantics of CSL)

$M, Q \models p$ iff $p \in \pi(q)$ for every $q \in Q$;

$M, Q \models \neg\varphi$ iff not $M, Q \models \varphi$;

$M, Q \models \varphi \land \psi$ iff $M, Q \models \varphi$ and $M, Q \models \psi$;

Definition 2.7 (Semantics of CSL)

$M, Q \models p$ iff $p \in \pi(q)$ for every $q \in Q$;

$M, Q \models \neg\varphi$ iff not $M, Q \models \varphi$;

$M, Q \models \varphi \wedge \psi$ iff $M, Q \models \varphi$ and $M, Q \models \psi$;

$M, Q \models \langle\langle A \rangle\rangle \gamma$ iff there exists S_A such that, for every $\lambda \in out(Q, S_A)$, we have that $M, \lambda[1] \models \varphi$;

$M, Q \models \mathcal{K}_A\varphi$ iff $M, q \models \varphi$ for every $q \in \mathrm{img}(Q, \sim_A^{\mathcal{K}})$ (where $\mathcal{K} = C, E, D$);

$M, Q \models \mathcal{K}_A\varphi$ iff $M, q \models \varphi$ for every $q \in \mathrm{img}(Q, \sim_A^{\mathcal{K}})$ (where $\mathcal{K} = C, E, D$);

$M, Q \models \hat{\mathcal{K}}_A\varphi$ iff $M, \mathrm{img}(Q, \sim_A^{\mathcal{K}}) \models \varphi$ (where $\hat{\mathcal{K}} = \mathbb{C}, \mathbb{E}, \mathbb{D}$ and $\mathcal{K} = C, E, D$, respectively).

Example: Simple Market

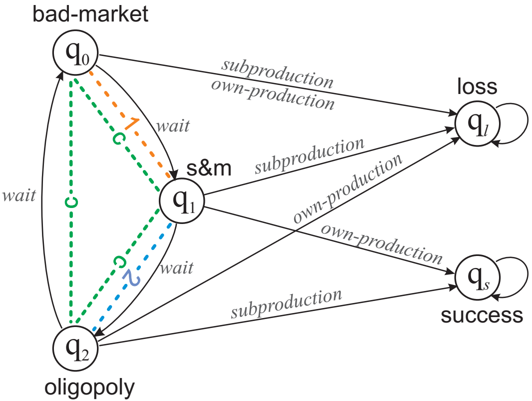

@ q_1 :

$\neg \mathbb{K}_c \langle\langle c \rangle\rangle \Diamond \text{success}$

Example: Simple Market

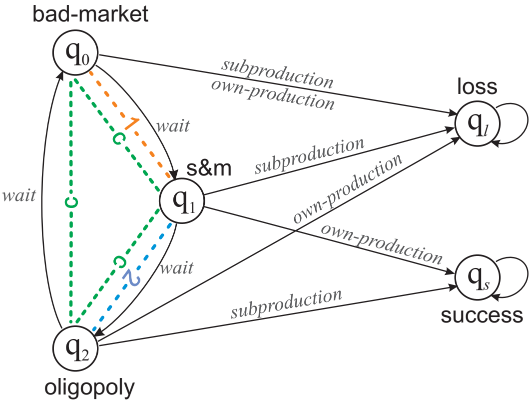

@ q_1 :

$\neg\mathbb{K}_c\langle\langle c\rangle\rangle\Diamond\text{success}$

$\neg\mathbb{E}_{\{1,2\}}\langle\langle c\rangle\rangle\Diamond\text{success}$

$\neg\mathbb{K}_1\langle\langle c\rangle\rangle\Diamond\text{success}$

$\neg\mathbb{K}_2\langle\langle c\rangle\rangle\Diamond\text{success}$

Example: Simple Market

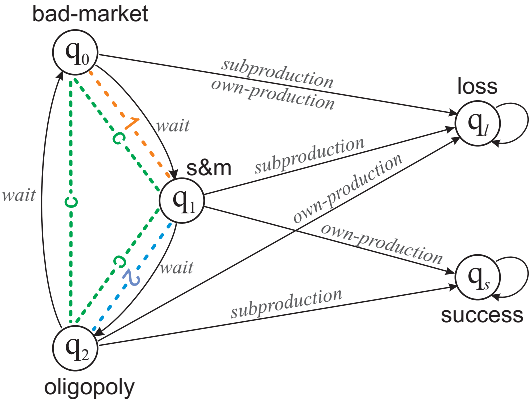

$@\ q_1$:

$\neg\mathbb{K}_c\langle\langle c\rangle\rangle\Diamond$success

$\neg\mathbb{E}_{\{1,2\}}\langle\langle c\rangle\rangle\Diamond$success

$\neg\mathbb{K}_1\langle\langle c\rangle\rangle\Diamond$success

$\neg\mathbb{K}_2\langle\langle c\rangle\rangle\Diamond$success

$\mathbb{D}_{\{1,2\}}\langle\langle c\rangle\rangle\Diamond$success

Theorem 2.8 (Expressivity)

CSL is strictly more expressive than most previous proposals.

Theorem 2.8 (Expressivity)

CSL is strictly more expressive than most previous proposals.

Theorem 2.9 (Verification complexity)

The complexity of model checking CSL is minimal.

2.5 Model Checking

Model Checking Formulae of CTL and ATL

- Model checking: Does φ hold in model M and state q?

Model Checking Formulae of CTL and ATL

- Model checking: Does φ hold in model M and state q?
- Natural for verification of existing systems; also during design (“prototyping”)
- Can be used for automated planning

function $plan(\varphi)$.
Returns a subset of St for which formula φ holds, together with a (conditional) plan to achieve φ. The plan is sought within the context of concurrent game structure $S = \langle \mathbb{A}\mathrm{gt}, St, \Pi, \pi, o\rangle$.

case $\varphi \in \Pi$ **:** return $\{\langle q, -\rangle \mid \varphi \in \pi(q)\}$
case $\varphi = \neg\psi$ **:** $P_1 := plan(\psi)$;
return $\{\langle q, -\rangle \mid q \notin states(P_1)\}$
case $\varphi = \psi_1 \vee \psi_2$ **:**
$P_1 := plan(\psi_1)$; $P_2 := plan(\psi_2)$;
return $\{\langle q, -\rangle \mid q \in states(P_1) \cup states(P_2)\}$
case $\varphi = \langle\langle A\rangle\rangle \bigcirc \psi$ **:** return $pre(A, states(plan(\psi)))$
case $\varphi = \langle\langle A\rangle\rangle \Box \psi$ **:**
$P_1 := plan(\mathbf{true})$; $P_2 := plan(\psi)$; $Q_3 := states(P_2)$;
while $states(P_1) \not\subseteq states(P_2)$
do $P_1 := P_2|_{states(P_1)}$; $P_2 := pre(A, states(P_1))|_{Q_3}$ **od**;
return $P_2|_{states(P_1)}$
case $\varphi = \langle\langle A\rangle\rangle \psi_1 \mathcal{U} \psi_2$ **:**
$P_1 := \emptyset$; $Q_3 := states(plan(\psi_1))$; $P_2 := plan(\mathbf{true})|_{states(plan(\psi_2))}$;
while $states(P_2) \not\subseteq states(P_1)$
do $P_1 := P_1 \oplus P_2$; $P_2 := pre(A, states(P_1))|_{Q_3}$ **od**;
return P_1
end case

Complexity od Model Checking ATL

Theorem (Alur, Kupferman & Henzinger 1998)

ATL model checking is P-complete, and can be done in time linear in the size of the model and the length of the formula.

Complexity od Model Checking ATL

Theorem (Alur, Kupferman & Henzinger 1998)

ATL model checking is P-complete, and can be done in time linear in the size of the model and the length of the formula.

So, let's model-check!

Complexity od Model Checking ATL

Theorem (Alur, Kupferman & Henzinger 1998)

ATL model checking is P-complete, and can be done in time linear in the size of the model and the length of the formula.

So, let's model-check!

Not as easy as it seems.

- Nice results: model checking ATL is tractable.

- Nice results: model checking ATL is tractable.
- But: the result is relative to the size of the model and the formula

- Nice results: model checking ATL is tractable.
- But: the result is relative to the size of the model and the formula
- Well known catch: size of models is exponential wrt a higher-level description

- Nice results: model checking ATL is tractable.
- But: the result is relative to the size of the model and the formula
- Well known catch: size of models is exponential wrt a higher-level description
- Another problem: transitions are labeled
- So: the number of transitions can be exponential in the number of agents.

3 agents/attributes, 12 states, 216 transitions

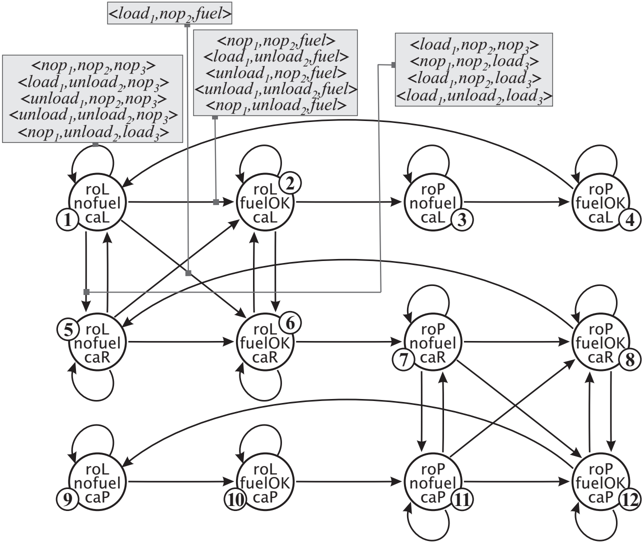

Model Checking Temporal & Strategic Logics

	m, l	n, k, l	n_{local}, k, l
CTL			
ATL			
CSL			

Model Checking Temporal & Strategic Logics

	m, l	n, k, l	n_{local}, k, l
CTL	P [1]	P [1]	PSPACE [2]
ATL			
CSL			

[1] Clarke, Emerson & Sistla (1986).
[2] Kupferman, Vardi & Wolper (2000).

Model Checking Temporal & Strategic Logics

	m, l	n, k, l	n_{local}, k, l
CTL	P [1]	P [1]	PSPACE [2]
ATL	P [3]	Δ_3^P [5,6]	EXPTIME [8,9]
CSL			

[1] Clarke, Emerson & Sistla (1986).
[2] Kupferman, Vardi & Wolper (2000).
[3] Alur, Henzinger & Kupferman (2002).

[5] Jamroga & Dix (2005).
[6] Laroussinie, Markey & Oreiby (2006).

[8] Hoek, Lomuscio & Wooldridge (2006).

Model Checking Temporal & Strategic Logics

	m, l	n, k, l	n_{local}, k, l
CTL	P [1]	P [1]	PSPACE [2]
ATL	P [3]	Δ_3^P [5,6]	EXPTIME [8,9]
CSL	Δ_2^P [4,7]	Δ_3^P [7]	PSPACE [9]

[1] Clarke, Emerson & Sistla (1986).
[2] Kupferman, Vardi & Wolper (2000).
[3] Alur, Henzinger & Kupferman (2002).
[4] Schobbens (2004).
[5] Jamroga & Dix (2005).
[6] Laroussinie, Markey & Oreiby (2006).
[7] Jamroga & Dix (2007).
[8] Hoek, Lomuscio & Wooldridge (2006).
[9] Jamroga & Ågotnes (2007).

Main message:

- Complexity is very sensitive to the context!

Main message:

- Complexity is very sensitive to the context!
- In particular, the way we define the input, and measure its size, is crucial.

Even if model checking appears very easy, it can be very hard.

Even if model checking appears very easy, it can be very hard.

Still, people do automatic model checking!

Even if model checking appears very easy, it can be very hard.

Still, people do automatic model checking!

- LTL: SPIN

Even if model checking appears very easy, it can be very hard.

Still, people do automatic model checking!

- LTL: SPIN
- CTL/ATL: MOCHA, MCMAS, VeriCS

Even if model checking appears very easy, it can be very hard.

Still, people do automatic model checking!

- LTL: SPIN
- CTL/ATL: MOCHA, MCMAS, VeriCS

Even if model checking is theoretically hard, it can be feasible in practice.

2.6 References

[Alur et al. 2002] R. Alur, T. A. Henzinger, and O. Kupferman.
Alternating-time Temporal Logic.
Journal of the ACM, 49:672–713, 2002.

[Emerson 1990] E. A. Emerson.
Temporal and modal logic.
Handbook of Theoretical Computer Science, volume B, 995–1072. Elsevier, 1990.

[Fisher 2006] Fisher, M..
Temporal Logics. Kluwer, 2006.

[Jamroga and Ågotnes 2007] W. Jamroga and T. Ågotnes.
Constructive knowledge: What agents can achieve under incomplete information.
Journal of Applied Non-Classical Logics, 17(4):423–475, 2007.