Alternating-time Temporal Logic

Thomas Ågotnes

University of Bergen, Norway

Introduction: Reasoning about Coalitional Ability

- This lecture will be about reasoning about coalitional ability in modal logic

- Will study different variants of logics with coalition operators of the form

$$\langle C \rangle \phi$$

  - where C is a coalition (= set of agents)

  - meaning: C has the ability to make phi true

Introduction: Reasoning about Coalitional Ability

- We will look at
  - different meanings of *ability*
  - different combinations with temporal, epistemic, public announcement, ..., operators

Introduction: Reasoning about Coalitional Ability

- Most common frameworks:
  - Pauly's Coalition Logic (CL):
    - extends propositional logic with coalition operators
    - interpreted in game structures: ability = the coalition can choose a joint action such that phi becomes true *no matter what the other agents do*
  - Alur et al.'s Alternating-time Temporal Logic (ATL):
    - can be seen as an extension of CL with temporal operators
    - ability = the coalition can choose a joint strategy such that phi becomes true no matter what the other agents do
- Others: van Benthem on forcing, Seeing-to-it-that (STIT) logics, ...

Confusion: is it a diamond or a box

- In CL and ATL: ability = the coalition can choose a joint action such that phi becomes true *no matter what the other agents do*
- *“exists... for all”*-pattern
- Notation that is sometimes used for this: $\langle\langle C \rangle\rangle \phi \quad [C]\phi$
- We will use the following notation: $\langle\![C]\!\rangle \phi$

Examples:

$$\langle\!\langle Xi, Obama \rangle\!\rangle \neg crisis \qquad \text{(CL)}$$

$$\langle\!\langle Thomas, Hans \rangle\!\rangle \diamond students\_happy \qquad \text{(ATL)}$$

Example: a simple social choice mechanism

Example: a simple social choice mechanism

- Two individuals, **a** and **b**, must choose between two outcomes, **p** and **q**. We want a mechanism that will allow them to choose which will satisfy the following requirements: we want an outcome to be possible – that is, we want the two agents to choose, collectively, either p or q. We do not want them to be able to bring about both outcomes simultaneously. Finally, we do not want either agent to be able to unilaterally dictate an outcome – we want them both to have “equal power”.

Example: a simple social choice mechanism

- Two individuals, **a** and **b**, must choose between two outcomes, **p** and **q**. We want a mechanism that will allow them to choose which will satisfy the following requirements: we want an outcome to be possible – that is, we want the two agents to choose, collectively, either p or q. We do not want them to be able to bring about both outcomes simultaneously. Finally, we do not want either agent to be able to unilaterally dictate an outcome – we want them both to have “equal power”.

- Specification in Coalition Logic:

Example: a simple social choice mechanism

- Two individuals, **a** and **b**, must choose between two outcomes, **p** and **q**. We want a mechanism that will allow them to choose which will satisfy the following requirements: we want an outcome to be possible – that is, we want the two agents to choose, collectively, either p or q. We do not want them to be able to bring about both outcomes simultaneously. Finally, we do not want either agent to be able to unilaterally dictate an outcome – we want them both to have “equal power”.

- Specification in Coalition Logic:

$$\langle\![a, b]\!\rangle p \quad \langle\![a, b]\!\rangle q \quad \neg\langle\![a, b]\!\rangle(p \wedge q)$$
$$\neg\langle\![a]\!\rangle p \quad \neg\langle\![b]\!\rangle p \quad \neg\langle\![a]\!\rangle q \quad \neg\langle\![b]\!\rangle q$$

ATL

- Alternating-time Temporal Logic was introduced by Alur et. al for strategic reasoning in game-like situations

- It can be viewed as an extension of both

  - Coalition Logic

  - Computation Tree Logic (CTL)

- CTL is a branching-time temporal logic, one of the most well-known temporal logics

Contents

- Branching-time temporal logic: CTL
- ATL
- Bisimulations and the role of memory
- Irrevocable strategies

Branching-time temporal logics

- Natural to view the possible unfoldings of events as a tree linear in the past, branching into the future.
- Branching corresponds to different ways in which non-determinism can be resolved.

Example

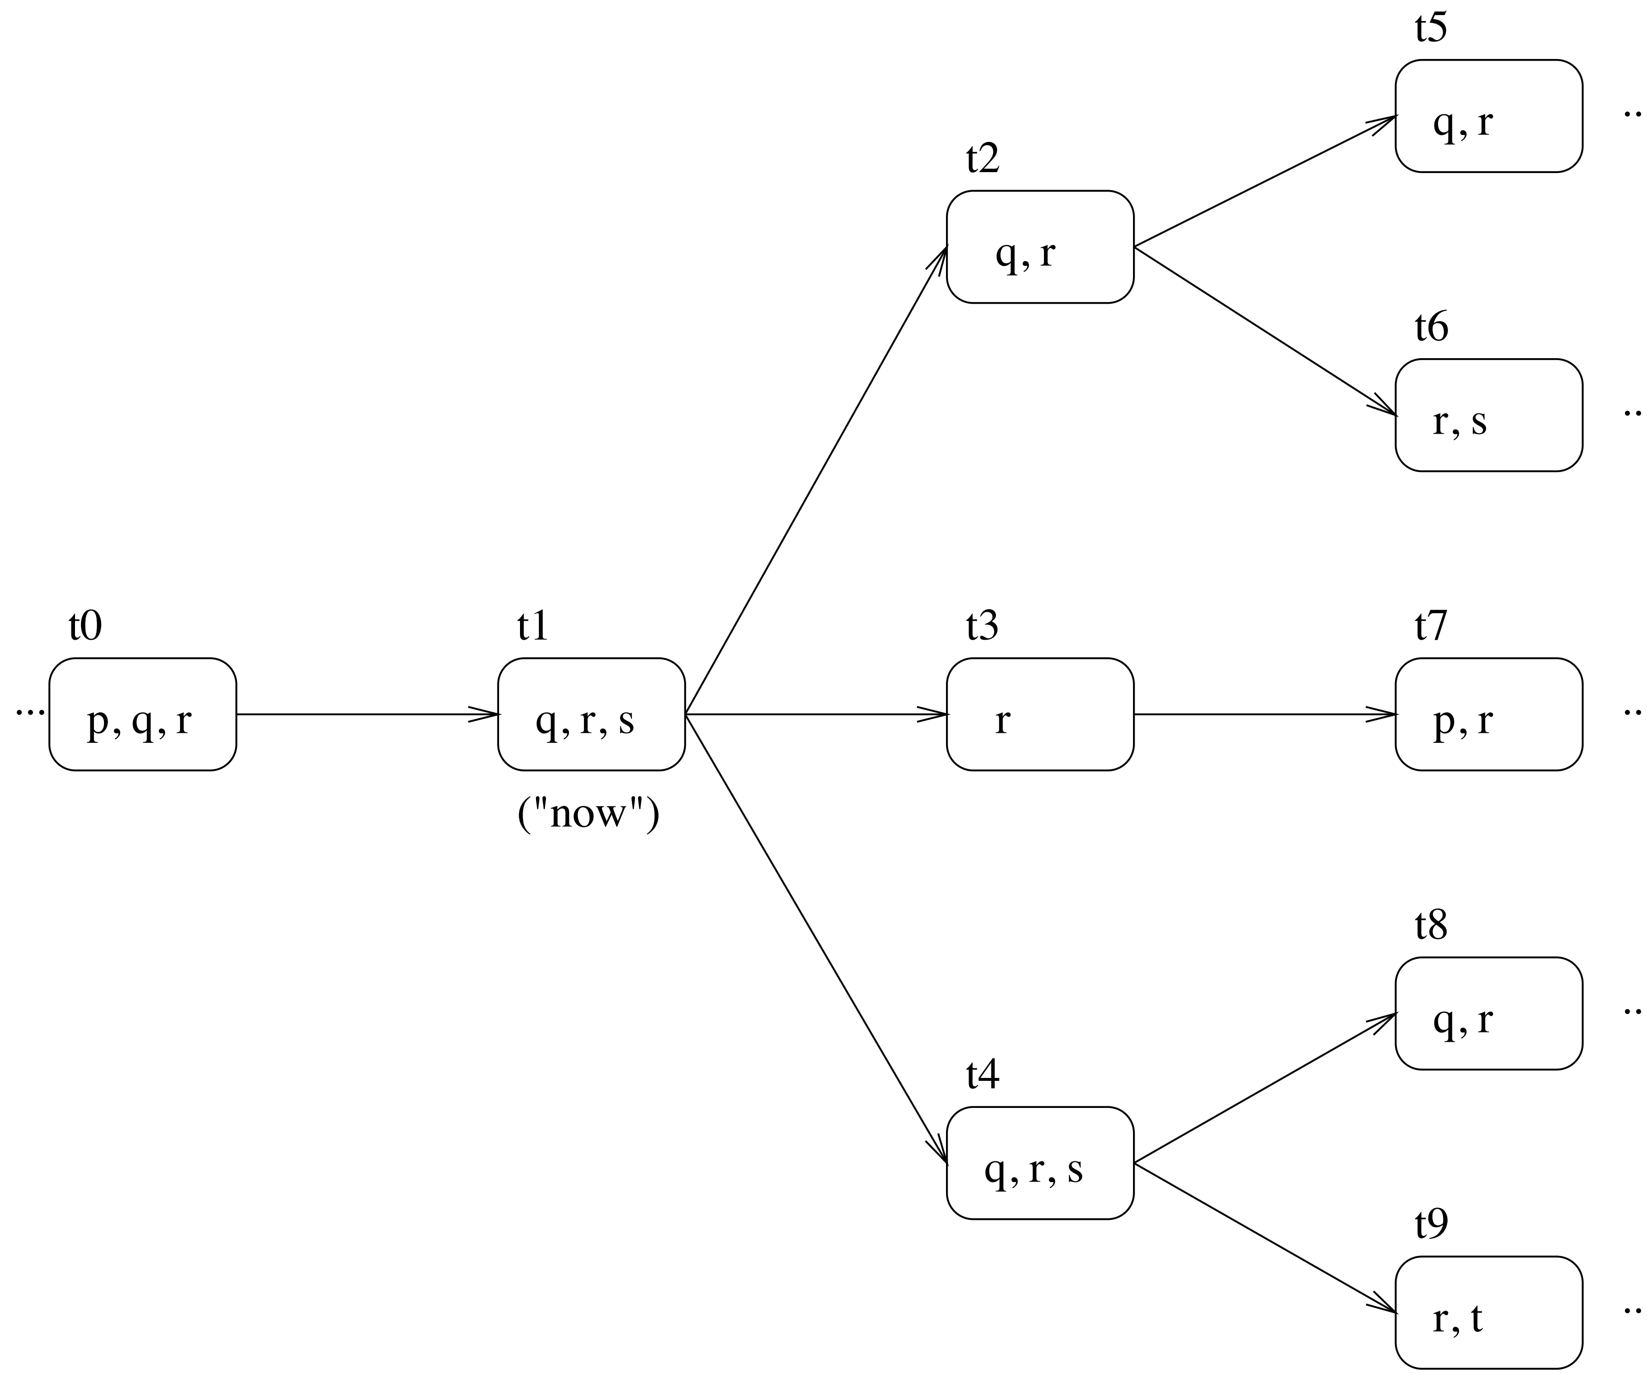

Computation Tree Logic (CTL)

- Extends propositional logic with
  - *path quantifiers* A, E
  - *tense modalities* $\bigcirc, \Diamond, \Box, \mathcal{U}$

CTL: syntax

$\mathsf{A}\bigcirc\phi$	"on all paths, ϕ is true next
$\mathsf{A}\Diamond\phi$	"on all paths, ϕ is eventually true
$\mathsf{A}\,\Box\,\phi$	"on all paths, ϕ is always true
$\mathsf{A}\phi\,\mathcal{U}\,\psi$	"on all paths, ϕ is true until ψ
$\mathsf{E}\bigcirc\phi$	"on some path, ϕ is true next
$\mathsf{E}\Diamond\phi$	"on some path, ϕ is eventually true
$\mathsf{E}\,\Box\,\phi$	"on some path, ϕ is always true
$\mathsf{E}\phi\,\mathcal{U}\,\psi$	"on some path, ϕ is true until ψ

CTL: models

Models for CTL are *Kripke structures*:

$$\langle S, R, \pi \rangle$$

where

- S is the set of possible system states
- $R \subseteq S \times S$ is a *next state* relation
- $\pi : S \to 2^{\Pi}$ says which propositions are true in each state.

The branches are obtained by *unwinding* this relation, giving *paths* through the structure.

Example 1

$M, t_1 \models \mathsf{A}\bigcirc r$

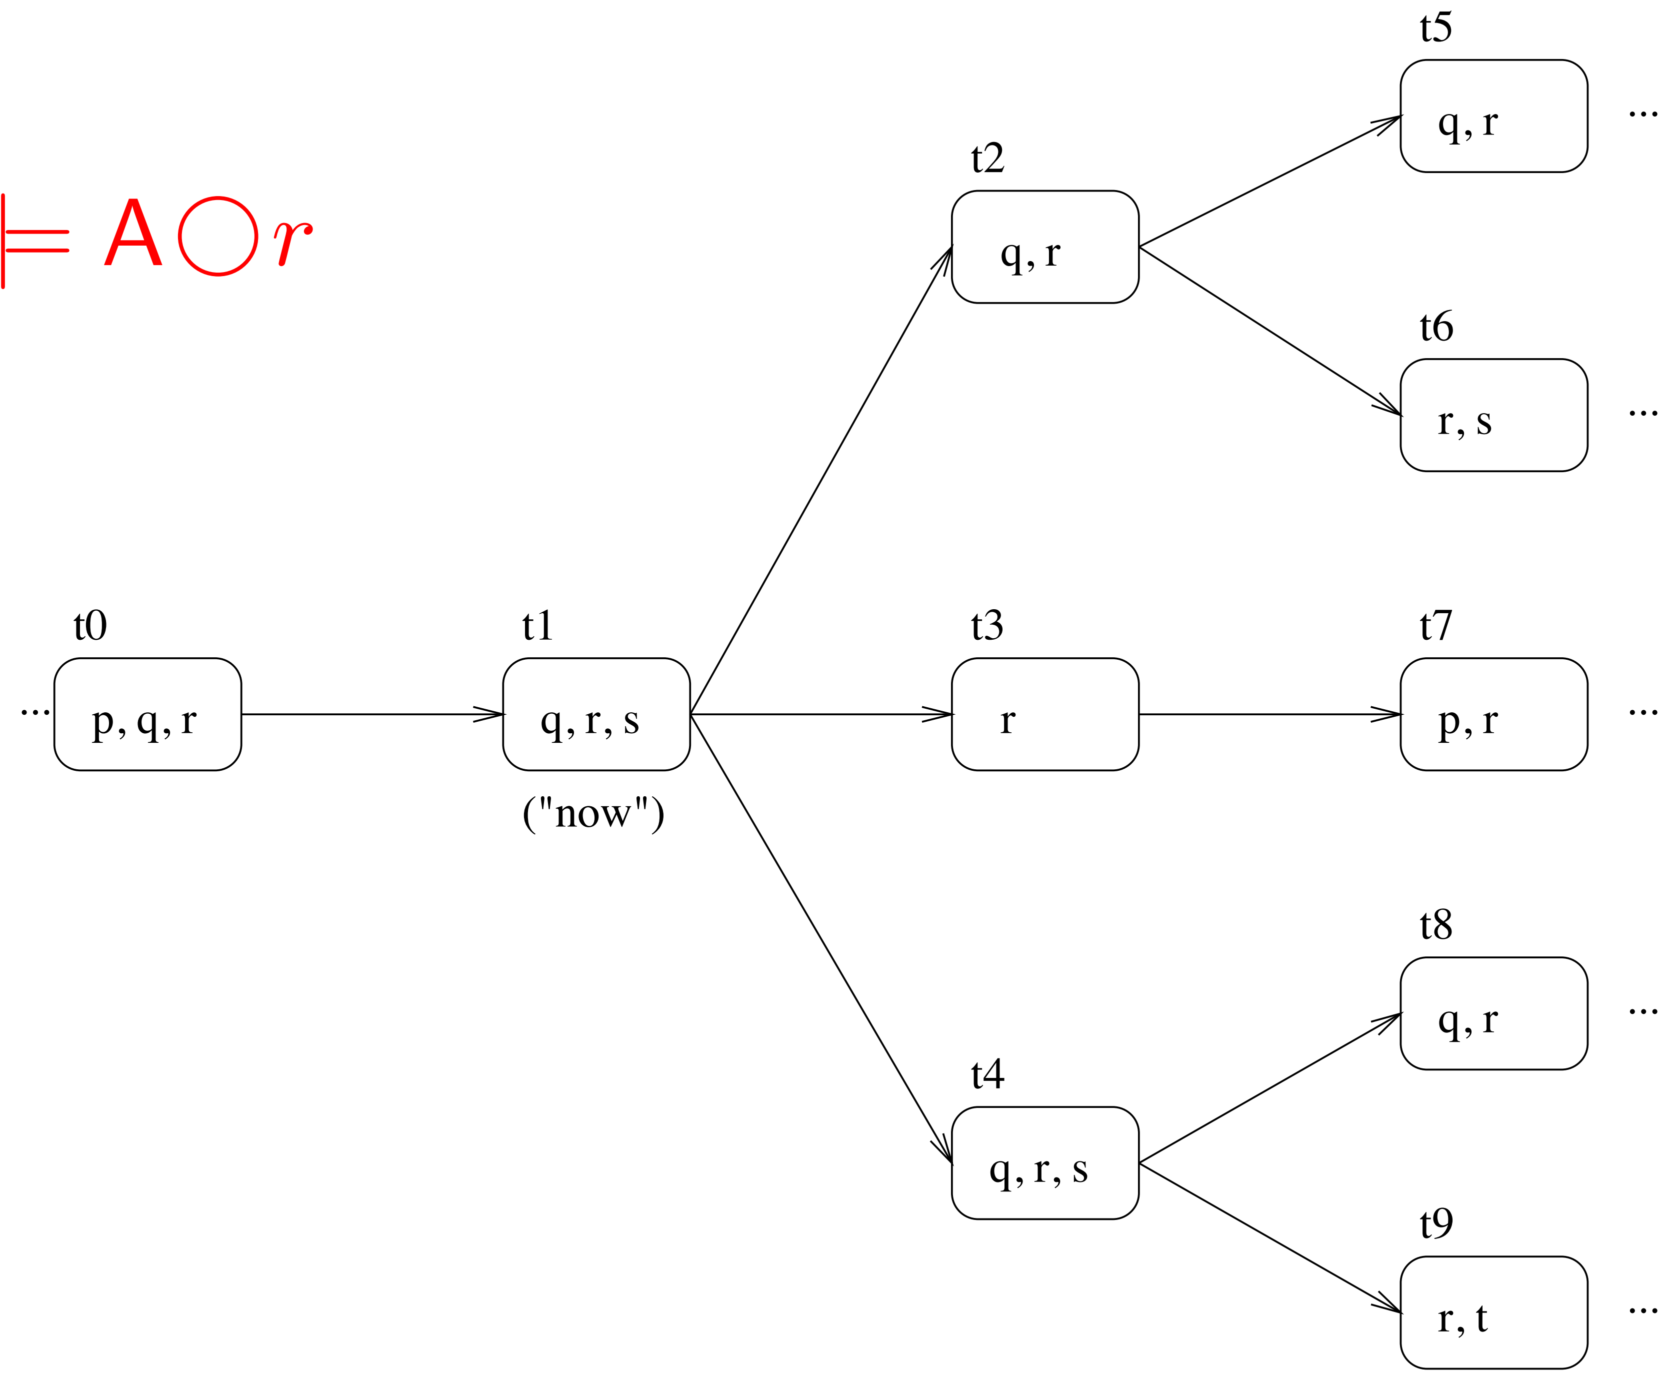

Example 2

$M, t_1 \models \mathsf{A}\,\Box r$

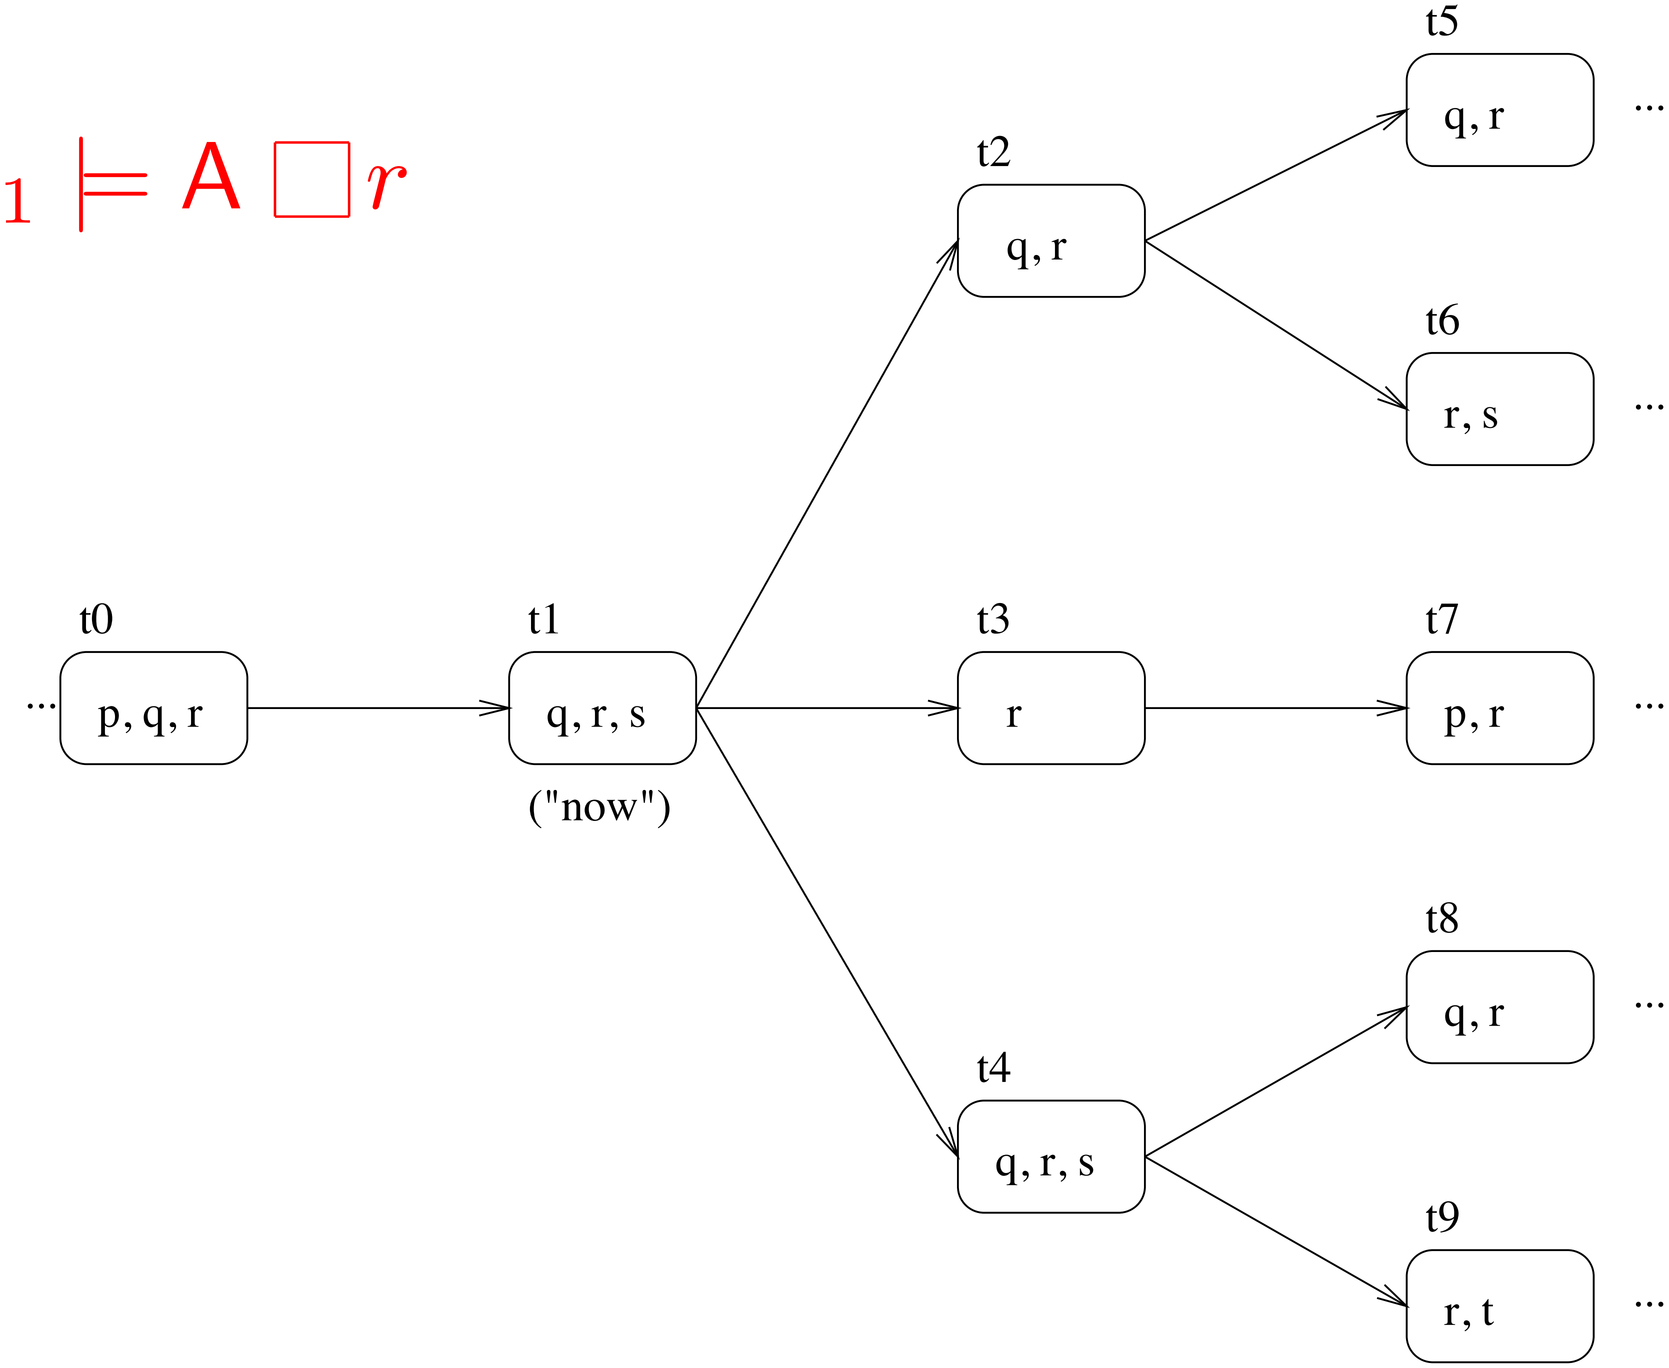

Example 3

$M, t_1 \models \mathsf{E}\,\Box q$

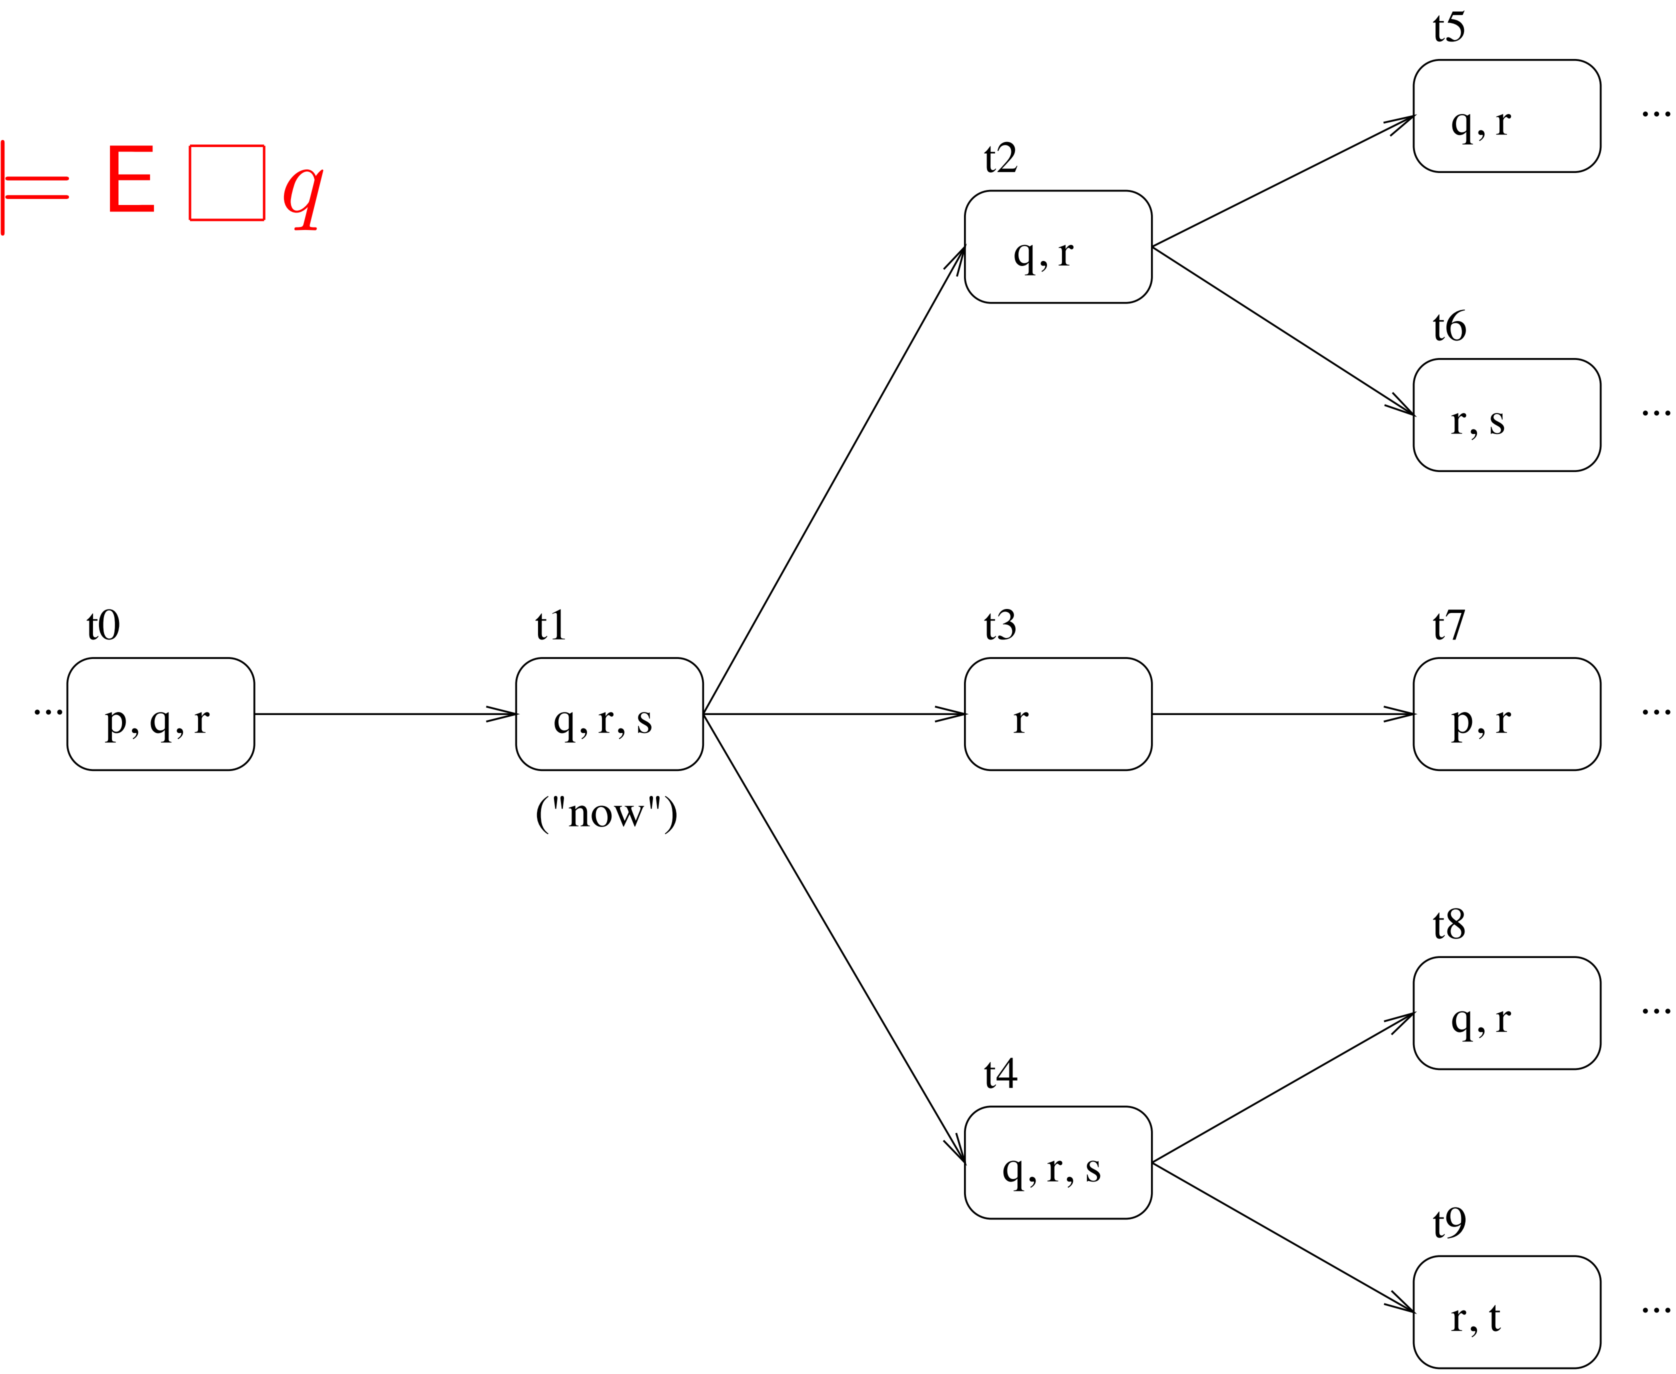

Example 4

$M, t_1 \models \mathsf{E}(s\,\mathcal{U}\,t)$

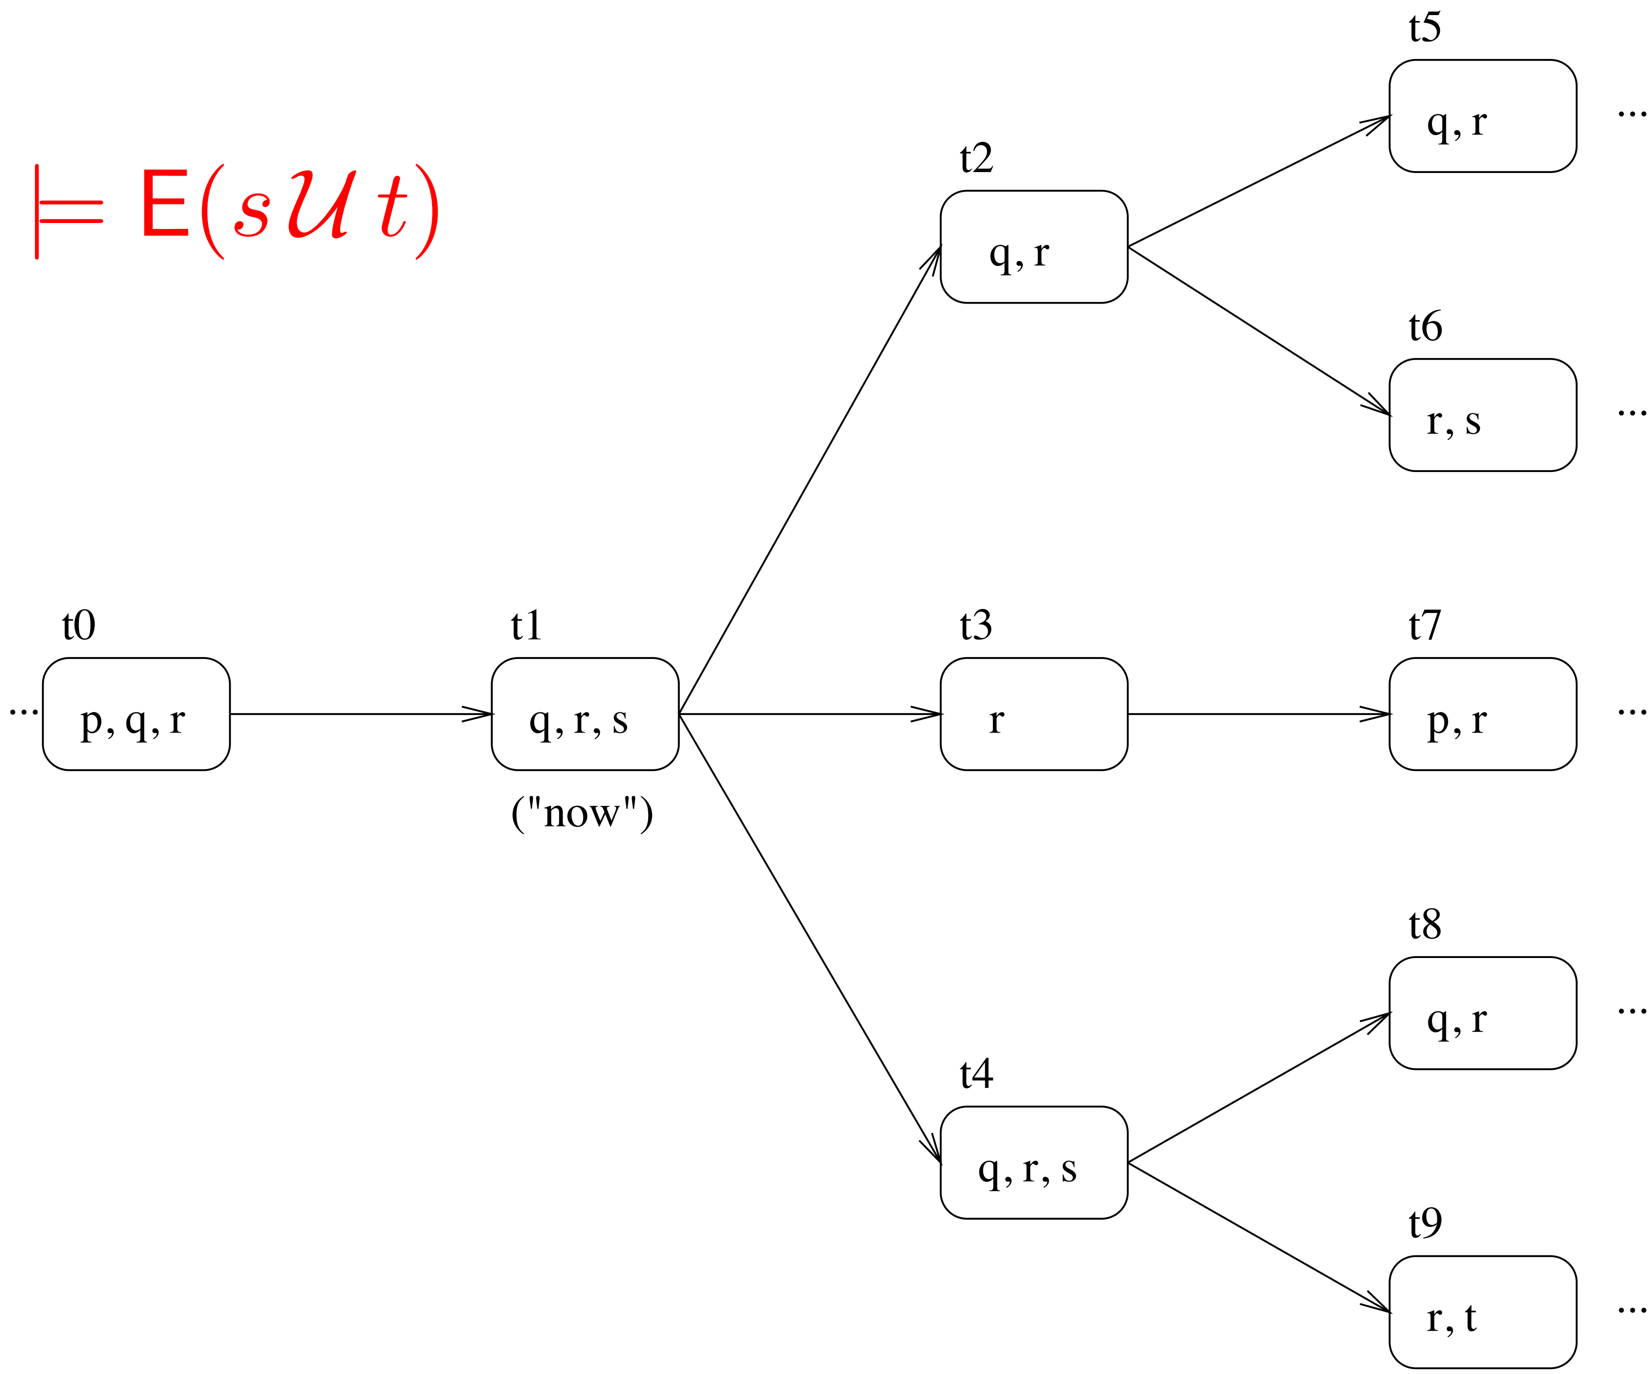

Contents

- Branching-time temporal logic: CTL
- ATL
- Bisimulations and the role of memory
- Irrevocable strategies

Alternating-time Temporal Logic (ATL)

- No notion of *agency* in CTL.
- In 1997, Alur, Henzinger & Kupferman proposed *Alternating-time Temporal Logic* (ATL).
- Branching used to model evolution of a system controlled by *agents*, which can affect the future by making *choices*.
- The particular future that will emerge depends on *combination* of choices that agents make.
- A temporal logic built on *agency*.

Coalition operators

In ATL the path quantifiers A, E are replaced by coalition operators:

$$\langle\!\langle G \rangle\!\rangle \phi$$

means

"group G has the ability to make ϕ true, no matter what the other agents do"

equivalently:

"G have a collective strategy to force ϕ"

ATL: syntax

Let N be set of all agents, Θ be set of atomic propositions:

$$
\begin{array}{rcll}
\phi & ::= & \top & \text{(truth constant)} \\
& | & p & \text{(primitive propositions)} \\
& | & \neg\phi & \text{(negation)} \\
& | & \phi \wedge \phi & \text{(conjunction)} \\
& | & \langle\!\langle C \rangle\!\rangle \bigcirc \phi & \text{(next)} \\
& | & \langle\!\langle C \rangle\!\rangle \Box \phi & \text{(always)} \\
& | & \langle\!\langle C \rangle\!\rangle \phi \mathcal{U} \phi & \text{(until)}
\end{array}
$$

where $C \subseteq N$ and $p \in \Theta$. Derived: $\langle\!\langle C \rangle\!\rangle \Diamond \phi \equiv \langle\!\langle C \rangle\!\rangle (\top \mathcal{U} \phi)$

Examples

$$\langle\!\llbracket thomas \rrbracket\!\rangle \Diamond\, bored\, audience$$

Examples

$$\langle\!\![thomas]\!\!\rangle \Diamond\, bored\, audience$$

$$\neg \langle\!\![thomas]\!\!\rangle\, \Box\, excited$$

Examples

$$\langle\!\langle thomas \rangle\!\rangle \Diamond bored\, audience$$

$$\neg \langle\!\langle thomas \rangle\!\rangle \, \Box \, excited$$

$$\langle\!\langle Thomas, Meiyun \rangle\!\rangle \Diamond students\_happy$$

Examples

$$\langle\!\langle thomas \rangle\!\rangle \Diamond bored\, audience$$

$$\neg \langle\!\langle thomas \rangle\!\rangle \, \Box \, excited$$

$$\langle\!\langle Thomas, Meiyun \rangle\!\rangle \Diamond students\_happy$$

$$\langle\!\langle 1 \rangle\!\rangle \neg enter\, \mathcal{U}\, permission$$

Examples

$$\langle\!\langle thomas \rangle\!\rangle \Diamond bored\,audience$$

$$\neg \langle\!\langle thomas \rangle\!\rangle \Box\, excited$$

$$\langle\!\langle Thomas, Meiyun \rangle\!\rangle \Diamond students\_happy$$

$$\langle\!\langle 1 \rangle\!\rangle \neg enter\, \mathcal{U}\; permission$$

$$\langle\!\langle Ann \rangle\!\rangle \Box \langle\!\langle Bob \rangle\!\rangle \Diamond win$$

ATL models: concurrent game structures

A concurrent game structure is a tuple $M = \langle N, S, \pi, Act, d, o \rangle$, where:

- N: a finite set of all agents
- S: a set of states
- π: a valuation of propositions
- Act: a finite set of (atomic) actions
- $d : N \times S \rightarrow \wp(Act)$ 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

CGS: example

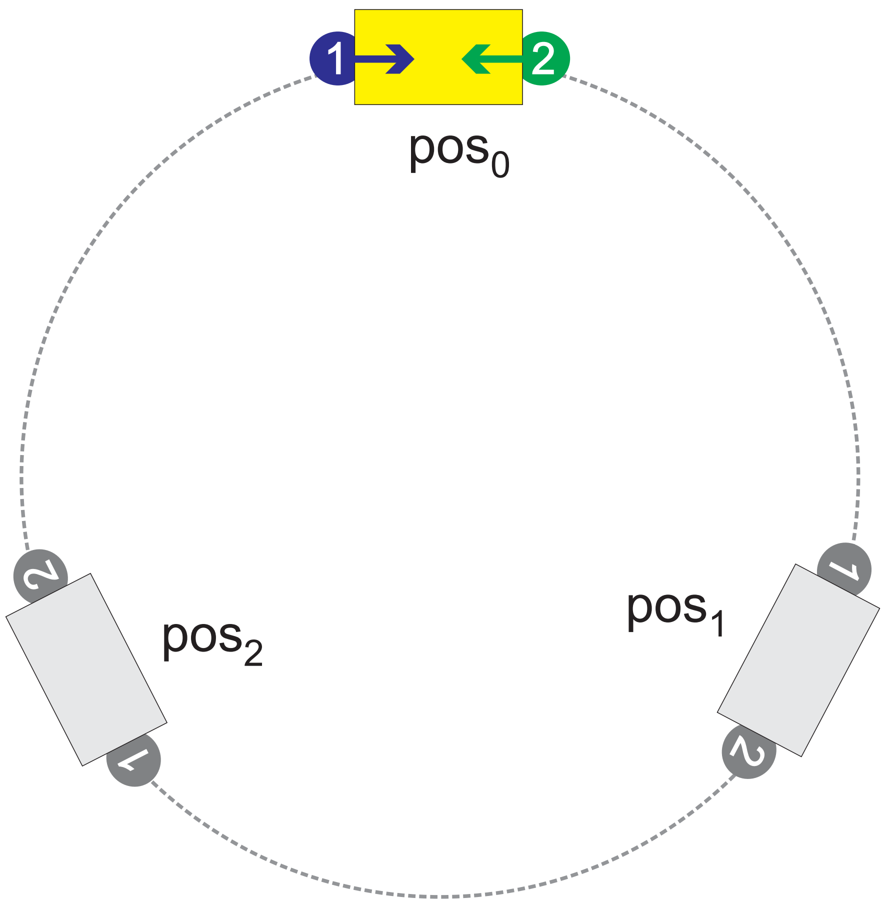

CGS: example

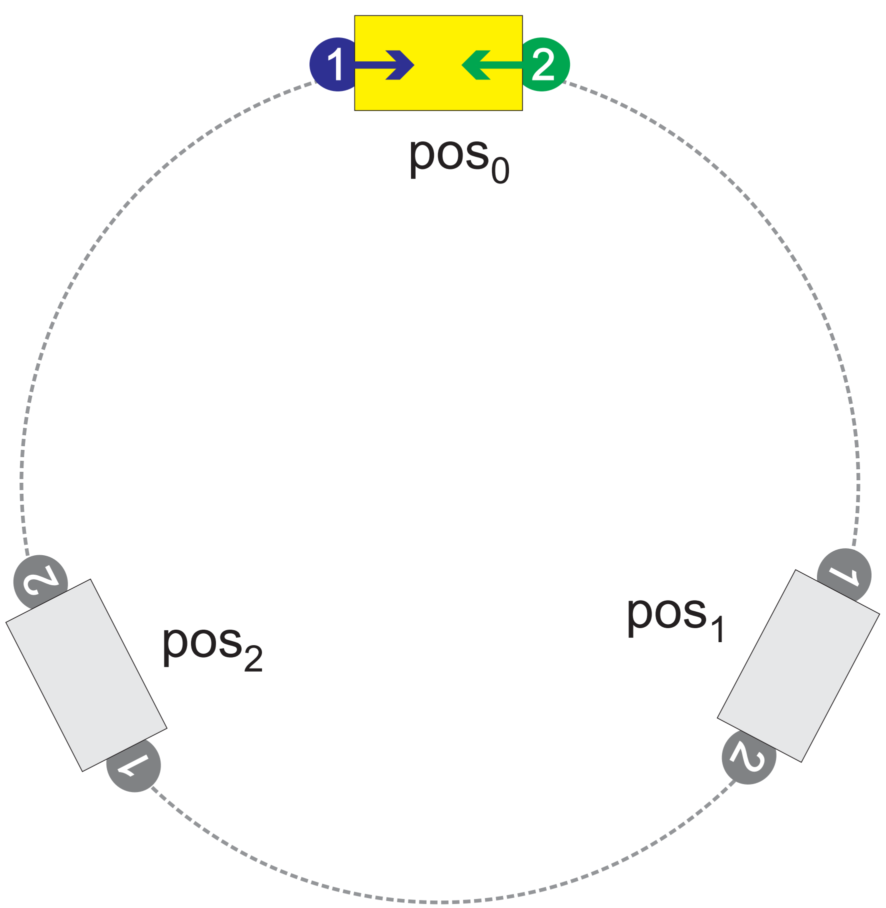
1
2
pos0
pos2
pos1

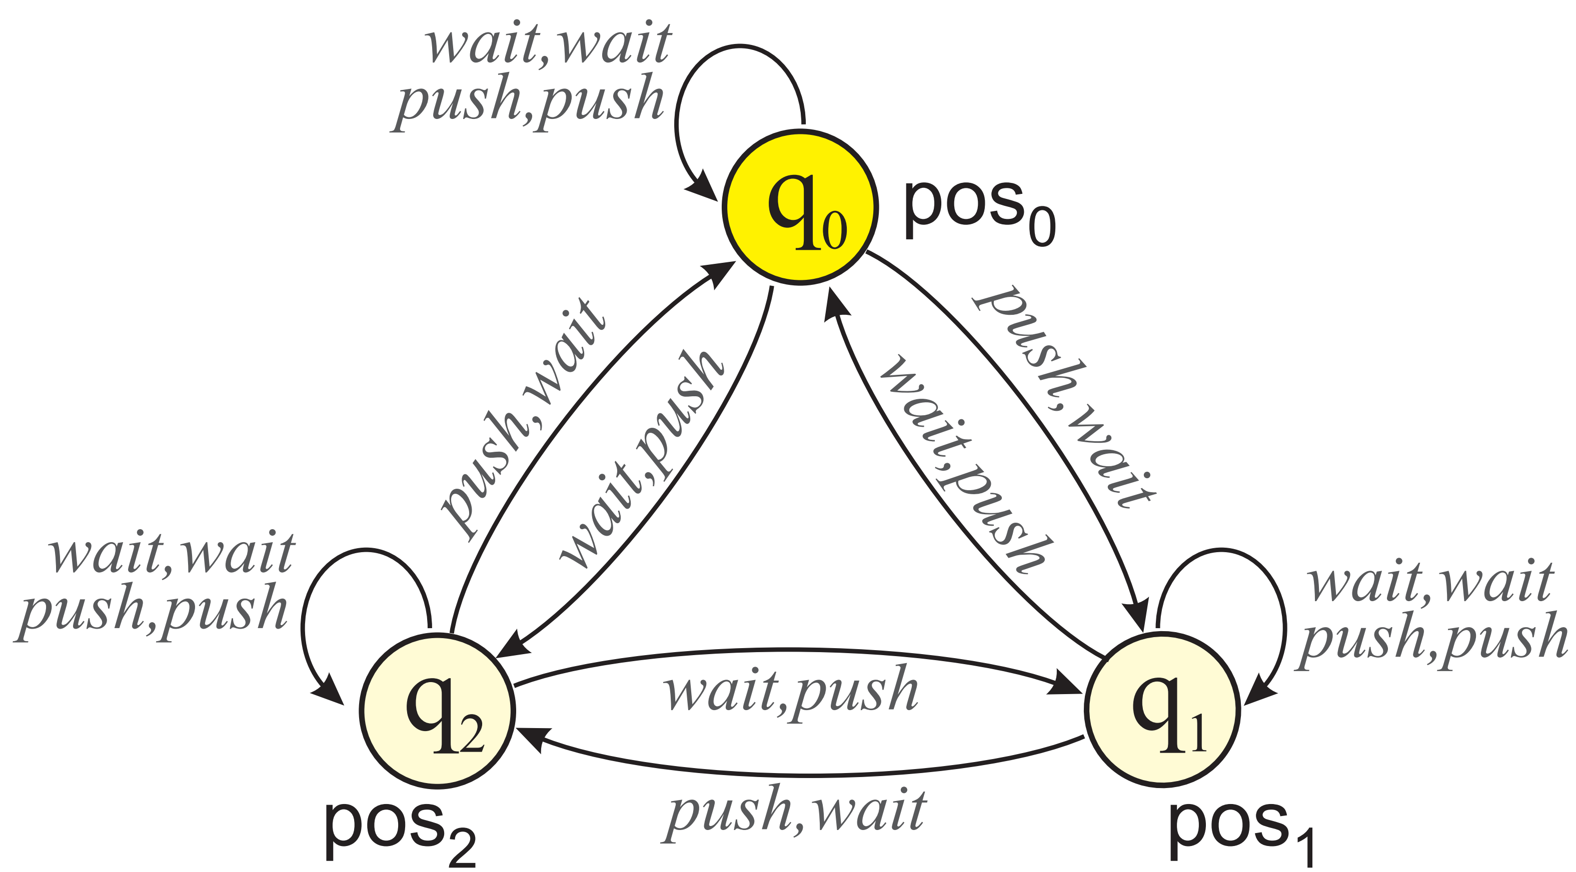
wait,wait
push,push
q0
pos0
push,wait
wait,push
wait,push
push,wait
wait,wait
push,push
q2
wait,push
q1
wait,wait
push,push
pos2
push,wait
pos1

Strategies and paths

A strategy for an agent a is a function

$$f_a : S \to Act$$

such that $f_a(s) \in d(a, s)$ for any state $s \in S$.

Strategies and paths

A strategy for an agent a is a function

$$f_a : S \to Act$$

such that $f_a(s) \in d(a, s)$ for any state $s \in S$.

A strategy for a coalition G is a set of one strategy for each agent in G

$$f_G = \{f_a : a \in G\}$$

Strategies and paths

A strategy for an agent a is a function

$$f_a : S \to Act$$

such that $f_a(s) \in d(a, s)$ for any state $s \in S$.

A strategy for a coalition G is a set of one strategy for each agent in G

$$f_G = \{f_a : a \in G\}$$

A path is an infinite sequence of states $s_1, s_2, s_3, \ldots$.

Strategies and paths

A strategy for an agent a is a function

$$f_a : S \to Act$$

such that $f_a(s) \in d(a, s)$ for any state $s \in S$.

A strategy for a coalition G is a set of one strategy for each agent in G

$$f_G = \{f_a : a \in G\}$$

A path is an infinite sequence of states $s_1, s_2, s_3, \ldots$.

$out(s, f_G)$ denotes the set of all possible paths starting in s where the agents in G uses the strategies in f_G.

ATL: semantics

$M, q \models \langle\!\langle A \rangle\!\rangle \Box \varphi$ iff there is f_A such that, for every $\lambda \in out(q, f_A)$, we have $M, \lambda[i] \models \varphi$ for all $i \geq 0$;

ATL: semantics

$M, q \models p$	iff p is in $\pi(q)$;
$M, q \models \neg\varphi$	iff $M, q \not\models \varphi$;
$M, q \models \varphi_1 \wedge \varphi_2$	iff $M, q \models \varphi_1$ and $M, q \models \varphi_2$;
$M, q \models \langle\!\langle A \rangle\!\rangle \bigcirc \varphi$	iff there is f_A such that, for every $\lambda \in out(q, f_A)$, we have $M, \lambda[1] \models \varphi$;
$M, q \models \langle\!\langle A \rangle\!\rangle \Box \varphi$	iff there is f_A such that, for every $\lambda \in out(q, f_A)$, we have $M, \lambda[i] \models \varphi$ for all $i \geq 0$;
$M, q \models \langle\!\langle A \rangle\!\rangle \varphi_1 \mathcal{U} \varphi_2$	iff there is f_A such that, for every $\lambda \in out(q, f_A)$, we have $M, \lambda[i] \models \varphi_2$ for some $i \geq 0$ and $M, \lambda[j] \models \varphi_1$ for all $0 \leq j \leq i$.

Example

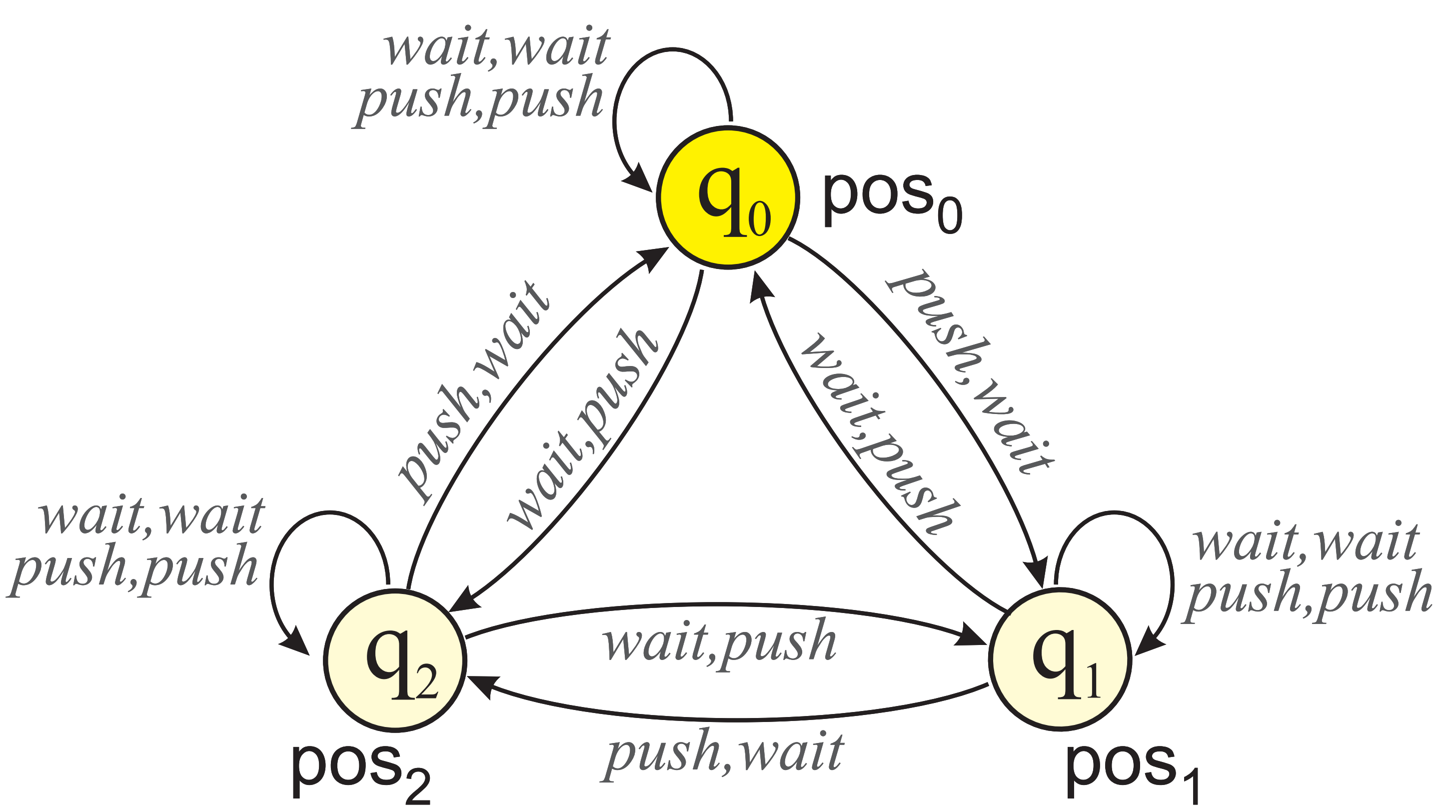

$$pos_0 \rightarrow \langle\!\langle 1 \rangle\!\rangle \Box \neg pos_1$$

Example

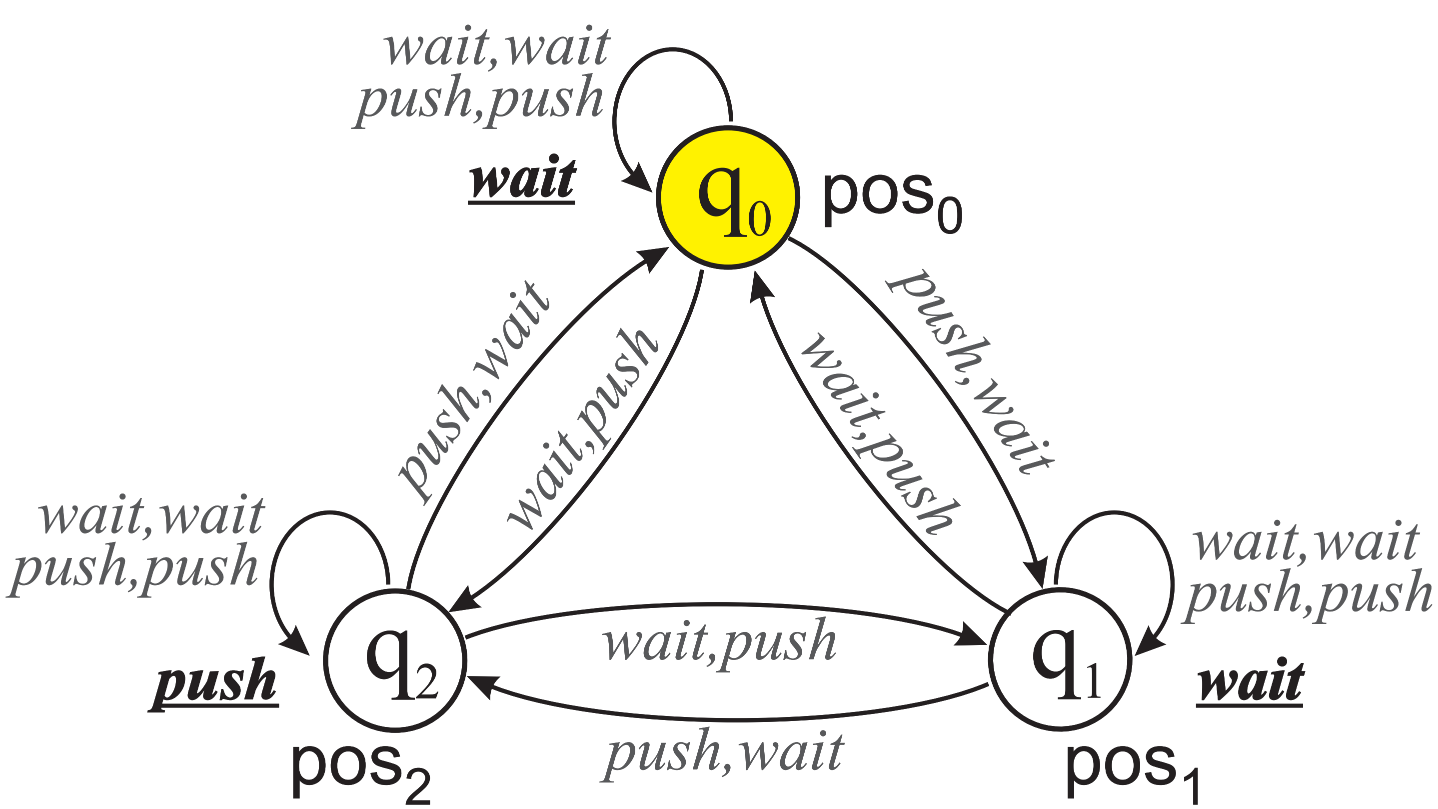

$$pos_0 \rightarrow \langle\!\langle 1 \rangle\!\rangle \Box \neg pos_1$$

Example

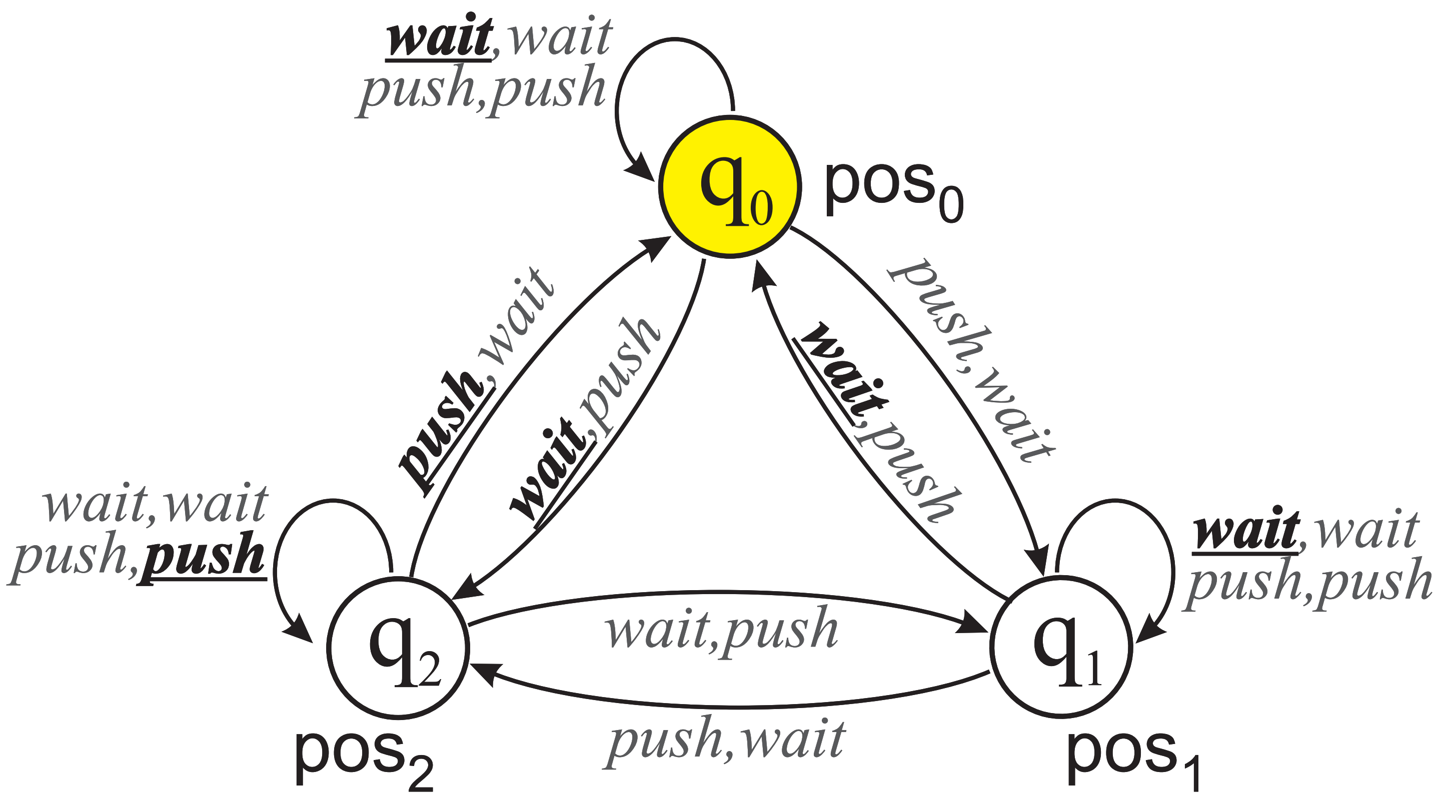

$$pos_0 \rightarrow \langle\!\langle 1 \rangle\!\rangle \Box \neg pos_1$$

Example

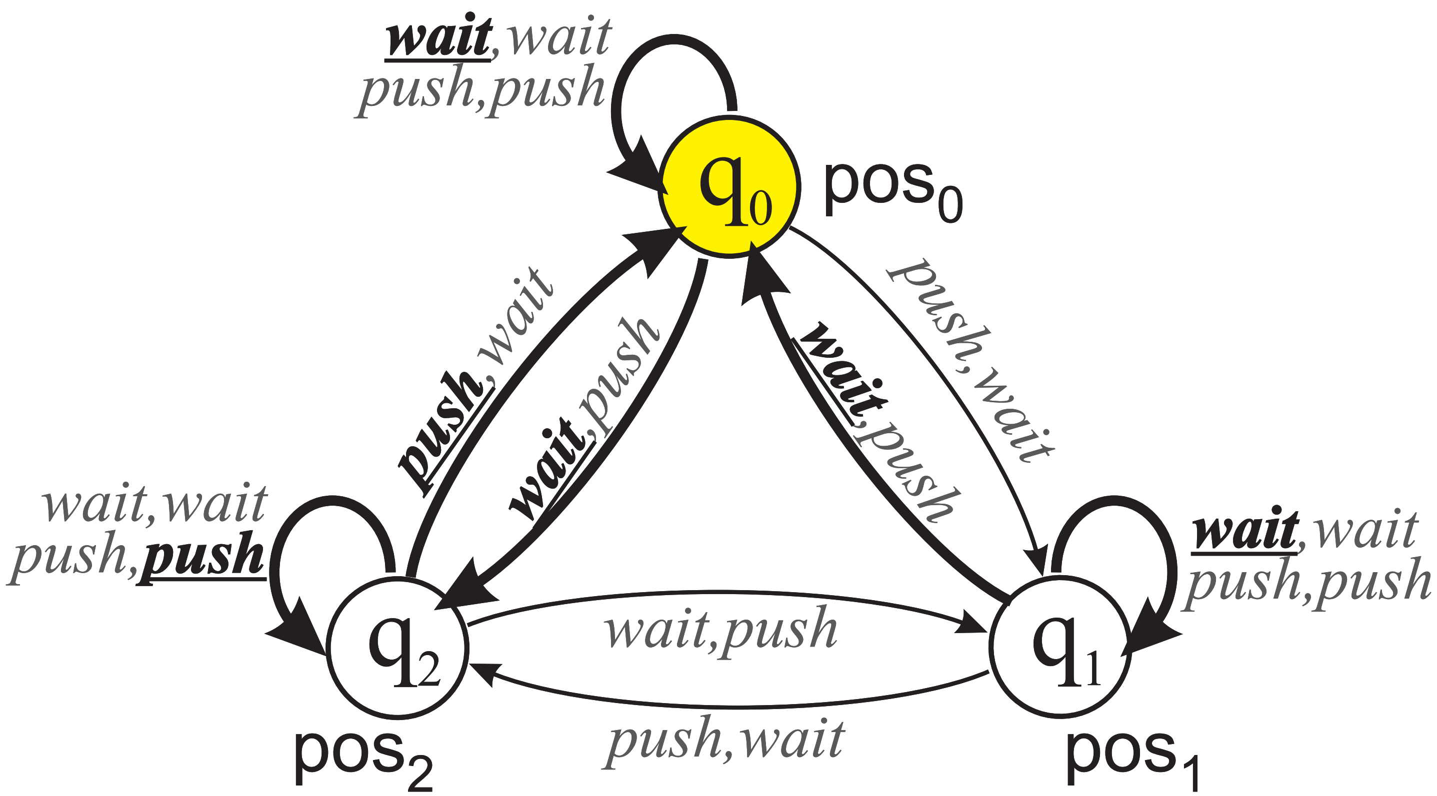

$$pos_0 \rightarrow \langle\!\langle 1 \rangle\!\rangle \Box \neg pos_1$$

Example

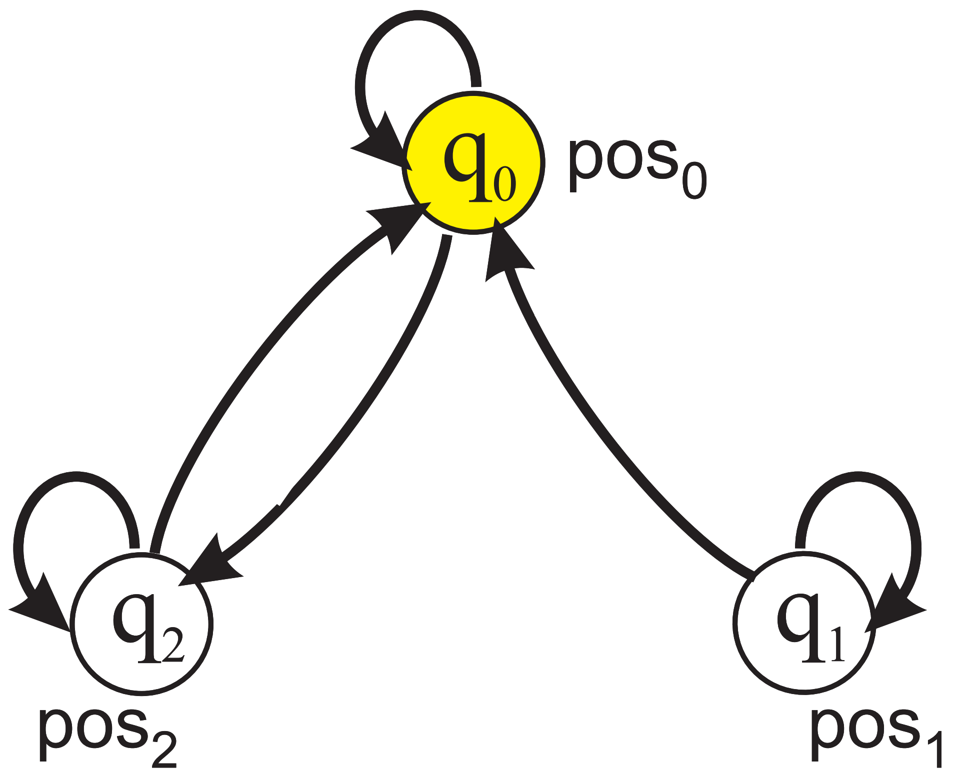

$$pos_0 \rightarrow \langle\!\langle 1 \rangle\!\rangle \Box \neg pos_1$$

Example

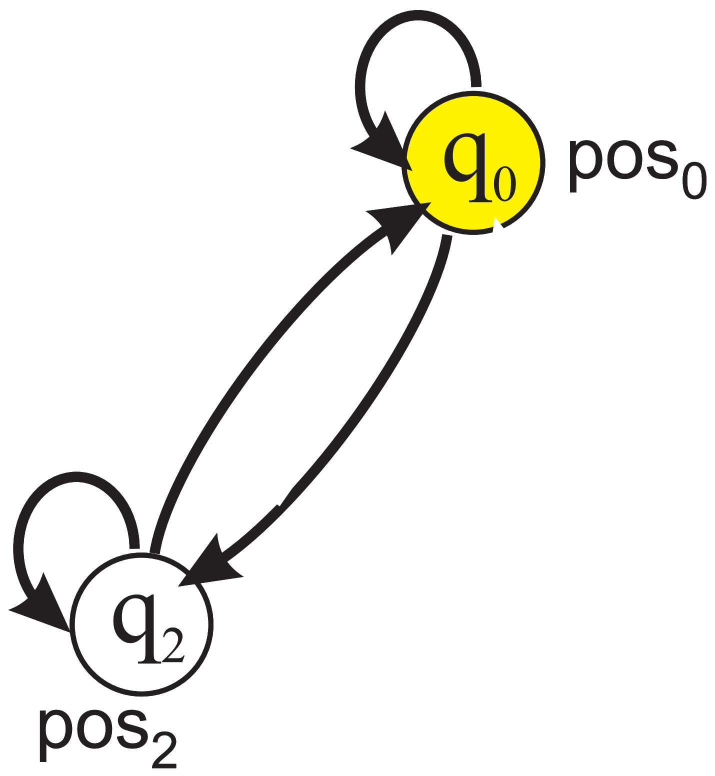

$$pos_0 \rightarrow \langle\!\langle 1 \rangle\!\rangle \,\Box \neg pos_1$$

ATL as an extension of CTL

- A $\equiv \langle\!\langle \emptyset \rangle\!\rangle$ $\quad$ ("for all paths")
  E $\equiv \langle\!\langle N \rangle\!\rangle$ $\quad$ ("there is a path")

ATL as an extension of CL

- $\langle\![G]\!\rangle \phi \equiv \langle\![G]\!\rangle \bigcirc \phi$

- Concurrent game structures are equivalent to game models

ATL and games

- Concurrent game structure:

$$\langle\!\langle G \rangle\!\rangle \phi$$

  - sequence of strategic form games
  - generalised extensive form game
- Coalition operator splits the players into proponents G and opponents N\G
  - True iff proponents have a winning strategy
  - Flexible and compact specification of winning conditions

ATL and games

- Model checking: finding a winning strategy

- Satisfiability checking: mechanism design

ATL*

ATL* is a generalisation of ATL where coalition operators and temporal operators can be mixed freely:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\!\langle A \rangle\!\rangle \gamma,$$
$$\gamma ::= \varphi \mid \neg\gamma \mid \gamma \wedge \gamma \mid \bigcirc\gamma \mid \Diamond\gamma \mid \;\Box\gamma \mid \gamma\,\mathcal{U}\,\gamma.$$

ATL*: example

$$\langle\!\langle producer, dealer \rangle\!\rangle \Box (carRequested \rightarrow \Diamond carDelivered)$$

ATL*: semantics

$M, q \models p$ iff p is in $\pi(q)$;
$M, q \models \neg\varphi$ iff $M, q \not\models \varphi$;
$M, q \models \varphi_1 \wedge \varphi_2$ iff $M, q \models \varphi_1$ and $M, q \models \varphi_2$;
$M, \lambda \models \neg\gamma$ iff $M, q \not\models \gamma$

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

$M, \lambda \models \bigcirc\gamma$ iff $M, \lambda[1..\infty] \models \gamma$;
$M, \lambda \models \Box\gamma$ iff $M, \lambda[i..\infty] \models \gamma$ for all $i \geq 0$;
$M, \lambda \models \gamma_1 \mathcal{U} \gamma_2$ iff $M, \lambda[i..\infty] \models \gamma_2$ for some $i \geq 0$, and $M, \lambda[j..\infty] \models \gamma_1$ forall $0 \leq j \leq i$.

Fixpoint properties

- $\langle\!\langle A \rangle\!\rangle \, \Box \varphi \quad \leftrightarrow \quad \varphi \wedge \langle\!\langle A \rangle\!\rangle \bigcirc \langle\!\langle A \rangle\!\rangle \, \Box \varphi$
- $\langle\!\langle A \rangle\!\rangle \varphi_1 \, \mathcal{U} \, \varphi_2 \quad \leftrightarrow \quad \varphi_2 \vee \varphi_1 \wedge \langle\!\langle A \rangle\!\rangle \bigcirc \langle\!\langle A \rangle\!\rangle \varphi_1 \, \mathcal{U} \, \varphi_2$

Fixpoint properties

- $\langle\!\langle A \rangle\!\rangle \, \Box \varphi \quad \leftrightarrow \quad \varphi \wedge \langle\!\langle A \rangle\!\rangle \bigcirc \langle\!\langle A \rangle\!\rangle \, \Box \varphi$
- $\langle\!\langle A \rangle\!\rangle \varphi_1 \, \mathcal{U} \, \varphi_2 \quad \leftrightarrow \quad \varphi_2 \vee \varphi_1 \wedge \langle\!\langle A \rangle\!\rangle \bigcirc \langle\!\langle A \rangle\!\rangle \varphi_1 \, \mathcal{U} \, \varphi_2$

ATL: valid

Fixpoint properties

- $\langle\!\langle A \rangle\!\rangle \,\Box\, \varphi \quad \leftrightarrow \quad \varphi \wedge \langle\!\langle A \rangle\!\rangle \bigcirc \langle\!\langle A \rangle\!\rangle \,\Box\, \varphi$
- $\langle\!\langle A \rangle\!\rangle \varphi_1 \,\mathcal{U}\, \varphi_2 \quad \leftrightarrow \quad \varphi_2 \vee \varphi_1 \wedge \langle\!\langle A \rangle\!\rangle \bigcirc \langle\!\langle A \rangle\!\rangle \varphi_1 \,\mathcal{U}\, \varphi_2$

ATL: valid

ATL*: not valid

ATL: axioms

$(\bot)$ $\neg\langle\!\langle C\rangle\!\rangle \bigcirc \bot$

$(\top)$ $\langle\!\langle C\rangle\!\rangle \bigcirc \top$

(N) $\neg\langle\!\langle \emptyset\rangle\!\rangle \bigcirc \neg\varphi \to \langle\!\langle N\rangle\!\rangle \bigcirc \varphi$

(S) $\langle\!\langle C_1\rangle\!\rangle \bigcirc \varphi_1 \wedge \langle\!\langle C_2\rangle\!\rangle \bigcirc \varphi_2 \to \langle\!\langle C_1 \cup C_2\rangle\!\rangle \bigcirc (\varphi_1 \wedge \varphi_2)$, $C_1 \cap C_2 = \emptyset$

$(FP_\Box)$ $\langle\!\langle C\rangle\!\rangle \Box\varphi \leftrightarrow \varphi \wedge \langle\!\langle C\rangle\!\rangle \bigcirc \langle\!\langle C\rangle\!\rangle \Box\varphi$

$(GFP_\Box)$ $\langle\!\langle \emptyset\rangle\!\rangle \Box(\theta \to (\varphi \wedge \langle\!\langle C\rangle\!\rangle \bigcirc \theta)) \to \langle\!\langle \emptyset\rangle\!\rangle \Box(\theta \to \langle\!\langle C\rangle\!\rangle \Box\varphi)$

$(FP_{\mathcal{U}})$ $\langle\!\langle C\rangle\!\rangle(\varphi_1 \,\mathcal{U}\, \varphi_2) \leftrightarrow \varphi_2 \vee (\varphi_1 \wedge \langle\!\langle C\rangle\!\rangle \bigcirc \langle\!\langle C\rangle\!\rangle(\varphi_1 \,\mathcal{U}\, \varphi_2))$

$(LFP_{\mathcal{U}})$ $\langle\!\langle \emptyset\rangle\!\rangle \Box((\varphi_2 \vee (\varphi_1 \wedge \langle\!\langle C\rangle\!\rangle \bigcirc \theta)) \to \theta) \to \langle\!\langle \emptyset\rangle\!\rangle \Box(\langle\!\langle C\rangle\!\rangle(\varphi_1 \,\mathcal{U}\, \varphi_2) \to \theta)$

$$\frac{\varphi_1, \varphi_1 \to \varphi_2}{\varphi_2}(MP) \quad \frac{\varphi_1 \to \varphi_2}{\langle\!\langle C\rangle\!\rangle \bigcirc \varphi_1 \to \langle\!\langle C\rangle\!\rangle \bigcirc \varphi_2}(Mon) \quad \frac{\varphi}{\langle\!\langle \emptyset\rangle\!\rangle \Box\varphi}(Nec)$$

ATL: axioms

Sound and complete

$(\bot)$ $\neg\langle\!\langle C\rangle\!\rangle\bigcirc\bot$

$(\top)$ $\langle\!\langle C\rangle\!\rangle\bigcirc\top$

(N) $\neg\langle\!\langle\emptyset\rangle\!\rangle\bigcirc\neg\varphi\rightarrow\langle\!\langle N\rangle\!\rangle\bigcirc\varphi$

(S) $\langle\!\langle C_1\rangle\!\rangle\bigcirc\varphi_1\wedge\langle\!\langle C_2\rangle\!\rangle\bigcirc\varphi_2\rightarrow\langle\!\langle C_1\cup C_2\rangle\!\rangle\bigcirc(\varphi_1\wedge\varphi_2),\ C_1\cap C_2=\emptyset$

$(FP_\Box)$ $\langle\!\langle C\rangle\!\rangle\Box\varphi\leftrightarrow\varphi\wedge\langle\!\langle C\rangle\!\rangle\bigcirc\langle\!\langle C\rangle\!\rangle\Box\varphi$

$(GFP_\Box)$ $\langle\!\langle\emptyset\rangle\!\rangle\Box(\theta\rightarrow(\varphi\wedge\langle\!\langle C\rangle\!\rangle\bigcirc\theta))\rightarrow\langle\!\langle\emptyset\rangle\!\rangle\Box(\theta\rightarrow\langle\!\langle C\rangle\!\rangle\Box\varphi)$

$(FP_{\mathcal{U}})$ $\langle\!\langle C\rangle\!\rangle(\varphi_1\,\mathcal{U}\,\varphi_2)\leftrightarrow\varphi_2\vee(\varphi_1\wedge\langle\!\langle C\rangle\!\rangle\bigcirc\langle\!\langle C\rangle\!\rangle(\varphi_1\,\mathcal{U}\,\varphi_2))$

$(LFP_{\mathcal{U}})$ $\langle\!\langle\emptyset\rangle\!\rangle\Box((\varphi_2\vee(\varphi_1\wedge\langle\!\langle C\rangle\!\rangle\bigcirc\theta))\rightarrow\theta)\rightarrow\langle\!\langle\emptyset\rangle\!\rangle\Box(\langle\!\langle C\rangle\!\rangle(\varphi_1\,\mathcal{U}\,\varphi_2)\rightarrow\theta)$

$$\frac{\varphi_1,\varphi_1\rightarrow\varphi_2}{\varphi_2}(MP)\qquad\frac{\varphi_1\rightarrow\varphi_2}{\langle\!\langle C\rangle\!\rangle\bigcirc\varphi_1\rightarrow\langle\!\langle C\rangle\!\rangle\bigcirc\varphi_2}(Mon)\qquad\frac{\varphi}{\langle\!\langle\emptyset\rangle\!\rangle\Box\varphi}(Nec)$$

Contents

- Branching-time temporal logic: CTL
- ATL
- Bisimulations and the role of memory
- Irrevocable strategies

Some definitions

$$D(q, C) = \times_{i \in C} d(i, q)$$

When $\vec{a}_C \in D(q, C)$ let

$$next_M(q, \vec{a}_C) = \{\delta(q, \vec{b}) : \vec{b} \in D(q), a_i = b_i \text{ for all } i \in C\}$$

denote the set of possible next states in CGS M when coalition C choose actions $\vec{a}_C$.

Bisimulation for CGSs

Given CGS $M_1 = (Q_1, \pi_1, Act_1, d_1, \delta_1)$; CGS $M_2 = (Q_2, \pi_2, Act_2, d_2, \delta_2)$; $\beta \subseteq Q_1 \times Q_2$.

$M_1 \rightleftarrows_\beta^C M_2$ (for some $C \subseteq N$): for all q_1, q_2, $q_1 \beta q_2$ implies that

> Local harmony $\pi_1(q_1) = \pi_2(q_2)$;
>
> Forth For all joint actions $\vec{a}_C^1 \in D_1(q_1, C)$ for C, there exists a joint action $\vec{a}_C^2 \in D_2(q_2, C)$ for C such that for all states $s_2 \in next_{M_2}(q_2, \vec{a}_C^2)$, there exists a state $s_1 \in next_{M_1}(q_1, \vec{a}_C^1)$ such that $s_1 \beta s_2$;
>
> Back Likewise, for 1 and 2 swapped.

$M_1 \rightleftarrows_\beta M_2$: $M_1 \rightleftarrows_\beta^C M_2$ for every $C \subseteq N$

Bisimulation: example

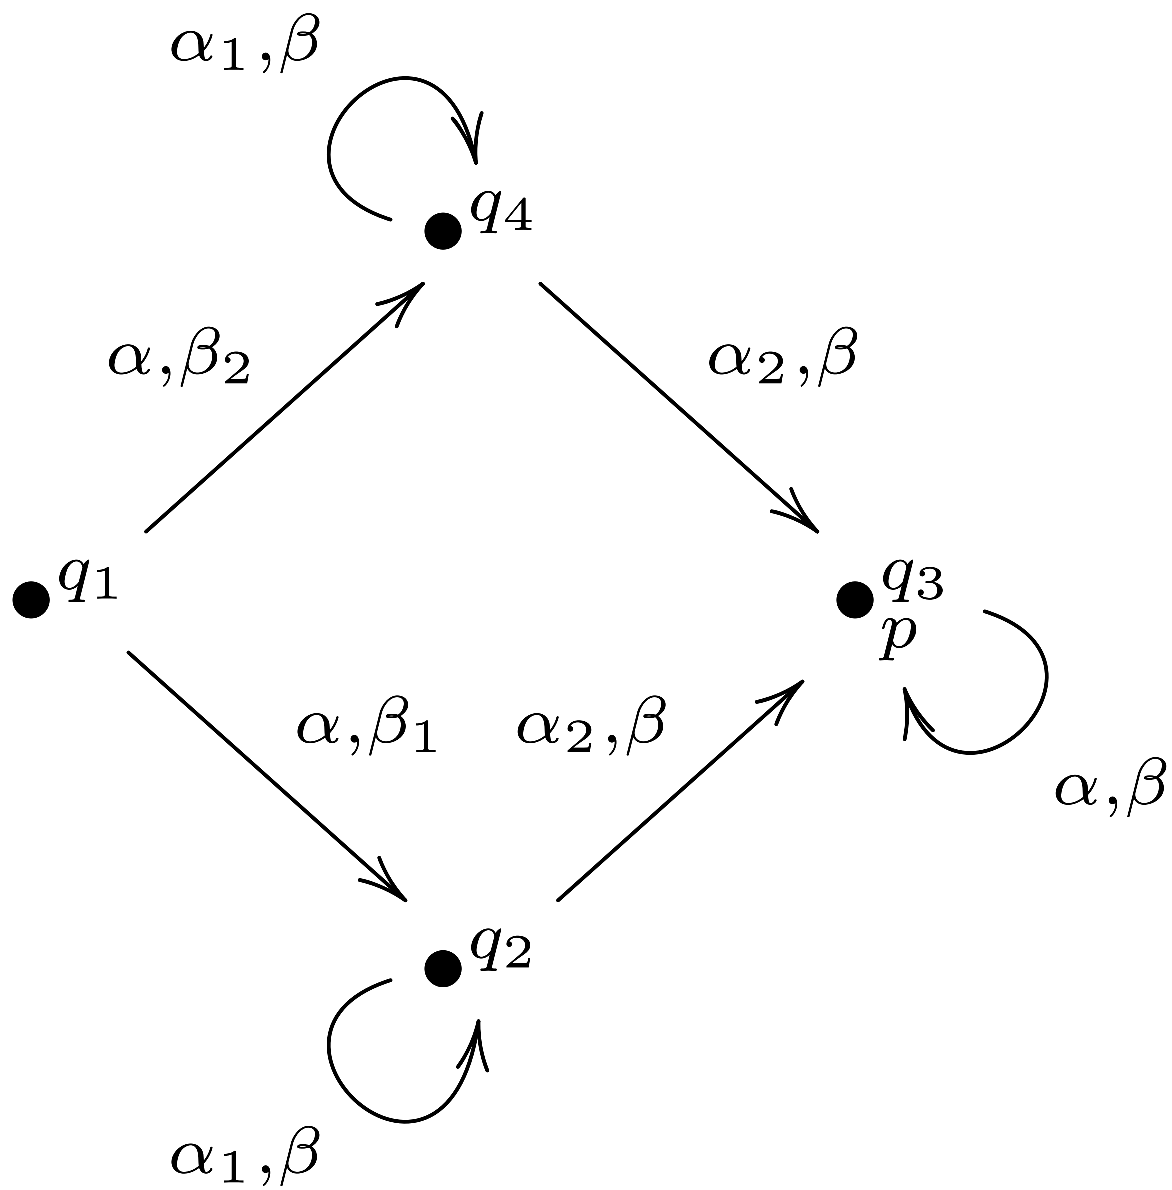

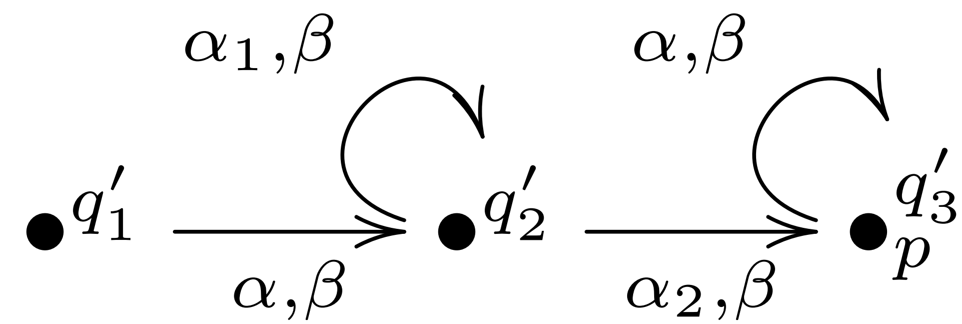

$$\beta = \{(q_1, q'_1), (q_2, q'_2), (q_4, q'_2), (q_3, q'_3)\}$$

Strategies and memory

Let us discern between two definitions of the satisfaction relation:

$\models_F$: perfect recall is assumed, all strategies

$$f : Q^+ \rightarrow Act$$

are allowed

$\models_L$: only memoryless strategies are allowed, i.e., strategies

$$f : Q \rightarrow Act$$

Invariance under bisimulation: the memoryless case

Theorem: *If $M_1 \rightleftarrows_\beta M_2$ and $s_1 \beta s_2$, then for every ATL formula φ:*

$$M_1, s_1 \models_L \varphi \quad \textit{iff} \quad M_2, s_2 \models_L \varphi$$

Tree-unfoldings

Let $fincomp_M(q)$ denote the set of finite prefixes of paths starting in q. Let $\ell(q_0 \cdots q_k) = q_k$.

Tree-unfoldings

Let $fincomp_M(q)$ denote the set of finite prefixes of paths starting in q. Let $\ell(q_0 \cdots q_k) = q_k$.

Given a CGS

$$M = (Q, \pi, Act, d, \delta)$$

and $q \in Q$, the tree-unfolding $T(M, q)$ of M from q is defined as follows:

$$T(M, q) = (Q^*, \pi^*, Act, d^*, \delta^*),$$

where $Q^* = fincomp_M(q)$; $\pi^*(\sigma) = \pi(\ell(\sigma))$; $d_i^*(\sigma) = d_i(\ell(\sigma))$; and $\delta^*(\sigma, \mathbf{a}) = \sigma\delta(\ell(\sigma), \mathbf{a})$.

More memory does not increase ability

Lemma: *For any M, q,*

$$T(M, q) \rightleftarrows_{\beta} M$$

where $\beta = \{(\sigma, \ell(\sigma)) \mid \sigma \in fincomp_M(q)\}$

More memory does not increase ability

Lemma: *For any M, q,*

$$T(M,q) \rightleftarrows_{\beta} M$$

where $\beta = \{(\sigma, \ell(\sigma)) \mid \sigma \in fincomp_M(q)\}$

Lemma: For any M, q and φ,

$$T(M,q), q \models_L \varphi \Leftrightarrow M, q \models_F \varphi$$

More memory does not increase ability

Lemma: *For any M, q,*

$$T(M,q) \rightleftarrows_\beta M$$

where $\beta = \{(\sigma, \ell(\sigma)) \mid \sigma \in fincomp_M(q)\}$

Lemma: For any M, q and φ,

$$T(M,q), q \models_L \varphi \Leftrightarrow M, q \models_F \varphi$$

Corollary: For any M, q and φ,

$$M, q \models_L \varphi \Leftrightarrow M, q \models_F \varphi$$

More memory does not increase ability

Lemma: *For any M, q,*

$$T(M,q) \rightleftarrows_\beta M$$

$where\ \beta = \{(\sigma, \ell(\sigma)) \mid \sigma \in fincomp_M(q)\}$

Lemma: For any M, q and φ,

Also: axiomatisation is sound and complete wrt. both semantics

Corollary: For any M, q and φ,

$$M, q \models_L \varphi \Leftrightarrow M, q \models_F \varphi$$

Invariance under bisimulation: the perfect recall case

Corollary: If $M_1 \rightleftarrows_\beta M_2$ and $s_1 \beta s_2$, then

$$M_1, s_1 \models_F \varphi \text{ iff } M_2, s_2 \models_F \varphi$$

for every ATL formula φ.

ATL* and memory

Unlike for ATL, for ATL* memory matters:

M: 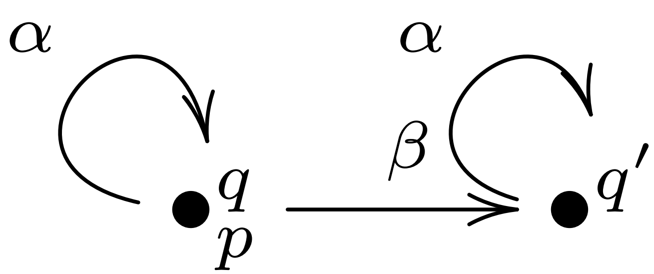

$$\varphi = \langle\!\langle a \rangle\!\rangle (\bigcirc p \wedge \bigcirc \bigcirc \neg p)$$

ATL* and memory

Unlike for ATL, for ATL* memory matters:

M: 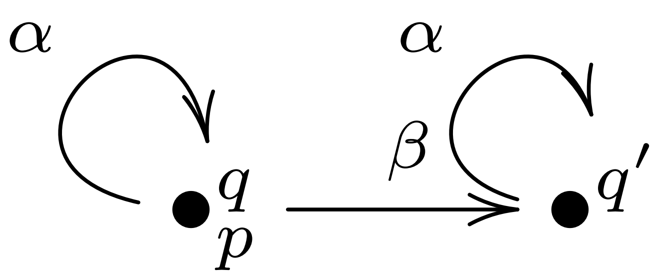

$$M, q \models_F \phi$$

$$\varphi = \langle\!\langle a \rangle\!\rangle (\bigcirc p \wedge \bigcirc \bigcirc \neg p)$$

ATL* and memory

Unlike for ATL, for ATL* memory matters:

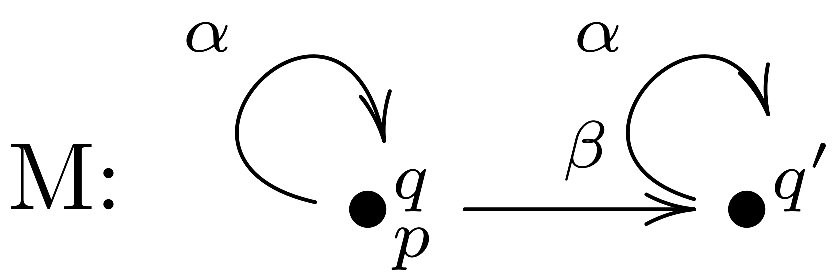

$$M, q \models_F \phi$$

$$M, q \not\models_L \phi$$

$$\varphi = \langle\!\langle a \rangle\!\rangle (\bigcirc p \wedge \bigcirc \bigcirc \neg p)$$

Contents

- Branching-time temporal logic: CTL
- ATL
- Bisimulations and the role of memory
- Irrevocable strategies

Example

- p: agent a controls the resource

Example

- p: agent a controls the resource
- $\langle\!\langle a \rangle\!\rangle \bigcirc p$: a has the ability to control the resource next

Example

- p: agent a controls the resource
- $\langle\!\langle a \rangle\!\rangle \bigcirc p$: a has the ability to control the resource next
- $\langle\!\langle a \rangle\!\rangle \Box \langle\!\langle a \rangle\!\rangle \bigcirc p$: a has the ability to ensure that $\langle\!\langle a \rangle\!\rangle \bigcirc p$ is always true

Example

- p: agent a controls the resource
- $\langle\!\langle a \rangle\!\rangle \bigcirc p$: a has the ability to control the resource next
- $\langle\!\langle a \rangle\!\rangle \Box \langle\!\langle a \rangle\!\rangle \bigcirc p$: a has the ability to ensure that $\langle\!\langle a \rangle\!\rangle \bigcirc p$ is always true

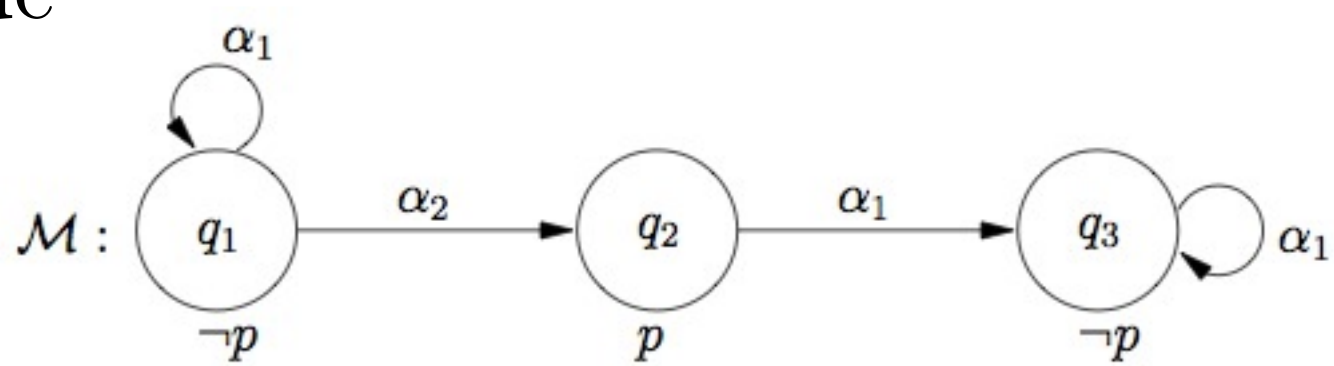

Example

- p: agent a controls the resource
- $\langle\!\langle a \rangle\!\rangle \bigcirc p$: a has the ability to control the resource next
- $\langle\!\langle a \rangle\!\rangle \Box \langle\!\langle a \rangle\!\rangle \bigcirc p$: a has the ability to ensure that $\langle\!\langle a \rangle\!\rangle \bigcirc p$ is always true

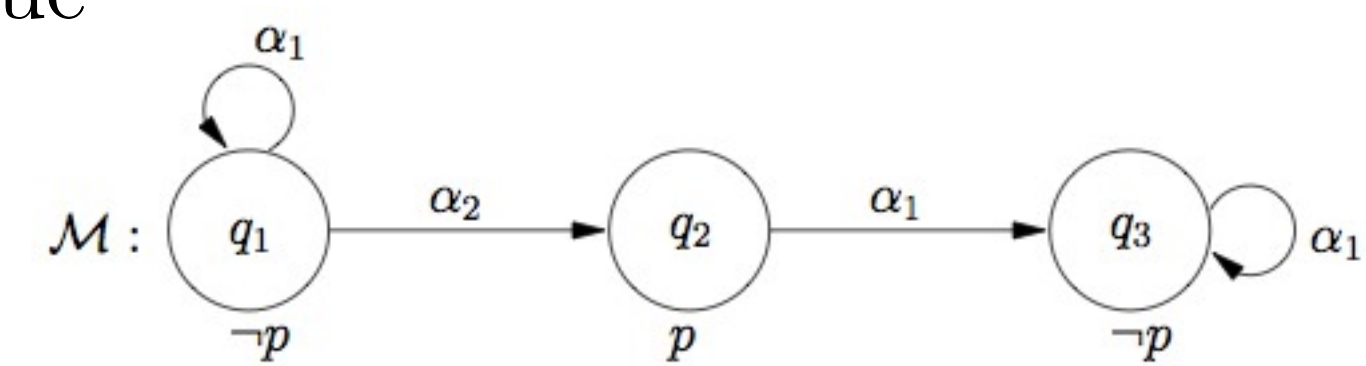

$$M, q_1 \models \langle\!\langle a \rangle\!\rangle \bigcirc p$$

Example

- p: agent a controls the resource
- $\langle\!\langle a \rangle\!\rangle \bigcirc p$: a has the ability to control the resource next
- $\langle\!\langle a \rangle\!\rangle \square \langle\!\langle a \rangle\!\rangle \bigcirc p$: a has the ability to ensure that $\langle\!\langle a \rangle\!\rangle \bigcirc p$ is always true

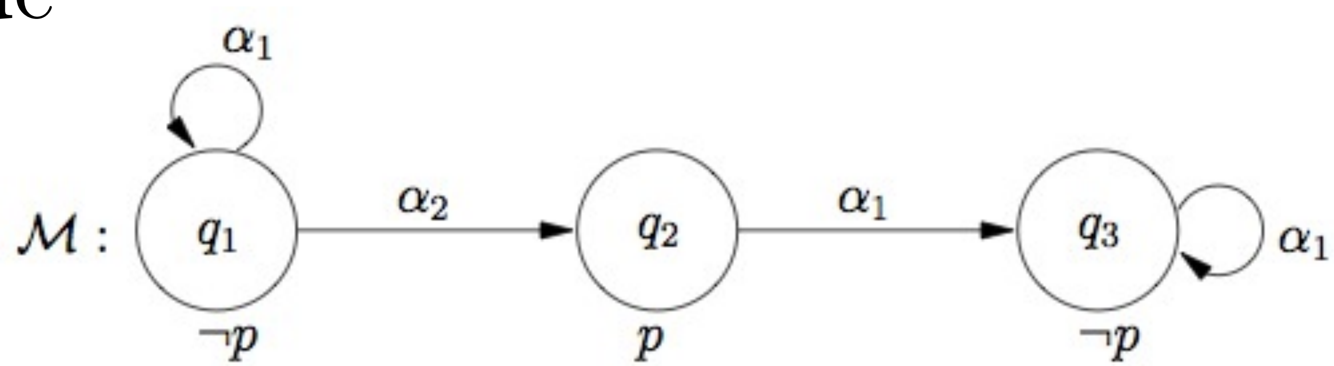

$M, q_1 \models \langle\!\langle a \rangle\!\rangle \bigcirc p$

$M, q_1 \models \langle\!\langle a \rangle\!\rangle \square \langle\!\langle a \rangle\!\rangle \bigcirc p$

Example

- p: agent a controls the resource
- $\langle\!\langle a \rangle\!\rangle \bigcirc p$: a has the ability to control the resource next
- $\langle\!\langle a \rangle\!\rangle \Box \langle\!\langle a \rangle\!\rangle \bigcirc p$: a has the ability to ensure that $\langle\!\langle a \rangle\!\rangle \bigcirc p$ is always true

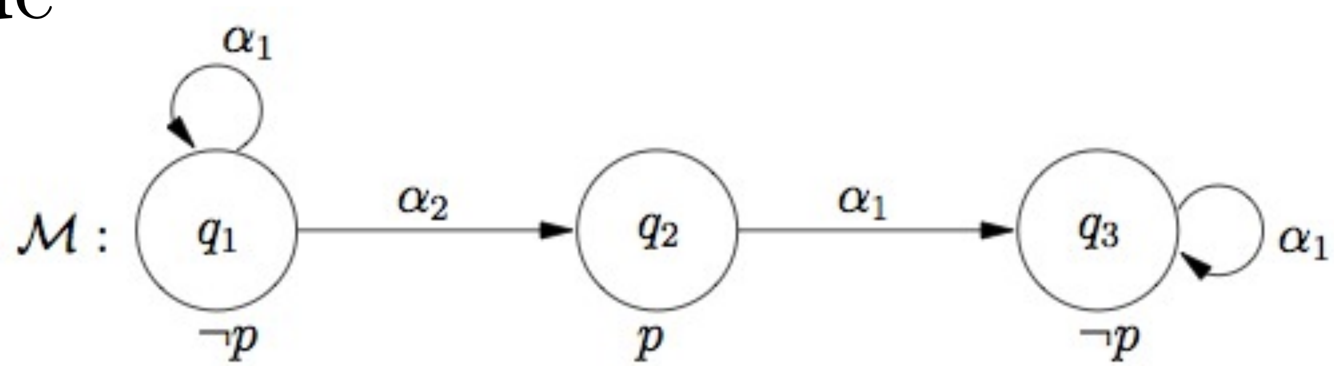

$M, q_1 \models \langle\!\langle a \rangle\!\rangle \bigcirc p$ $\qquad$ $M, q_1 \models \langle\!\langle a \rangle\!\rangle \Box \langle\!\langle a \rangle\!\rangle \bigcirc p$

Paradox? a has the ability to ensure that she can always access the resource - but only by never actually accessing it

Revocability of strategies in ATL

- In the evaluation of a formula such as $\langle\!\langle a \rangle\!\rangle \Box \varphi$, when the goal φ is evaluated the agent (a) is no longer restricted by the strategy she chose in order to get to the state where the goal is evaluated (as the example illustrates)
- In this sense, strategies in ATL are revocable
- In some contexts, it would be more natural to reason about strategies which are *not* revocable and completely specify the future behaviour of the agent

Alternative: irrevocable strategies

Irrevocable strategies can be modelled by using model updates in the semantics.

Alternative: irrevocable strategies

Irrevocable strategies can be modelled by using model updates in the semantics.

Assume memoryless strategies (for now).

Alternative: irrevocable strategies

Irrevocable strategies can be modelled by using model updates in the semantics.

Assume memoryless strategies (for now).

Let M be a CGS, C a coalition, and f_C a memoryless strategy for C. The update of M by f_C, denoted $M \dagger f_C$, is the same as M, except that the choices of each agent $i \in C$ are fixed by the strategy f_i:

$$d_i(q) = \{f_i(q)\}$$

for each state q.

Model update: example

Model update: example

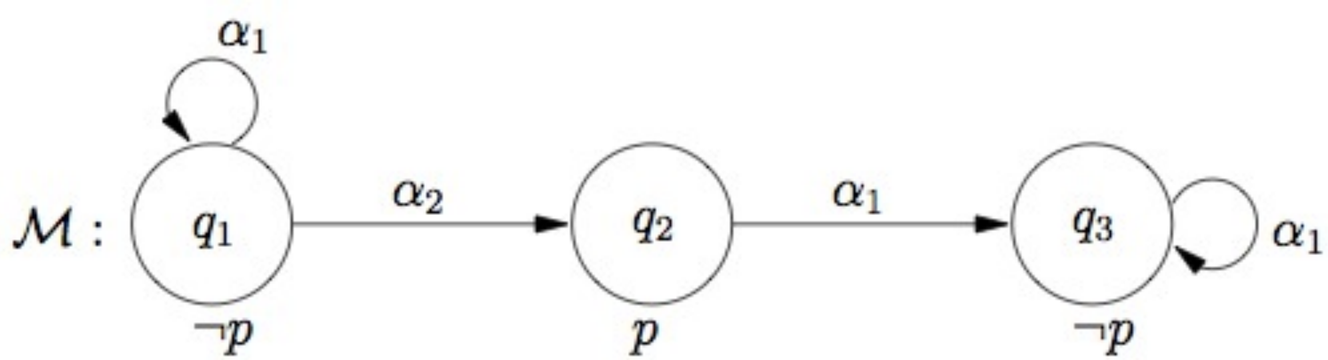

Model update: example

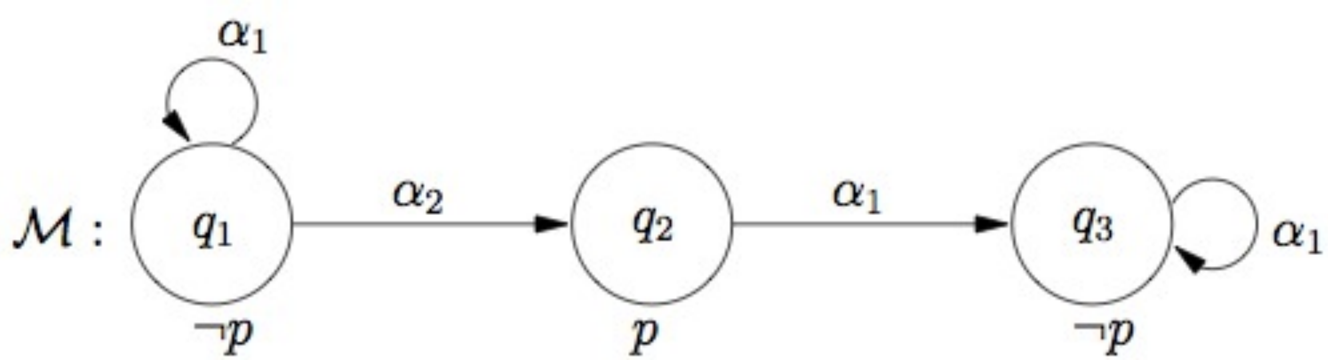

$$f_1 = \{q_1 \mapsto \alpha_1, q_2 \mapsto \alpha_1, q_3 \mapsto \alpha_1\}$$

Model update: example

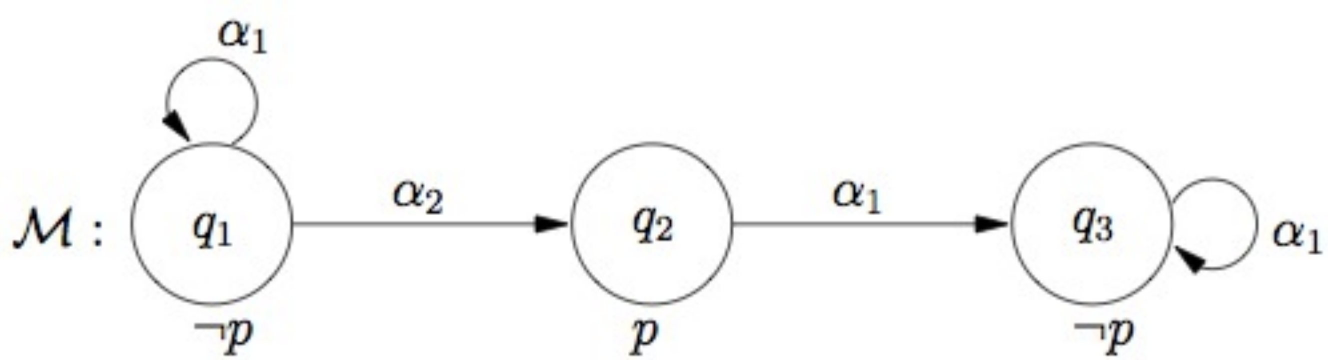

$$f_1 = \{q_1 \mapsto \alpha_1, q_2 \mapsto \alpha_1, q_3 \mapsto \alpha_1\}$$

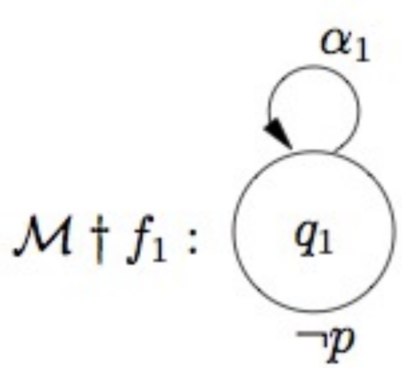

Satisfiability under irrevocable semantics

We can now define a new variant of the satisfiability relation:

$$
\begin{array}{lcl}
M, q \models_i \langle\!\langle C \rangle\!\rangle \bigcirc \phi & \Leftrightarrow & \exists f_C \forall \lambda \in out_{M \dagger f_C}(q, f_C) \\
 & & (M \dagger f_C, \lambda[1] \models_i \phi) \\
M, q \models_i \langle\!\langle C \rangle\!\rangle \Box \phi & \Leftrightarrow & \exists f_C \forall \lambda \in out_{M \dagger f_C}(q, f_C) \\
 & & \forall j \geq 0 (M \dagger f_C, \lambda[j] \models_i \phi) \\
M, q \models_i \langle\!\langle C \rangle\!\rangle (\phi_1 \,\mathcal{U}\, \phi_2) & \Leftrightarrow & \exists f_C \forall \lambda \in out_{M \dagger f_C}(q, f_C) \\
 & & \exists j \geq 0 (M \dagger f_C, \lambda[j] \models_i \phi_2 \text{ and} \\
 & & \forall 0 \leq k < j (M \dagger f_C, \lambda[k] \models_i \phi_1))
\end{array}
$$

Example (continued)

Example (continued)

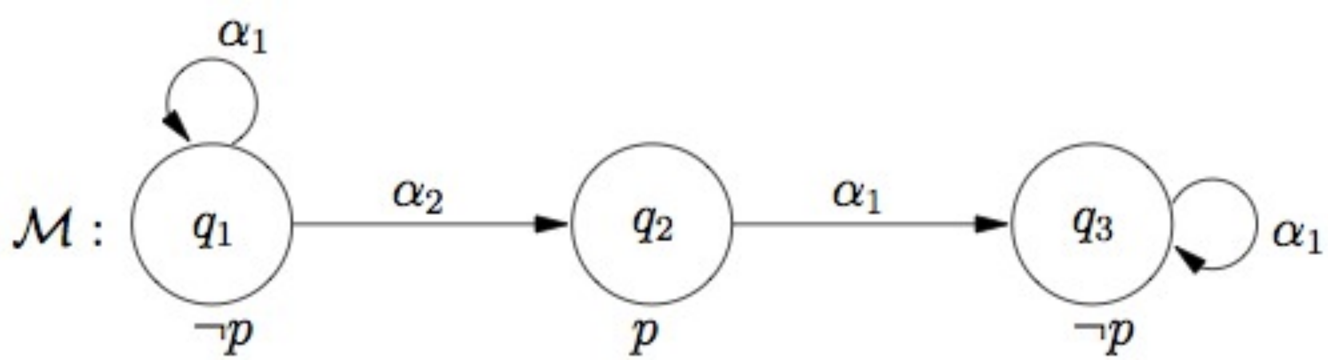

Example (continued)

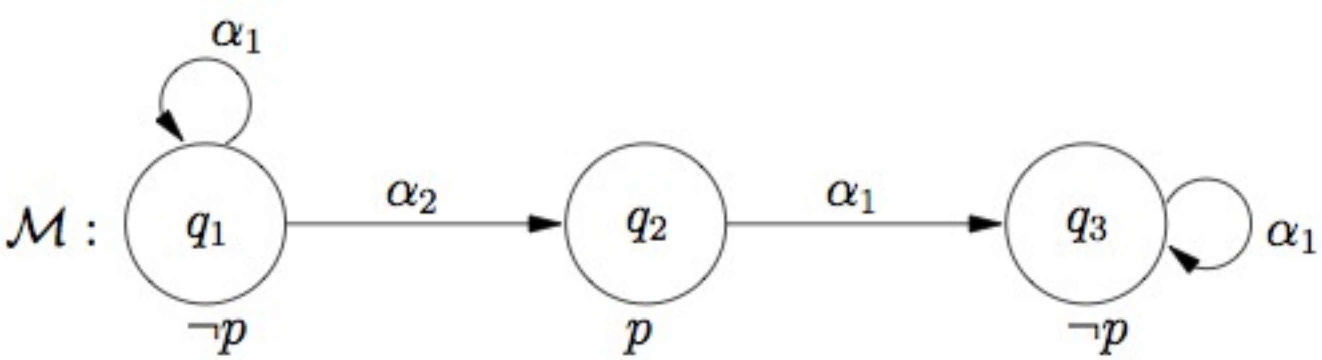

$$f_1 = \{q_1 \mapsto \alpha_1, q_2 \mapsto \alpha_1, q_3 \mapsto \alpha_1\}$$
$$f_2 = \{q_1 \mapsto \alpha_2, q_2 \mapsto \alpha_1, q_3 \mapsto \alpha_1\}$$

Example (continued)

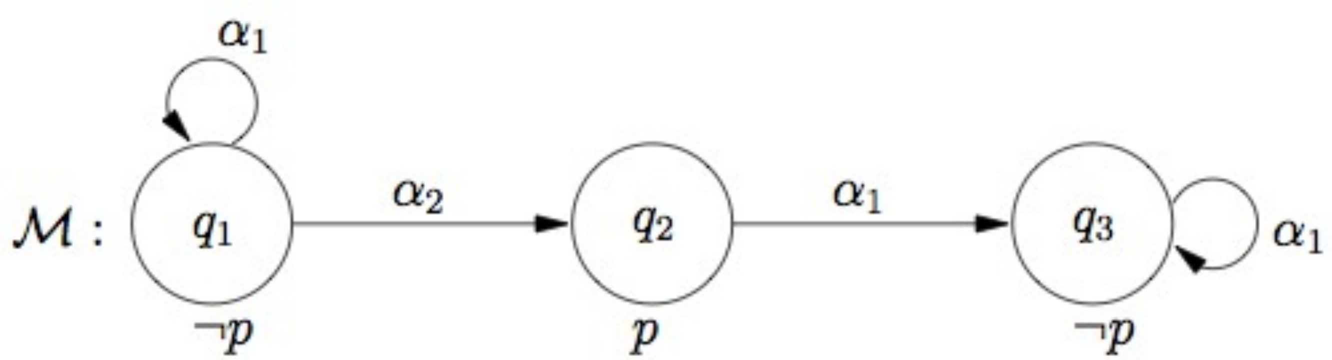

$$f_1 = \{q_1 \mapsto \alpha_1, q_2 \mapsto \alpha_1, q_3 \mapsto \alpha_1\}$$
$$f_2 = \{q_1 \mapsto \alpha_2, q_2 \mapsto \alpha_1, q_3 \mapsto \alpha_1\}$$

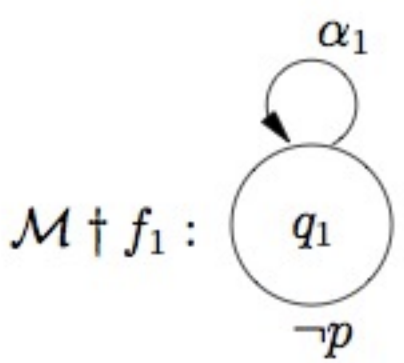

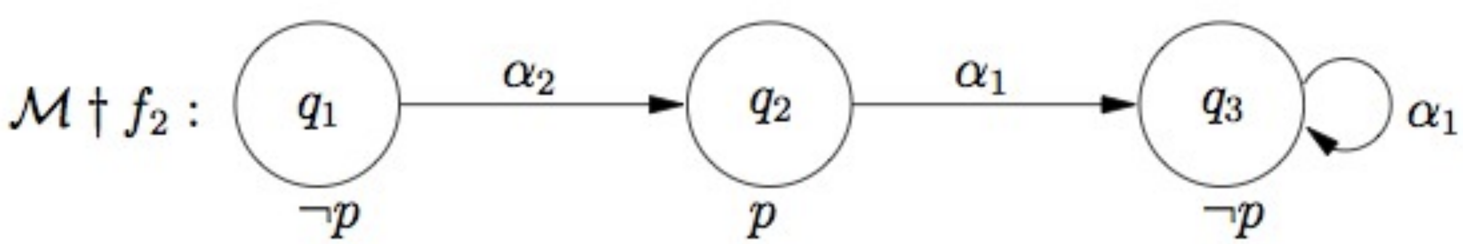

Example (continued)

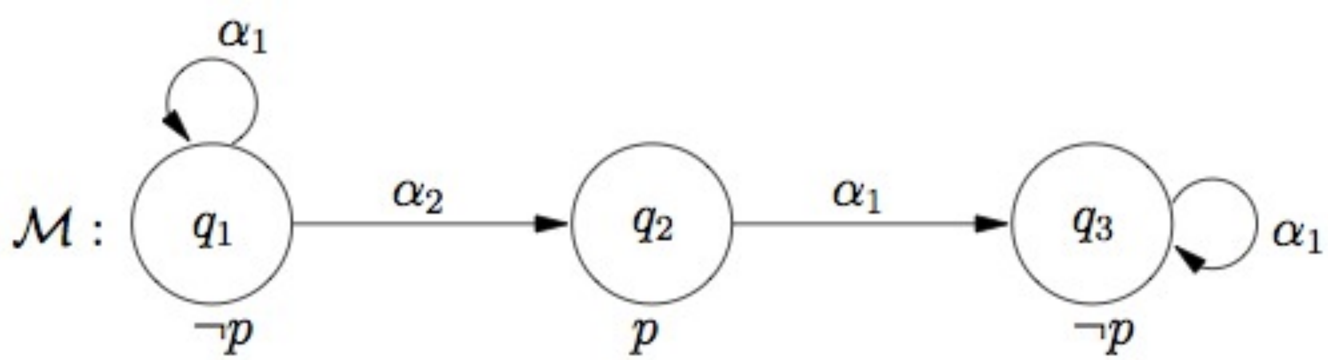

$$f_1 = \{q_1 \mapsto \alpha_1, q_2 \mapsto \alpha_1, q_3 \mapsto \alpha_1\}$$
$$f_2 = \{q_1 \mapsto \alpha_2, q_2 \mapsto \alpha_1, q_3 \mapsto \alpha_1\}$$

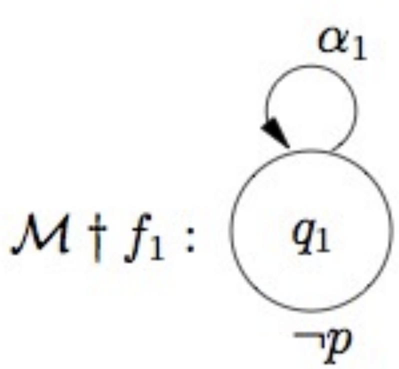

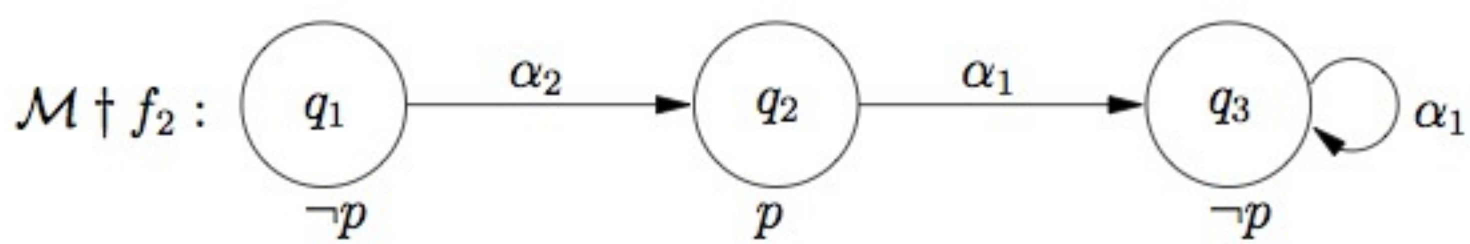

$M, q_1 \models \langle\!\langle a \rangle\!\rangle \Box \langle\!\langle a \rangle\!\rangle \bigcirc p$ (standard definition)

$M, q_1 \not\models_i \langle\!\langle a \rangle\!\rangle \Box \langle\!\langle a \rangle\!\rangle \bigcirc p$ (with irrevocable strategies)

Irrevocable strategies and bisimulation

With irrevocable strategies, truth of formulae is not invariant under bisimulations:

Irrevocable strategies and bisimulation

With irrevocable strategies, truth of formulae is not invariant under bisimulations:

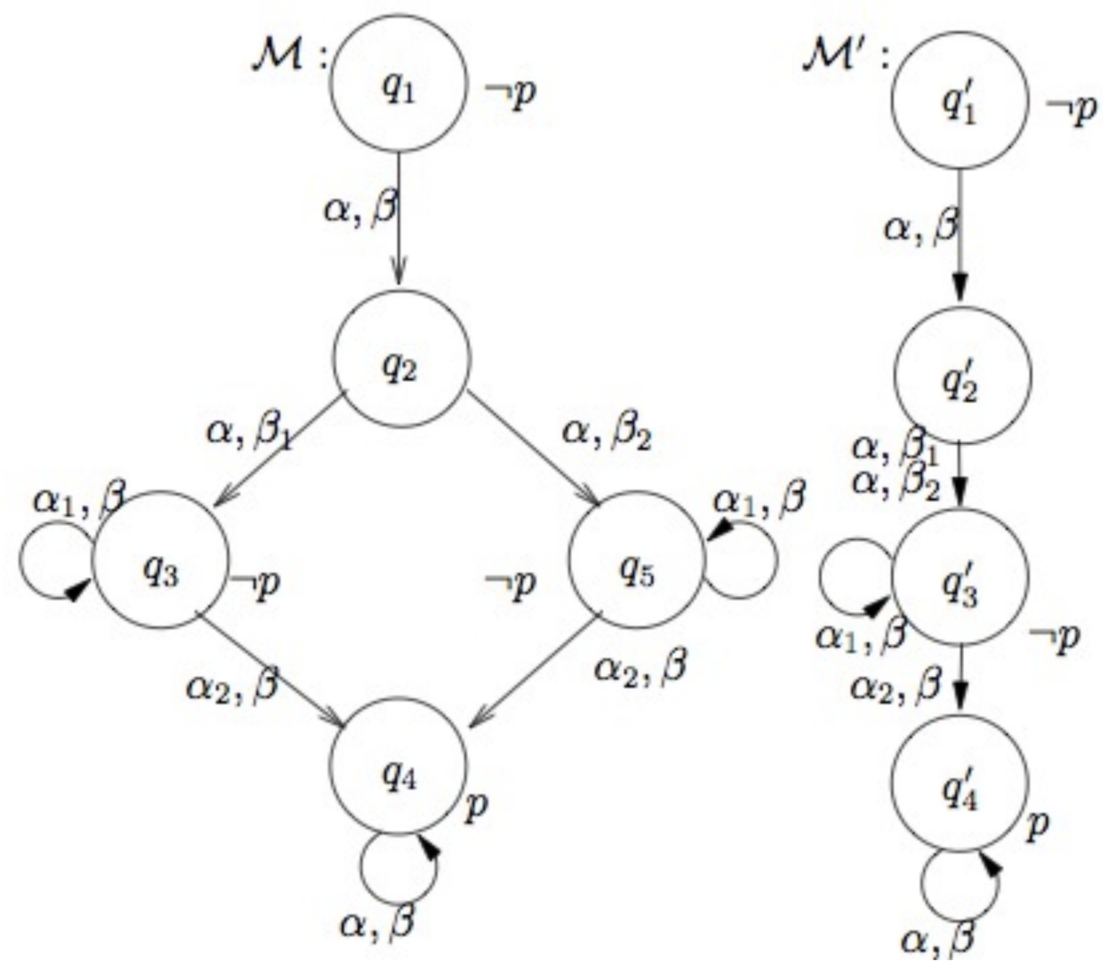

Irrevocable strategies and bisimulation

With irrevocable strategies, truth of formulae is not invariant under bisimulations:

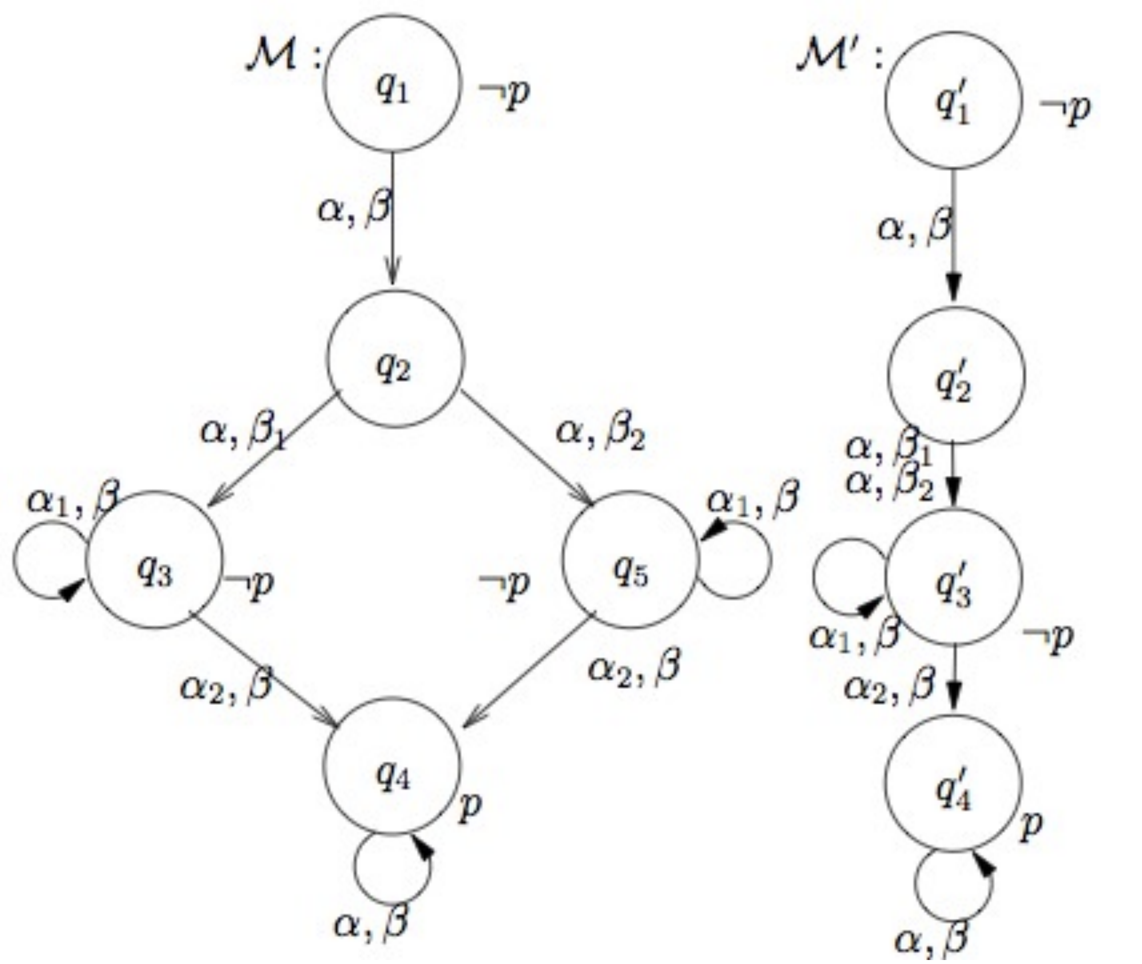

$$M, q_1 \models_i \langle\!\langle 1 \rangle\!\rangle \bigcirc ((\langle\!\langle 2 \rangle\!\rangle \bigcirc \langle\!\langle \emptyset \rangle\!\rangle \bigcirc \neg p) \wedge \langle\!\langle 2 \rangle\!\rangle \bigcirc \langle\!\langle \emptyset \rangle\!\rangle \bigcirc p)$$

(strategies: $\{q_3 \mapsto \alpha_1, q_5 \mapsto \alpha_2\}$; $\{q_2 \mapsto \beta_1\}$; $\{q_2 \mapsto \beta_2\}$)

Irrevocable strategies and bisimulation

With irrevocable strategies, truth of formulae is not invariant under bisimulations:

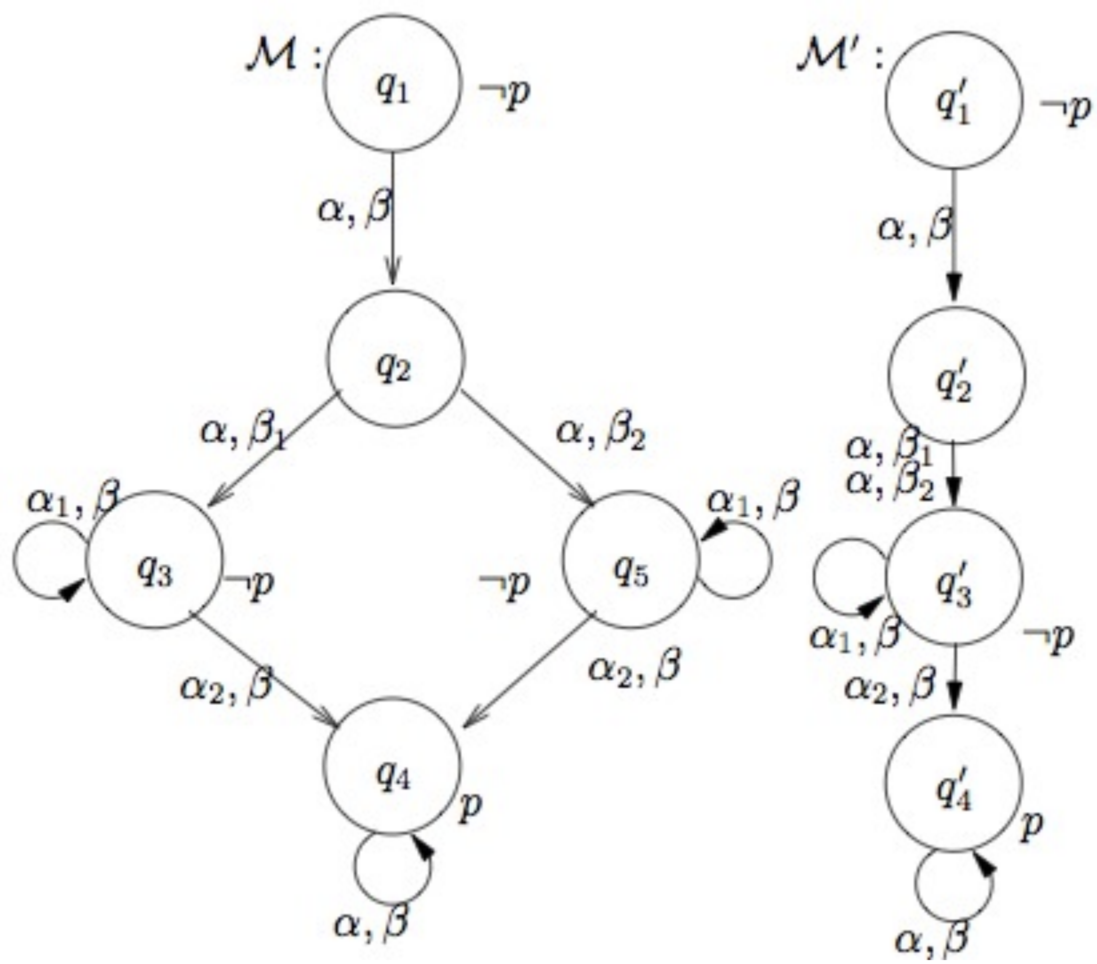

$$M, q_1 \models_i \langle\!\langle 1 \rangle\!\rangle \bigcirc ((\langle\!\langle 2 \rangle\!\rangle \bigcirc \langle\!\langle \emptyset \rangle\!\rangle \bigcirc \neg p) \wedge \langle\!\langle 2 \rangle\!\rangle \bigcirc \langle\!\langle \emptyset \rangle\!\rangle \bigcirc p)$$

(strategies: $\{q_3 \mapsto \alpha_1, q_5 \mapsto \alpha_2\}$; $\{q_2 \mapsto \beta_1\}$; $\{q_2 \mapsto \beta_2\}$)

$$M', q_1 \not\models_i \langle\!\langle 1 \rangle\!\rangle \bigcirc ((\langle\!\langle 2 \rangle\!\rangle \bigcirc \langle\!\langle \emptyset \rangle\!\rangle \bigcirc \neg p) \wedge \langle\!\langle 2 \rangle\!\rangle \bigcirc \langle\!\langle \emptyset \rangle\!\rangle \bigcirc p)$$

On valid reasoning about irrevocable strategies

- Formulae valid under the standard definition is not necessarily valid under irrevocable strategies. For example, the principle of uniform substitution does not hold. The ATL axiom

$$\neg\langle\!\langle\emptyset\rangle\!\rangle\bigcirc\neg p \to \langle\!\langle N\rangle\!\rangle\bigcirc p$$

  is still valid with irrevocable strategies, but the result of substituting

$$\langle\!\langle N\rangle\!\rangle\bigcirc p \wedge \langle\!\langle N\rangle\!\rangle\bigcirc\neg p$$

  for p in it is not valid.

On valid reasoning about irrevocable strategies

- Formulae valid under the standard definition is not necessarily valid under irrevocable strategies. For example, the principle of uniform substitution does not hold. The ATL axiom

$$\neg\langle\!\langle\emptyset\rangle\!\rangle\bigcirc\neg p \to \langle\!\langle N\rangle\!\rangle\bigcirc p$$

  is still valid with irrevocable strategies, but the result of substituting

$$\langle\!\langle N\rangle\!\rangle\bigcirc p \wedge \langle\!\langle N\rangle\!\rangle\bigcirc\neg p$$

  for p in it is not valid.

- Formulae valid under irrevocable strategies are not necessarily valid under the standard definition. Example:

$$\langle\!\langle C\rangle\!\rangle\bigcirc\langle\!\langle C\rangle\!\rangle\bigcirc\phi \leftrightarrow \langle\!\langle C\rangle\!\rangle\bigcirc\langle\!\langle\emptyset\rangle\!\rangle\bigcirc\phi$$

  for $C \neq \emptyset$.

Perfect recall

With perfect recall strategies, we cannot update the model directly. Instead, unwind it first, and recall that a perfect recall strategy in M is equivalent to a memoryless strategy in $T(M, q)$:

$$M, q \models_{mi} \varphi \Leftrightarrow^{def} T(M, q), q \models_{i} \varphi$$

Perfect recall

$$M, q \models_{mi} \varphi \Leftrightarrow^{def} T(M, q), q \models_i \varphi$$

We get that:

- Still non-invariant under bisimulation

Perfect recall

$$M, q \models_{mi} \varphi \Leftrightarrow^{def} T(M, q), q \models_i \varphi$$

We get that:

- Still non-invariant under bisimulation

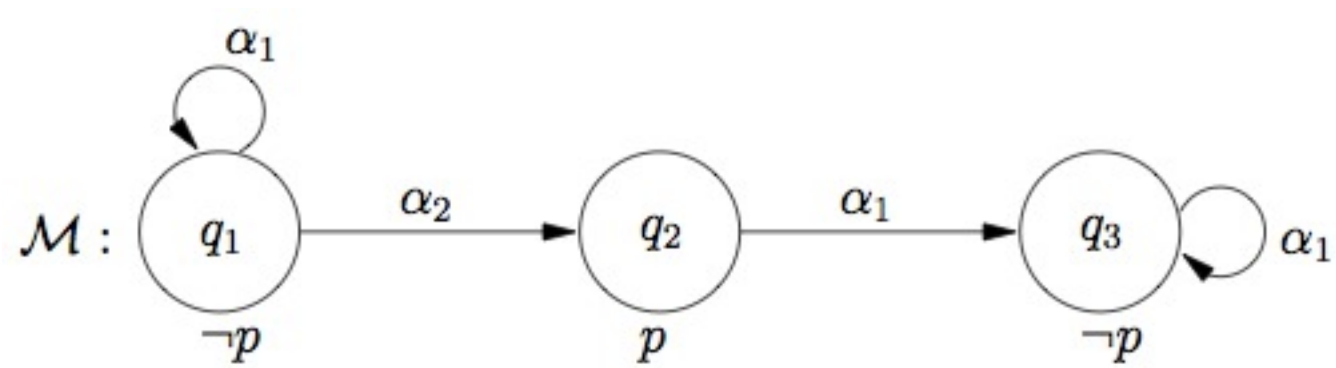

Perfect recall

$$M, q \models_{mi} \varphi \Leftrightarrow^{def} T(M, q), q \models_i \varphi$$

We get that:

- Still non-invariant under bisimulation
- With irrevocable strategies (unlike under the standard definition), memory matters:

$$M, q_1 \models_{mi} \langle\!\langle a \rangle\!\rangle \bigcirc \langle\!\langle a \rangle\!\rangle \bigcirc p$$
$$M, q_1 \not\models_{i} \langle\!\langle a \rangle\!\rangle \bigcirc \langle\!\langle a \rangle\!\rangle \bigcirc p$$

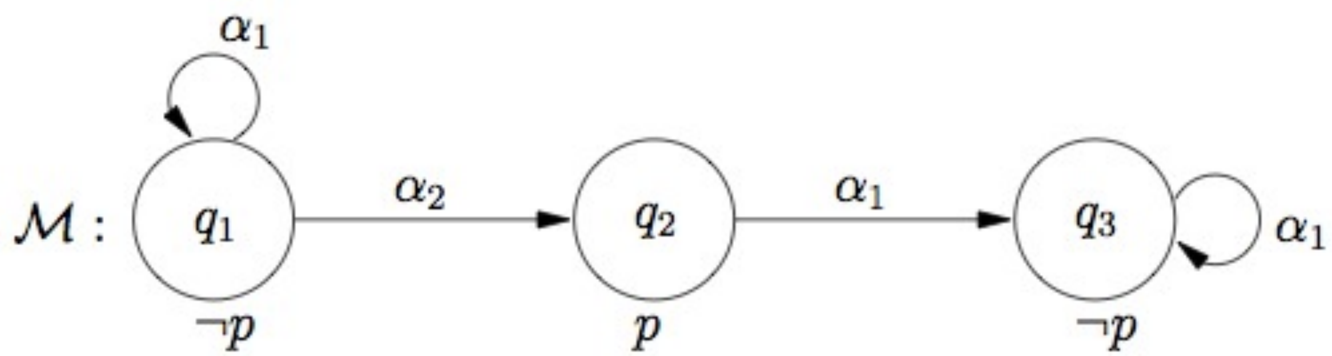

Summary

- Introduced ATL as an extension of both CL and CTL
- ATL*: more expressive
- The role of memory: do you have to remember the past?
  - ATL: no
  - ATL*: yes
  - Irrevocable ATL: yes

ATL and epistemic logic can be combined to allow strategic reasoning under imperfect information

- We extend CGSs with indistinguishability relations $\sim_a$, one per agent
- We add epistemic operators to ATL

$\rightsquigarrow$ Problems!

ATL and epistemic logic can be combined to allow strategic reasoning under imperfect information

- We extend CGSs with indistinguishability relations $\sim_a$, one per agent
- We add epistemic operators to ATL

$\rightsquigarrow$ Problems!

Combining Dimensions

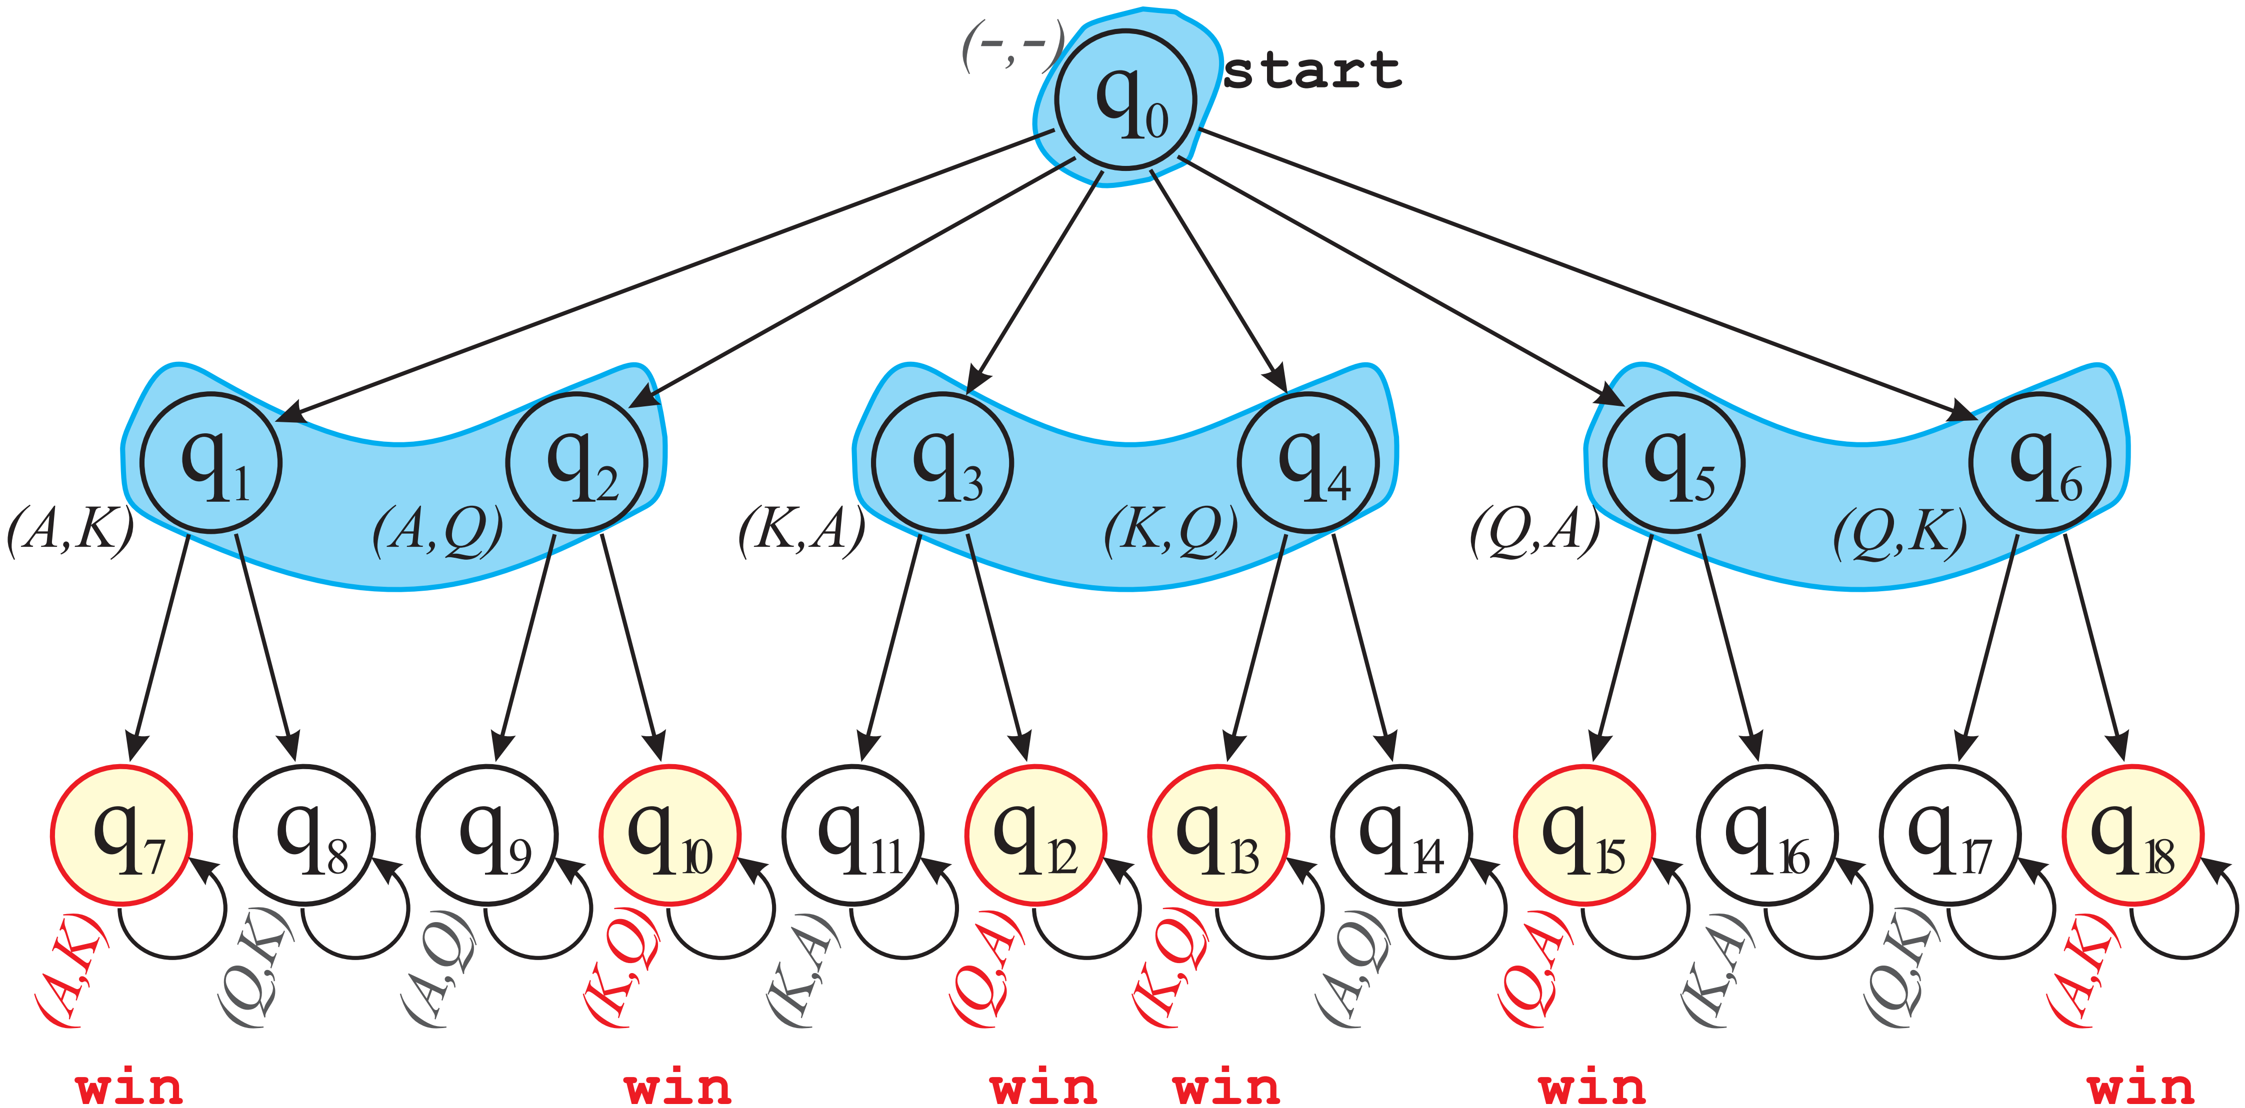

Combining Dimensions

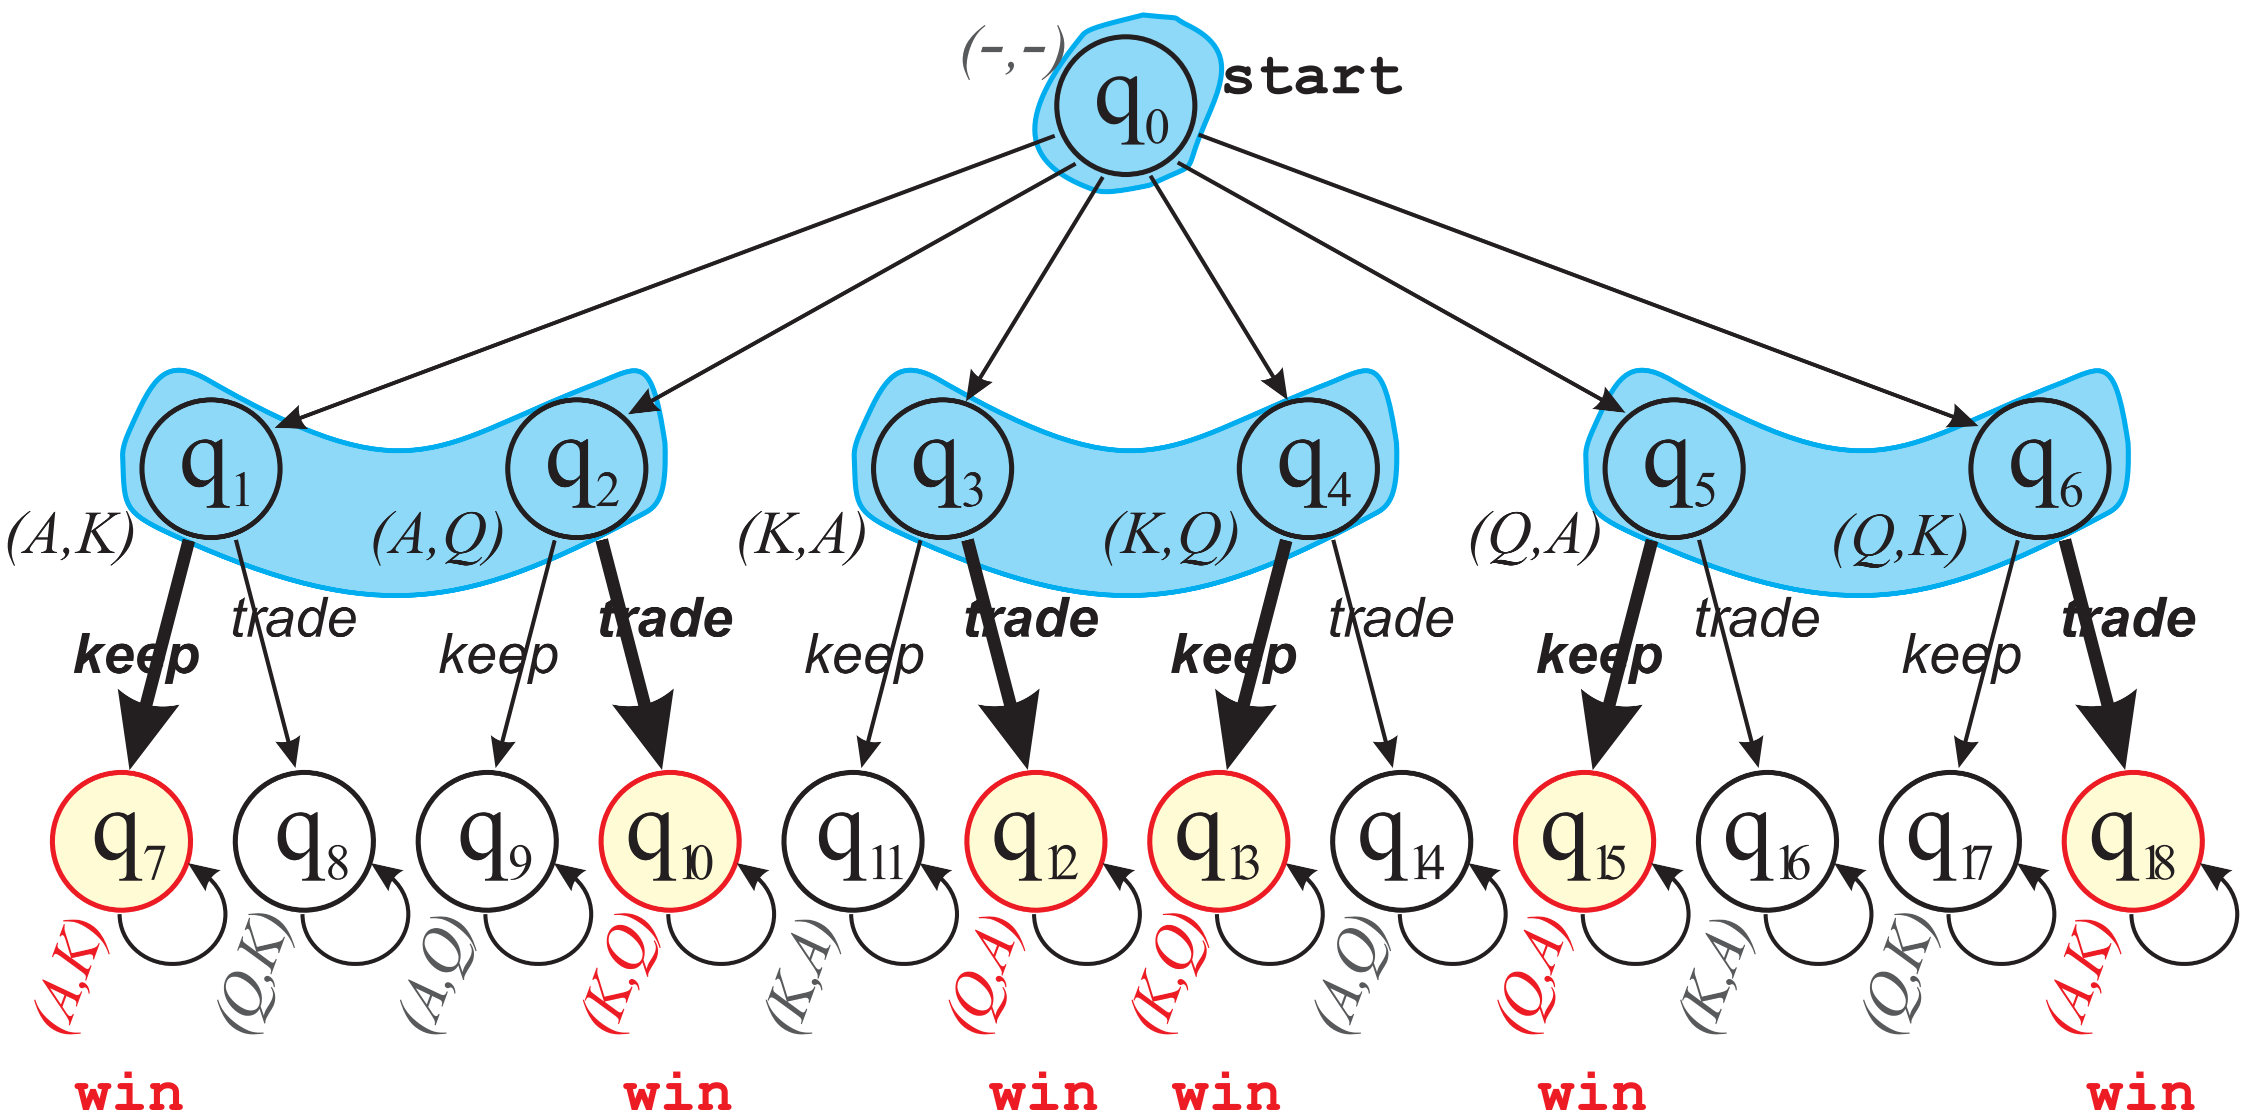

Combining Dimensions

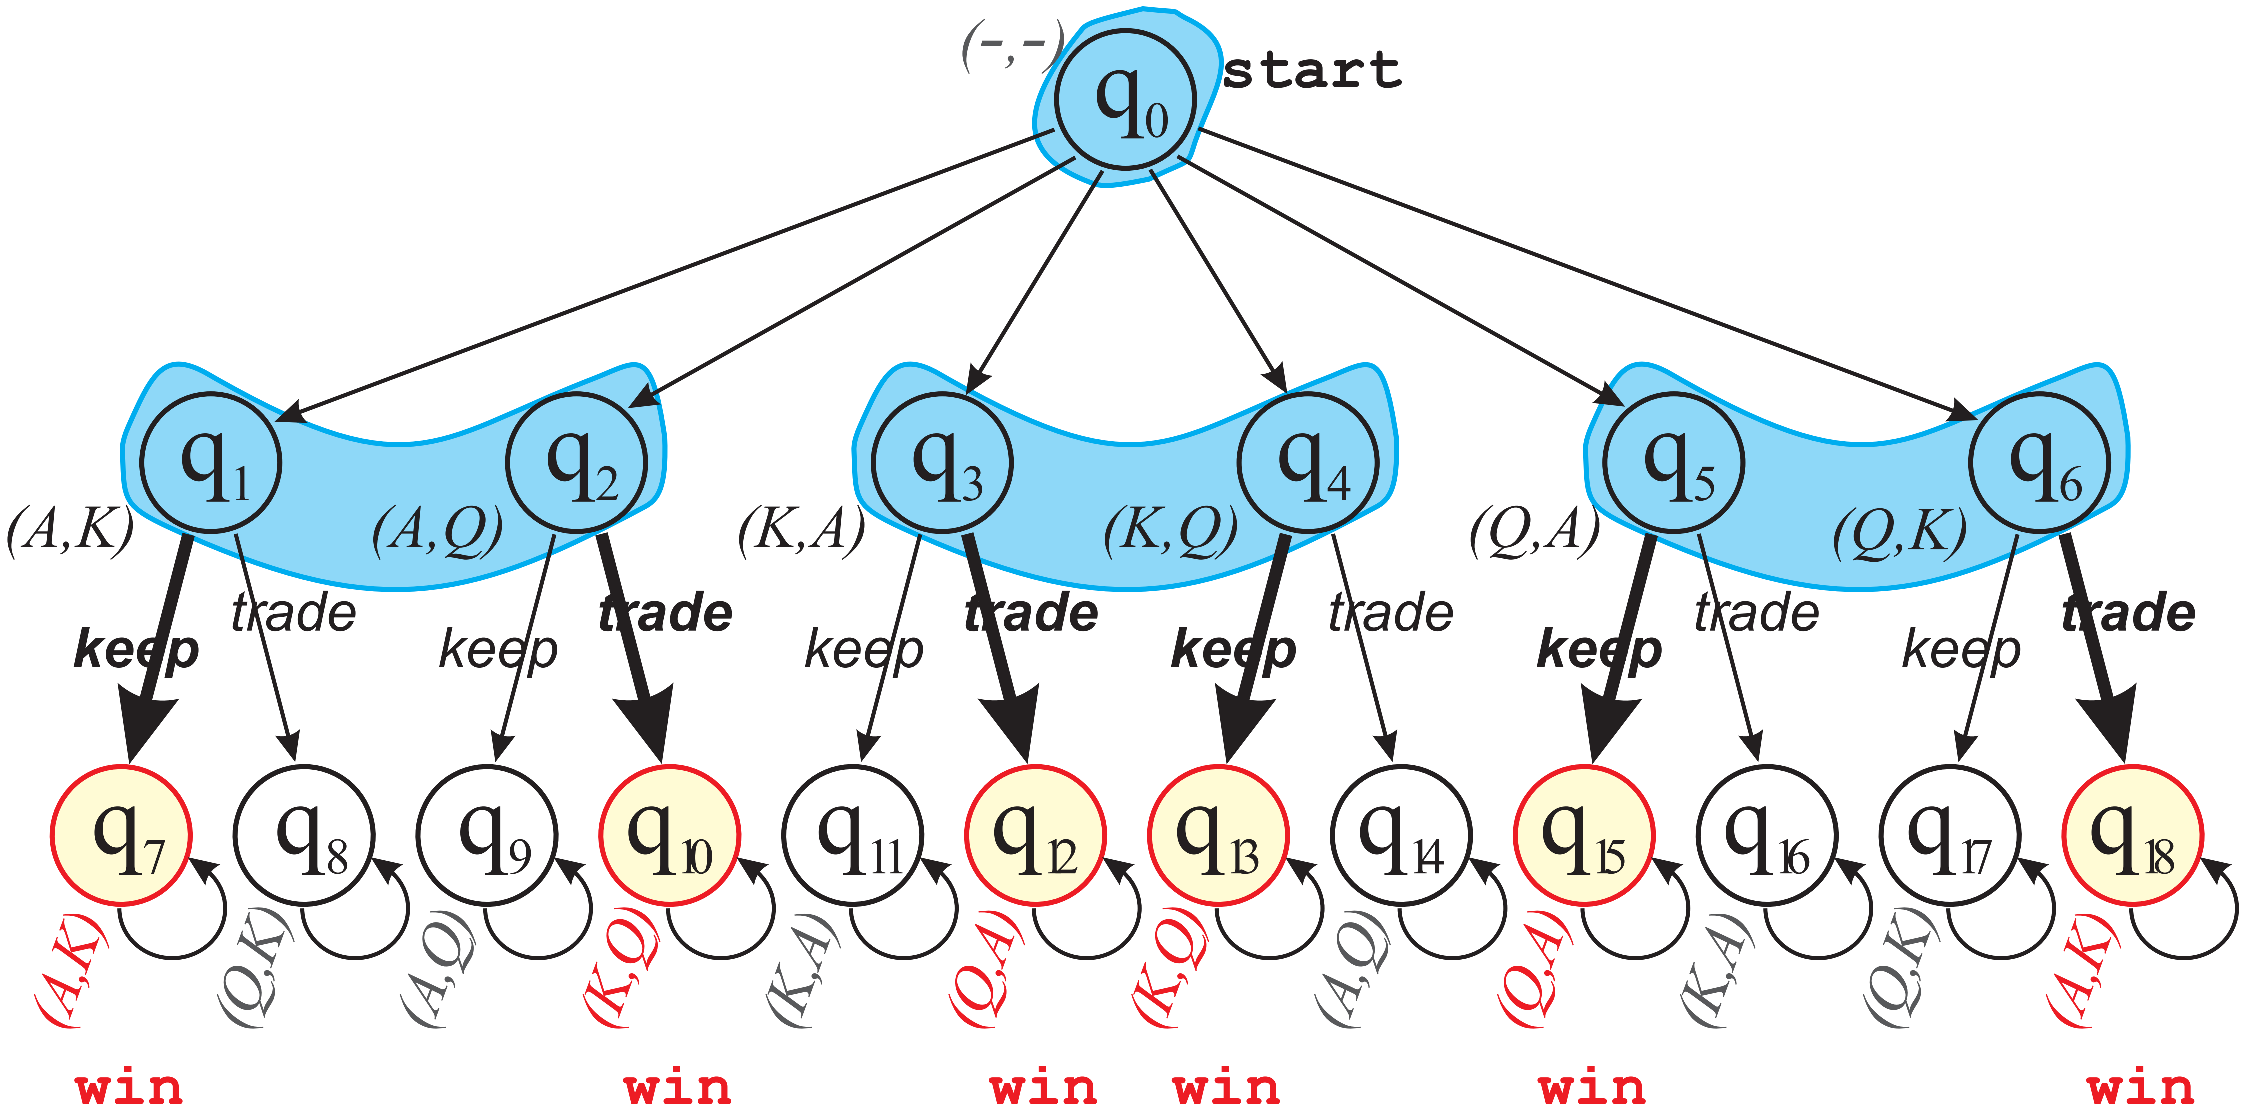

$$start \to \langle\langle a \rangle\rangle \Diamond win$$

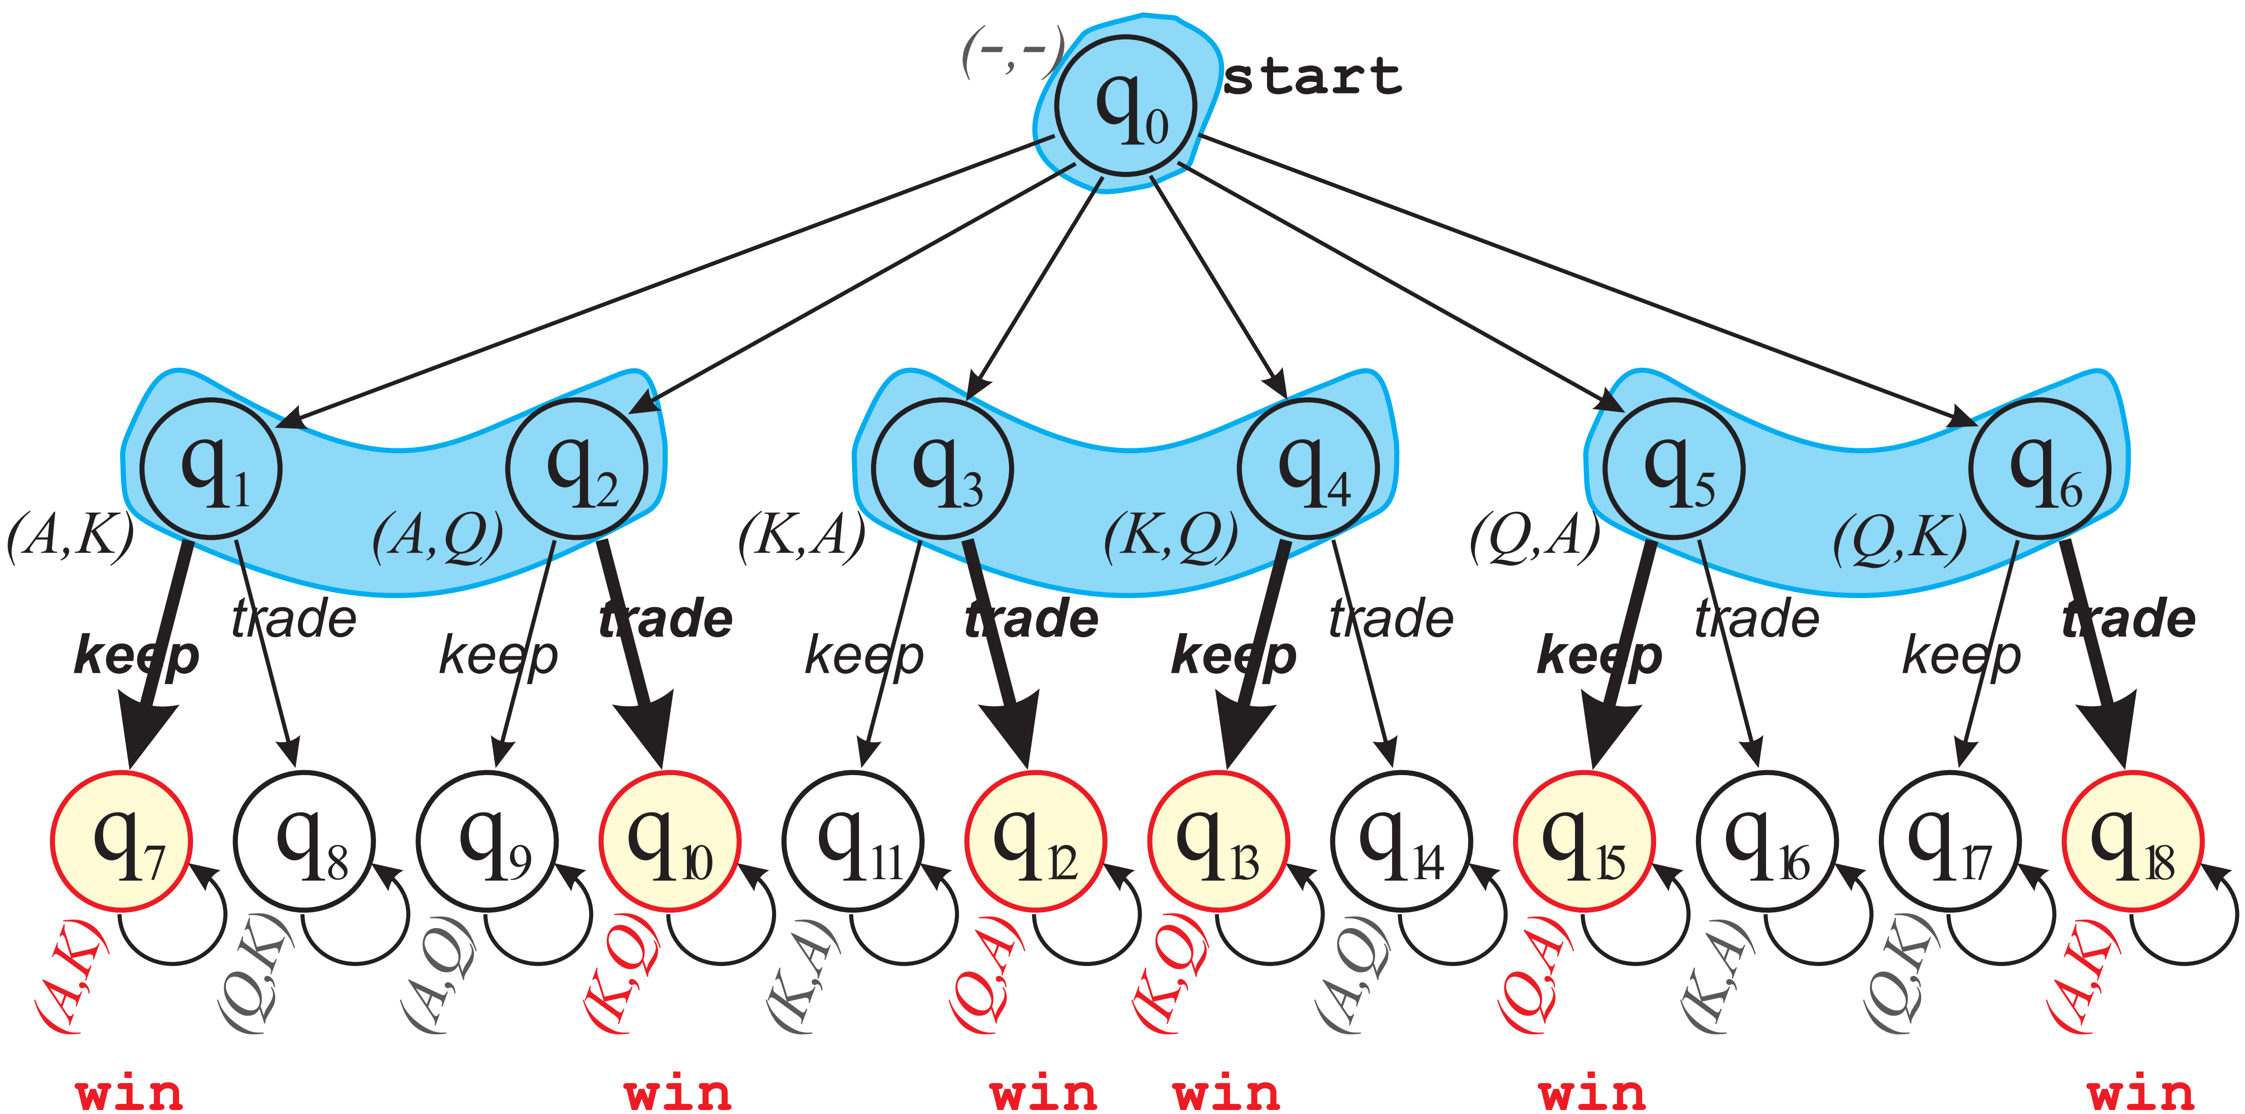

$$start \to \langle\langle a \rangle\rangle \Diamond win$$

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

Combining Dimensions

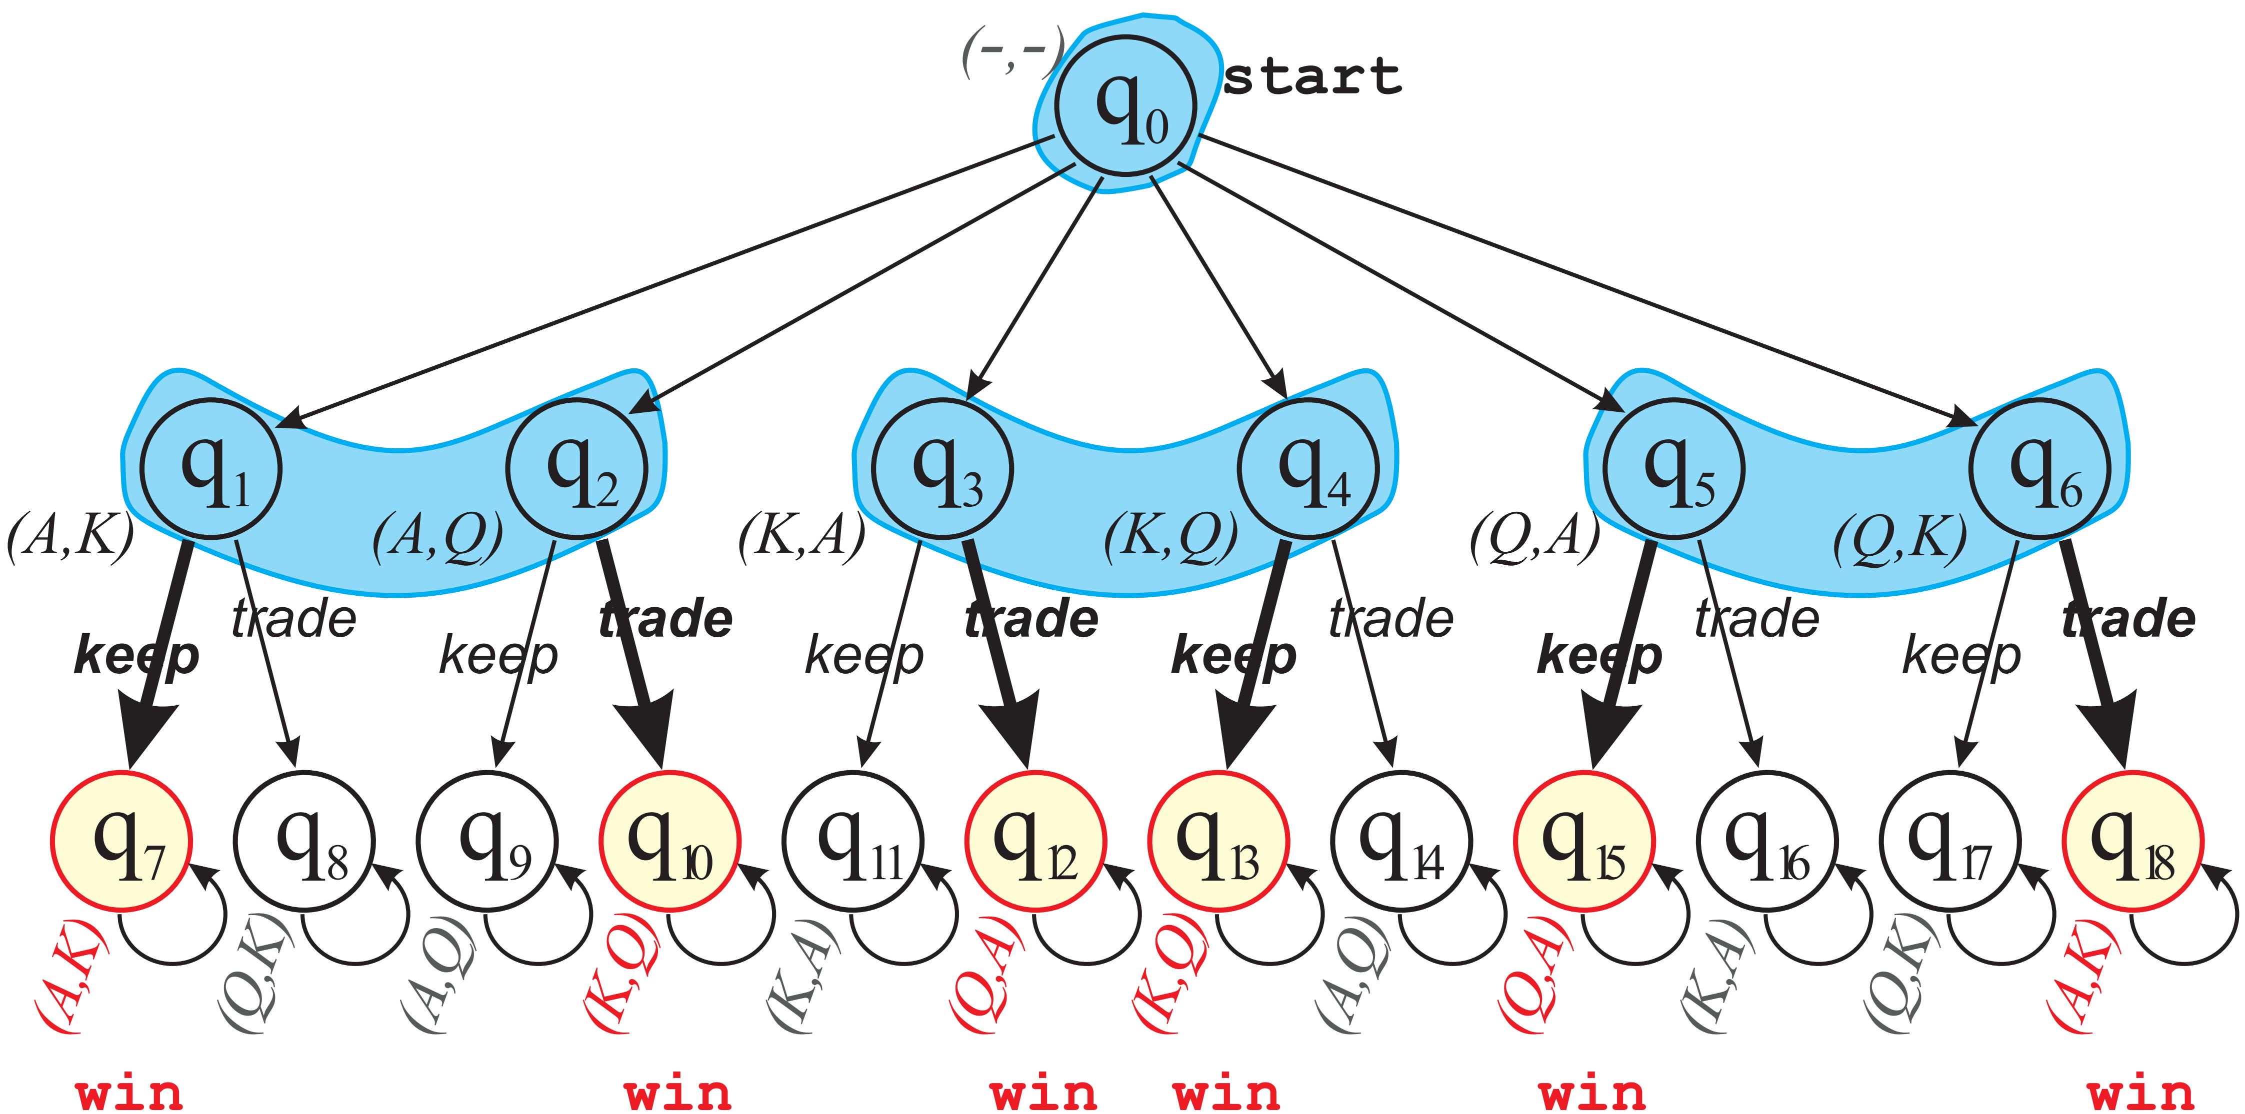

$$start \to \langle\langle a \rangle\rangle \Diamond win$$

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

Does it make sense?

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

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

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

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 (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.

Definition (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!

Example

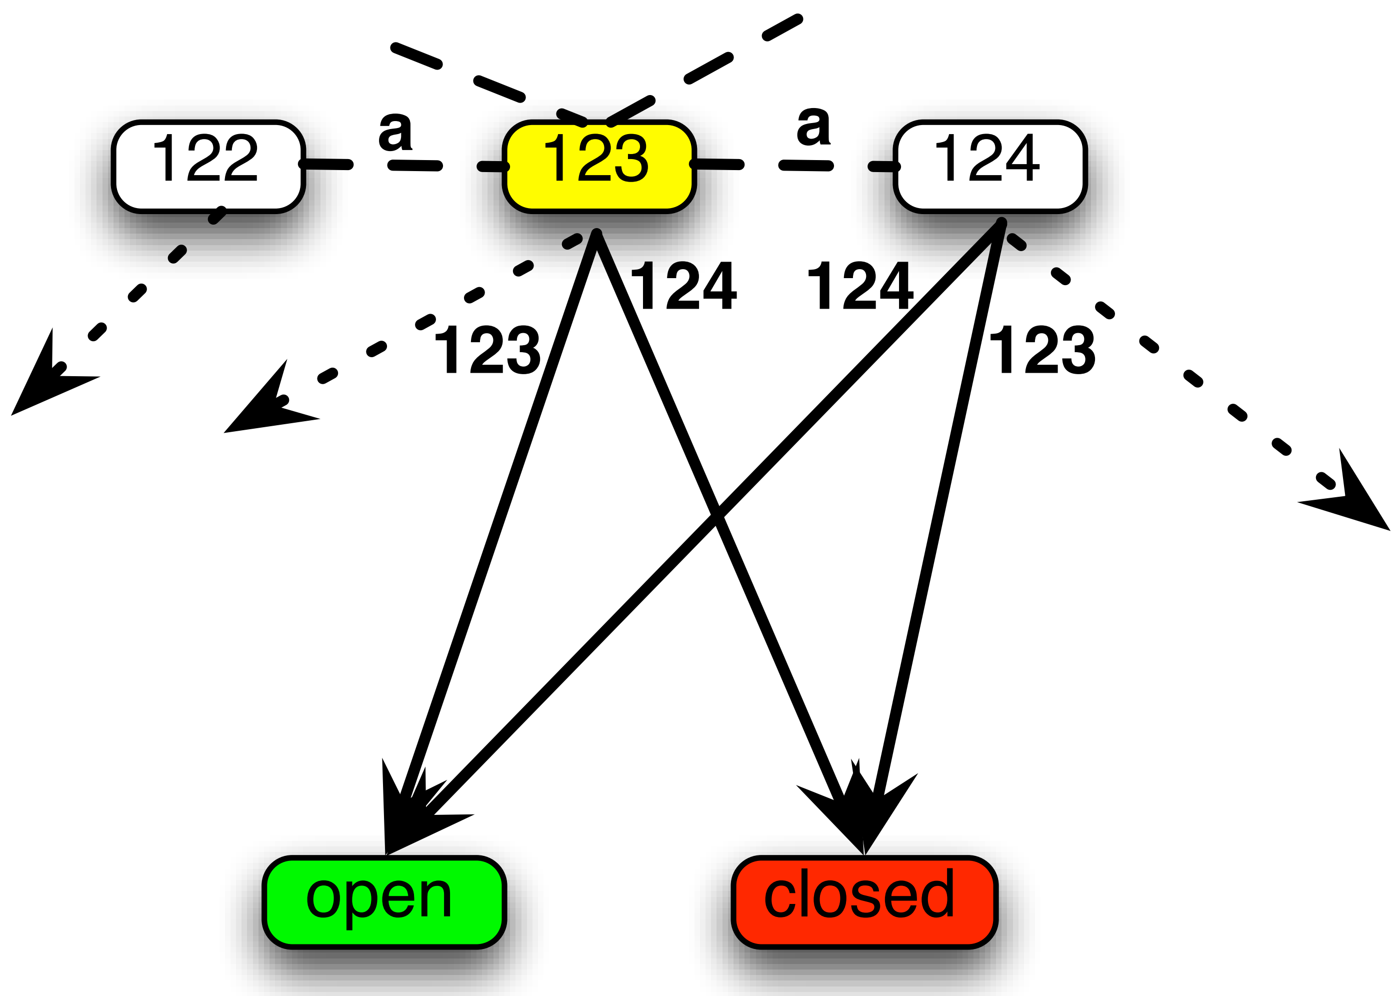

$\langle\langle a \rangle\rangle \bigcirc open$

$K_a \langle\langle a \rangle\rangle \bigcirc open$

Example

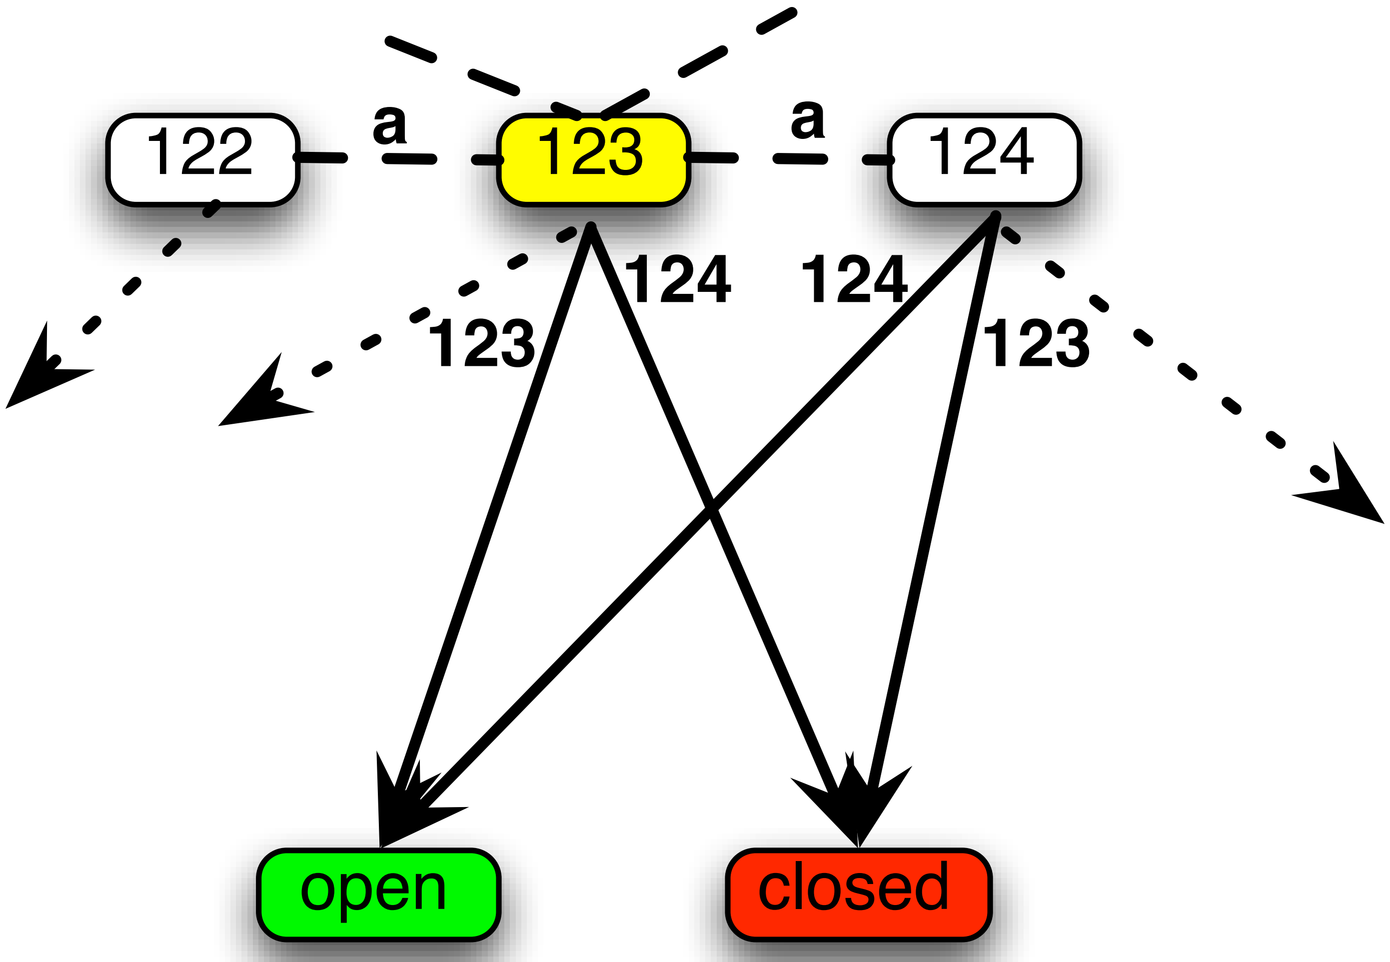

$$\langle\langle a \rangle\rangle \bigcirc open$$

$$K_a \langle\langle a \rangle\rangle \bigcirc open$$

Example

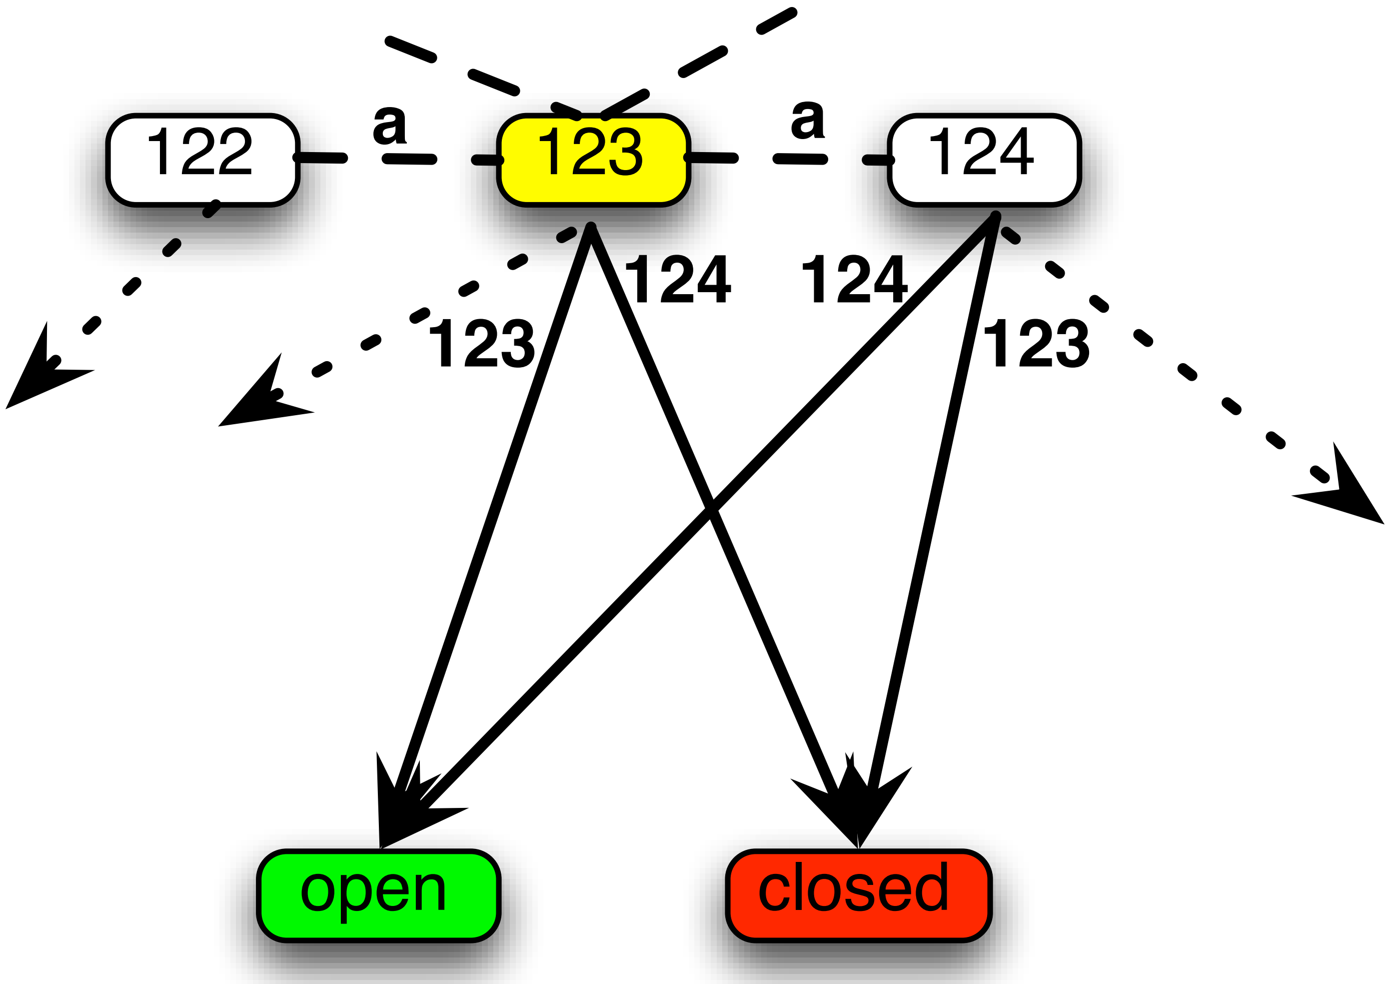

$$\langle\langle a \rangle\rangle \bigcirc open$$

$$K_a \langle\langle a \rangle\rangle \bigcirc open$$

Example

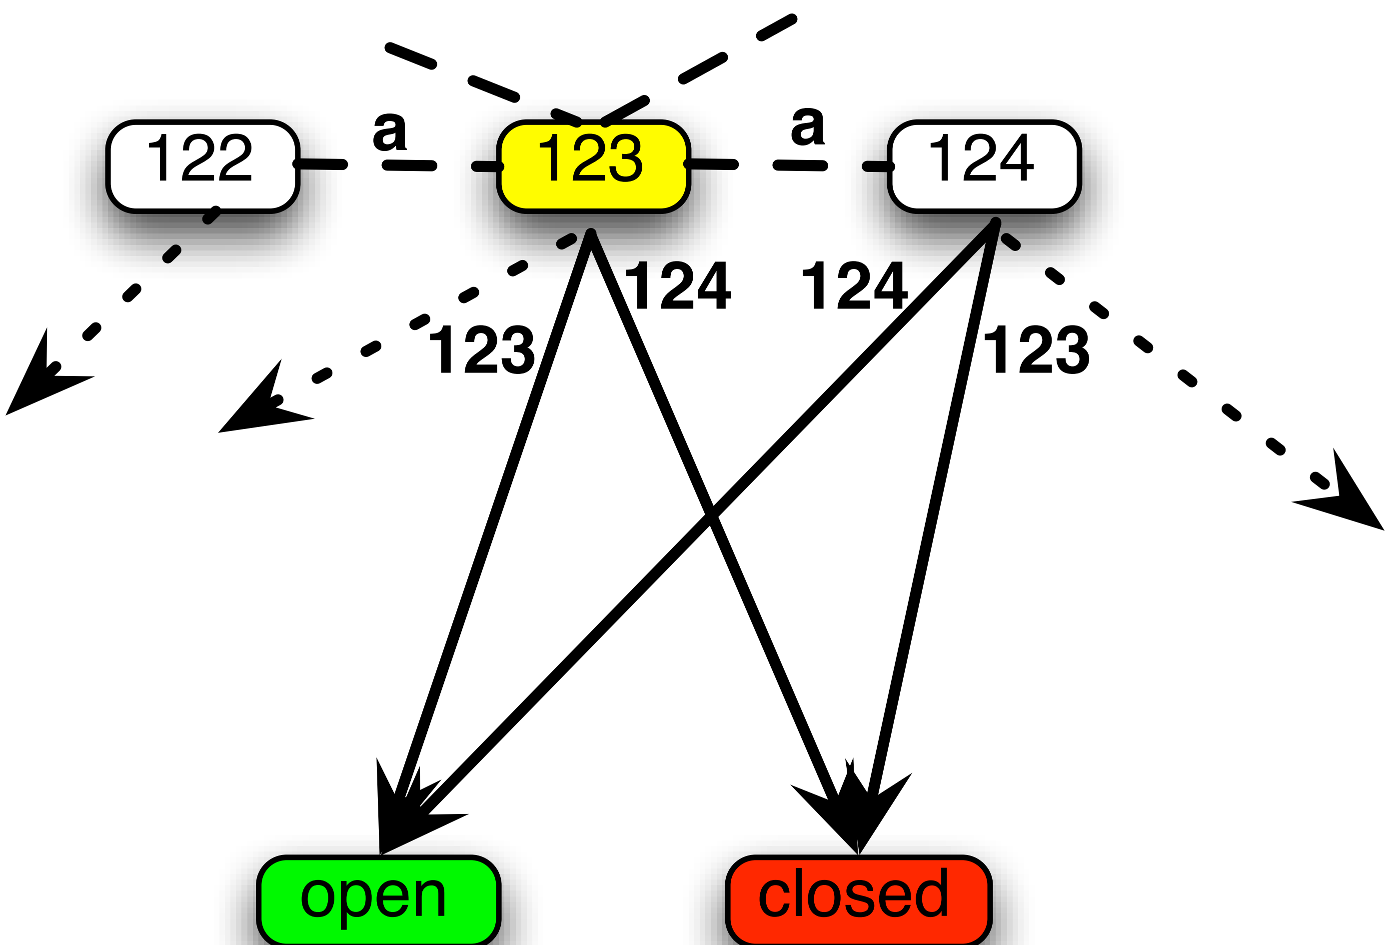

$$\langle\langle a \rangle\rangle \bigcirc open$$

$$K_a \langle\langle a \rangle\rangle \bigcirc open$$

Note:

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

Levels of Strategic Ability

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

1. There is σ (not necessarily executable!) such that, for every execution of σ, Φ holds
2. There is a uniform σ such that, for every execution of σ, Φ holds
3. A know that there is a uniform σ such that, for every execution of σ, Φ holds
4. There is a uniform σ such that A know that, for every execution of σ, Φ holds

From now on, we restrict our discussion to uniform memoryless strategies (unless explicitly stated otherwise).

Levels of Strategic Ability

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

1. There is σ (not necessarily executable!) such that, for every execution of σ, Φ holds
2. There is a uniform σ such that, for every execution of σ, Φ holds
3. A know that there is a uniform σ such that, for every execution of σ, Φ holds
4. There is a uniform σ such that A know that, for every execution of σ, Φ holds

From now on, we restrict our discussion to uniform memoryless strategies (unless explicitly stated otherwise).

Levels of Strategic Ability

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

1. There is σ (not necessarily executable!) such that, for every execution of σ, Φ holds
2. There is a uniform σ such that, for every execution of σ, Φ holds
3. A know that there is a uniform σ such that, for every execution of σ, Φ holds
4. There is a uniform σ such that A know that, for every execution of σ, Φ holds

From now on, we restrict our discussion to uniform memoryless strategies (unless explicitly stated otherwise).

Levels of Strategic Ability

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

1. There is σ (not necessarily executable!) such that, for every execution of σ, Φ holds
2. There is a uniform σ such that, for every execution of σ, Φ holds
3. *A* know that there is a uniform σ such that, for every execution of σ, Φ holds
4. There is a uniform σ such that *A* know that, for every execution of σ, Φ holds

From now on, we restrict our discussion to uniform memoryless strategies (unless explicitly stated otherwise).

Levels of Strategic Ability

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

1. There is σ (not necessarily executable!) such that, for every execution of σ, Φ holds
2. There is a uniform σ such that, for every execution of σ, Φ holds
3. *A* know that there is a uniform σ such that, for every execution of σ, Φ holds
4. There is a uniform σ such that *A* know that, for every execution of σ, Φ holds

From now on, we restrict our discussion to uniform memoryless strategies (unless explicitly stated otherwise).

Levels of Strategic Ability

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

1. There is σ (not necessarily executable!) such that, for every execution of σ, Φ holds
2. There is a uniform σ such that, for every execution of σ, Φ holds
3. A know that there is a uniform σ such that, for every execution of σ, Φ holds
4. There is a uniform σ such that A know that, for every execution of σ, Φ holds

From now on, we restrict our discussion to uniform memoryless strategies (unless explicitly stated otherwise).

Knowing how to play

- It turns out that knowledge of ability *de re* is not expressible in the language
- In *Constructive strategic logic (CSL)* (Jamroga and Ågotnes, 2007) ATL is extended with *constructive knowledge operators* such that

$$\mathbb{K}_a \langle\langle a \rangle\rangle \phi$$

means that a knows *de re* that she can achieve the goal

Constructive Strategic Logic: key idea

1. Interpret ability modalities in *sets* of states:
   - $M, Q \models \langle\langle a \rangle\rangle \phi$: there exists some strategy such that if a follows it *from any of the states in the set Q*, ϕ is guaranteed to be true
2. Introduce new *constructive knowledge* operators:
   - $M, q \models \mathbb{K}_a \phi \Leftrightarrow M, [q]_{\sim_a} \models \phi$

We get that:

$$M, q \models \mathbb{K}_a \langle\langle a \rangle\rangle \phi \Leftrightarrow M, [q]_{\sim_a} \models \langle\langle a \rangle\rangle \phi \Leftrightarrow$$

there exists some strategy such that if a follows it *from any of the states she considers possible*, ϕ is guaranteed to be true

Constructive Strategic Logic: key idea

1. Interpret ability modalities in *sets* of states:
   - $M, Q \models \langle\langle a \rangle\rangle \phi$: there exists some strategy such that if a follows it *from any of the states in the set Q*, ϕ is guaranteed to be true
2. Introduce new *constructive knowledge* operators:
   - $M, q \models \mathbb{K}_a \phi \Leftrightarrow M, [q]_{\sim_a} \models \phi$

We get that:

$$M, q \models \mathbb{K}_a \langle\langle a \rangle\rangle \phi \Leftrightarrow M, [q]_{\sim_a} \models \langle\langle a \rangle\rangle \phi \Leftrightarrow$$

there exists some strategy such that if a follows it *from any of the states she considers possible*, ϕ is guaranteed to be true

Constructive Strategic Logic: key idea

1. Interpret ability modalities in *sets* of states:
   - $M, Q \models \langle\langle a \rangle\rangle \phi$: there exists some strategy such that if a follows it *from any of the states in the set Q*, ϕ is guaranteed to be true
2. Introduce new *constructive knowledge* operators:
   - $M, q \models \mathbb{K}_a \phi \Leftrightarrow M, [q]_{\sim_a} \models \phi$

We get that:

$$M, q \models \mathbb{K}_a \langle\langle a \rangle\rangle \phi \Leftrightarrow M, [q]_{\sim_a} \models \langle\langle a \rangle\rangle \phi \Leftrightarrow$$

there exists some strategy such that if a follows it *from any of the states she considers possible*, ϕ is guaranteed to be true

Example

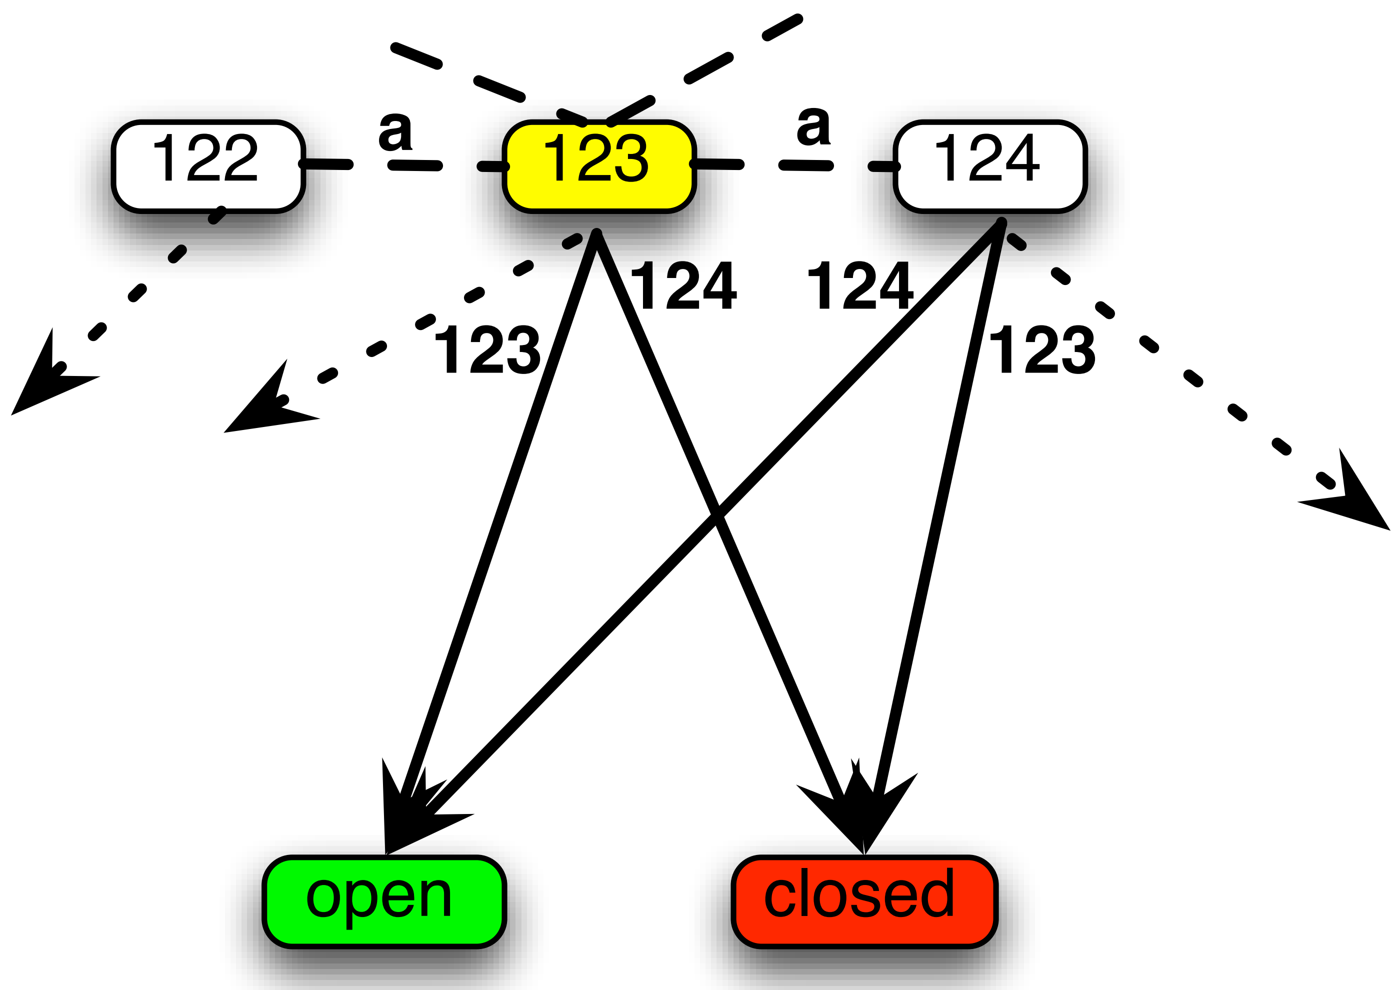

$$\langle\langle a\rangle\rangle \bigcirc open$$

$$K_a\langle\langle a\rangle\rangle \bigcirc open$$

$$\neg ck_a\langle\langle a\rangle\rangle \bigcirc open$$

Example

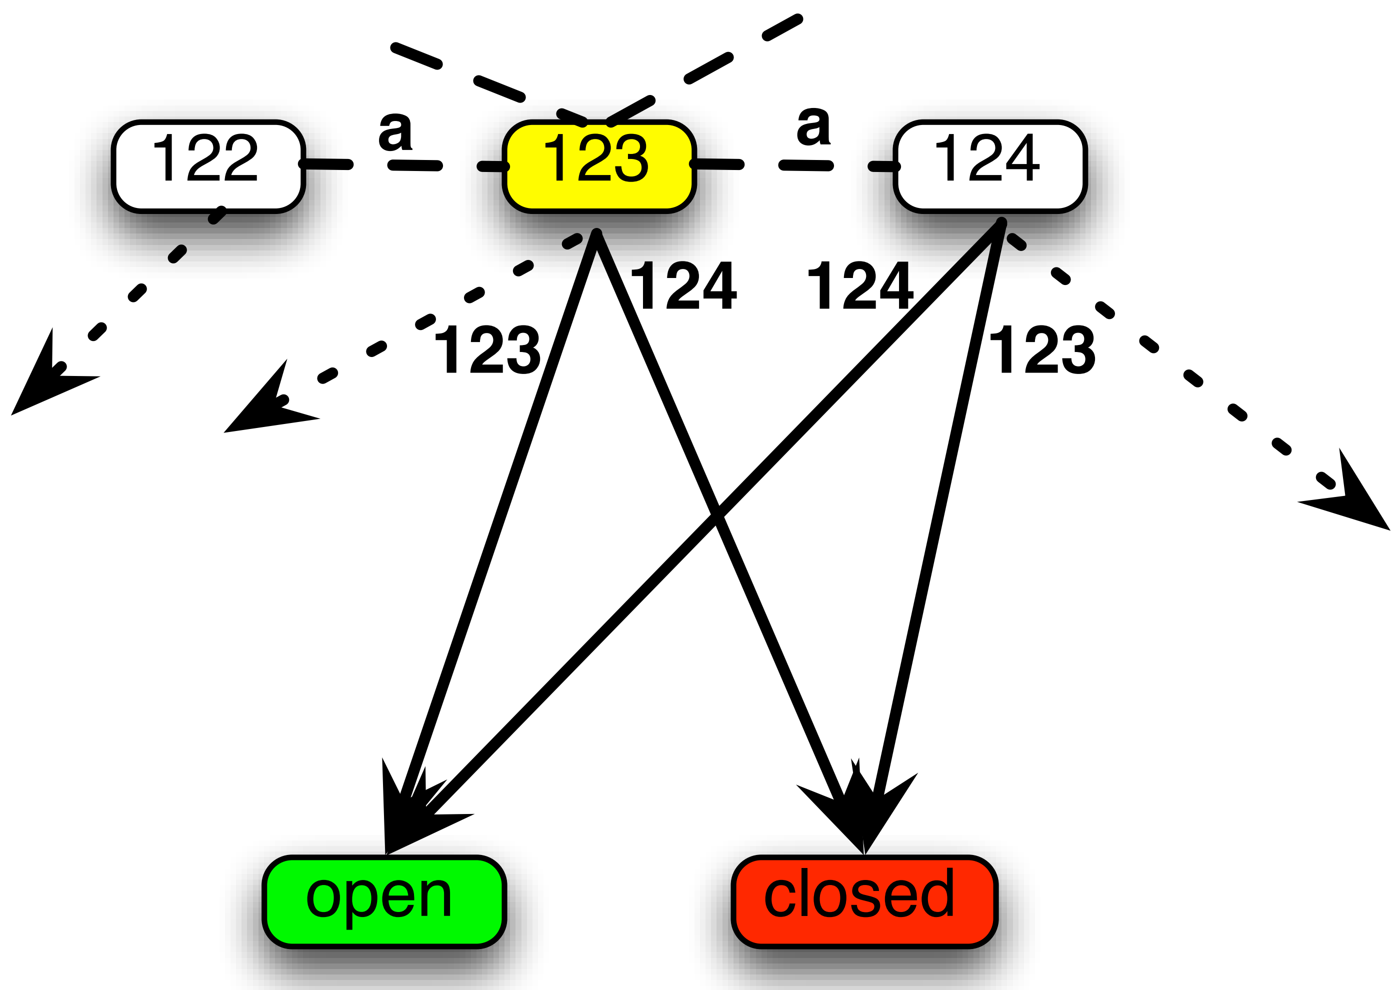

$$\langle\langle a \rangle\rangle \bigcirc open$$

$$K_a \langle\langle a \rangle\rangle \bigcirc open$$

$$\neg ck_a \langle\langle a \rangle\rangle \bigcirc open$$

Knowing how to Play

- Single agent case: we take into account the paths starting from indistinguishable states
- What about coalitions? In what sense should they know the strategy? Common knowledge (C_A), mutual knowledge (E_A), distributed knowledge (D_A)...?
- Other options also make sense!

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

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
- "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

- 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

- 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

- 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...

$\leadsto$ Solution: (general) constructive knowledge operators

Many subtle cases...

$\rightsquigarrow$ Solution: (general) constructive knowledge operators

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: $\mathbb{C}_A, \mathbb{E}_A, \mathbb{D}_A, \ldots$

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: $\mathbb{C}_A, \mathbb{E}_A, \mathbb{D}_A, \ldots$

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
- $M, Q \models \langle\langle A \rangle\rangle \gamma$: A have a single strategy to enforce γ 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 $M, \{q\} \models \varphi$

Non-standard semantics:

- Formulae are evaluated in sets of states
- $M, Q \models \langle\langle A \rangle\rangle \gamma$: A have a single strategy to enforce γ 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 $M, \{q\} \models \varphi$

Non-standard semantics:

- Formulae are evaluated in sets of states
- $M, Q \models \langle\langle A \rangle\rangle \gamma$: A have a single strategy to enforce γ 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 (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 $M, \lambda \models \gamma$;

$M, Q \models \mathcal{K}_A \varphi$ 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).

Definition (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 $M, \lambda \models \gamma$;

$M, Q \models \mathcal{K}_A \varphi$ 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).

Definition (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 $M, \lambda \models \gamma$;

$M, Q \models \mathcal{K}_A\varphi$ 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).

Definition (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 $M, \lambda \models \gamma$;

$M, Q \models \mathcal{K}_A\varphi$ 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).

Definition (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 $M, \lambda \models \gamma$;

$M, Q \models \mathcal{K}_A\varphi$ 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).

Definition (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 $M, \lambda \models \gamma$;

$M, Q \models \mathcal{K}_A \varphi$ 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).