Planning Over Multi-Agent Epistemic States: A Classical Planning Approach

Christian Muise*, Vaishak Belle†, Paolo Felli*, Sheila McIlraith†
Tim Miller*, Adrian R. Pearce*, Liz Sonenberg*

*Department of Computing and Information Systems, University of Melbourne
†Department of Computer Science, University of Toronto
*{christian.muise,paolo.felli,tmiller,adrianrp,l.sonenberg}@unimelb.edu.au, †{vaishak,sheila}@cs.toronto.edu

Contribution

We formally characterize a notion of multi-agent epistemic planning, and demonstrate how to solve a rich subclass of these problems using classical planning techniques.

General Approach

1. Model how the actions update both the state of the world and the agents' belief of that state
2. Assume a single nesting of belief is a fluent and convert to a classical planning problem
3. Reformulate the problem to maintain desired properties
4. "Project" and reason *as if* we were another agent

Classical Planning

Planning Problem $\langle F, G, I, O \rangle$ where,

F: set of fluent atoms

G: set of fluents describing the goal condition

I: setting of the fluents describing the initial state

O: set of operators of the form $\langle Pre, \mathit{eff}^+, \mathit{eff}^- \rangle$

Pre: set of fluents for the precondition

eff^+: set of conditional effects that add a fluent

eff^-: set of conditional effects that delete a fluent

$(\langle \mathcal{C}^+, \mathcal{C}^- \rangle \rightarrow l)$: conditional effect that fires when $\mathcal{C}^+$ holds and $\mathcal{C}^-$ does not hold

E.g., PickupBlock

- If the agent is strong and the block is not slippery, then the agent holds the block: eff^+ contains $(\langle \{strong\}, \{slippery\} \rangle \rightarrow holding\_block)$
- If the block is big, then the agent's hand will no longer be free (i.e., we should delete the $hand\_free$ fluent): eff^- contains $(\langle \{big\_block\}, \emptyset \rangle \rightarrow hand\_free)$

Note: We distinguish between $\mathcal{C}^+/\mathcal{C}^-$ and $\mathit{eff}^+/\mathit{eff}^-$ so that our encoding is more legible

Multi-Agent Epistemic Planning

- State represents our belief about the world
- Our belief includes the nested belief of others
- Action precondition / effects can mention belief

Encoded fluents are Restricted Modal Literals (RMLs):

$$\phi ::= p \mid B_{ag}\phi \mid \neg\phi$$

- $ag \in Ag$: A particular agent
- $p \in \mathcal{P}$: An original fluent without belief
- E.g., $B_{Sue}raining$: "Sue believes it is raining"

Key Issue: How do we maintain properties on the state of the world, such as believing logical deductions or never believing contradictory information? *Additional effects*

Example: Grapevine

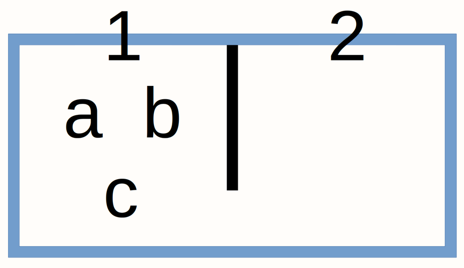

Agents each have their own secret to (possibly) share with one another and start knowing only their own secret. They can move freely between a pair of rooms, and broadcast any secret they currently believe to everyone in the room.

Actions: $share(i, secret_j, room_k)$ (i and j may differ), $move\_left(i)$, and $move\_right(i)$

Goal: Misconception – one agent believes another does not know their secret, when in fact they do.

Solution: Consider a goal of $\{B_a secret_b, \neg B_b B_a secret_b\}$

$$[\, move\_right(a), share(b, secret_b, 1), move\_right(c), share(c, secret_b, 2)\,]$$

Example Effect

E.g., Consider the action $share(c, secret_b, room_1)$
Precondition: $\{at\_c\_room_1,\ B_c secret_b\}$
One effect: If a is in the room, they will learn the secret:
$(\langle \{at\_a\_room_1\}, \emptyset \rangle \rightarrow B_a secret_b) \in \mathit{eff}^+$

Ancillary Conditional Effects

Idea: Compile new conditional effects from existing ones in order to ensure certain properties hold

Negation Removal

Delete the negation of any added RML

$$(\langle \mathcal{C}^+, \mathcal{C}^- \rangle \rightarrow l) \in \mathit{eff}^+$$
$$\Rightarrow (\langle \mathcal{C}^+, \mathcal{C}^- \rangle \rightarrow \neg l) \in \mathit{eff}^-$$

E.g., $(\langle \{at\_a\_room_1\}, \emptyset \rangle \rightarrow B_a secret_b) \in \mathit{eff}^+$
$\Rightarrow (\langle \{at\_a\_room_1\}, \emptyset \rangle \rightarrow \neg B_a secret_b) \in \mathit{eff}^-$

Uncertain Firing

If we are uncertain if an effect fires, we should be uncertain about the original outcome of the effect

E.g., $(\langle \{at\_a\_room_1\}, \emptyset \rangle \rightarrow B_a secret_b) \in \mathit{eff}^+$
$\Rightarrow (\langle \emptyset, \{\neg at\_a\_room_1\} \rangle \rightarrow \neg B_a secret_b) \in \mathit{eff}^-$

$KD45_n$ Closure

The agent's belief should remain deductively closed under the logic of $KD45_n$ (e.g., $B_i p \vdash_{KD45} \neg B_i \neg p$)

E.g., $(\langle \{at\_a\_room_1\}, \emptyset \rangle \rightarrow B_a secret_b) \in \mathit{eff}^+$
$\Rightarrow (\langle \{at\_a\_room_1\}, \emptyset \rangle \rightarrow \neg B_a \neg secret_b) \in \mathit{eff}^+$

Conditioned Mutual Awareness

Idea: Given a condition μ_i for agent i to witness an action, add the conditional effects for us to update our belief about agent i

Examples for μ_i:

- in_room_i: Agent i observes the effect if they are in the room (i.e., physically present).
- $True$: Agent i always observes the effect
- $False$: Agent i never observes the effect

E.g., $(\langle \emptyset, \{\neg at\_a\_room_1\} \rangle \rightarrow \neg B_a secret_b) \in \mathit{eff}^-$
$\Rightarrow (\langle \{\neg B_c \neg at\_a\_room_1\}, \emptyset \rangle \rightarrow \neg B_c \neg B_a secret_b) \in \mathit{eff}^+$

Agent Projection

Idea: Project our beliefs about the world and action effects to reason *as if* we were another agent.

- Can be repeated to depths greater than 1
- Represents our view of how another will reason
- Works in concert with Conditioned Mutual Awareness

$Proj(s, \vec{Ag})$: Projecting state s for agent sequence $\vec{Ag}$:

$$\begin{cases} \{\phi \mid B_i\phi \in s\} & \text{if } \vec{Ag} = [i] \\ Proj(Proj(s, [i]), \vec{Ag}') & \text{if } \vec{Ag} = [i] + \vec{Ag}' \end{cases}$$

Projecting effects for agent i works in a similar fashion, but can only be done if the effect is *uniform in* i:
$\mathcal{C}^- = \emptyset$ and all RMLs in the effect begin with B_i.

Preliminary Evaluation

Problem	$\lvert Ag \rvert$	d	$\lvert F \rvert$	$\lvert\vec{\sigma}\rvert$	Time (s) Plan	Time (s) Total
Corridor	3	1	70	5	0.01	0.11
	7	1	150	5	0.01	0.21
	3	3	2590	5	0.05	6.85
Grapevine	4	1	216	10	0.04	0.27
	3	2	774	4	0.09	1.84
	4	2	1752	4	0.70	6.61

Table: Results for various Corridor and Grapevine problems. Ag: agents, d: max depth, F: compiled fluents, and $\vec{\sigma}$: found plan. Plan time is the time spent solving the encoded problem, while Total time additionally includes the encoding and parsing phases.

Summary

- Multi-agent planning settings often require us to model the nested belief of agents
- We leveraged a tractable fragment of epistemic reasoning to maintain consistency of agents' belief
- Realized an automated planning system that deals with the nested belief in a multi-agent setting

www.haz.ca/research/pdkb
Paper · Slides · Code · Demo