Cooperative Epistemic Multi-Agent Planning with Implicit Coordination

Thorsten Engesser[1], Thomas Bolander[2], Robert Mattmüller[1], Bernhard Nebel[1]

[1]University of Freiburg, [2]Technical University of Denmark

Why Epistemic Planning?

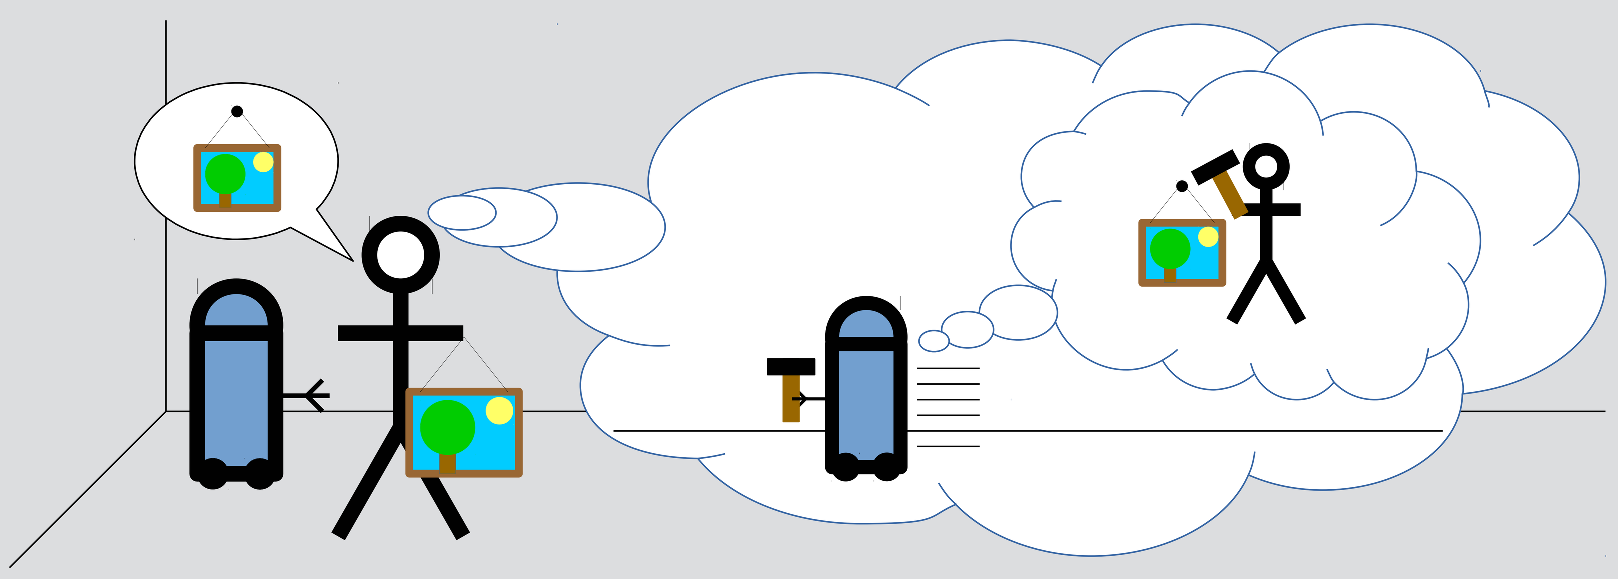

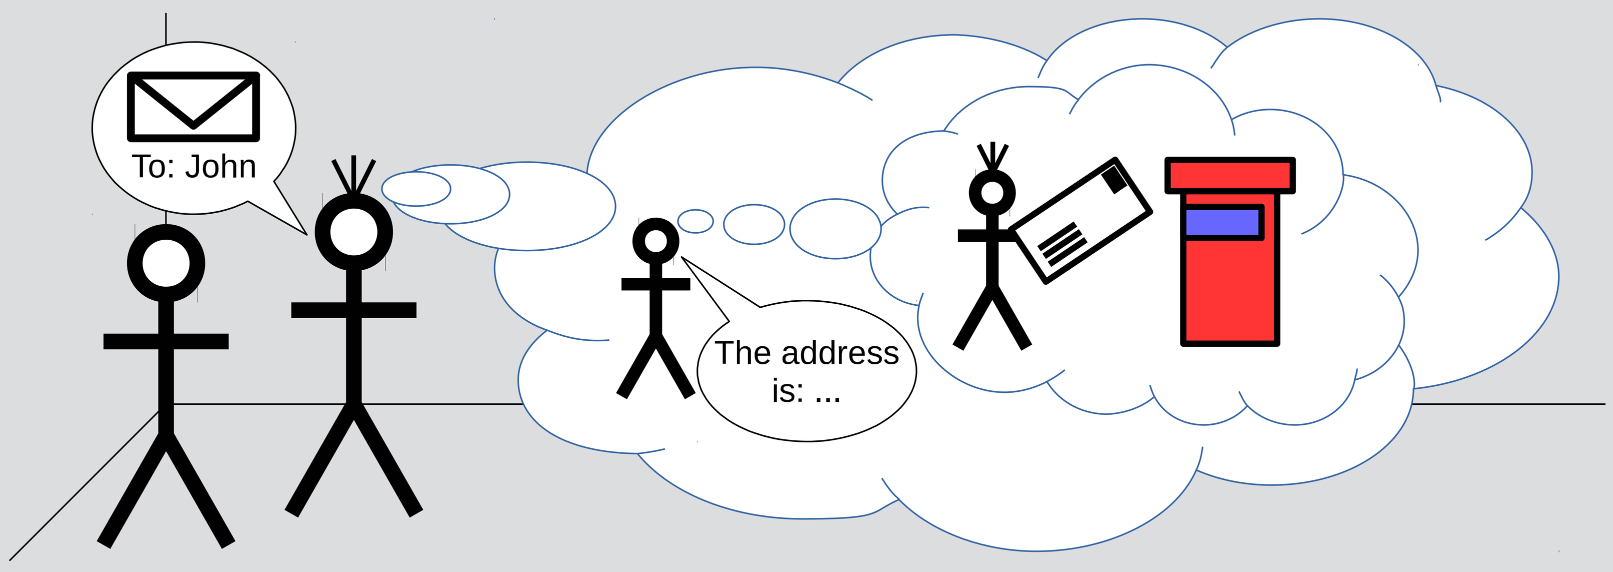

- Implicit Coordination by projecting into the other agents' state of knowledge
- Epistemic action preconditions and goals desired for certain problems

- **Collaborative, decentralized** setting: Agents **individually plan and decide**, when and how to act under **incomplete/distributed knowledge**.
- Idea: **Make observational and communicative actions part of the plan!**

Dynamic Epistemic Logic

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi \mid C\varphi \mid (\!(a)\!)\varphi$$

- Allows reasoning about knowledge and action consequences
- $K_i(\!(a)\!)\varphi$: Agent i knows a is applicable and leads to a situation where φ holds
- Interpretation over standard **S5**$_n$ **Kripke models**

$\mathcal{M} =$ $w_1 : p$ —$1,2$— $w_2 : \neg p$

- $\mathcal{M}, w_1 \models p$ and $\mathcal{M}, w_2 \models \neg p$
- $\mathcal{M}, w_1 \models \neg K_1 p \wedge \neg K_1 \neg p$

Epistemic States and Actions

- ***Epistemic States*** $(\mathcal{M}, W_d)$ correspond to the **belief states** $\{(\mathcal{M}, w) \mid w \in W_d\}$
- *Global states* $(\mathcal{M}, \{w\})$ contain **complete information** about the situation
- *Local states* (with W_d closed under R_i) match the **perspective of an agent** i
- **Shifting perspective** to the ***associated local state*** of another agent possible
- ***Epistemic Actions*** $(\mathcal{E}, E_d)$ as multi-pointed Kripke structures on possible events with additional per-event pre-/postconditions
- Successor states via the ***Product Update*** $(\mathcal{M}', W_d') = (\mathcal{M}, W_d) \otimes (\mathcal{E}, E_d)$

Example: Observational and Communicative Actions

- **Sensing action** `sense` for agent 1, checking whether p or $\neg p$:

$w_1 : p$ —$1,2,3$— $w_2 : \neg p$ $\otimes$ $e_1 : \langle p, \top\rangle$ —$2,3$— $e_2 : \langle \neg p, \top\rangle$ $=$ $w_1, e_1 : p$ —$2,3$— $w_2, e_2 : \neg p$

- From the global perspective s_0 (given w_1 is the *actual world*):

$w_1 : p$ —$1,2,3$— $w_2 : \neg p$ $\otimes$ $e_1 : \langle p, \top\rangle$ —$2,3$— $e_2 : \langle \neg p, \top\rangle$ $=$ $w_1, e_1 : p$ —$2,3$— $w_2, e_2 : \neg p$

- Agent 1 will be able to identify the actual world at run-time.
- $s_0 \models (\!(\texttt{sense})\!) K_1 p$, but $s_0 \not\models K_1 (\!(\texttt{sense})\!) p$.
- $s_0 \models K_1 (\!(\texttt{sense})\!) (K_1 p \vee K_1 \neg p)$

- **Communication action** `inform2p` for agent 1, informing agent 2 of p:

$w_1 : p$ —$2,3$— $w_2 : \neg p$ $\otimes$ $e_1 : \langle p, \top\rangle$ —3— $e_2 : \langle \neg p, \top\rangle$ $=$ $w_1, e_1 : p$ —3— $w_2, e_2 : \neg p$

- From the perspective of agent 3 (using his associated local state/action):

$w_1 : p$ —$2,3$— $w_2 : \neg p$ $\otimes$ $e_1 : \langle p, \top\rangle$ —3— $e_2 : \langle \neg p, \top\rangle$ $=$ $w_1, e_1 : p$ —3— $w_2, e_2 : \neg p$

- For agent 3, `inform2p` is indistinguishable to agent 1 informing agent 2 of $\neg p$.
- $s_0 \otimes \texttt{sense} \models K_3 (\!(\texttt{inform2p})\!) p$, but $s_0 \otimes \texttt{sense} \not\models (\!(\texttt{inform2p})\!) K_3 p$
- $s_0 \otimes \texttt{sense} \models K_3 (\!(\texttt{inform2p})\!) (K_2 p \vee K_2 \neg p)$

Cooperative Epistemic Planning Problem

A ***Cooperative Epistemic Planning Problem*** consists of

- An **initial state** s_0, an **action set** A and a **goal formula** φ (see also [1])
- An **owner function** ω, determining the acting agent $\omega(a)$ for each action a

Solution Concepts

- ***Centralized Plan*** $a_1, a_2, \ldots, a_n \in A^n$, iff (concept from [1])

$$s_0 \models (\!(a_1)\!)(\!(a_2)\!) \ldots (\!(a_n)\!)\varphi$$

- Issue: Owner of an action may **not even know** that the action **is applicable**, let alone that it makes **progress toward the goal**

- ***Implicitly-Coordinated Plan*** $a_1, a_2, \ldots, a_n \in A^n$, iff

$$s_0 \models K_{\omega(a_1)}(\!(a_1)\!) K_{\omega(a_2)}(\!(a_2)\!) \ldots K_{\omega(a_n)}(\!(a_n)\!)\varphi_g$$

The owner of the first action a_1 knows that a_1 is initially applicable and will lead to a situation where the owner of the second action a_2 knows that a_2 is applicable and will lead to a situation where... the owner of the nth action a_n knows that a_n is applicable and will lead to the goal being satisfied.

- Allows the agents to **self-coordinate implicitly** during the plan execution
- **Forward search** using product updates, **shifting perspective** in each step
- We generalized this concept also for **conditional plans**

Example: Russian Card Game [2]

- Problem: Seven cards 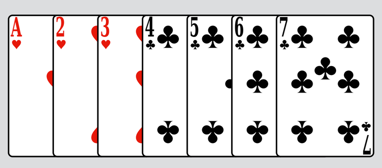 randomly dealt to Alice, Bob & Eve

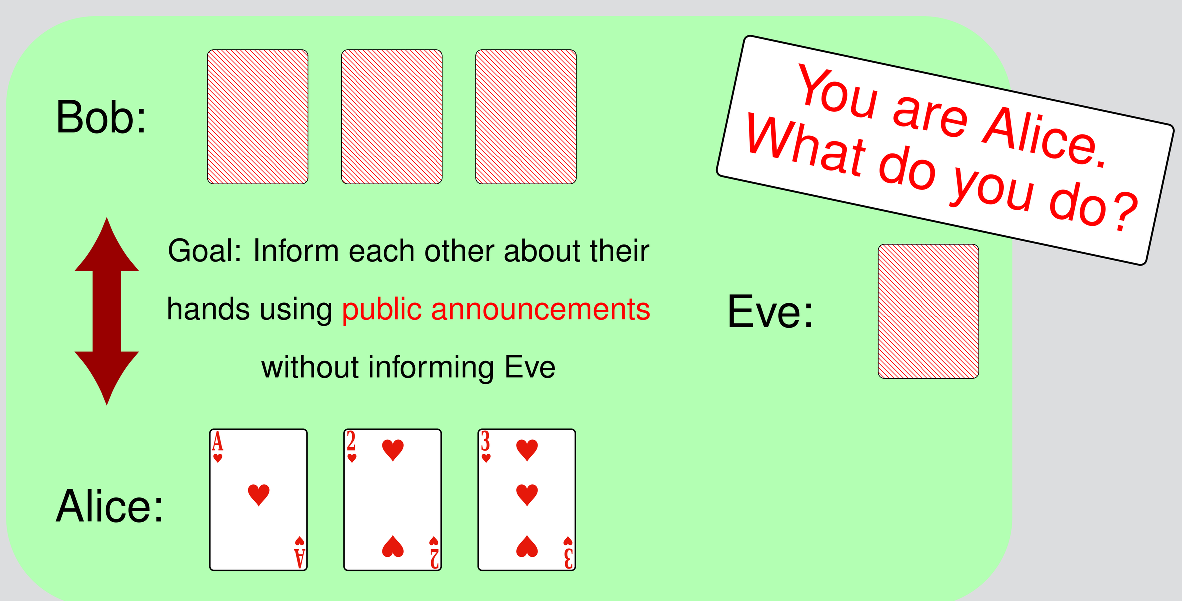

- Announce: "I have one of 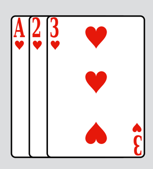 , 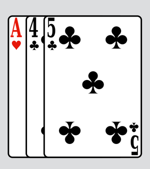 , 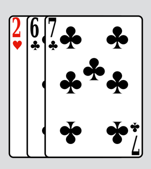 , 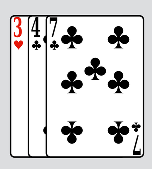 , and 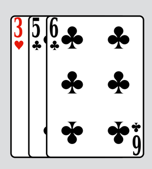 !"
- Wait for Bob announcing Eve's true card

Conclusion and Future Work

- Synthesis of epistemic plans, with **coordination as part of the plan**
- Practical Issues: **Scheduling & livelock avoidance**
- Applications, e.g., in **Human-Robot Collaboration**

[1] Thomas Bolander and Mikkel Birkegaard Andersen. Epistemic planning for single and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21(1):9–34, 2011.

[2] Hans van Ditmarsch. The Russian Cards Problem. *Studia Logica*, 75(1):31–62, 2003.