Epistemic 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**):

- One of the most central subfields of **artificial intelligence**.
- 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: Get A on B and B on C.

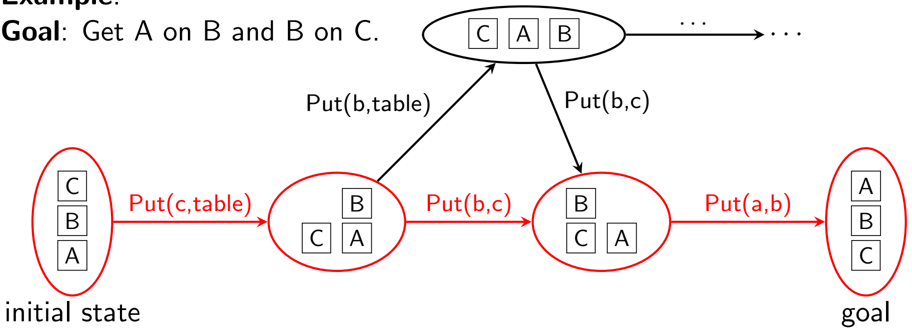

Everything in three slides

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

Example. Restack $\begin{array}{|c|}\hline C\\\hline B\\\hline A\\\hline\end{array}$ as $\begin{array}{|c|}\hline A\\\hline B\\\hline C\\\hline\end{array}$.

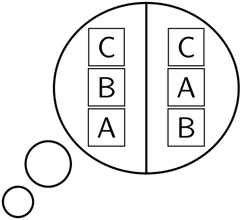

Classical planning	Planning w/belief states	Multi-agent planning
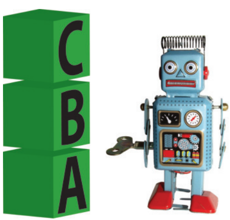	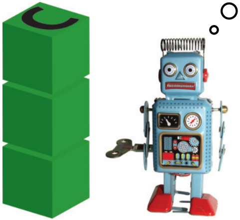	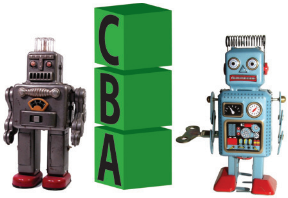

Why multi-agent planning?

Efficient planning in the presence of other agents requires the planning agent to have a model of the other agents' knowledge (a **Theory of Mind**).

Example. My wife wants to check her webmail at the Hotel. Only I know the password for the wifi, only she knows the password for her webmail account (distributed knowledge).

Case 1: I want her to check her webmail. I say: "Please check your webmail. The wifi password is xyz123."

My model of her knowledge tells me that she only needs the wifi password to be able to achieve the goal of having checked her webmail.

Case 2: She wants to check her webmail. She say: "Thomas, what's the wifi password?"

Her model of my knowledge tells her that I (might) know the wifi password. Her model of our kids' knowledge tells her that neither of them knows.

Classical planning

Definition (Planning problem). *Consists of:*

1. **States** *(including an* **initial state***): Models of propositional logic.*
2. **Goal formula***: Formula of propositional logic.*
3. *A set of possible* **actions***: mapping states to states.*

Definition (Solution). *A solution is a sequence of actions leading from the initial state to a state satisfying the goal formula.*

Planning based on DEL

Definition (Planning problem). *Consists of:*

1. **States** *(including an* **initial state** s_0*): Models of multi-agent epistemic logic.*
2. **Goal formula** ϕ_g*: Formula of multi-agent epistemic logic.*
3. *A set A of possible* **actions***: event models (action models) of DEL.*

Definition (Solution). *A solution is a sequence of actions $a_1, \ldots, a_n \in A$ s.t.*

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

Advantages of planning based on DEL

Increase in **expressive power**:

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

Natural **generalisation** of standard types of planning:

- DEL-based planning with singleton models and actions generalises **classical planning**.
- DEL-based planning with 1 agent generalises **planning with partial observability using belief state**.

DEL by example: Hidden coin toss

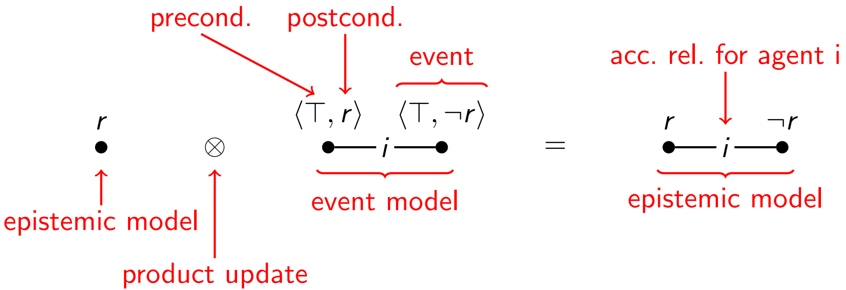

- **Epistemic models**: *Finite* multi-agent *S5* models. Reflexive edges omitted. Elements of domain called **worlds**.
- **Event models**: Both pre- and post-conditions as in [van Ditmarsch and Kooi, 2008] (allows ontic actions). Ours differ only in the definition of **postconditions**: conjunctions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

Planning interpretation of DEL

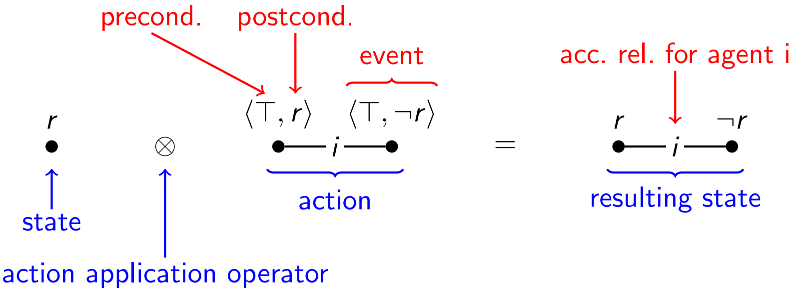

- **States**: Epistemic models.
- **Actions**: Event models.
- **Result of applying an action in a state**: Product update of state with action.

Epistemic planning problems

Definition (Epistemic planning problem). *An* **epistemic planning problem** *consists of:*

- **States** *(including an* **initial state** s_0*): Models of multi-agent epistemic logic.*
- *A* **Goal formula** ϕ_g*: Formula of multi-agent epistemic logic.*
- *A set A of possible* **actions***: Event models.*

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

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

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

Coin toss followed by sensing

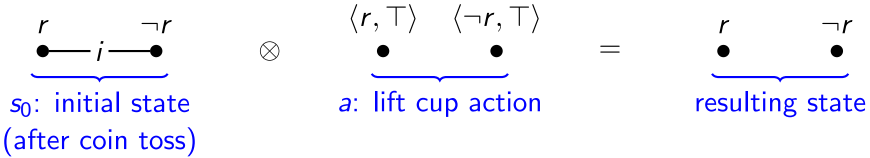

Epistemic planning (and **knowledge-based planning** in general) is about:

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.

Coin toss followed by sensing (cont'd)

$$\underbrace{\overset{r}{\bullet} \;\text{—}\; i \;\text{—}\; \overset{\neg r}{\bullet}}_{\substack{s_0\text{: initial state} \\ \text{(after coin toss)}}} \quad \otimes \quad \underbrace{\overset{\langle r, \top \rangle}{\bullet} \qquad \overset{\langle \neg r, \top \rangle}{\bullet}}_{a\text{: lift cup action}} \quad = \quad \underbrace{\overset{r}{\bullet} \qquad \overset{\neg r}{\bullet}}_{\text{resulting state}}$$

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

Answer (provisional): 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!

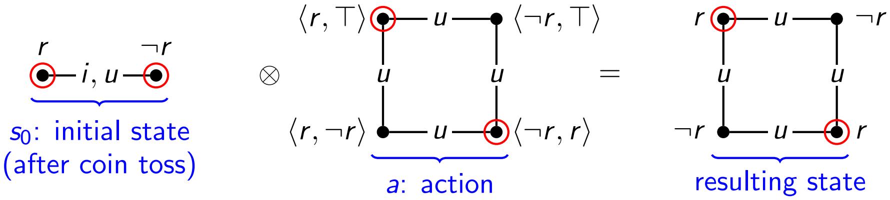

- Action a: I look at the coin and either flip it or not. You see the action, but not the result.
- I might choose to flip iff it's $\neg r$, thereby enforcing r.
- But then afterwards, I'll know r. How do I see this in the resulting state?
- **Solution**: Use **designated** worlds and events: ⊙ (gives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

Multiple agents and designated worlds (cont'd)

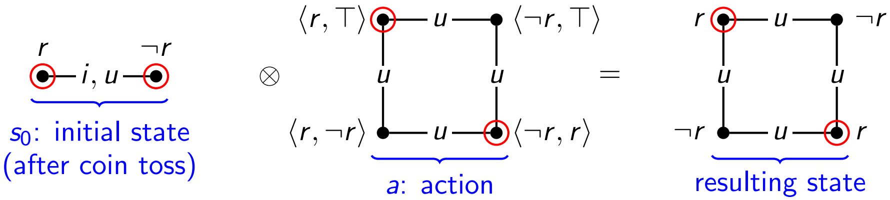

Recall question: In which world(s) in the resulting state do we evaluate a goal formula?

Answer (final): In the **designated** worlds.

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

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.*

Distinct from the standard **external perspective**, where an actual world is always pointed out.

Applicability

Consider this **plan** (sequence of actions): *Drive to bank*, *Get cash at bank*.

Drive to bank:

$\langle atH \wedge carOK, atB \wedge \neg atH \rangle$

$\langle atH \wedge \neg carOK, \top \rangle$

Get cash at bank: $\langle atB, haveC \rangle$

Driving to the bank:

$atH \wedge carOK \wedge \neg haveC$

i

$atH \wedge \neg carOK \wedge \neg haveC$

$\otimes$

$\langle atH \wedge carOK, atB \wedge \neg atH \rangle$

$\langle atH \wedge \neg carOK, \top \rangle$

$=$

$atB \wedge carOK \wedge \neg haveC$

$atH \wedge \neg carOK \wedge \neg haveC$

Applicability (cont'd)

Getting the cash:

$$\begin{array}{l} \odot\ atB \land carOK \land \neg haveC \\ \\ \odot\ atH \land \neg carOK \land \neg haveC \end{array} \quad \otimes \quad \odot \langle atB, haveC \rangle \; = \; \odot\ atB \land carOK \land haveC$$

Problem: I can now, incorrectly, conclude that after having executed *Drive to bank*, *Get cash from bank*, **I know** I have cash.

Solution: Concept of applicability.

Definition (Applicability). An action a is said to be **applicable** in a state s if:

> *for each designated world in s there is a designated event in a having its precondition satisfied in the world.*

In other words: For each world the agent considers possible, the action specifies at least one applicable event.

Redefine concept of **solution** accordingly.

Main results

Theorem

Plan existence in single-agent epistemic planning is decidable.

Proof sketch.

1. If a connected component of a state contains two worlds making the same propositions true, these worlds are bisimilar and can be collapsed into one.
2. Thus: There can only be finitely many distinct bisimulation minimal states (we are assuming there is only finitely many propositional symbols).
3. Plan existence then becomes s-t connectivity in a finite graph.

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 sketch. **Reduction to Halting problem**: Given any Turing machine M we can construct an epistemic planning problem P_M that has a solution iff M halts.

Instantaneous description (ID):

$$x_1 \cdots x_{n-2} x_{n-1} q_s x_n x_{n+1} \cdots x_m$$

encoded by:

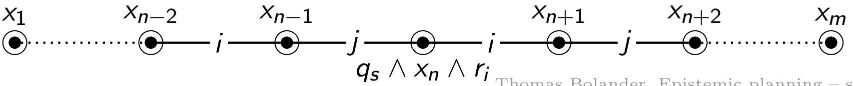

Some related work

DEL planning and some tractable cases [Löwe et al., 2011]:

- Concurrent, independent work on very similar ideas.
- **Differs by**: Having external perspective, having no ontic events, not having full generality.
- Focus on specialised **tractable cases**.

Tractable Multiagent Planning for Epistemic Goals [Hoek and Wooldridge, 2002]:

- Planning as model checking in ATEL.
- **Differs by**: Having no internal structure on actions (actions specified by transition function on states), being less expressive (but decidable).
- **Note**: Tractability is in the size of the state space, not the size of the planning problem.

Current and future work

From **epistemic models** to **plausibility models**:

- Using the framework of [Baltag and Smets, 2008].
- Different types of **applicability** gives different strengths of planning.
- Plausibility applicability gives "**defeasible planning**".
- Can deal efficiently with **exogenous events**.

Taking and giving **instructions**:

- Inspired by [Benotti, 2010].
- Planning-based inference of conversational implicature.

Reducing model sizes:

- Partial epistemic models?
- Depth-limited models?

Learning:

- Learning of facts is automatically included.
- Learning of actions?

Summing up

- Modelling the knowledge of other agents is essential to efficient planning and interaction in multi-agent settings (cf. webmail example).
- Presented a **planning framework based on DEL** (with ontic actions).
- Using **internal perspective**: The planning agent is the modeler.
- **Uniform generalisation** of several "classical" planning frameworks.
- Single agent planning is **decidable**, multi-agent planning is **undecidable**.

References

Baltag, A. and Smets, S. (2008).
A Qualitative Theory of Dynamic Interactive Belief Revision.
In Logic and the Foundations of Game and Decision Theory (LOFT7), (Bonanno, G., van der Hoek, W. and Wooldridge, M., eds), vol. 3, of Texts in Logic and Games pp. 13–60, Amsterdam University Press.

Benotti, L. (2010).
Implicature as an Interactive Process.
PhD thesis, Université Henri Poincaré, INRIA Nancy Grand Est, Francia.

Hoek, W. V. D. and Wooldridge, M. (2002).
Tractable Multiagent Planning for Epistemic Goals.
In In Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2002 pp. 1167–1174, ACM Press.

Löwe, B., Pacuit, E. and Witzel, A. (2011).
DEL planning and some tractable cases.
In LORI 2011 (to appear), (Hans van Ditmarsch, Jérôme Lang, S. J., ed.), vol. 6953, of Lecture Notes in Artificial Intelligence pp. 179–192, Springer.

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.