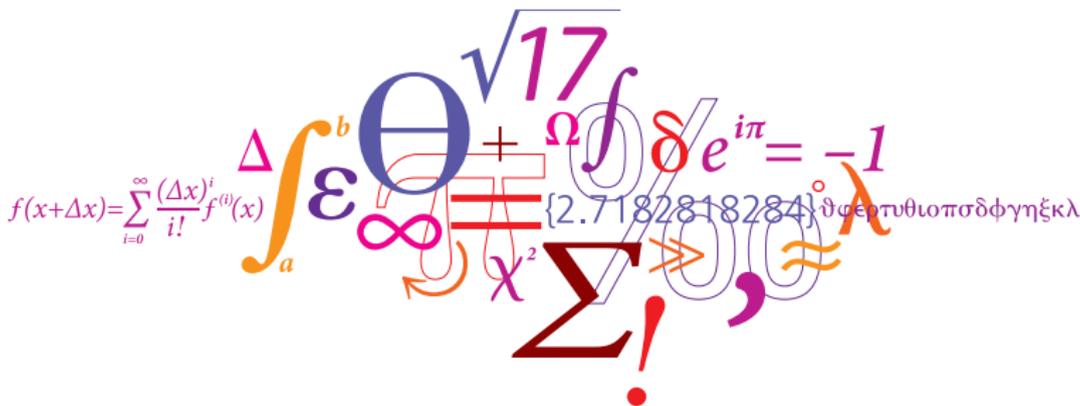


Epistemic and doxastic planning for single- and multi-agent systems

Thomas Bolander, DTU Informatics

Joint work with Mikkel Birkegaard Andersen and Martin Holm Jensen



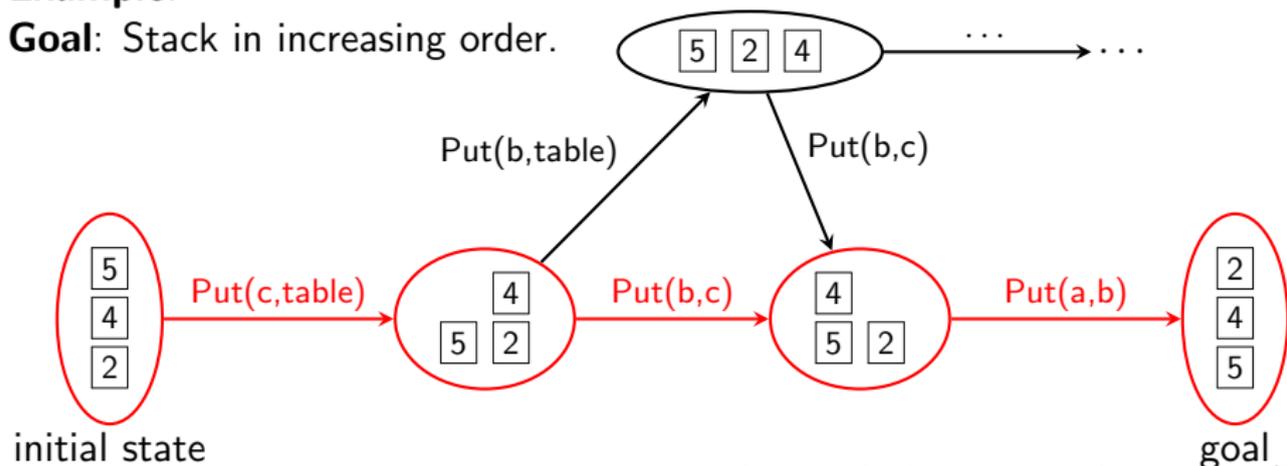
Automated planning

Automated planning (or, simply, **planning**):

- A central subfield of **artificial intelligence** (AI).
- Aims at generating **plans** (**sequences of actions**) leading to desired outcomes.
- More precisely: Given a **goal formula**, an **initial state** and some **possible actions**, an **automated planner** outputs a plan that leads from the initial state to a state satisfying the goal formula.

Example.

Goal: Stack in increasing order.



Main idea of our work

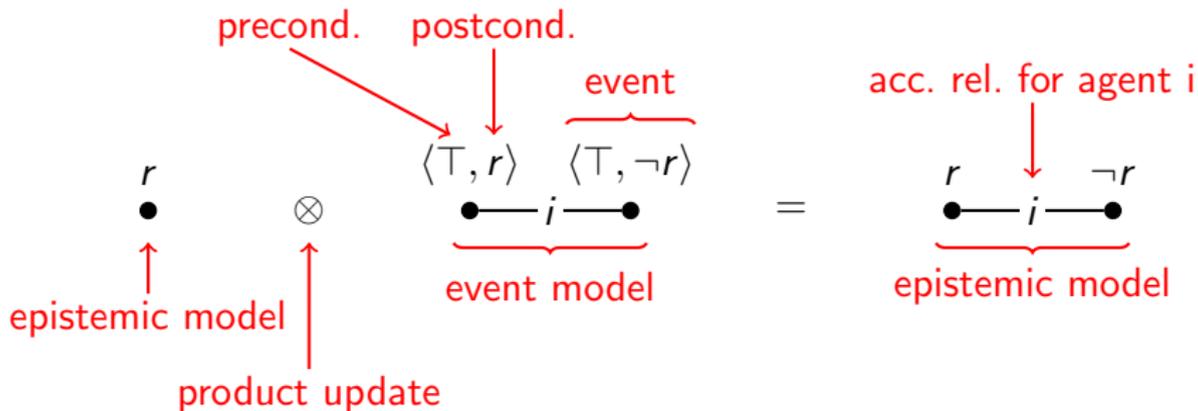
Essentially: A transition from **classical planning** based on propositional logic to planning based on **Dynamic Epistemic Logic (DEL)**.

	Classical	DEL-based
States	models of prop. logic	models of MA epist. logic
Goal formula	formula of prop. logic	formula of MA epist. logic
Actions	induced by action schemas	event models of DEL

Advantages: Generalises classical planning by allowing

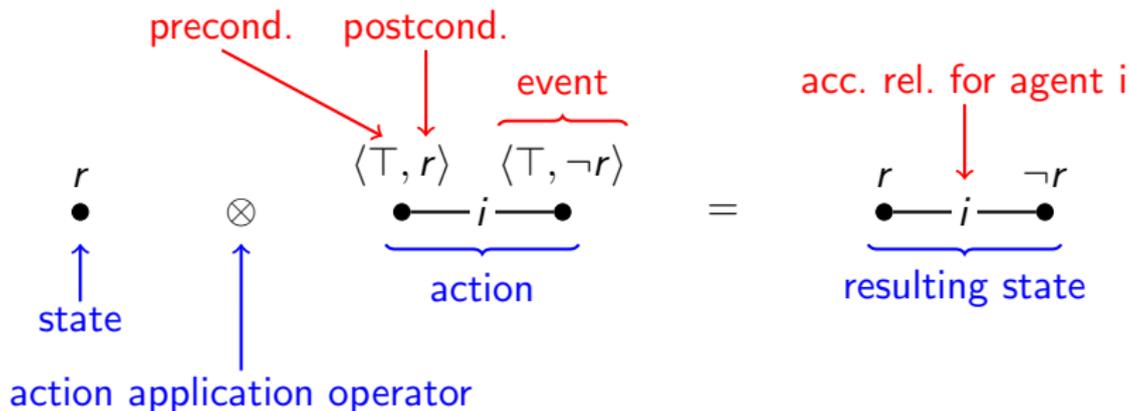
- Planning under partial observability and/or non-determinism with sensing actions.
- Planning including reasoning about other agents (essential to agent communication and collaboration).

DEL by example: Hidden coin toss



- **Epistemic models:** *Finite* multi-agent *S5* models. Reflexive edges omitted. Elements of domain called **worlds**.
- **Event models:** Both pre- and post-conditions as in [van Ditmarsch and Kooi, 2008] (allows ontic actions). Ours differ only in the definition of **postconditions**: conjunctions of propositional literals (as in classical planning). Same expressivity.
- **Product update:** As in [van Ditmarsch and Kooi, 2008].

Planning interpretation of DEL



- **States:** Epistemic models.
- **Actions:** Event models.
- **Result of applying an action in a state:** Product update of state with action.

Epistemic planning problems

Definition. An epistemic planning problem consists of:

- An **initial state** s_0 : Finite model of multi-agent epistemic logic.
- A **goal formula** ϕ_g : Formula of multi-agent epistemic logic.
- A finite set A of available **actions**: Finite event models.

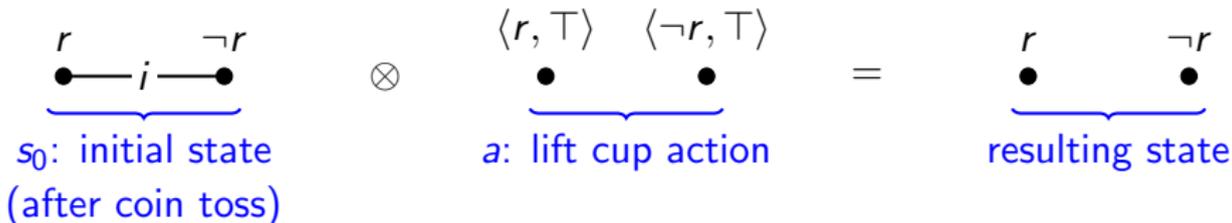
Definition. A **solution** to an epistemic planning problem is a sequence of actions $a_1, \dots, a_n \in A$ such that

$$s_0 \otimes a_1 \otimes \dots \otimes a_n \models \phi_g.$$

We then also say that a_1, \dots, a_n is a **plan** for achieving ϕ_g from s_0 .

But wait! In which world(s) is ϕ_g evaluated?...

Planning: hypothesising about the future



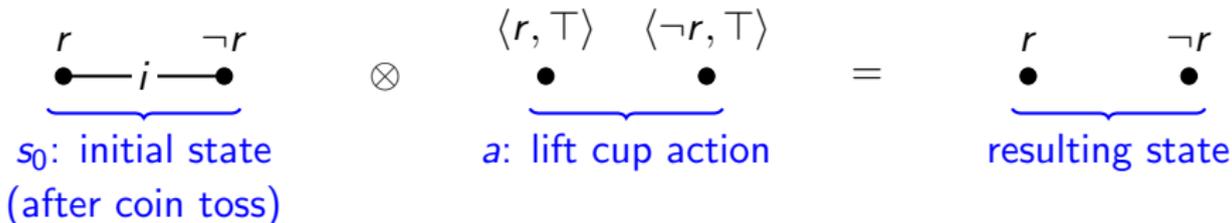
Epistemic planning (and **knowledge-based planning** in general) is about:

hypothesising about the possible outcomes of your actions.

The models (states) represent what the planning agent knows at **plan time** (*a priori*) about the knowledge it will achieve at **run time** (*a posteriori*).

In the example above: The agent will at **run time** (after the action has been performed) **come to know** whether r holds. But at **plan time** (before the action has been performed), it can't point out which of r or $\neg r$ it'll be.

Where is the goal formula evaluated?



Question: So in which world(s) in the resulting state do we evaluate a goal formula?

1st suggestion: Goal formula has to hold globally in the model.

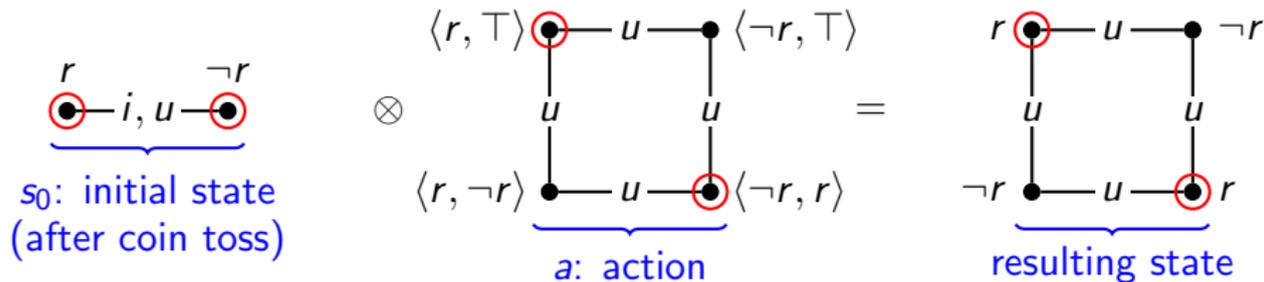
Examples. i is the planning agent.

1. $s_0 \otimes a \models K_i r \vee K_i \neg r$. Thus performing a in s_0 is a plan for achieving knowledge of **whether** r .
2. $s_0 \otimes a \not\models K_i r$. Performing a in s_0 is **not** a plan for achieving the knowledge that r .
3. $s_0 \otimes a \not\models K_i \neg r$. Performing a in s_0 is **not** a plan for achieving the knowledge that $\neg r$.

Multiple agents and designated worlds

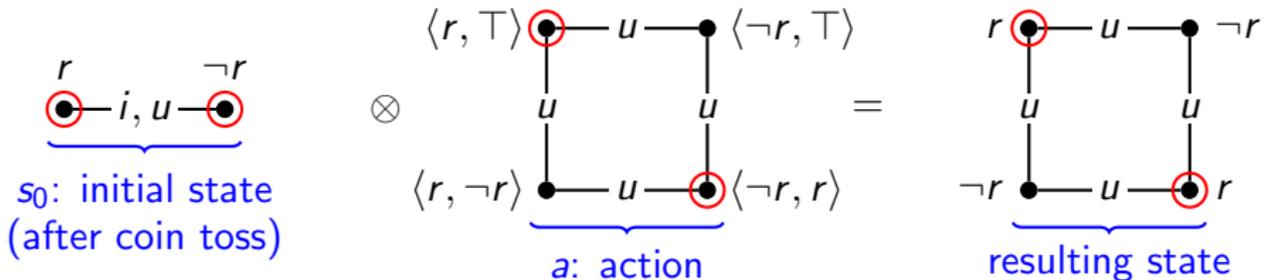
In the multi-agent case things get slightly more complicated.

Let i be I and u be you!



- Action a : I look at the coin and either flip it or not. You see the action, but not the result.
- I might choose to flip iff it's $\neg r$, thereby enforcing r .
- Afterwards I'll know r and you won't. How do I see this in the resulting state?
- **Solution**: Use **designated** worlds and events: \odot (gives **multi-pointed** epistemic models and event models).

Multiple agents and designated worlds (cont'd)



Recall question: In which world(s) in the resulting state do we evaluate a goal formula?

2nd suggestion (final): In the **designated** worlds.

Example. Applying a in s_0 achieves the goal of me knowing r but not you.

Redefinitions.

- **State:** *Multi-pointed epistemic model.*
- **Action:** *Multi-pointed event model.*
- $s \models \phi$ means ϕ holds in all the **designated** worlds of the state s .

Modelling the internal perspective

Multi-pointed models provide an **internal perspective**:

The planning agent can not always himself point out the actual world, but can point out the subset of worlds he considers possible.

A slight generalisation of the standard **external perspective**, where an actual world is always pointed out.

Main results

Theorem

Plan existence in single-agent epistemic planning is decidable.

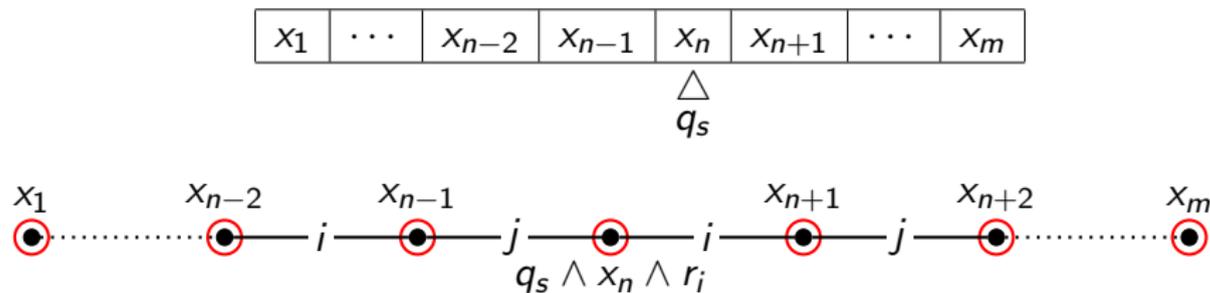
Proof idea: The number of propositional symbols is assumed to be finite. Hence there can only be finitely many distinct single-agent epistemic models (S5 models) up to bisimulation.

Theorem

Plan existence in multi-agent epistemic planning is undecidable in each of the following cases:

- There are at least 3 agents.
- There are at least 2 agents, and the epistemic language includes the common knowledge modality.
- There is at least 1 agent, and we allow arbitrary frames (not only S5).

Proof idea: Reduction to Halting problem. States (epistemic models) encode IDs of TM, actions (event models) encode transitions of TM.



Current work

We are currently extending the framework in two ways:

- **Conditional planning.** Build a language of conditional plans on top of the atomic actions (event models):

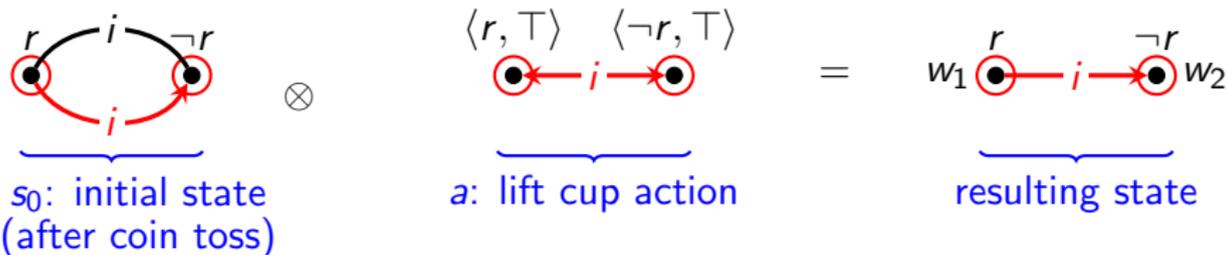
$$\pi ::= \text{skip} \mid a \mid \text{if } K_i\phi \text{ then } \pi \text{ else } \pi \mid \pi; \pi$$

where $a \in A$ (an available action), and ϕ is a formula of DEL.

- **Plausibility planning.** A transition from DEL-based planning to planning based on **epistemic plausibility models** [Baltag and Smets, 2006]. Agents can do **plausibility planning** where only the n most plausible layers of plausibility are taken into account in the planning phase. (Defeasible planning).

Plausibility planning example

Example. Tossing a biased coin.

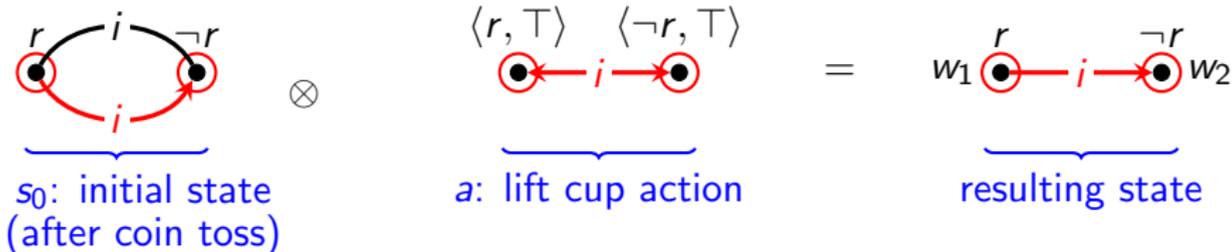


The resulting state represents the agents **plan time** knowledge about the possible outcomes of executing the plan:

- There are two possible outcomes, r and $\neg r$.
- When the plan has been executed, it will be known which it is (no epistemic link between the two).
- Currently (at plan time), it is considered most plausible that it will be $\neg r$.

Note: Our plausibility relation is the **a priori plausibility relation** (“beliefs about some virtual state”), not the **local plausibility relation** (“beliefs about the actual, current state”) [Baltag and Smets, 2006].

Example cont'd



- *Toss coin, lift cup* is a 1-strong plausibility plan for achieving $K_i \neg r$.
- *Toss coin, lift cup* is not a 2-strong plausibility plan for $K_i \neg r$.

Summing up

- Presented a **planning framework based on DEL** (with ontic actions): partial observability, non-determinism, multiple agents.
- Single agent planning is **decidable**, multi-agent planning is **undecidable**.
- The framework is currently generalised to **conditional planning** and **plausibility planning**.

References



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.