Cooperative Epistemic Multi-Agent Planning With Implicit Coordination

Albert-Ludwigs-Universität Freiburg

Robert Mattmüller[1]
(joint work with
Thorsten Engesser[1], Thomas Bolander[2], and Bernhard Nebel[1])
[1]University of Freiburg, Germany
[2]Technical University of Denmark, Copenhagen, Denmark
August 20, 2015

Classical planning on one slide:

- Given:
  - Initial world state
  - Goal description
  - Available actions
- Wanted:
  - Plan leading from initial state to goal state
- Assumptions:
  - Single agent
  - Full observability
  - Deterministic actions
  - Static and discrete environment
  - Reachability goal
  - ...

From Classical to Epistemic Planning

Classical

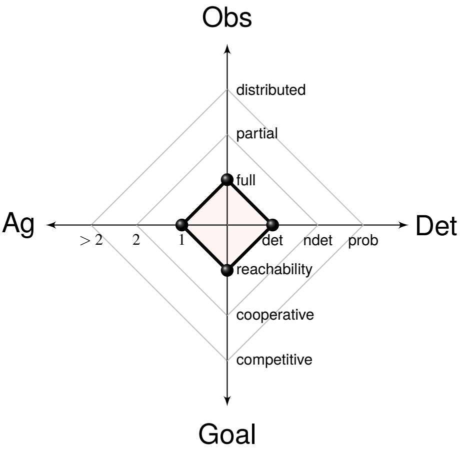

From Classical to Epistemic Planning

Classical, FOND

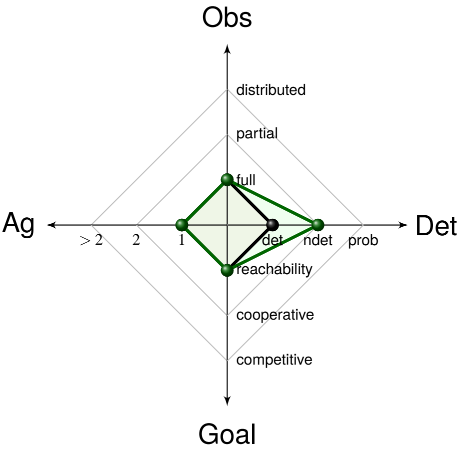

From Classical to Epistemic Planning

Classical, FOND, POND

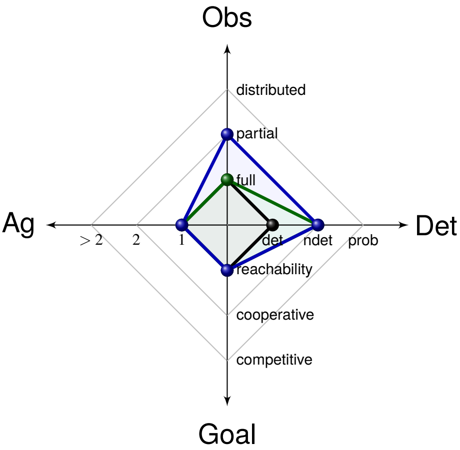

From Classical to Epistemic Planning

Classical, FOND, POND, epistemic planning, ...

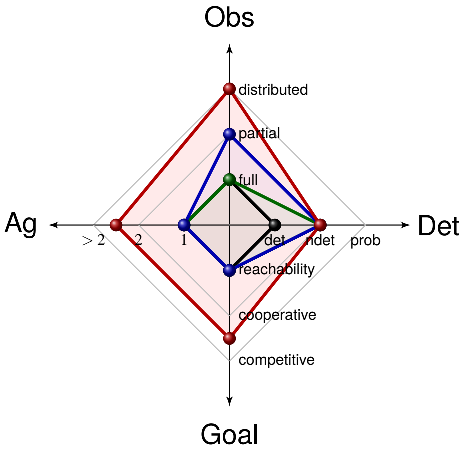

From Classical to Epistemic Planning

Classical, FOND, POND, epistemic planning, ...

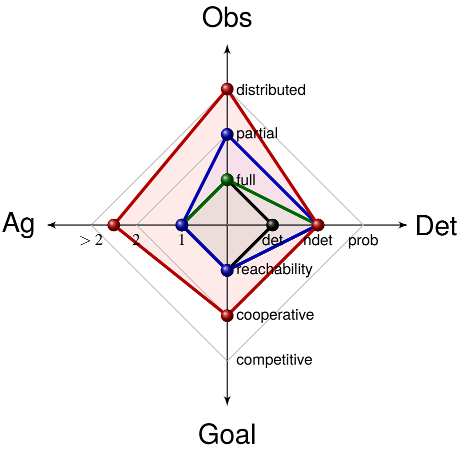

From Classical to Epistemic Planning

UNI FREIBURG

Algorithmic techniques successful in (satisficing) classical planning:

Mainly state-space search

- guided by goal-distance heuristics
- based on delete relaxation,
- abstractions, and
- landmarks,
- enhanced with pruning techniques (helpful actions, commutativity, symmetry),
- as well as invariants, causal relationships, decoupling techniques, ...

From Classical to Epistemic Planning

- Question: How well do they serve us in epistemic planning?
- Attempt at answer:
  - Start with simple state-space search.
  - Later try to add in other techniques step by step.
  - Contrast: Compilation to classical planning (cf. Muise, Belle, McIlraith, et al.).

Example: Robot Collaborating with Human

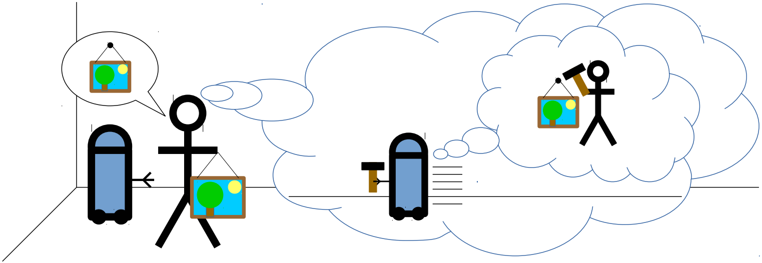

Example: Robot Collaborating with Human

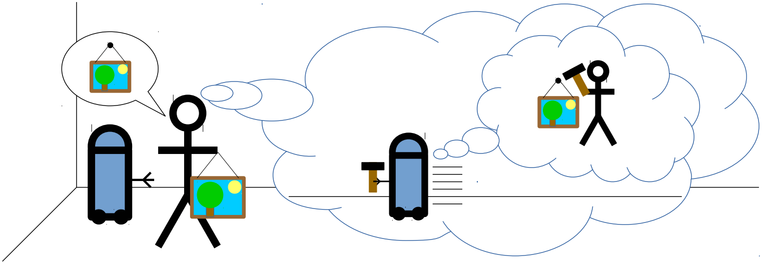

- Epistemic planning useful if we want the agents to coordinate implicitly

Cooperative Epistemic Planning

Cooperative epistemic planning:

- Task: Collaboratively reach joint goal
- Challenge: Required knowledge and capabilities distributed among agents
- Idea: Communication / coordination as part of the plan

Cooperative Epistemic Planning

Cooperative epistemic planning:

- Task: Collaboratively reach joint goal
- Challenge: Required knowledge and capabilities distributed among agents
- Idea: Communication / coordination as part of the plan

This talk:

- Cooperative epistemic planning: the problem
- Some solution concepts and their properties

Epistemic Language and Models

Reasoning about knowledge:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi \mid C\varphi$$

- $K_i\varphi$: Agent i knows φ

Epistemic Language and Models

Reasoning about knowledge and actions:

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

- $K_i\varphi$: Agent i knows φ
- $(\!(a)\!)\varphi$: a is applicable, leads to a state where φ holds

Epistemic Language and Models

Reasoning about knowledge and actions:

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

- $K_i\varphi$: Agent i knows φ
- $(\!(a)\!)\varphi$: a is applicable, leads to a state where φ holds

Epistemic formulas without $(\!(\cdot)\!)$ interpreted over standard $\mathbf{S5}_n$ Kripke models $\mathcal{M} = \langle W, R_1, \ldots, R_n, V \rangle$.

Epistemic Language and Models

Reasoning about knowledge and actions:

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

- $K_i\varphi$: Agent i knows φ
- $(\!(a)\!)\varphi$: a is applicable, leads to a state where φ holds

Epistemic formulas without $(\!(\cdot)\!)$ interpreted over standard **S5**$_n$ Kripke models $\mathcal{M} = \langle W, R_1, \ldots, R_n, V \rangle$.

$\mathcal{M} =$ 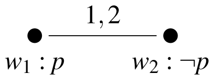

Epistemic Language and Models

Reasoning about knowledge and actions:

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

- $K_i\varphi$: Agent i knows φ
- $(\!(a)\!)\varphi$: a is applicable, leads to a state where φ holds

Epistemic formulas without $(\!(\cdot)\!)$ interpreted over standard **S5**$_n$ Kripke models $\mathcal{M} = \langle W, R_1, \dots, R_n, V\rangle$.

$\mathcal{M} =$ 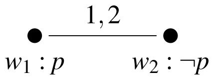

- $\mathcal{M}, w_1 \models p$ and $\mathcal{M}, w_2 \models \neg p$

Epistemic Language and Models

Reasoning about knowledge and actions:

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

- $K_i\varphi$: Agent i knows φ
- $(\!(a)\!)\varphi$: a is applicable, leads to a state where φ holds

Epistemic formulas without $(\!(\cdot)\!)$ interpreted over standard **S5**$_n$ Kripke models $\mathcal{M} = \langle W, R_1, \ldots, R_n, V\rangle$.

$\mathcal{M} =$ 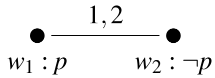

- $\mathcal{M}, w_1 \models p$ and $\mathcal{M}, w_2 \models \neg p$
- $\mathcal{M}, w_1 \models \neg K_1 p \wedge \neg K_1 \neg p$

Epistemic States

$\mathcal{M} =$

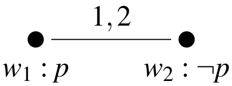

Epistemic States

$(\mathcal{M}, \{w_1\}) =$

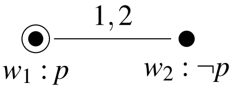

Global epistemic state $s = (\mathcal{M}, \{w\})$:

- Epistemic model $\mathcal{M}$
- World w designates actual world

Epistemic States

$(\mathcal{M}, \{w_1, w_2\}) =$

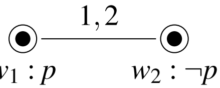

Global epistemic state $s = (\mathcal{M}, \{w\})$:

- Epistemic model $\mathcal{M}$
- World w designates actual world

Local epistemic state $s = (\mathcal{M}, W_d)$ for agent i:

- Epistemic model $\mathcal{M}$
- Worlds $W_d \subseteq W$ considered possible by agent i
- $(\mathcal{M}, W_d) \models \varphi$ iff $\mathcal{M}, w \models \varphi$ for all $w \in W_d$

Epistemic Actions

Event models $\langle E, Q_1, \ldots, Q_n, \mathrm{pre}, \mathrm{post} \rangle$ are $\mathbf{S5}_n$ Kripke frames with additional

- precondition function pre and
- postcondition function post

assigning formulas to events $e \in E$.

Epistemic Actions

Event models $\langle E, Q_1, \dots, Q_n, \mathrm{pre}, \mathrm{post} \rangle$ are $\mathbf{S5}_n$ Kripke frames with additional

- precondition function pre and
- postcondition function post

assigning formulas to events $e \in E$.

An epistemic action $(\mathcal{E}, E_d)$ consists of an event model $\mathcal{E}$ and a set $E_d \subseteq E$ of designated events.

Epistemic Actions

Event models $\langle E, Q_1, \ldots, Q_n, \mathrm{pre}, \mathrm{post} \rangle$ are $\mathbf{S5}_n$ Kripke frames with additional

- precondition function pre and
- postcondition function post

assigning formulas to events $e \in E$.

An epistemic action $(\mathcal{E}, E_d)$ consists of an event model $\mathcal{E}$ and a set $E_d \subseteq E$ of designated events.

E.g. partially observable / sensing / nondeterministic actions

Action application and successor states defined using product update.

For epistemic state s and epistemic action a, the product update $s \otimes a$ is the product of the two Kripke structures, with

- world-event pairs (w,e) eliminated if the precondition of e is violated in w and
- the valuation function updated according to the postcondition function.

Epistemic Actions

Semantics

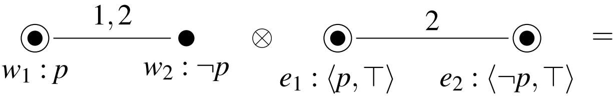

Epistemic Actions

Semantics

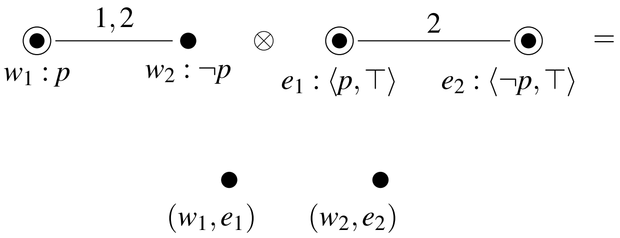

Epistemic Actions
Semantics

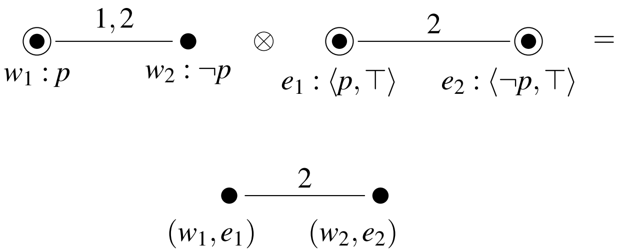

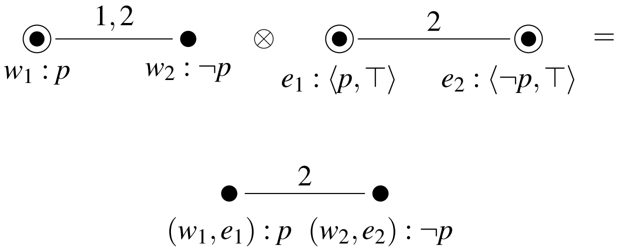

1, 2
$w_1 : p$
$w_2 : \neg p$
$\otimes$
2
$e_1 : \langle p, \top \rangle$
$e_2 : \langle \neg p, \top \rangle$
=
2
$(w_1, e_1) : p$
$(w_2, e_2) : \neg p$

Epistemic Actions

Semantics

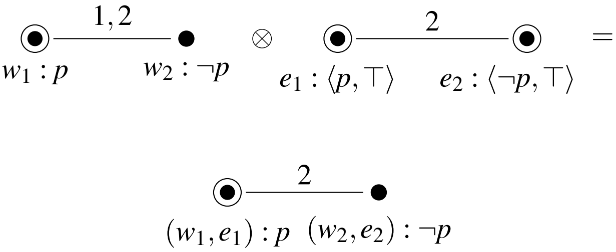

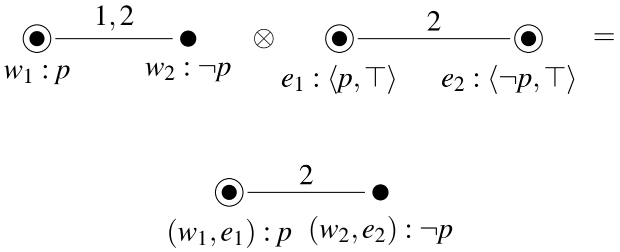

Action $(\mathcal{E}, E_d)$ is applicable in $(\mathcal{M}, W_d)$ iff for all possible situations $w \in W_d$ an outcome is defined, i.e., there is $e \in E_d$ such that $\mathcal{M}, w \models \text{pre}(e)$.

Epistemic Actions

Semantics

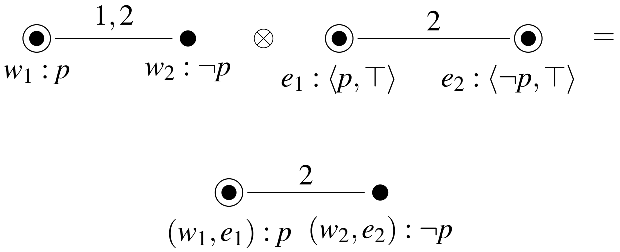

Action $(\mathcal{E}, E_d)$ is applicable in $(\mathcal{M}, W_d)$ iff for all possible situations $w \in W_d$ an outcome is defined, i.e., there is $e \in E_d$ such that $\mathcal{M}, w \models \text{pre}(e)$.

$s \models (\!(a)\!)\varphi$ iff a is applicable in s and $s \otimes a \models \varphi$

Cooperative Epistemic Planning Problem

A cooperative epistemic planning problem $\Pi = \langle s_0, A, \omega, \gamma \rangle$ consists of

- an initial epistemic state s_0,
- a finite, set A of epistemic actions,
- an owner function ω assigning agents to actions, and
- a goal formula γ such that

each action a is local for $\omega(a)$.

The action set is common knowledge among all agents.

A centralized plan $\pi = a_1, a_2, \ldots, a_n$ for Π with goal γ is a sequence of actions such that

$$s_0 \models (\!(a_1)\!)(\!(a_2)\!)\ldots(\!(a_n)\!)\gamma \quad .$$

[Bolander and Andersen, 2011]

A centralized plan $\pi = a_1, a_2, \ldots, a_n$ for Π with goal γ is a sequence of actions such that

$$s_0 \models (\!(a_1)\!)(\!(a_2)\!) \ldots (\!(a_n)\!)\gamma \quad .$$

[Bolander and Andersen, 2011]

Issue with centralized plans: Agent whose turn it is to act may not even know that the supposed action is applicable!

Sequential Planning

Centralized Plans

Example: Agents $1, 2$, propositions p, q, r, goal $\gamma = r$, initial state $s_0 = \underset{w_1:}{\odot}$, and these actions:

action	owner	pre	post	observability

Example: Agents $1, 2$, propositions p, q, r, goal $\gamma = r$, initial state $s_0 = \underset{w_1:}{\odot}$, and these actions:

action	owner	pre	post	observability
$setP$	1	$\top$	p	Indistinguishable by agent 2 at execution time
$setQ$	1	$\top$	q	

Sequential Planning

Centralized Plans

Example: Agents $1, 2$, propositions p, q, r, goal $\gamma = r$, initial state $s_0 = \underset{w_1:}{\odot}$, and these actions:

action	owner	pre	post	observability
setP	1	$\top$	p	$setP = \underset{e_1 : \langle \top, p \rangle}{\odot} \overset{2}{\text{———}} \underset{e_2 : \langle \top, q \rangle}{\bullet}$
setQ	1	$\top$	q	

Example: Agents $1, 2$, propositions p, q, r, goal $\gamma = r$, initial state $s_0 = \underset{w_1:}{\odot}$, and these actions:

action	owner	pre	post	observability
$setP$	1	$\top$	p	$setQ = \underset{e_1 : \langle \top, p \rangle}{\bullet} \overset{2}{\text{———}} \underset{e_2 : \langle \top, q \rangle}{\odot}$
$setQ$	1	$\top$	q	

Sequential Planning

Centralized Plans

Example: Agents $1, 2$, propositions p, q, r, goal $\gamma = r$, initial state $s_0 = \underset{w_1:}{\odot}$, and these actions:

action	owner	pre	post	observability
$setP$	1	$\top$	p	$setP = \underset{e_1 : \langle \top, p \rangle}{\odot} \overset{2}{\text{———}} \underset{e_2 : \langle \top, q \rangle}{\bullet}$
$setQ$	1	$\top$	q	
$setR$	2	p	r	$setR = \odot\ e_1 : \langle p, r \rangle$

Example: Agents $1, 2$, propositions p, q, r, goal $\gamma = r$, initial state $s_0 = \odot_{w_1:}$, and these actions:

action	owner	pre	post	observability
$setP$	1	$\top$	p	$setP = \odot\ e_1 : \langle \top, p \rangle \overset{2}{\text{———}} \bullet\ e_2 : \langle \top, q \rangle$
$setQ$	1	$\top$	q	
$setR$	2	p	r	$setR = \odot\ e_1 : \langle p, r \rangle$

Let $s_1 = s_0 \otimes setP$ and $s_2 = s_1 \otimes setR$. Then

state	remark

Example: Agents $1, 2$, propositions p, q, r, goal $\gamma = r$, initial state $s_0 =$ 

, and these actions:

action	owner	pre	post	observability
setP	1	$\top$	p	$setP =$ 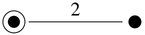
setQ	1	$\top$	q	
setR	2	p	r	$setR =$ 

Let $s_1 = s_0 \otimes setP$ and $s_2 = s_1 \otimes setR$. Then

state	remark
$s_1 =$ 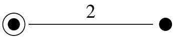	$s_1 \models p$, but $s_1 \not\models K_2 p$. 2 does not know he can apply *setR*.

Sequential Planning

Centralized Plans

Example: Agents $1,2$, propositions p,q,r, goal $\gamma = r$, initial state $s_0 = \underset{w_1:}{\odot}$, and these actions:

action	owner	pre	post	observability
$setP$	1	$\top$	p	$setP = \underset{e_1:\langle\top,p\rangle}{\odot} \overset{2}{—} \underset{e_2:\langle\top,q\rangle}{\bullet}$
$setQ$	1	$\top$	q	
$setR$	2	p	r	$setR = \odot\ e_1 : \langle p,r\rangle$

Let $s_1 = s_0 \otimes setP$ and $s_2 = s_1 \otimes setR$. Then

state	remark
$s_1 = \underset{(w_1,e_1):p}{\odot} \overset{2}{—} \underset{(w_1,e_2):q}{\bullet}$	$s_1 \models p$, but $s_1 \not\models K_2 p$. 2 does not know he can apply $setR$.
$s_2 = \odot\ (w_1,e_1,e_1) : p,r$	$s_2 \models r$. Goal is achieved.

More formally:

- $s_0 \models ((setP))((setR))r \quad \Rightarrow \quad (setP, setR)$ centralized plan

More formally:

- $s_0 \models (\!(setP)\!)(\!(setR)\!)r \quad \Rightarrow \quad (setP, setR)$ centralized plan
- $s_0 \models (\!(setP)\!)\neg K_2(\!(setR)\!)r$

More formally:

- $s_0 \models (\!(setP)\!)(\!(setR)\!)r \quad \Rightarrow \quad (setP, setR)$ centralized plan
- $s_0 \models (\!(setP)\!)\neg K_2(\!(setR)\!)r$

Motivation for different concept of plans:

If there is no central instance, then

- agents should coordinate themselves, and
- agents whose turn it is to act should know that the supposed action (a) is applicable and (b) makes progress to the goal.

Implicitly Coordinated Plans

An implicitly coordinated plan $\pi = a_1, a_2, \ldots, a_n$ for Π with goal γ is a sequence of actions such that

$$s_0 \models K_{\omega(a_1)} (\!(a_1)\!) K_{\omega(a_2)} (\!(a_2)\!) \ldots K_{\omega(a_n)} (\!(a_n)\!) \gamma \quad .$$

Implicitly Coordinated Plans

An implicitly coordinated plan $\pi = a_1, a_2, \ldots, a_n$ for Π with goal γ is a sequence of actions such that

$$s_0 \models K_{\omega(a_1)} (\!(a_1)\!) K_{\omega(a_2)} (\!(a_2)\!) \ldots K_{\omega(a_n)} (\!(a_n)\!) \gamma \quad .$$

Example: Agent 1 has to tell agent 2 that (as a consequence of his action *setP*) the proposition p is now true.

Implicitly Coordinated Plans

Example (ctd.):

action	owner	pre	post	observability
...	...	...	...	...
tellP	1	p	$\top$	fully observable Agent 2 receives message p

Example (ctd.):

action	owner	pre	post	observability
…	…	…	…	…
tellP	1	p	$\top$	$tellP = \odot\ e_1 : \langle p, \top \rangle$

Let $s_1 = s_0 \otimes setP$, $s_2 = s_1 \otimes tellP$, and $s_3 = s_2 \otimes setR$. Then

state	remark
$s_1 =$ 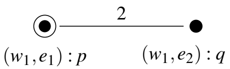	$s_1 \models p$, but $s_1 \not\models K_2 p$. 2 does not know he can apply *setR*.

Implicitly Coordinated Plans

Example (ctd.):

action	owner	pre	post	observability
...	...	...	...	...
tellP	1	p	$\top$	$tellP = \odot\ e_1 : \langle p, \top \rangle$

Let $s_1 = s_0 \otimes setP$, $s_2 = s_1 \otimes tellP$, and $s_3 = s_2 \otimes setR$. Then

state	remark
$s_1 =$ 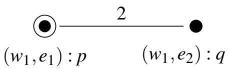	$s_1 \models p$, but $s_1 \not\models K_2 p$. 2 does not know he can apply $setR$.
$s_2 =$ 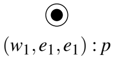	$s_2 \models p$, and $s_2 \models K_2 p$. 2 now knows that he can apply $setR$.

Motivation
Theoretical Background
Cooperative Epistemic Planning
Planning Tasks
Sequential Planning
Conditional Planning
Agent Types and Plan Executions
Conclusion

Implicitly Coordinated Plans

Example (ctd.):

action	owner	pre	post	observability
...	...	...	...	...
tellP	1	p	$\top$	$tellP = \odot\, e_1 : \langle p, \top \rangle$

Let $s_1 = s_0 \otimes setP$, $s_2 = s_1 \otimes tellP$, and $s_3 = s_2 \otimes setR$. Then

state	remark
$s_1 =$ $\odot$ $(w_1, e_1) : p$ —2— $\bullet$ $(w_1, e_2) : q$	$s_1 \models p$, but $s_1 \not\models K_2 p$. 2 does not know he can apply *setR*.
$s_2 =$ $\odot$ $(w_1, e_1, e_1) : p$	$s_2 \models p$, and $s_2 \models K_2 p$. 2 now knows that he can apply *setR*.
$s_3 = \odot\ (w_1, e_1, e_1, e_1) : p, r$	$s_3 \models r$. Goal is achieved.

Implicitly Coordinated Plans

More formally:

- $s_0 \models K_1 (\!(setP)\!) K_1 (\!(tellP)\!) K_2 (\!(setR)\!) r$
- $(setP, tellP, setR)$ is an implicitly coordinated plan for Π.

Implicitly Coordinated Plans

Searching for implicitly coordinated plans:

- Forward search in space of epistemic states using product update.
- In each step, perform a perspective shift to the agent whose action is considered, by considering its associated local state.

Implicitly Coordinated Conditional Plans

Conditional plans:

- Often, sequential plans are not sufficient to solve a task.
- One can also apply an AND-OR search to find conditional (branching) plans.

Implicitly Coordinated Conditional Plans

Conditional plans:

- Often, sequential plans are not sufficient to solve a task.
- One can also apply an AND-OR search to find conditional (branching) plans.

Remark:

- Needed, e.g., to solve Russian card games problem (initial state uncertainty necessitates branching)

Implicitly Coordinated Conditional Plans

The Russian Card Game Problem

Seven cards randomly dealt to Alice, Bob & Eve:

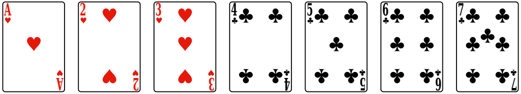

Alice:

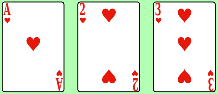

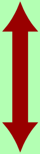

Goal: Inform each other about their hands using public announcements without informing Eve

Eve:

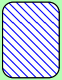

Bob:

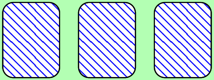

Conditional Plans

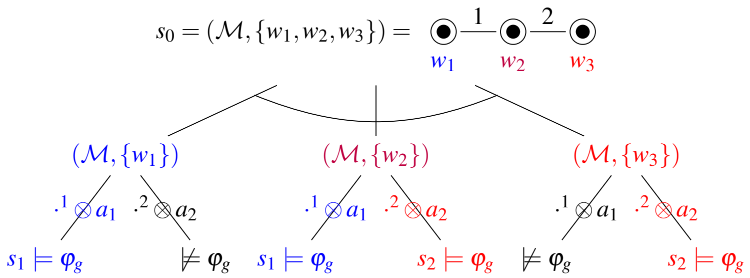

AND-OR Search

- AND: Solve an arbitrary state $(\mathcal{M}, W_d)$ by solving all global states $(\mathcal{M}, w)$ with $w \in W_d$
- OR: Solve a global state $(\mathcal{M}, w)$ by finding an agent i and an action a with $\omega(a) = i$, and solving $(\mathcal{M}, w)^i \otimes a$

Conditional Plans

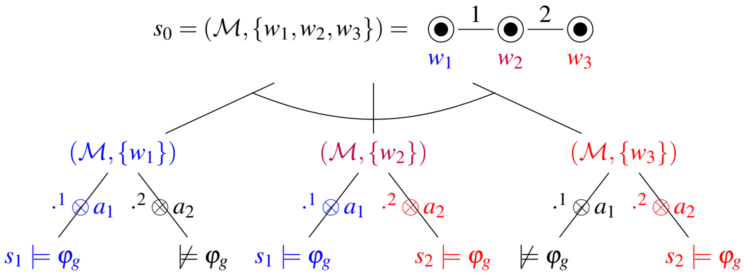

Global policy

- $\pi((\mathcal{M}, \{w_1\})) = \{a_1\}$ with $\omega(a_1) = 1$
- $\pi((\mathcal{M}, \{w_3\})) = \{a_2\}$ with $\omega(a_2) = 2$
- $\pi((\mathcal{M}, \{w_2\})) = \{a_1, a_2\}$

Agents in a Decentralized System

- Each agent plans and decides for himself when/how to act
- No imposed agent/action precedence:
  First agent that decides to act updates the system

⇒ Agents may have to replan.

Agent Types: Lazy Agents

An agent is called lazy if he choses another agents' action whenever allowed (= it is part of a strong policy).

Example problem: Who gets the door?

The goal, for Jim and John, is to go to the door and let Sarah in. Both agents are perfectly capable of doing so in one action.

What happens if both agents are lazy?

Agent Types: Naively Eager Agents

An agent is called naively eager if he choses an action owned by himself whenever allowed (= it is part of a strong policy).

Example problem: Pulling the lever (I)

The goal, for **L**ewis and **R**alph, is to pull the lever either fully to the left (-2), or to the right (2). Lewis can only pull left while Ralph can only pull right (both in steps of 1).

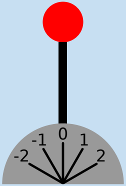

What happens if both agents are naively eager?

Agent Types: Intelligently Eager Agents

An agent is called intelligently eager if he choses an action owned by himself whenever this action is part of a strong policy of minimal depth.

Example problem: Pulling the lever (II)

Same problem as before, but **L**ewis only knows about -2 being a goal position, while **R**alph only knows about 2 being one.

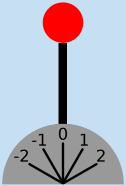

What happens if both agents are intelligently eager?

A More Interesting Problem. . .

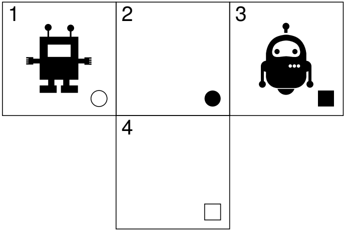

(robot icons made by Freepik and SimpleIcon from www.flaticon.com)

- Round and Square Robot have to pass each other.
- The corridor is narrow, only one agent per cell is allowed.
- Each agent is uncertain about the other's destination.

Interesting Questions

Livelocks and Deadlocks? Successful Plan Executions?

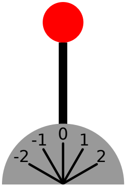

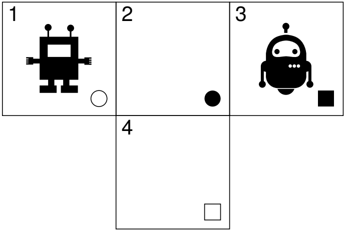

- When do we need which communicative actions?
- What about meta-reasoning?

Conclusion

Summary:

- Synthesis of epistemic plans/strategies
- Centralized vs. implicitly coordinated planning
- Communication modeled as epistemic actions
- Coordination becomes part of the plan
- Relies on the agents' ability to shift perspective