Implicit coordination for epistemic planning and human-robot collaboration

Thomas Bolander, DTU Compute

LOFT 2024

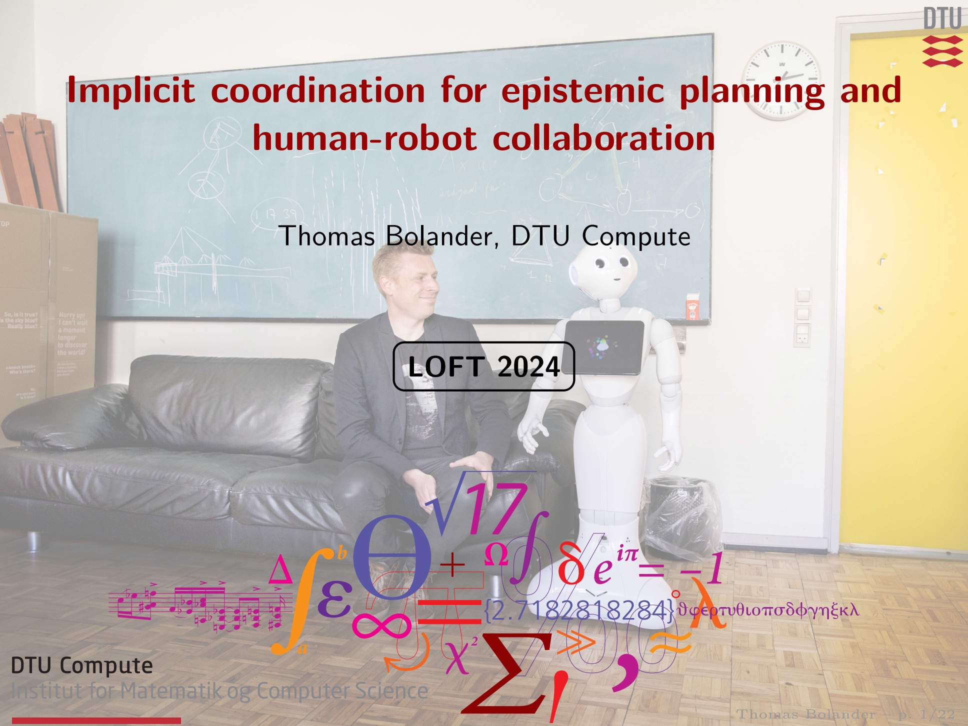

DTU Compute
Institut for Matematik og Computer Science

Epistemic planning $=$

automated *planning* (AI) $+$ *epistemic* reasoning (epistemic logic)

Aim: To compute plans that can take the mental states of other agents into account.

Essentially: (Decentralised) **multi-agent planning** in environments with (potentially higher-order) **information asymmetry**.

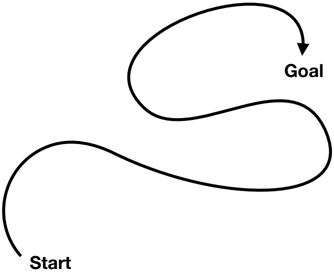

Automated planning

$+$

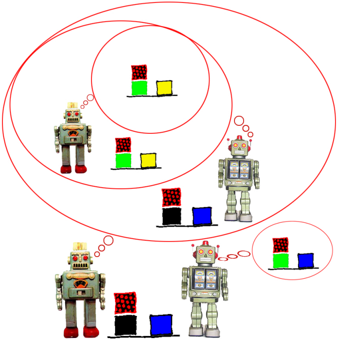

Epistemic reasoning about the mental states of others

Classical automated planning: state space search and domain descriptions

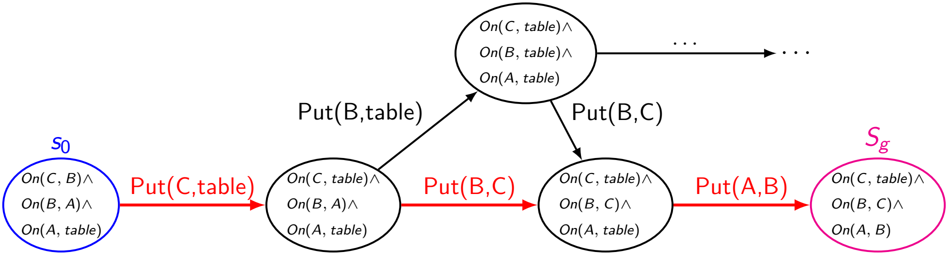

Action schema describing the $Put(x, y)$ action for put object x on top of object y:

$$\begin{array}{|l|}\hline \textsc{Action} : Put(x,y) \\ \textsc{Precondition} : On(x,z) \wedge \cdots \\ \textsc{Effect} : On(x,y) \wedge \neg On(x,z) \\ \hline \end{array}$$

STRIPS/PDDL

$$\begin{array}{|ll|}\hline \text{pre}: & On(x,z) \wedge \cdots \\ \hline \text{post}: & \begin{array}{l} On(x,y) := \top \\ On(x,z) := \bot \end{array} \\ \hline \end{array}$$

Dynamic Epistemic Logic (DEL)

[Ghallab et al., 2004, Baltag et al., 1998, van Ditmarsch and Kooi, 2008]

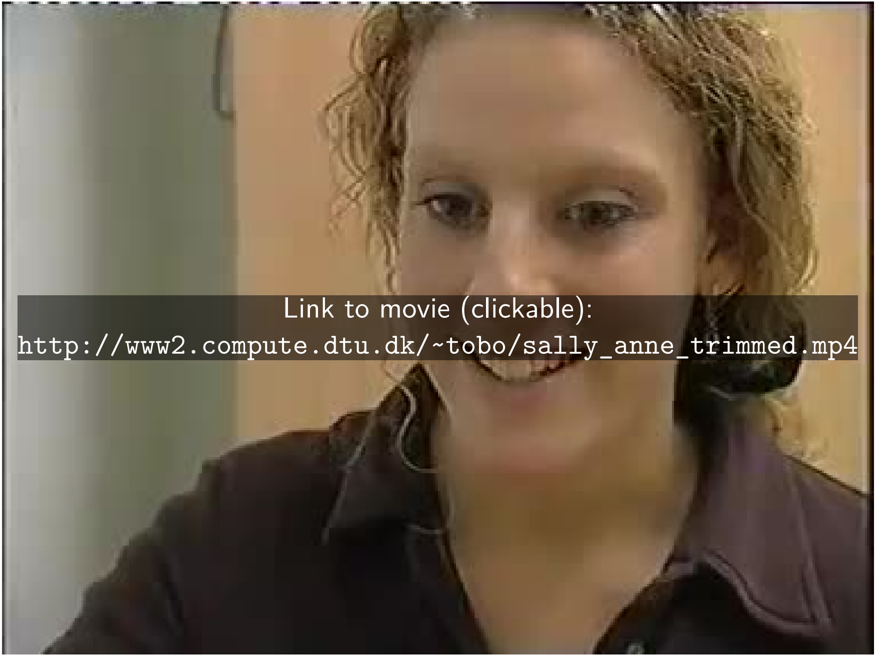

Link to movie (clickable):
http://www2.compute.dtu.dk/~tobo/sally_anne_trimmed.mp4

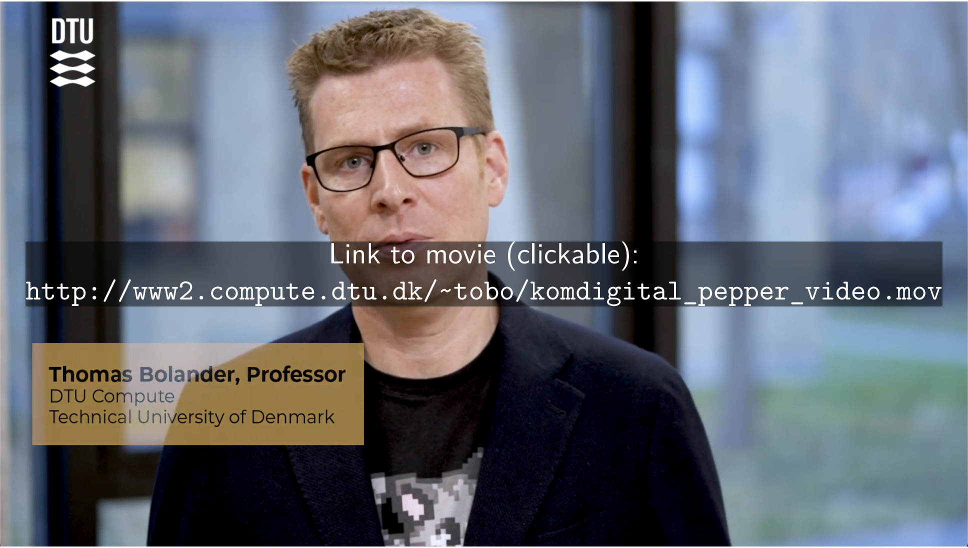

KomDigital: R2DTU – A Pepper robot, 25 November 2020 [?]

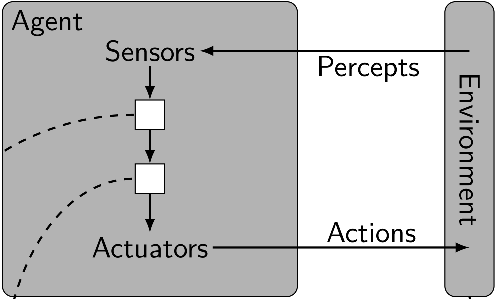

perception layer:
deep neural networks (subsymbolic)

cognition layer:
DEL + planning (symbolic)

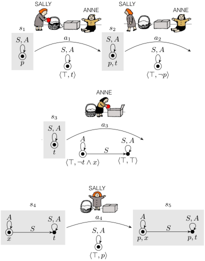

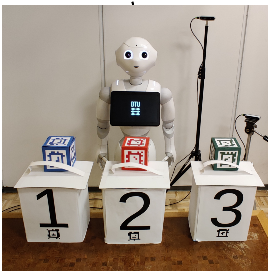

But...

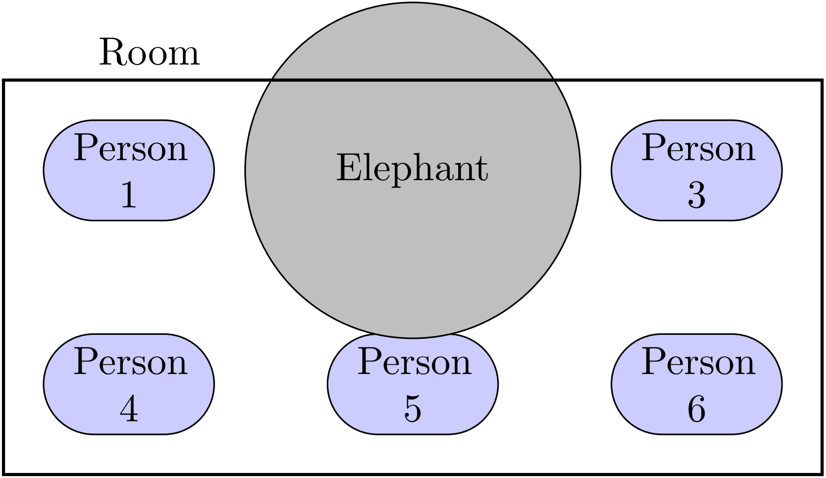

(poor person 2)

Sally-Anne with a twist (in GPT-4 = ChatGPT+)

USER

Sally and Anne are playing with a marble. Sally puts the marble into her basket. Then she goes outside to play. While she is outside, Anne drinks a coke. Now Sally comes back. Sally sees Anne taking the marble from the basket and putting it into her box. Where will Sally look for her marble?

ASSISTANT

Sally will look for her marble in her basket. Since that was where she originally placed it before leaving, and she did not see Anne move the marble, she would expect it to still be there.

Multiagent case: States as S5 Kripke models

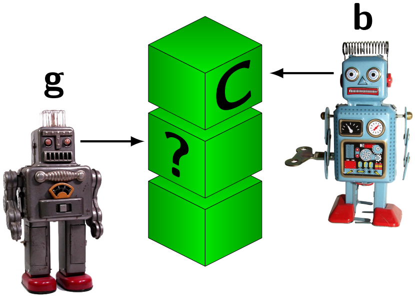

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.
Designated worlds: ▢ (those considered possible by planning agent).

Multiagent case: States as S5 Kripke models

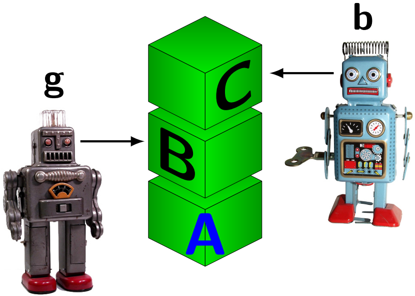

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.
Designated worlds: (those considered possible by planning agent).

Multiagent case: States as S5 Kripke models

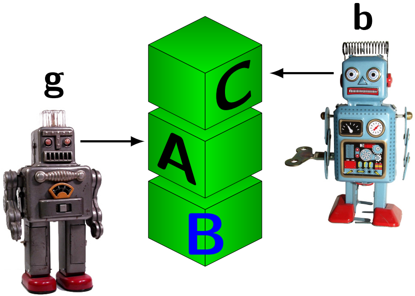

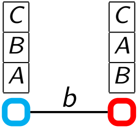

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.
Designated worlds: (those considered possible by planning agent).

Multiagent case: States as S5 Kripke models

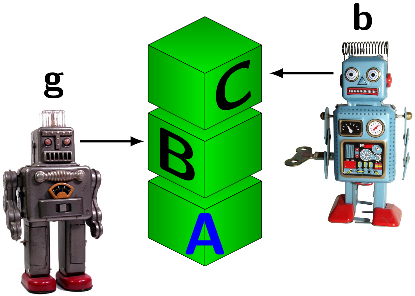

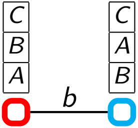

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.
Designated worlds: (those considered possible by planning agent).

Multiagent case: States as S5 Kripke models

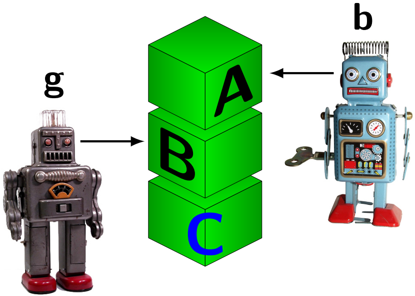

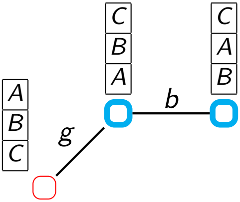

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.
Designated worlds: (those considered possible by planning agent).

Multiagent case: States as S5 Kripke models

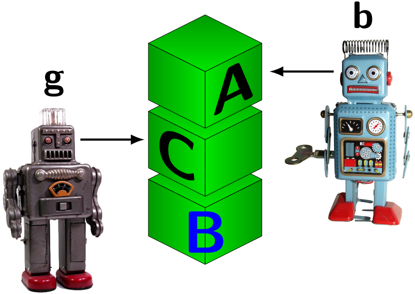

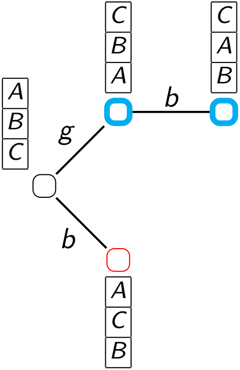

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.
Designated worlds: (those considered possible by planning agent).

Multiagent case: States as S5 Kripke models

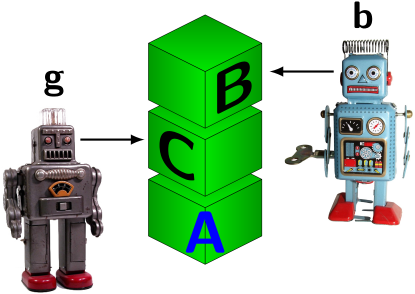

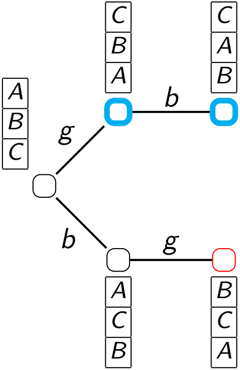

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.
Designated worlds: (those considered possible by planning agent).

Multiagent case: States as S5 Kripke models

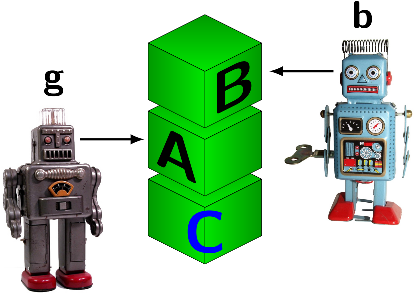

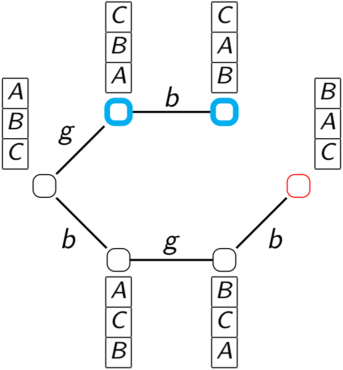

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.
Designated worlds: (those considered possible by planning agent).

Multiagent case: States as S5 Kripke models

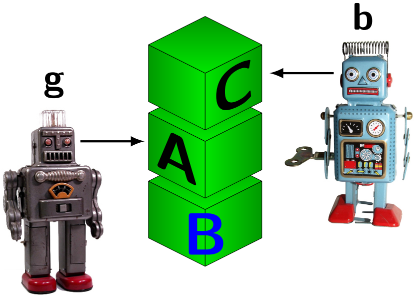

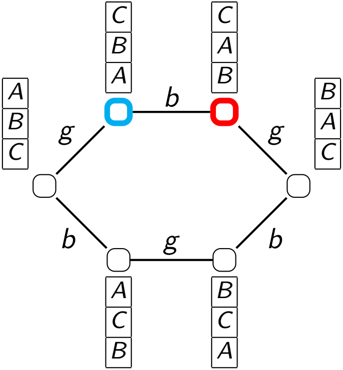

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.
Designated worlds: (those considered possible by planning agent).

Multiagent case: States as S5 Kripke models

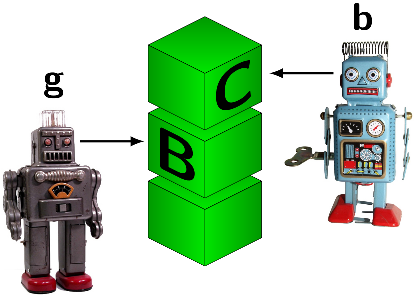

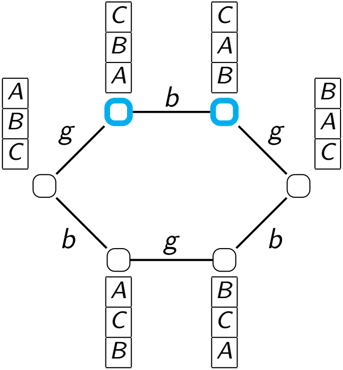

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.
Designated worlds: (those considered possible by planning agent).

Multiagent case: States as S5 Kripke models

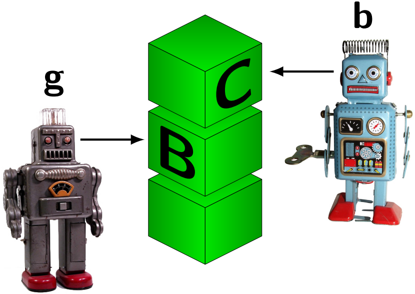

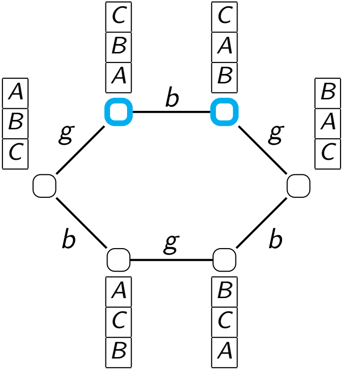

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.
Designated worlds: (those considered possible by planning agent).

Agent *b*: "Which letter does the middle block have?"

Multiagent case: States as S5 Kripke models

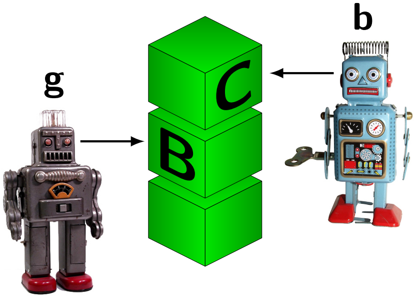

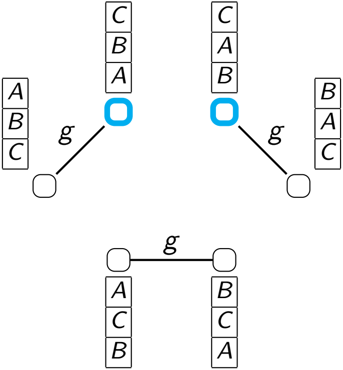

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.
Designated worlds: (those considered possible by planning agent).

Agent *b*: "Which letter does the middle block have?"

Dynamic epistemic logic (DEL) by example: product update

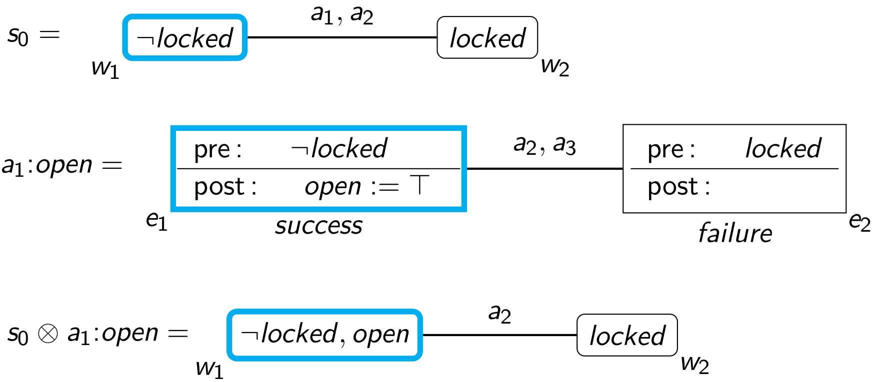

$a_1{:}open$ is an **event model** (representing an action). In these, nodes are **events**, and each event has a **precondition** (epistemic formula) and **postconditions** for all atoms (also epistemic formulas).

[Baltag et al., 1998, van Ditmarsch and Kooi, 2008]

Dynamic epistemic logic (DEL) by example: product update

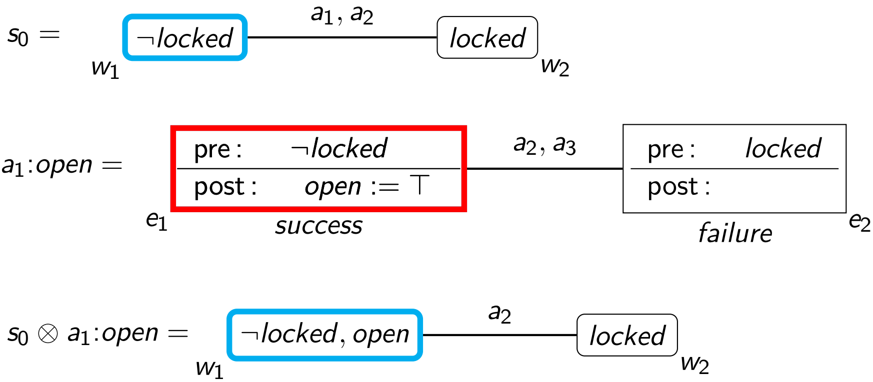

$a_1{:}open$ is an **event model** (representing an action). In these, nodes are **events**, and each event has a **precondition** (epistemic formula) and **postconditions** for all atoms (also epistemic formulas).

[Baltag et al., 1998, van Ditmarsch and Kooi, 2008]

Dynamic epistemic logic (DEL) by example: product update

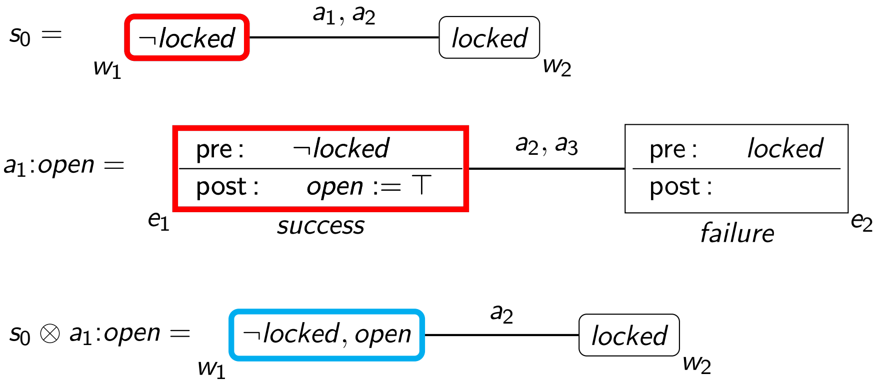

$a_1{:}open$ is an **event model** (representing an action). In these, nodes are **events**, and each event has a **precondition** (epistemic formula) and **postconditions** for all atoms (also epistemic formulas).

[Baltag et al., 1998, van Ditmarsch and Kooi, 2008]

Dynamic epistemic logic (DEL) by example: product update

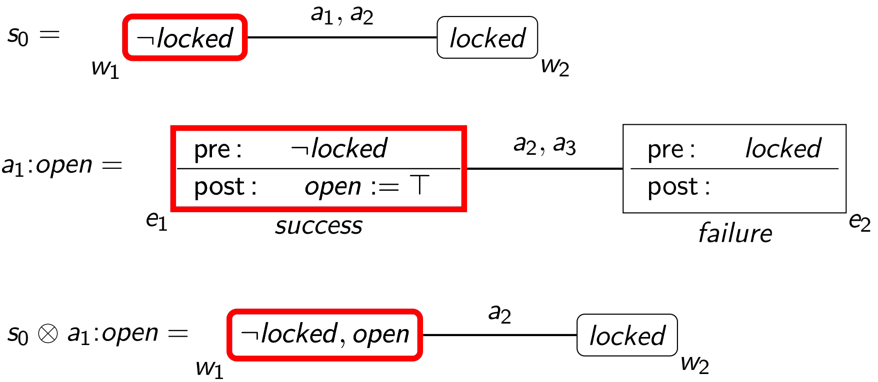

$a_1{:}open$ is an **event model** (representing an action). In these, nodes are **events**, and each event has a **precondition** (epistemic formula) and **postconditions** for all atoms (also epistemic formulas).

[Baltag et al., 1998, van Ditmarsch and Kooi, 2008]

Dynamic epistemic logic (DEL) by example: product update

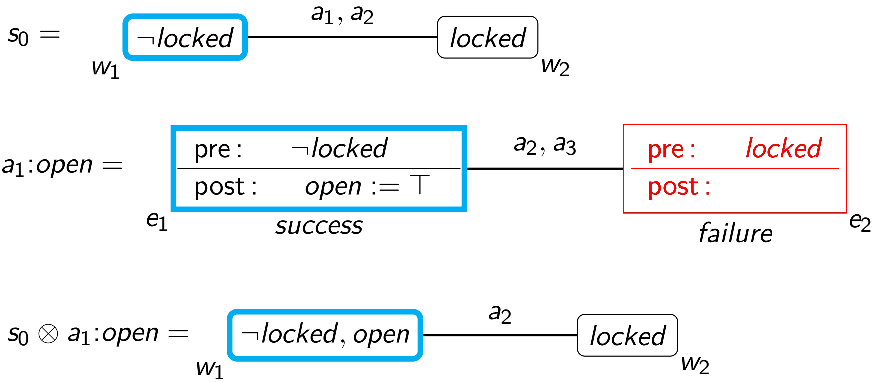

$a_1{:}open$ is an **event model** (representing an action). In these, nodes are **events**, and each event has a **precondition** (epistemic formula) and **postconditions** for all atoms (also epistemic formulas).

[Baltag et al., 1998, van Ditmarsch and Kooi, 2008]

Dynamic epistemic logic (DEL) by example: product update

$$s_0 = \begin{tikzcd} \underset{w_1}{\boxed{\neg \mathit{locked}}} \arrow[rr, dash, "a_1, a_2"] & & \underset{w_2}{\boxed{\mathit{locked}}} \end{tikzcd}$$

$$a_1{:}\mathit{open} = \begin{tikzcd} \underset{e_1 \quad \mathit{success}}{\boxed{\begin{array}{ll} \text{pre}: & \neg \mathit{locked} \\ \hline \text{post}: & \mathit{open} := \top \end{array}}} \arrow[rr, dash, "a_2, a_3"] & & \underset{\mathit{failure} \quad e_2}{\boxed{\begin{array}{ll} \text{pre}: & \mathit{locked} \\ \hline \text{post}: & \end{array}}} \end{tikzcd}$$

$$s_0 \otimes a_1{:}\mathit{open} = \begin{tikzcd} \underset{w_1}{\boxed{\neg \mathit{locked}, \mathit{open}}} \arrow[rr, dash, "a_2"] & & \underset{w_2}{\boxed{\mathit{locked}}} \end{tikzcd}$$

$a_1{:}\mathit{open}$ is an **event model** (representing an action). In these, nodes are **events**, and each event has a **precondition** (epistemic formula) and **postconditions** for all atoms (also epistemic formulas).

[Baltag et al., 1998, van Ditmarsch and Kooi, 2008]

Dynamic epistemic logic (DEL) by example: product update

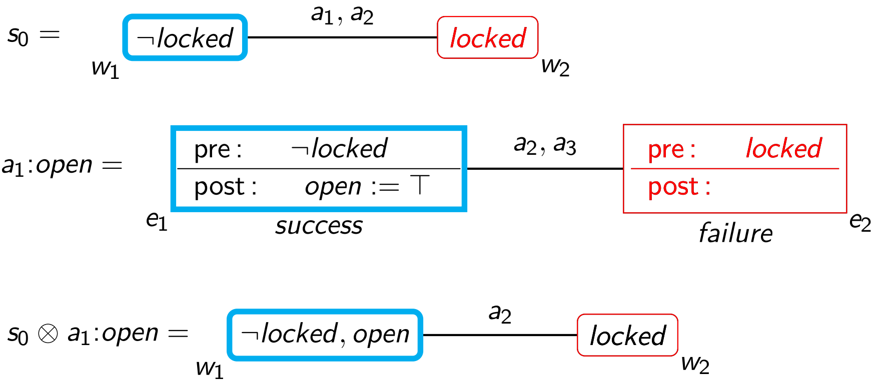

$a_1{:}open$ is an **event model** (representing an action). In these, nodes are **events**, and each event has a **precondition** (epistemic formula) and **postconditions** for all atoms (also epistemic formulas).

[Baltag et al., 1998, van Ditmarsch and Kooi, 2008]

Dynamic epistemic logic (DEL) by example: product update

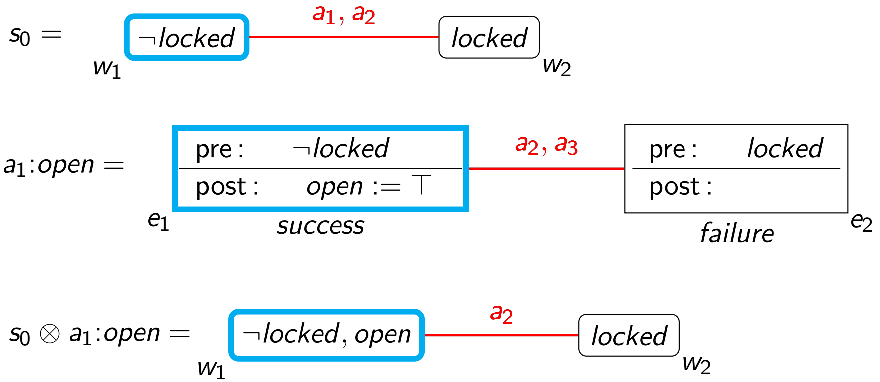

$a_1{:}open$ is an **event model** (representing an action). In these, nodes are **events**, and each event has a **precondition** (epistemic formula) and **postconditions** for all atoms (also epistemic formulas).

[Baltag et al., 1998, van Ditmarsch and Kooi, 2008]

s_0

A, B

$A \lhd B, B \lhd A$

α_1

A:
$+In(cube_{red}, box_1)$

s_1

A, B

$A \lhd B, B \lhd A,$
$In(cube_{red}, box_1)$

α_2

A:
$-A \lhd B, -B \lhd A$

s_2

A, B

$In(cube_{red}, box_1)$

α_3

B:
$-In(cube_{red}, box_1)$

Atom $In(c_1, c_2)$: obj. c_1 is in container c_2.

Atom $a \lhd b$ (a sees b): agent a currently observes the actions of agent b.

Action a:X: X is a list of assignments ($+p$ and $-p$ where p is an atom).

s_3

B $\quad$ A, B

$w_l \xrightarrow{A} w_r$

$In(cube_{red}, box_1)$

α_4

B:
$+In(cube_{red}, box_2)$

s_4

B $\quad$ A, B

$w_l \xrightarrow{A} w_r$

$In(cube_{red}, box_2)$ $\quad$ $In(cube_{red}, box_1)$

Second-order false-belief tasks

Full formalisation of second-order chocolate task:
$boy{:}{+}In(choc, drawer)$; $boy{:}{-}boy \triangleleft girl, -girl \triangleleft boy$; $boy{:}{+}boy \triangleleft girl$; $girl{:}{-}In(choc, drawer)$; $girl{:}{+}In(choc, box)$.

True in resulting state: $B_{girl}\, box \wedge B_{boy}\, box \wedge B_{girl} B_{boy}\, drawer$.

Planning based on DEL: epistemic planning tasks

Definition. An **(epistemic) planning task** $T = (s_0, A, \varphi_g)$ consists of

- A multipointed Kripke model s_0 called the **initial state**.
- A finite set of multipointed event models A called **actions**.
- A **goal formula** φ_g of epistemic logic.

Definition. A (sequential) **plan** for a planning task $T = (s_0, A, \varphi_g)$ is a sequence of actions $\alpha_1, \alpha_2, \ldots, \alpha_n$ from A such that for all $1 \leq i \leq n$, α_i is applicable in $s_0 \otimes \alpha_1 \otimes \cdots \otimes \alpha_{i-1}$ and

$$s_0 \otimes \alpha_1 \otimes \alpha_2 \otimes \cdots \otimes \alpha_n \models \varphi_g.$$

Defining $(\!(\alpha)\!)\varphi := \langle\alpha\rangle\top \wedge [\alpha]\varphi$, this can be reformulated as

$$s_0 \models (\!(\alpha_1)\!)(\!(\alpha_2)\!)\cdots(\!(\alpha_n)\!)\varphi_g.$$

Definition. A plan $i_1{:}\alpha_1, \ldots i_n{:}\alpha_n$ (using notation $i{:}\alpha$ for agent i performing action α) is **implicitly coordinated** if it furthermore holds that :

$$s_0 \models K_{i_1}(\!(i_1{:}\alpha_1)\!)K_{i_2}(\!(i_2{:}\alpha_2)\!)\cdots K_{i_n}(\!(i_n{:}\alpha_n)\!)\varphi_g.$$

[Bolander and Andersen, 2011, Engesser et al., 2017, Bolander et al., 2020]

Conflicting implicitly coordinated plans: Move to (nondeterministic) policies

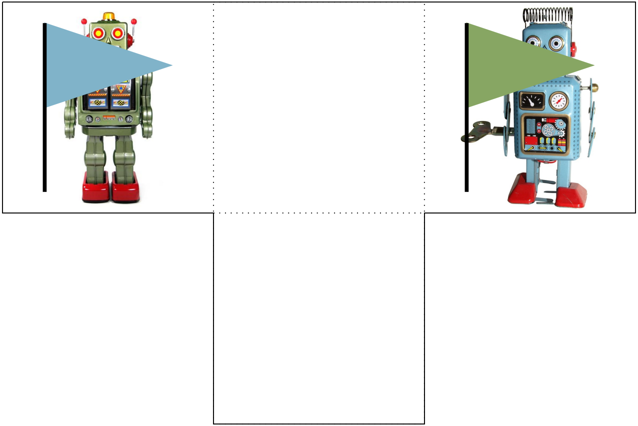

[Bolander et al., 2018]

Conflicting implicitly coordinated plans: Move to (nondeterministic) policies

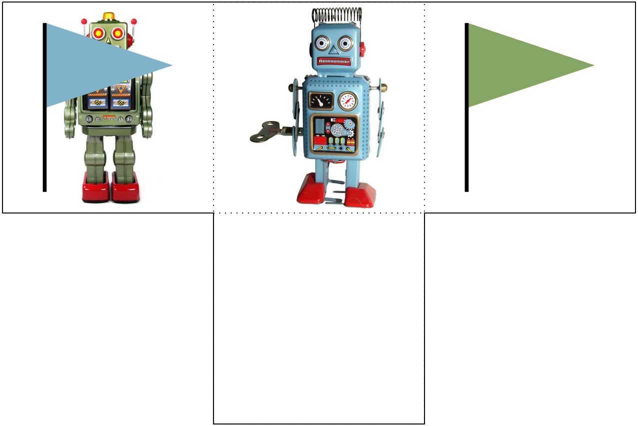

[Bolander et al., 2018]

Conflicting implicitly coordinated plans: Move to (nondeterministic) policies

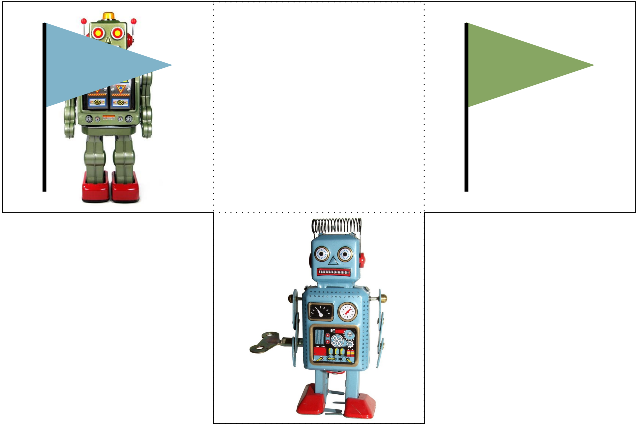

[Bolander et al., 2018]

Conflicting implicitly coordinated plans: Move to (nondeterministic) policies

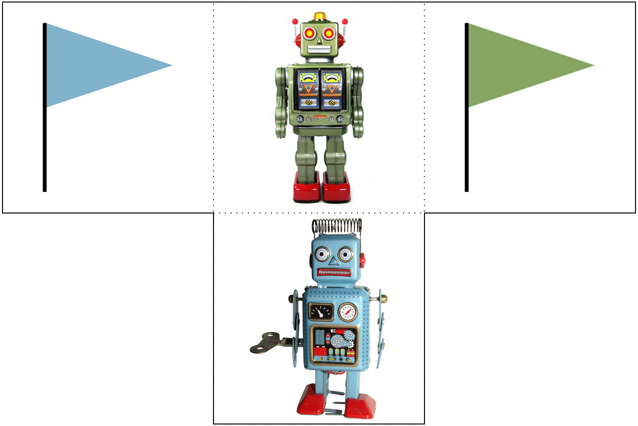

[Bolander et al., 2018]

Conflicting implicitly coordinated plans: Move to (nondeterministic) policies

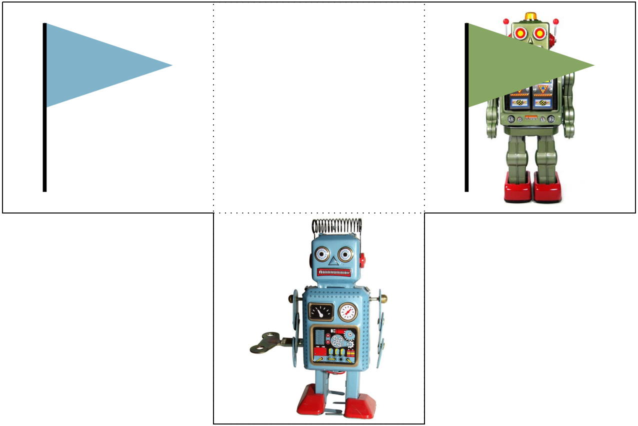

[Bolander et al., 2018]

Conflicting implicitly coordinated plans: Move to (nondeterministic) policies

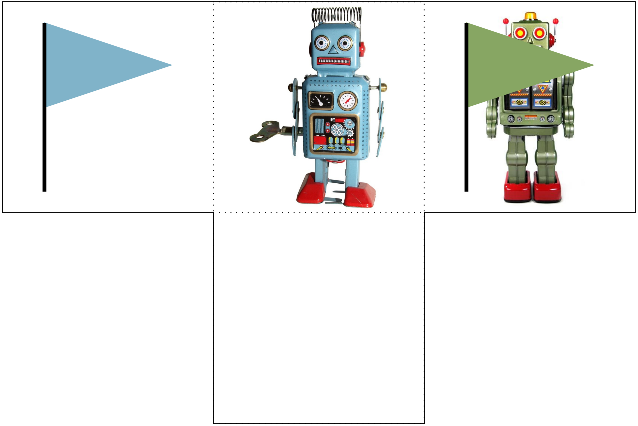

[Bolander et al., 2018]

Conflicting implicitly coordinated plans: Move to (nondeterministic) policies

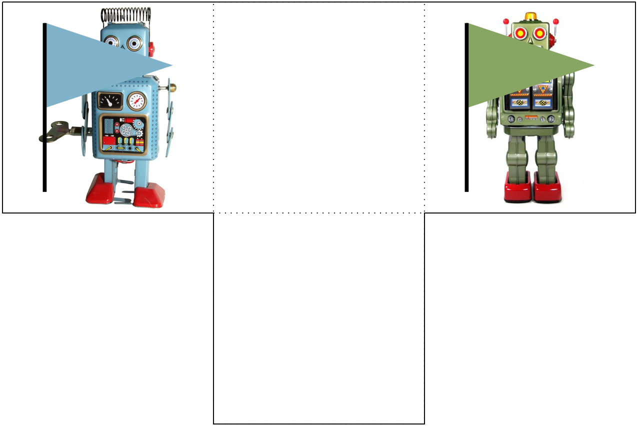

[Bolander et al., 2018]

Implicit coordination: multi-agent pathfinding with destination uncertainty

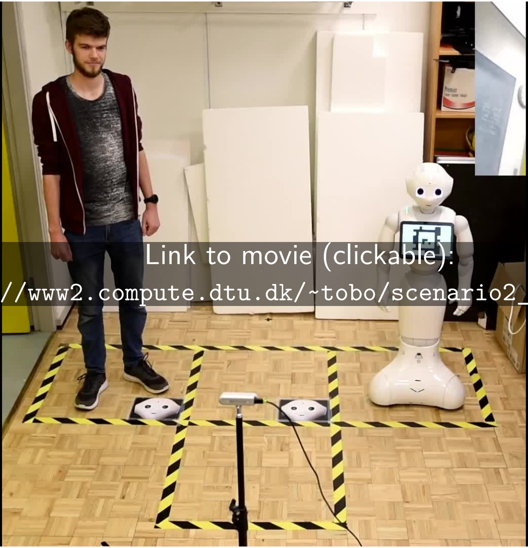

Link to movie (clickable):
http://www2.compute.dtu.dk/~tobo/scenario2_double.mp4

[Nebel et al., 2019, Bolander et al., 2021]

Towards forward induction in implicit coordination

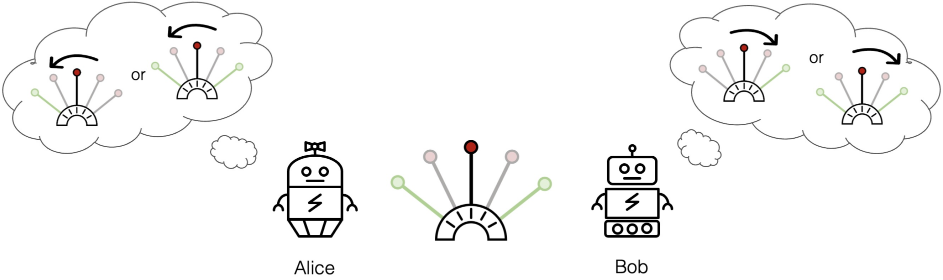

Atomic propositions:

- l: There's a goal in the leftmost position (pos. 1)
- r: There's a goal is the rightmost position (pos. 5)
- p_i: The lever is at position i

Goal: $\varphi_g = (l \wedge p_1) \vee (r \wedge p_5)$

Alice (a) can only pull left ($a{:}L$), Bob (b) can only pull right ($b{:}R$).

Implicit coordination with (simple) forward induction

Each agent plans from their local perspective, i.e., we close the set of designated worlds under the accessibility relation of that agent. Below the perspective of *a*.

Implicit coordination with (simple) forward induction

Each agent plans from their local perspective, i.e., we close the set of designated worlds under the accessibility relation of that agent. Below the perspective of a.

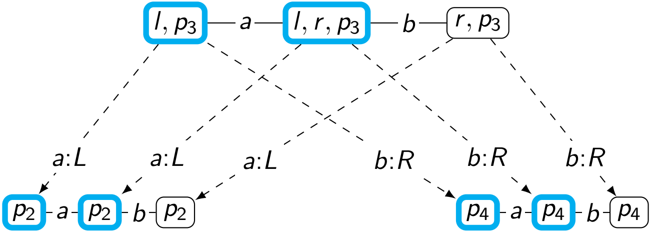

Implicit coordination with (simple) forward induction

Each agent plans from their local perspective, i.e., we close the set of designated worlds under the accessibility relation of that agent. Below the perspective of a.

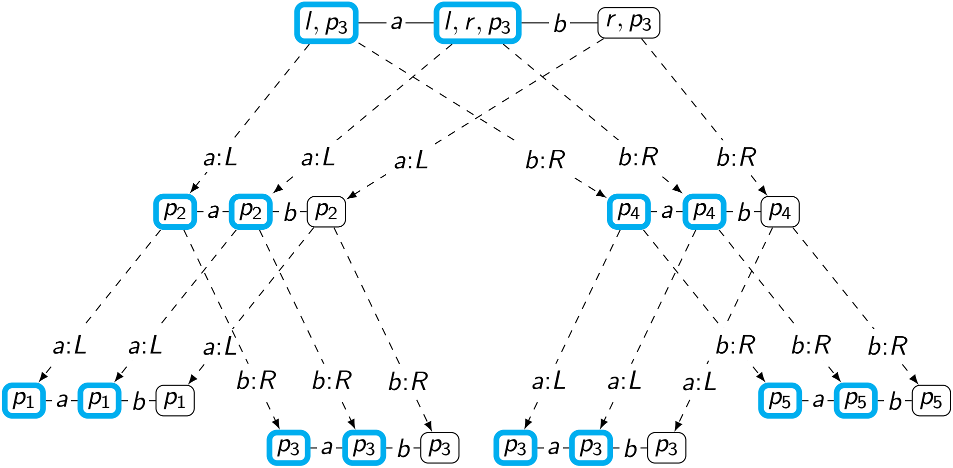

Implicit coordination with (simple) forward induction

Each agent plans from their local perspective, i.e., we close the set of designated worlds under the accessibility relation of that agent. Below the perspective of a.

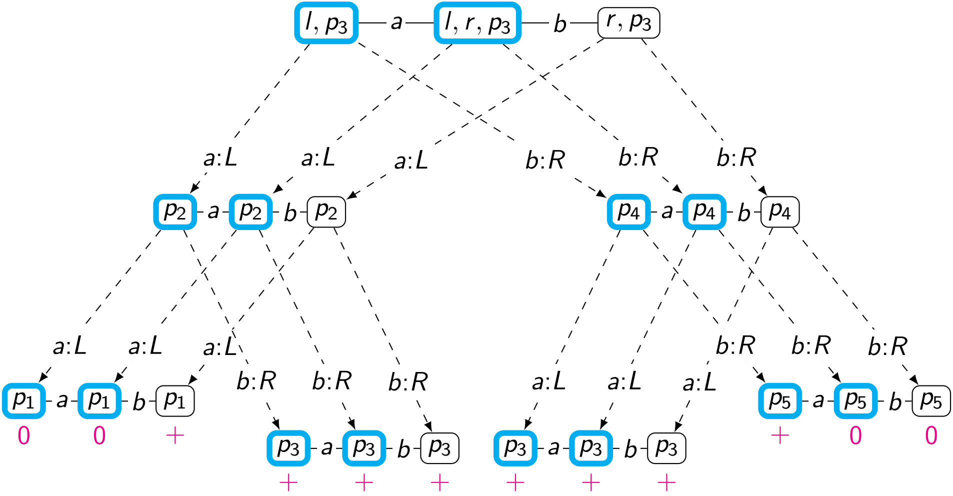

Implicit coordination with (simple) forward induction

Each agent plans from their local perspective, i.e., we close the set of designated worlds under the accessibility relation of that agent. Below the perspective of a.

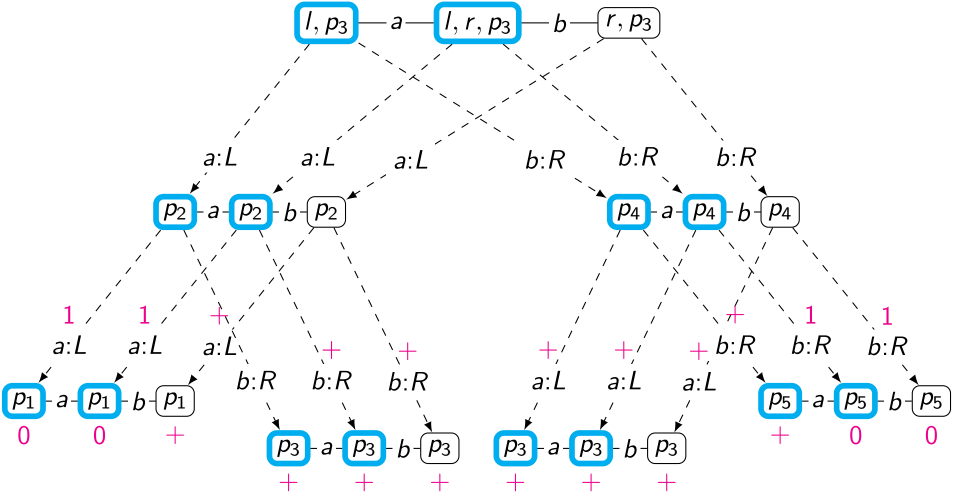

Implicit coordination with (simple) forward induction

Each agent plans from their local perspective, i.e., we close the set of designated worlds under the accessibility relation of that agent. Below the perspective of a.

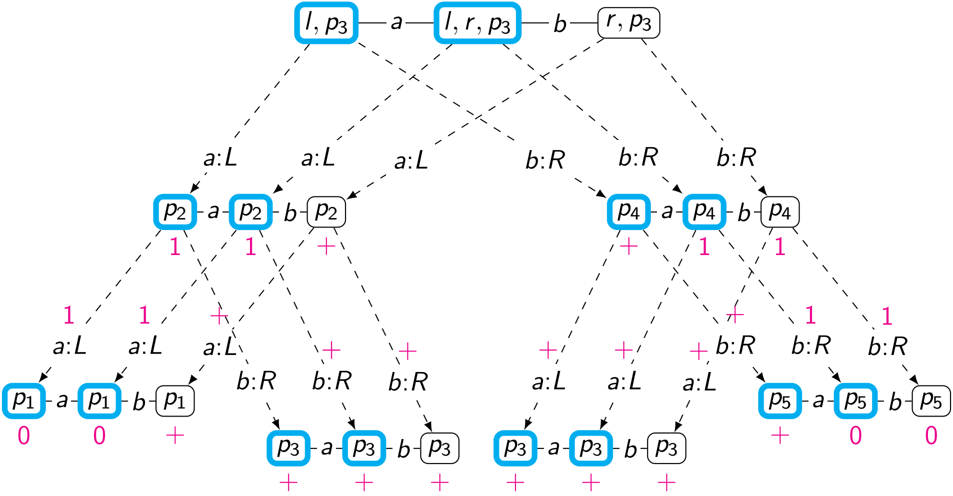

Implicit coordination with (simple) forward induction

Each agent plans from their local perspective, i.e., we close the set of designated worlds under the accessibility relation of that agent. Below the perspective of a.

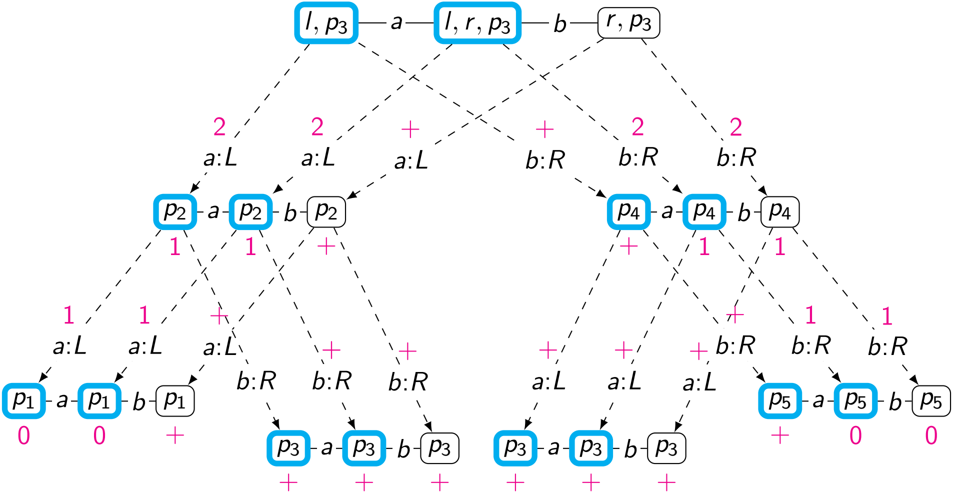

Implicit coordination with (simple) forward induction

Each agent plans from their local perspective, i.e., we close the set of designated worlds under the accessibility relation of that agent. Below the perspective of a.

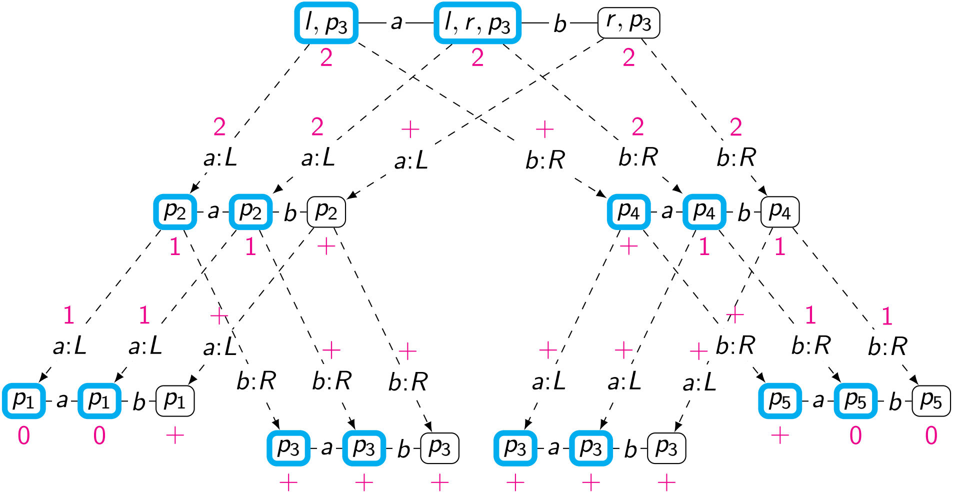

Implicit coordination with (simple) forward induction

Each agent plans from their local perspective, i.e., we close the set of designated worlds under the accessibility relation of that agent. Below the perspective of a.

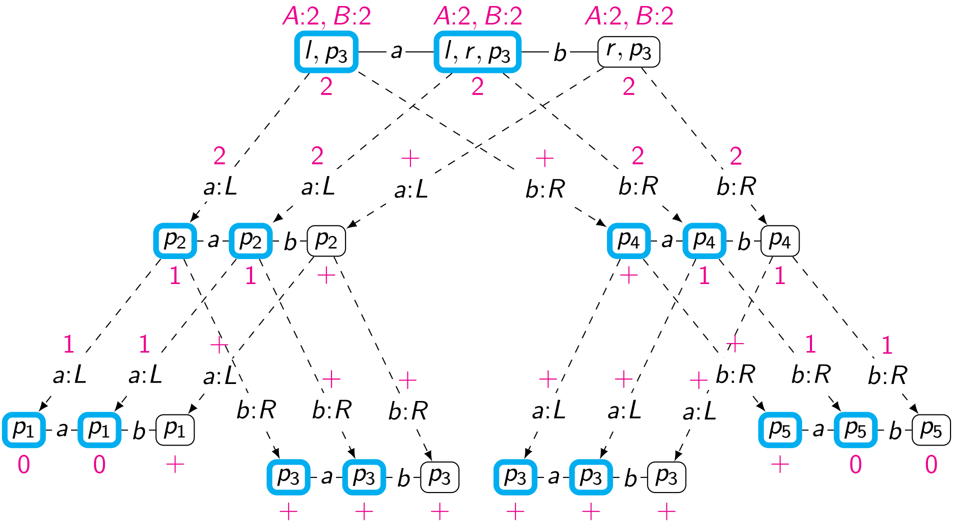

Implicit coordination with (simple) forward induction

Each agent plans from their local perspective, i.e., we close the set of designated worlds under the accessibility relation of that agent. Below the perspective of a.

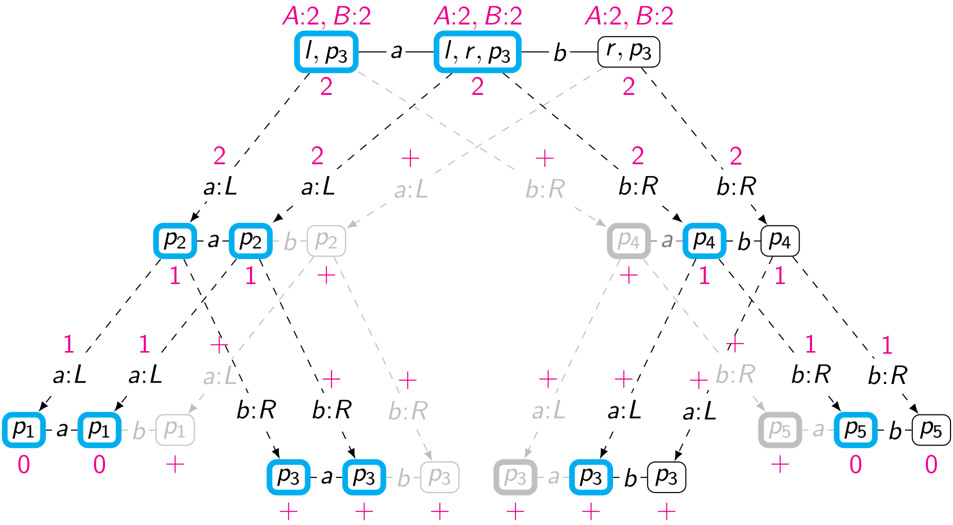

Implicit coordination with (simple) forward induction

Each agent plans from their local perspective, i.e., we close the set of designated worlds under the accessibility relation of that agent. Below the perspective of a.

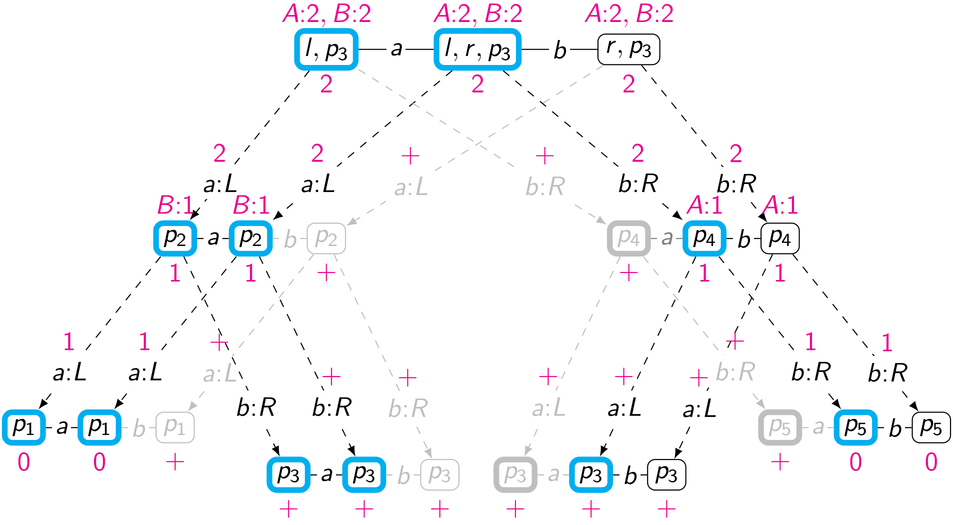

Implicit coordination with (simple) forward induction

Each agent plans from their local perspective, i.e., we close the set of designated worlds under the accessibility relation of that agent. Below the perspective of a.

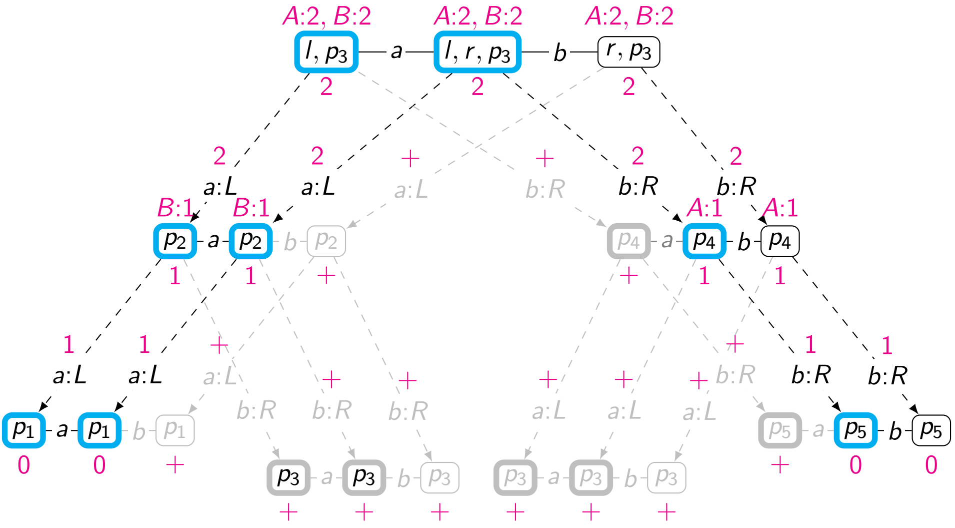

Implicit coordination with (simple) forward induction

Each agent plans from their local perspective, i.e., we close the set of designated worlds under the accessibility relation of that agent. Below the perspective of a.

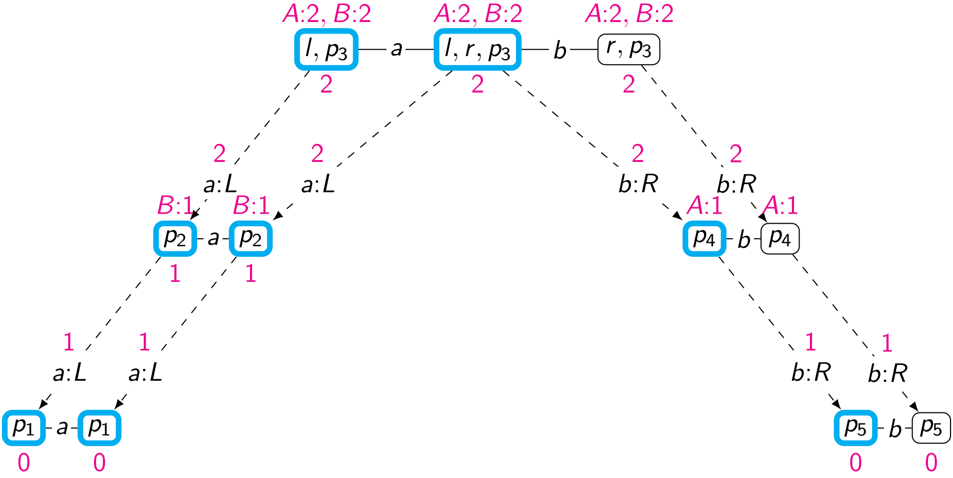

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat "realistic" solution concept for multiagent systems.

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat "realistic" solution concept for multiagent systems.

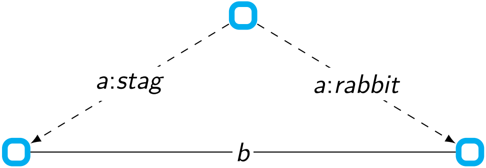

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat "realistic" solution concept for multiagent systems.

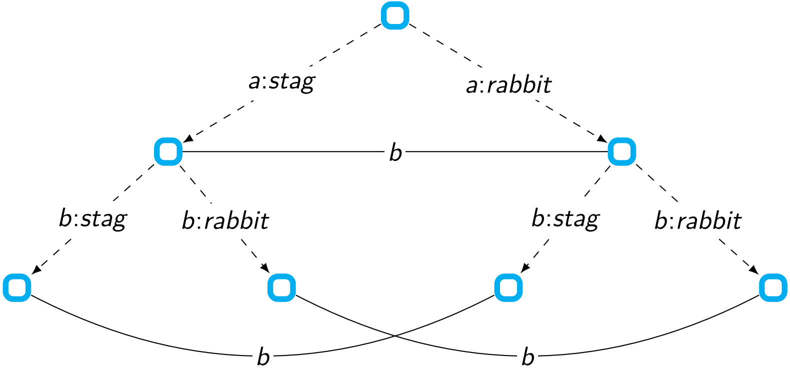

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat "realistic" solution concept for multiagent systems.

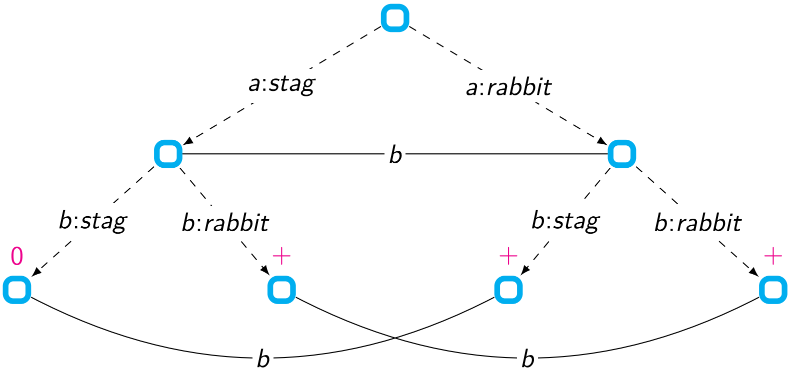

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat "realistic" solution concept for multiagent systems.

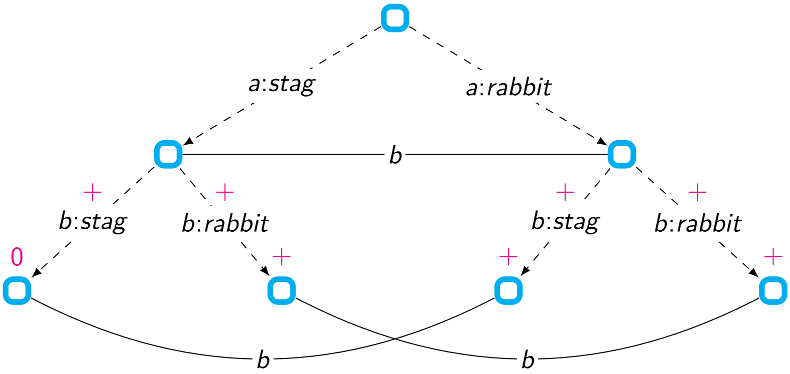

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat "realistic" solution concept for multiagent systems.

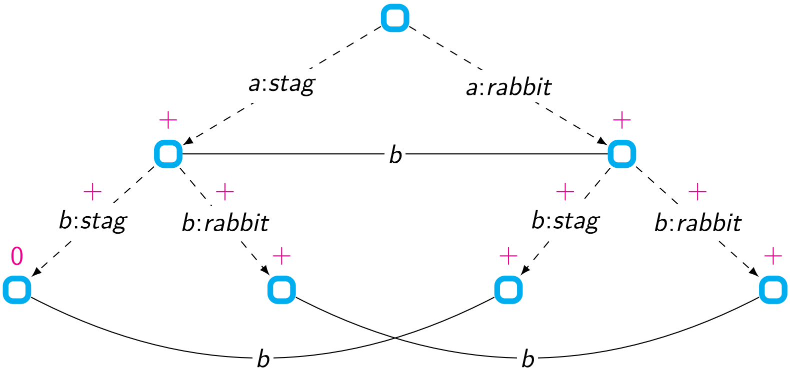

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat "realistic" solution concept for multiagent systems.

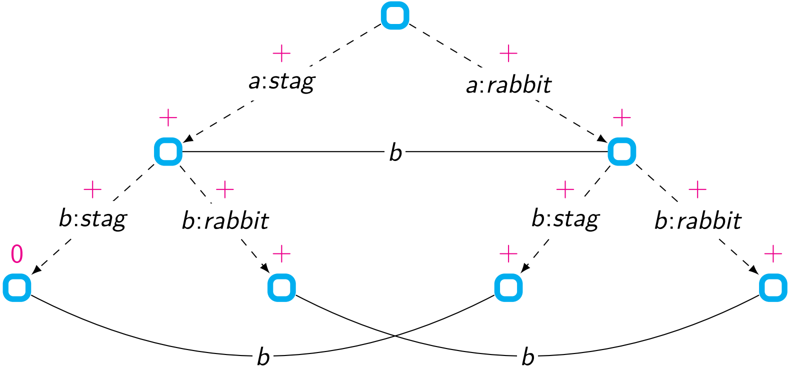

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat "realistic" solution concept for multiagent systems.

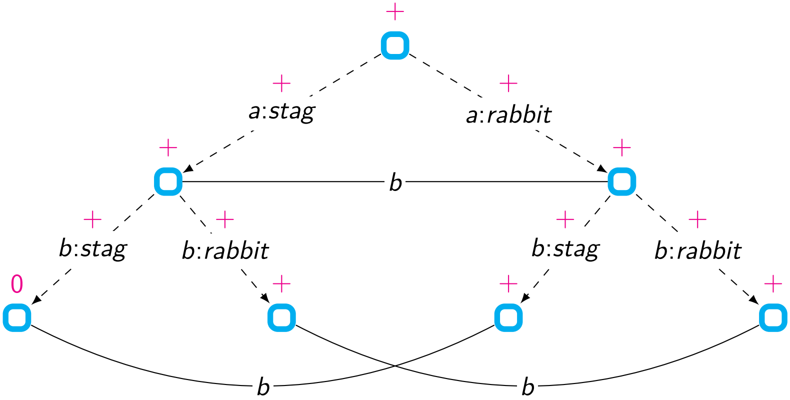

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat “realistic” solution concept for multiagent systems.

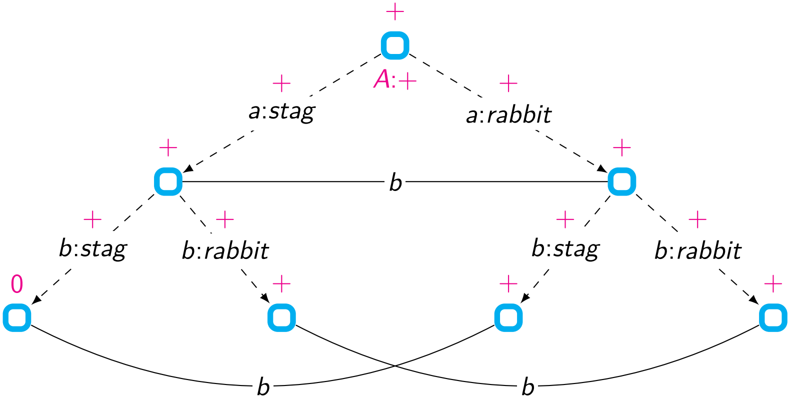

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat "realistic" solution concept for multiagent systems.

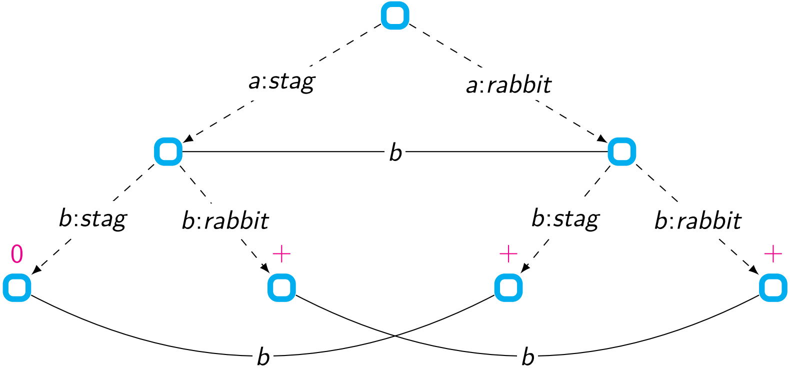

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat "realistic" solution concept for multiagent systems.

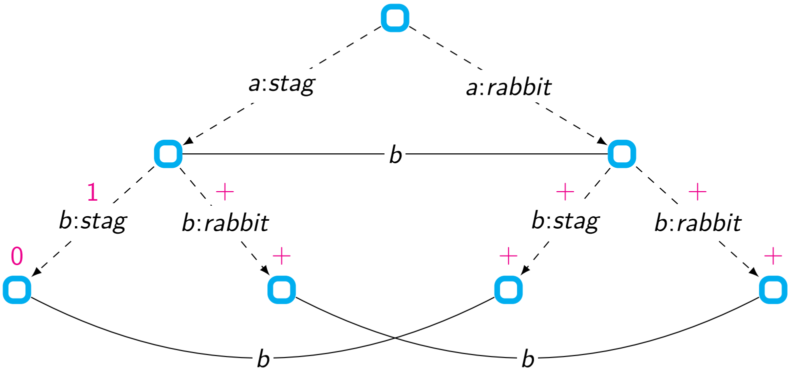

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat “realistic” solution concept for multiagent systems.

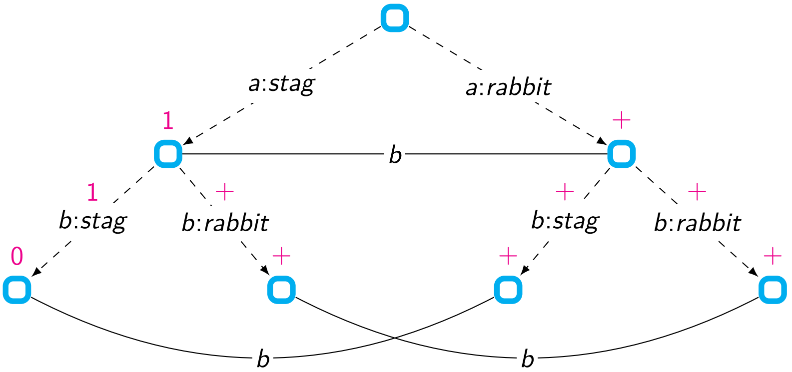

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat "realistic" solution concept for multiagent systems.

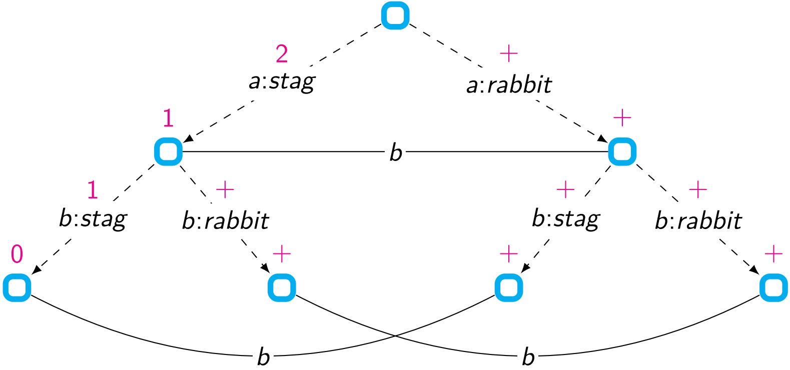

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat "realistic" solution concept for multiagent systems.

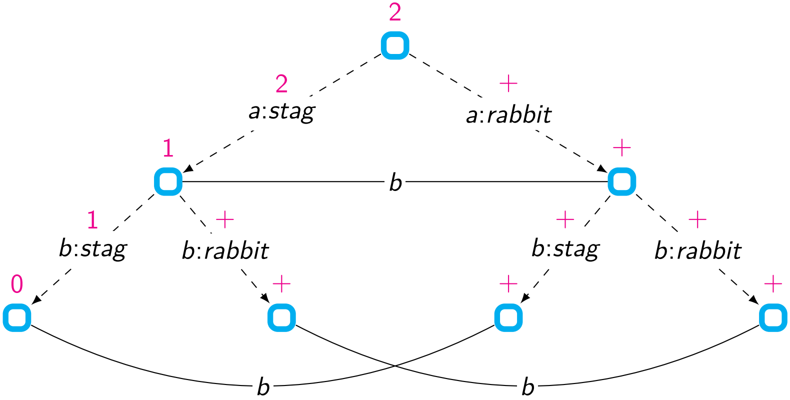

From global to local costs

Previous solution concept not strong enough to capture e.g. goal recognition. Solution: Compute first global costs, and prune indistinguishability edges. Goal of our solution concept:

- No exponential time complexity in the tree size (no quantification over all policies).
- Gives a somewhat "realistic" solution concept for multiagent systems.

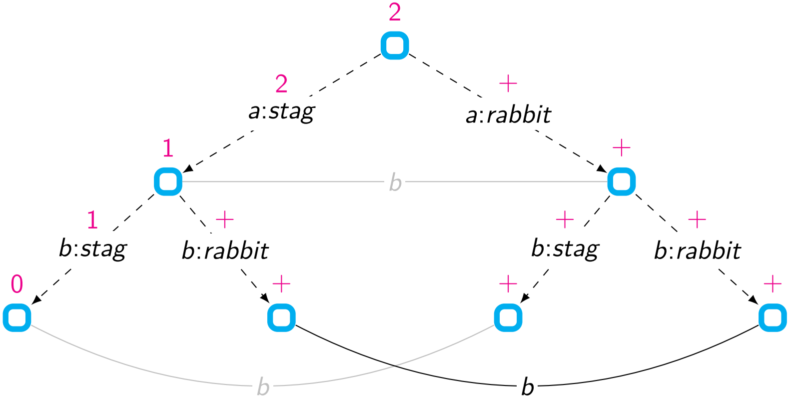

Implicit coordination in multi-agent pathfinding revisited

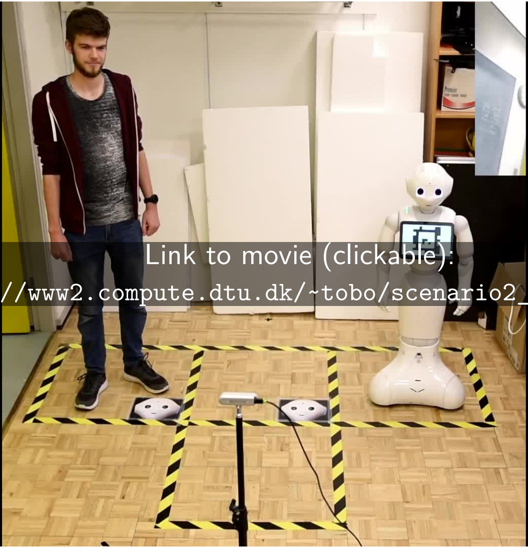

Link to movie (clickable):

http://www2.compute.dtu.dk/~tobo/scenario2_double.mp4

Towards depth-bounded reasoning in epistemic planning: k-bisimulation contractions

k-**bisimilarity**, $s \leftrightarroweq_k t$: Models s and t satisfy the back and forth conditions of bisimilarity up to distance k from the designated world(s).

Theorem [**Blackburn** et al., **2001**]. If $s \leftrightarroweq_k t$ then s and t are modally equivalent to modal depth k (agree on all formulas up to modal depth k).

k-**bisimulation contraction**, $\lfloor s \rfloor_k$: Quotient of s wrt. $\leftrightarroweq_k$.

Note: A k-bisimulation contraction is not necessarily minimal among models preserving modal equivalence to depth k. A recent revised notion of k-contraction remedies this [Bolander and Burigana, 2024].

Theorem [Bolander and Lequen, 2023]. *If $s \,\underline{\leftrightarrow}_k\, t$ then $s \otimes \alpha \,\underline{\leftrightarrow}_{k-md(\alpha)}\, t \otimes \alpha$, where $md(\alpha)$ is the maximal modal depth of any pre- or post-condition of α.*

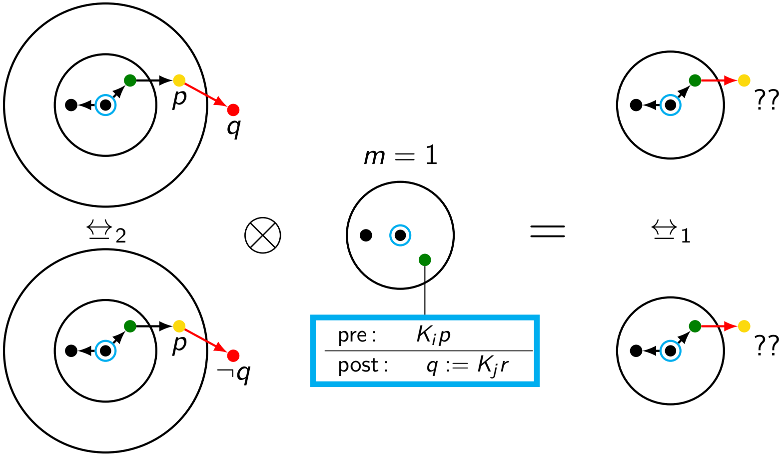

Theorem. *The plan existence problem in epistemic planning is decidable when all actions have propositional pre- and postconditions (i.e., $m = 0$). (Orig. proof by [Yu et al., 2013])*

Epistemic planning with depth-bounded reasoning

Paper currently under submission. We present a planning algorithm $\textsc{BoundedSearch}(T, b)$ with bound b (simplified version):

- "Approximate" initial state s_0 with $\lfloor s_0 \rfloor_b$. Define *bound* $b(s_0) = b$.
- For each computed product update $s \otimes \alpha$, let $b(s \otimes \alpha) = b(s) - md(\alpha)$ (cf. theorem on previous slide).
- If $b(s) < \textit{modal-depth}(\varphi_g)$, delete s.

Parameters of planning task T (we study parameterised complexity).

a: # agents	p: # propositional variables
o: modal depth of goal formula	u: maximal length of plan
c: max. modal depth of action preconditions	

Theorem (Soundness and completeness). $\textsc{BoundedSearch}(T, b)$ is sound, and if T is solvable, a solution will be found when $b \geq \mathrm{cu} + \mathrm{o}$.

Theorem (Complexity). $\textsc{BoundedSearch}(T, b)$ runs in time $|T|^{O(1)} \exp_2^{b+1} O(\mathrm{a} + \mathrm{p})$. (So fixed-parameter tractable in $\{\mathrm{a}, \mathrm{c}, \mathrm{o}, \mathrm{p}, \mathrm{u}\}$)

For any proper subset of the parameters $\{\mathrm{a}, \mathrm{c}, \mathrm{o}, \mathrm{p}, \mathrm{u}\}$, even plan verification is fixed-parameter intractable. [Bolander and Lequen, 2023]

Appendix: References I

Baltag, A., Moss, L. S. and Solecki, S. (1998).
The Logic of Public Announcements and Common Knowledge and Private Suspicions.
In Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98), (Gilboa, I., ed.), pp. 43–56, Morgan Kaufmann.

Blackburn, P., de Rijke, M. and Venema, Y. (2001).
Modal Logic, vol. 53, of Cambridge Tracts in Theoretical Computer Science.
Cambridge University Press, Cambridge, UK.

Bolander, T. and Andersen, M. B. (2011).
Epistemic Planning for Single- and Multi-Agent Systems.
Journal of Applied Non-Classical Logics *21*, 9–34.

Bolander, T. and Burigana, A. (2024).
Better Bounded Bisimulation Contractions.
In Proceedings of Advances in Modal Logic AiML '24.

Bolander, T., Charrier, T., Pinchinat, S. and Schwarzentruber, F. (2020).
DEL-based Epistemic Planning: Decidability and Complexity.
Artificial Intelligence *287*, 1–34.

Bolander, T., Dissing, L. and Herrmann, N. (2021).
DEL-based Epistemic Planning for Human-Robot Collaboration: Theory and Implementation.
In Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning (KR 2021).

Bolander, T., Engesser, T., Mattmüller, R. and Nebel, B. (2018).
Better Eager Than Lazy? How Agent Types Impact the Successfulness of Implicit Coordination.
In Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR 2018) AAAI Press.

Appendix: References II

Bolander, T. and Lequen, A. (2023).
Parameterized complexity of dynamic belief updates: A complete map.
Journal of Logic and Computation *33*, 1270–1300.

Dissing, L. and Bolander, T. (2020).
Implementing Theory of Mind on a Robot Using Dynamic Epistemic Logic.
In Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI).

Engesser, T., Bolander, T., Mattmüller, R. and Nebel, B. (2017).
Cooperative Epistemic Multi-Agent Planning for Implicit Coordination.
In Proceedings of Methods for Modalities number 243 in Electronic Proceedings in Theoretical Computer Science pp. 75–90,.

Ghallab, M., Nau, D. S. and Traverso, P. (2004).
Automated Planning: Theory and Practice.
Morgan Kaufmann.

Nebel, B., Bolander, T., Engesser, T. and Mattmüller, R. (2019).
Implicitly Coordinated Multi-Agent Path Finding under Destination Uncertainty: Success Guarantees and Computational Complexity.
Journal of Artificial Intelligence Research *64*, 497–527.

van Ditmarsch, H. and Kooi, B. (2008).
Semantic Results for Ontic and Epistemic Change.
In Logic and the Foundation of Game and Decision Theory (LOFT 7), (Bonanno, G., van der Hoek, W. and Wooldridge, M., eds), Texts in Logic and Games 3 pp. 87–117, Amsterdam University Press.

Appendix: References III

Yu, Q., Wen, X. and Liu, Y. (2013).
Multi-agent epistemic explanatory diagnosis via reasoning about actions.
In Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI) pp. 27–33,.