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

To Be Announced! Synthesis of Epistemic Protocols

August 19, 2015

Christian Muise

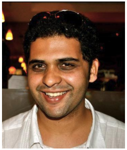

Vaishak Belle

Sheila McIlraith

Paolo Felli

Tim Miller

Adrian Pearce

Liz Sonenberg

50+ Years of Automated Planning

2015 marks the 50th anniversary of the **Shakey Project**, which resulted in such innovations as A* search and STRIPS planning.

We've seen **substantial advances** in AI Automated Planning:

- compact state and transition system encodings
- advances in heuristic search and SAT
- plans with hundreds of actions
- ... hundreds of objects
- ... 2^{1000} states
- plan computation ranging from milliseconds to a few hours

50+ Years of Automated Planning

2015 marks the 50th anniversary of the **Shakey Project**, which resulted in such innovations as A* search and STRIPS planning.

We've seen **substantial advances** in AI Automated Planning:

- compact state and transition system encodings
- advances in heuristic search and SAT
- plans with hundreds of actions
- ... hundreds of objects
- ... 2^{1000} states
- plan computation ranging from milliseconds to a few hours

Top-performing planners exploit **heuristic search**, though planning-tailored **SAT-based planners** are closing the gap.

Most of these advances have been in **classical planning**.

The Task

Synthesize plans, from the perspective of a single agent, that necessitate reasoning about the (nested) beliefs of other agents.

Desiderata:

- planning is at the belief level,
- goals and actions involve nested beliefs and respect KD45,
- agents have differing beliefs and differing capabilities, including the ability to perform communication actions and actions that are not public,
- an agent can reason *as if it were* another agent.

Example Goal: Deception

Make Bob believe that Sue believes the switch is off, when in fact Bob believes that it is on:

$$\{B_{Bob}B_{Sue}\neg switch\_on, B_{Bob}switch\_on\}$$

Example Action: Gossiping

Pre-condition for $share(Bob, secret, roomA)$ includes that Bob believes the secret:

$$B_{Bob}secret$$

Effects who perceives the gossip (and who is aware of this):

$$in(Sue, roomA) \rightarrow B_{Sue}secret,$$
$$B_{Joe}in(Sue, roomA) \land in(Joe, roomA) \rightarrow B_{Joe}B_{Sue}secret, \ldots$$

Contribution

We formally characterize a notion of Multi-agent Epistemic Planning (MEP), and demonstrate how to solve a **rich subclass** of these problems using state-of-the-art **classical planning** techniques.

Outline

1. How To Perform Epistemic Planning?
2. Background & Notation
3. Restricted Perspectival Multi-Agent Epistemic Planning
4. Projecting to Others
5. Implementation
6. Summary

Outline

1 How To Perform Epistemic Planning?

2 Background & Notation

3 Restricted Perspectival Multi-Agent Epistemic Planning

4 Projecting to Others

5 Implementation

6 Summary

How To Perform Epistemic Planning?

Option 1: Construct a Customized Planner

- (pro) Can directly handle epistemic reasoning
- (con) Rigid (forced to use one planner)

How To Perform Epistemic Planning?

Option 1: Construct a Customized Planner

- (pro) Can directly handle epistemic reasoning
- (con) Rigid (forced to use one planner)

Option 2: Cleverly Encode as Classical Planning

- (con) Must devise an encoding that ensures certain properties
  - Consistent belief
  - Logical closure
  - Etc...
- (pro) Planner agnostic – leverage every advance made!

How To Perform Epistemic Planning?

Option 1: Construct a Customized Planner

- (pro) Can directly handle epistemic reasoning
- (con) Rigid (forced to use one planner)

Option 2: Cleverly encode as Classical Planning

- (~~con~~) Must devise an encoding that ensures certain properties
  - Consistent belief
  - Logical closure
  - Etc...
- (pro) Planner agnostic – leverage every advance made!

Outline

1 How To Perform Epistemic Planning?

2 Background & Notation

3 Restricted Perspectival Multi-Agent Epistemic Planning

4 Projecting to Others

5 Implementation

6 Summary

Classical Planning

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

F: set of fluent atoms
G: set of fluents describing the goal condition
I: set of fluents describing the initial state
O: set of operators of the form $\langle Pre, \mathit{eff}^+, \mathit{eff}^- \rangle$

Operator elements

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: i.e., eff^+ contains

$$(\langle\{\mathit{strong}\}, \{\mathit{slippery}\}\rangle \rightarrow \mathit{holding\_block})$$

If the block is big, then the agent's hand will no longer be free (i.e., we should delete the *hand_free* fluent): i.e., eff^- contains

$$(\langle\{\mathit{big\_block}\}, \emptyset\rangle \rightarrow \mathit{hand\_free})$$

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

Outline

1. How To Perform Epistemic Planning?
2. Background & Notation
3. Restricted Perspectival Multi-Agent Epistemic Planning
4. Projecting to Others
5. Implementation
6. Summary

Restricted Perspectival MEP

Two restrictions:

- No disjunctive beliefs
- Bound depth of nested belief – $KD45_n$

General Approach

1. Restrict state, conditions and goals to *Proper Doxastic Knowledge Bases*

General Approach

1. Restrict state, conditions and goals to *Proper Doxastic Knowledge Bases*
2. Enforce $\mathsf{KD45}_n$ and realize belief revision by appealing to a solution to the *Ramification Problem*

General Approach

1. Restrict state, conditions and goals to *Proper Doxastic Knowledge Bases*
2. Enforce $\mathrm{KD45}_n$ and realize belief revision by appealing to a solution to the *Ramification Problem*
3. Encode conditions for mutual awareness as conditional effects

General Approach

1. Restrict state, conditions and goals to *Proper Doxastic Knowledge Bases*
2. Enforce $KD45_n$ and realize belief revision by appealing to a solution to the *Ramification Problem*
3. Encode conditions for mutual awareness as conditional effects

Result: A classical planning problem we can input to any off-the-shelf solver.

Theorem*: The approach is always sound and complete for problems that do not require disjunction

*- stated informally

1. Restrict to Proper Doxastic Knowledge Bases

Definition: Proper Doxastic Knowledge Base (PDKB)*

A set or conjunction of *Restricted Modal Literals*.

No disjunction!

1. Restrict to Proper Doxastic Knowledge Bases

Definition: Proper Doxastic Knowledge Base (PDKB)*

A set or conjunction of *Restricted Modal Literals*.

Definition: Restricted Modal Literal (RML)

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

$ag \in Ag$: A particular agent
$p \in \mathcal{P}$: An original fluent without belief

E.g., "Sue believes it is raining": $B_{Sue} raining$

1. Restrict to Proper Doxastic Knowledge Bases

Definition: Proper Doxastic Knowledge Base (PDKB)*

A set or conjunction of *Restricted Modal Literals*.

Definition: Restricted Modal Literal (RML)

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

$ag \in Ag$: A particular agent
$p \in \mathcal{P}$: An original fluent without belief

E.g., "Sue believes it is raining": $B_{Sue} raining$

Encoding

- Create a fluent for every RML (bounded depth and agents)
- States, preconditions, effects, & goals are PDKBs.

* - PEKB [Lakemeyer & Lespérance, 2012]

2. Enforce $KD45_n$ and Belief Revision via Ramifications

Treat $KD45_n$ axioms and belief revisions as indirect effects of actions.

2. Enforce $KD45_n$ and Belief Revision via Ramifications

Treat $KD45_n$ axioms and belief revisions as indirect effects of actions.

Potential Solutions to the Ramification Problem:

2. Enforce KD45$_n$ and Belief Revision via Ramifications

Treat KD45$_n$ axioms and belief revisions as indirect effects of actions.

Potential Solutions to the Ramification Problem:

Option 1: Derived Predicates

- (pro) Supported by some planners
- (con) Not a general solution

2. Enforce $KD45_n$ and Belief Revision via Ramifications

Treat $KD45_n$ axioms and belief revisions as indirect effects of actions.

Potential Solutions to the Ramification Problem:

Option 1: Derived Predicates

- (pro) Supported by some planners
- (con) Not a general solution

Option 2: Ramification Actions

- (pro) Heuristic Guidance
- (con) Not supported by planners

2. Enforce $KD45_n$ and Belief Revision via Ramifications

Treat $KD45_n$ axioms and belief revisions as indirect effects of actions.

Potential Solutions to the Ramification Problem:

Option 1: Derived Predicates

- (pro) Supported by some planners
- (con) Not a general solution

Option 2: Ramification Actions

- (pro) Heuristic Guidance
- (con) Not supported by planners

Option 3: Compile into Ancillary Conditional Effects

- (pro) Standard Approach [Pinto, 1999]
- (con) Size for non-binary constraints

Anatomy of an Ancillary Effect

If eff^+ (or eff^-) contains some conditional effect cond_1 then we will add to eff^+ (or eff^-) another effect cond_2:

$$\mathit{cond}_1 = (\langle \mathcal{C}_1^+, \mathcal{C}_1^- \rangle \rightarrow l_1)$$
$$\mathit{cond}_2 = (\langle \mathcal{C}_2^+, \mathcal{C}_2^- \rangle \rightarrow l_2)$$

$$(\langle \mathcal{C}_1^+, \mathcal{C}_1^- \rangle \rightarrow l_1) \in \mathit{eff}^+ (\text{or } \mathit{eff}^-)$$
$$\Rightarrow$$
$$(\langle \mathcal{C}_2^+, \mathcal{C}_2^- \rangle \rightarrow l_2) \in \mathit{eff}^+ (\text{or } \mathit{eff}^-)$$

Negation Removal

If you add belief, forget about its negation (belief update).

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

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

KD45 Logical Closure

The state should always be deductively closed.

$$
\begin{gathered}
(\langle \mathcal{C}^{+}, \mathcal{C}^{-} \rangle \rightarrow l) \in \mathit{eff}^{+} \\
\Rightarrow \\
\forall l' \in \mathit{Closure}(l), \;\; (\langle \mathcal{C}^{+}, \mathcal{C}^{-} \rangle \rightarrow l') \in \mathit{eff}^{+}
\end{gathered}
$$

$$
\begin{aligned}
&(\langle \{at\_a\_room_1\}, \emptyset \rangle \rightarrow B_a secret_b) \in \mathit{eff}^{+} \\
\Rightarrow\; &(\langle \{at\_a\_room_1\}, \emptyset \rangle \rightarrow \neg B_a \neg secret_b) \in \mathit{eff}^{+}
\end{aligned}
$$

KD45 Logical (un)Closure

Removing belief should also keep the state deductively closed.

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

$$(\langle \emptyset, \emptyset \rangle \rightarrow \neg B_a \mathit{secret}) \in \mathit{eff}^-$$
$$\Rightarrow (\langle \emptyset, \emptyset \rangle \rightarrow B_a \neg \mathit{secret}) \in \mathit{eff}^-$$

Uncertain Firing

Be aware of any effect that *may* have fired.

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

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

Uncertain Firing

Be aware of any effect that *may* have fired.

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

$$\begin{aligned} &(\langle \{at\_a\_room_1\}, \emptyset \rangle \rightarrow B_a secret_b) \in \mathit{eff}^+ \\ \Rightarrow\ &(\langle \emptyset, \{\neg at\_a\_room_1\} \rangle \rightarrow \neg B_a secret_b) \in \mathit{eff}^- \end{aligned}$$

Similar to the conformant planning technique:

$$p \rightarrow q \Rightarrow\ !K!p \rightarrow\ !K!q$$

3. Encode Conditioned Mutual Awareness as Conditional Effects

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.

3. Encode Conditioned Mutual Awareness as Conditional Effects

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.

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

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

Conditioned Mutual Awareness: Examples

Possible values 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

$$\mu_c = True$$

$$(\langle \emptyset, \{\neg at\_a\_room_1\}\rangle \rightarrow \neg B_a secret_b) \in eff^-$$

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

Outline

1 How To Perform Epistemic Planning?

2 Background & Notation

3 Restricted Perspectival Multi-Agent Epistemic Planning

4 Projecting to Others

5 Implementation

6 Summary

Projecting to Others

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

Projecting to Others

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

Example: *"Does Sue believe that Bob believes that Sue believes ..."*

Sue is the root agent. [Bob, Sue] is the *virtual agent* the root agent will reason about. We want Sue reasoning as if she was Bob reasoning as if he was Sue.

Projecting a PDKB state

$Proj(s, \vec{Ag})$: 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 to Others

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

Example: *"Does Sue believe that Bob believes that Sue believes ..."*

Sue is the root agent. [Bob, Sue] is the *virtual agent* the root agent will reason about. We want Sue reasoning as if she was Bob reasoning as if he was Sue.

Projecting a PDKB state

Proj

Projecting an Action

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.

Outline

1 How To Perform Epistemic Planning?

2 Background & Notation

3 Restricted Perspectival Multi-Agent Epistemic Planning

4 Projecting to Others

5 Implementation

6 Summary

Example Action Encoding: Sharing a Secret

```
(:action share_bob_secret3_room2

  :derive-condition   at_?_room2

  :precondition (and
                  (at_bob_room2)
                  [bob](secret3))

  :effect       (and
                  (when (at_sue_room2)
                        [sue](secret3))
                  (when (at_ann_room2)
                        [ann](secret3))))
```

Preliminary Evaluation

- Restricted Perspectival MEP problems specified in custom format and compiled for input to classical planners.
- Exploits the SIW^{+} and BFS_{f} planners [Lipovetzky & Geffner, 2012] (LAPKT planning library [Ramirez et al., 2015]).
- Simulates execution of an action sequence or generates a plan.
- Tested on several domains, including:
  - *Thief* domain model verified and pre-existing queries trivially solved (e.g., [Löwe, Pacuit, and Witzel, 2011])
  - *Corridor* domain [Kominis & Geffner, 2015]
  - *Grapevine* domain (combination of *Corridor* and *Gossip*)
  - Ask Christian about others!
- Time to generate plans in these more challenging domains remains small. Encoding and parsing times dominates. The max depth of nested belief, d, produces the bottleneck since the number of newly introduced fluents is exponential in d.

Corridor Domain [Kominis & Geffner, 2015]

Let a, b, and c be three agents in a corridor of four rooms (p_1, p_2, p_3, p_4) from left to right. The agents can move from a room to a contiguous room, and when agent i communicates some information, all the agents that are in the same room or a contiguous room, will hear what was said. E.g., if i expresses his knowledge about q in room p_3, all agents in rooms p_2, p_3 and p_4 will come to know it. a, b and c are initially in rooms p_1, p_2 and p_3, respectively, and a has to find out the truth value of a proposition q and let c know without b learning it.

Grapevine Domain (*Corridor* + *Gossip*)

There are two rooms with all agents starting in the first room. Every agent believes their own secret to begin with, and the agents can either move between rooms or broadcast a secret they believe. Movement is always observed by all, but the sharing of a secret is only observed by those in the same room. Goals include private communication and also misconception (e.g., $\{B_a s_b, B_b \neg B_a s_b\}$)

Experiments

Problem	$\lvert Ag \rvert$	d	$\lvert F \rvert$	$\lvert \vec{o} \rvert$	**Time (s)** **Plan**	**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

Ag: The number of agents included
d: Maximum depth of nested belief
F: Number of compiled fluents
$\vec{o}$: Computed plan

Demo Site: pdkb.haz.ca

Multi-agent Epistemic Planning With Proper Doxastic Knowledge Bases

Home | Demo | Examples | Paper | Code | Help

Agents: bob,sue,mike

Max Depth: 2

RML Editor: Enter Pi's, Bi's, and (!)p's — Delete

Positive Condition +

Negative Condition +

Condition Effect *light_on*

Mutual Awareness *in_room_?*

Effect Type
- (•) Positive
- ( ) Negative

Compile Effects

Positive Effects

ID	Pos Cond	Neg Cond	Effect	Source
89631	$\Box_{mike}\ in\_room\_mike$ $\Box_{mike}\Box_{sue}\ in\_room\_sue$		$\Diamond_{mike}\Diamond_{sue}\ light\_on$	(40604) Closure
11672	$\Box_{bob}\ in\_room\_bob$ $\Diamond_{bob}\Box_{mike}\ in\_room\_mike$		$\Diamond_{bob}\Diamond_{mike}\ light\_on$	(63000) Commonly Known
49907	$\Box_{mike}\ in\_room\_mike$ $\Diamond_{mike}\Box_{sue}\ in\_room\_sue$		$\Diamond_{mike}\Diamond_{sue}\ light\_on$	(81220) Commonly Known
48577	$\Box_{bob}\ in\_room\_bob$		$\Diamond_{bob}\ light\_on$	(42873) Closure

Negative Effects

ID	Pos Cond	Neg Cond	Effect	Source
49988	$\Box_{sue}\ in\_room\_sue$		$\Diamond_{sue}\neg\ light\_on$	(61097) Remove Negations
12292	$\Box_{mike}\Box_{bob}\ in\_room\_bob$ $\Box_{mike}\ in\_room\_mike$		$\Box_{mike}\Diamond_{bob}\neg\ light\_on$	(67269) Remove Negations
68604	$\Box_{bob}\ in\_room\_bob$ $\Box_{bob}\Box_{sue}\ in\_room\_sue$		$\Box_{bob}\Box_{sue}\neg\ light\_on$	(34766) Remove

Demo Site: pdkb.haz.ca

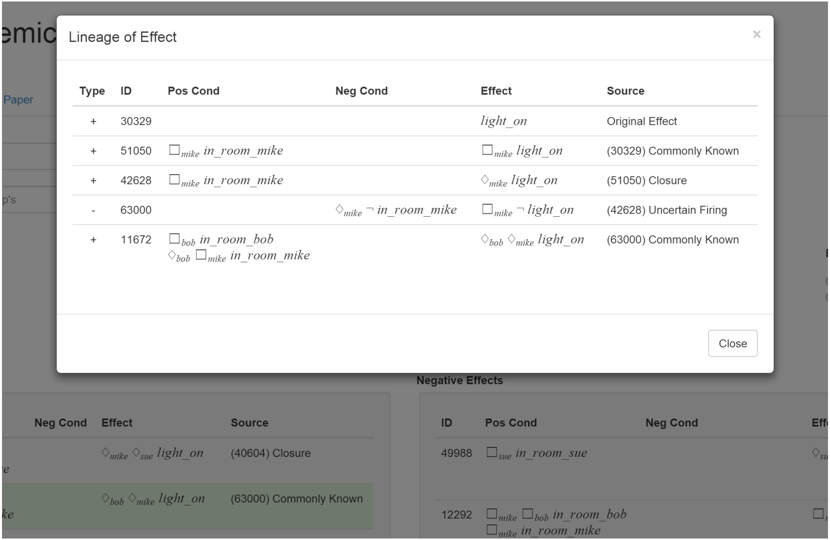

Lineage of Effect

Type	ID	Pos Cond	Neg Cond	Effect	Source
+	30329			$light\_on$	Original Effect
+	51050	$\Box_{mike}\ in\_room\_mike$		$\Box_{mike}\ light\_on$	(30329) Commonly Known
+	42628	$\Box_{mike}\ in\_room\_mike$		$\Diamond_{mike}\ light\_on$	(51050) Closure
-	63000		$\Diamond_{mike} \neg in\_room\_mike$	$\Box_{mike} \neg light\_on$	(42628) Uncertain Firing
+	11672	$\Box_{bob}\ in\_room\_bob$ $\Diamond_{bob} \Box_{mike}\ in\_room\_mike$		$\Diamond_{bob} \Diamond_{mike}\ light\_on$	(63000) Commonly Known

Close

Outline

1. How To Perform Epistemic Planning?
2. Background & Notation
3. Restricted Perspectival Multi-Agent Epistemic Planning
4. Projecting to Others
5. Implementation
6. Summary

Conclusion

- 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 effectively with the nested belief in a multi-agent setting

Future Work

Lots!

Any Questions?

Demo, benchmarks, code, and slides available at:
`http://www.haz.ca/research/pdkb/`

Some References and Follow-on Work I

P. Felli, T. Miller, C. Muise, A. R. Pearce, and L. Sonenberg.
Computing social behaviours using agent models.
In International Joint Conference on Artificial Intelligence, IJCAI 2015, 2015.

F. Kominis and H. Geffner.
Beliefs in multiagent planning: From one agent to many.
In Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling, ICAPS 2015, Jerusalem, Israel, June 7-11, 2015., 147–155, 2015.

G. Lakemeyer and Y. Lespérance.
Efficient reasoning in multiagent epistemic logics.
In ECAI 2012 - 20th European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS-2012) System Demonstrations Track, Montpellier, France, August 27-31 , 2012, 498–503, 2012.

N. Lipovetzky and H. Geffner.
Width and serialization of classical planning problems.
In 20th European Conference on Artificial Intelligence, 540–545, 2012.

B. Löwe, E. Pacuit, and A. Witzel.
DEL planning and some tractable cases.
In Logic, Rationality, and Interaction - Third International Workshop, LORI 2011, Guangzhou, China, October 10-13, 2011. Proceedings, 179–192, 2011.

C. Muise, V. Belle, P. Felli, S. A. McIlraith, T. Miller, A. R. Pearce, and L. Sonenberg.
Planning over multi-agent epistemic states: A classical planning approach.
In The 29th AAAI Conference on Artificial Intelligence, 2015.

Some References and Follow-on Work II

C. Muise, V. Belle, P. Felli, S. A. McIlraith, T. Miller, A. R. Pearce, and L. Sonenberg.
Planning over multi-agent epistemic states: A classical planning approach (amended version).
In Workshop on Distributed and Multi-Agent Planning (DMAP'15), 2015.

C. Muise, F. Dignum, P. Felli, T. Miller, A. R. Pearce, and L. Sonenberg.
Towards team formation via automated planning.
In International Workshop on Coordination, Organisation, Institutions and Norms in Multi-Agent Systems, 2015.

C. Muise, P. Felli, T. Miller, A. R. Pearce, and L. Sonenberg.
Leveraging fond planning technology to solve multi-agent planning problems.
In Workshop on Distributed and Multi-Agent Planning (DMAP'15), 2015.

C. Muise, T. Miller, P. Felli, A. R. Pearce, and L. Sonenberg.
Efficient reasoning with consistent proper epistemic knowledge bases.
In The International Conference on Autonomous Agents and Multiagent Systems, 2015.

J. Pinto.
Compiling ramification constraints into effect axioms.
Computational Intelligence, 15:280–307, 1999.

M. Ramirez, N. Lipovetzky, and C. Muise.
Lightweight Automated Planning ToolKiT.
`http://lapkt.org/`, 2015.
Accessed: 2015-01-17.