Multi-Agent Epistemic Planning with Inconsistent Beliefs, Trust and Lies

Francesco Fabiano[1], Alessandro Burigana[2], Agostino Dovier[1], Enrico Pontelli[3]

[1]University of Udine - Udine, Italy [2]Free University of Bozen-Bolzano - Bolzano, Italy

[3]New Mexico State University - Las Cruces, NM, U.S.A.

Multi-Agent Epistemic Reasoning

Reasoning about actions and information change has been one of the prominent interests since the beginning of the AI [4]. As discussed in [5], "***information*** **is something that is relative to a subject who has a certain perspective on the world, called an *agent*, and that is meaningful as a whole, not just loose bits and pieces. This makes us call it *knowledge* and, to a lesser extent, *belief*.**"

Initially formalized by logicians in the early sixties, epistemic reasoning rapidly evolved into *Dynamic Epistemic Logic* (DEL), a formalism used to reason both on the state of the world and *information change* in dynamic domains.

Automated planning and DEL has been recently merged together in a new framework called *Multi-agent Epistemic Planning* (MEP) [1]. Importantly, and differently from most approaches, the *epistemic states* (e-states) during the planning process must contain not only the state of the world (factual information), but also the *knowledge* or *beliefs* of the agents (epistemic information). In our setting, we focus on beliefs.

An Alternative Epistemic-State Representation

The traditional framework for MEP is built around the well known *Kripke models* formalism. In our previous work [2] we considered an alternative representation of e-states, namely ***Possibilities***.

Possibilities (first introduced in [3]) are *non-well-founded* objects that encode both factual and epistemic information. As shown in [3], they provide us with a more compact representation w.r.t. the traditional Kripke models.

Possibilities

Let $\mathcal{AG}$ be a set of agents and $\mathcal{F}$ a set of propositional variables:

- A *possibility* u is a function that assigns to each propositional variable $f \in \mathcal{F}$ a truth value $u(f) \in \{0, 1\}$ and to each agent $ag \in \mathcal{AG}$ an information state $u(ag) = \sigma$.
- An *information state* σ is a (non-well-founded) set of possibilities.

Each possibility u encodes a possible world. Specifically, u contains both the *interpretation of the world*, given by the component $u(f)$, and the *beliefs* of each agent, given by $u(ag)$. Intuitively, $u(ag)$ corresponds to the set of possibilities that ag considers to be possible in u.

A possibility

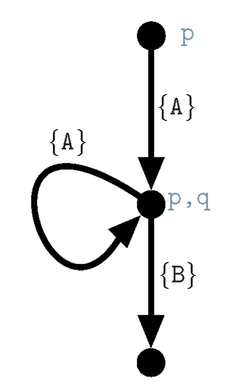

Its system of equation

$$\begin{cases} w(p) = 1 & w(q) = 0 \\ v(p) = 1 & v(q) = 1 \\ u(p) = 0 & u(q) = 0 \\ w(A) = \{v\} & w(B) = \{\emptyset\} \\ v(A) = \{v\} & v(B) = \{u\} \\ u(A) = \{\emptyset\} & u(B) = \{\emptyset\} \end{cases}$$

Corresponding K-Structure

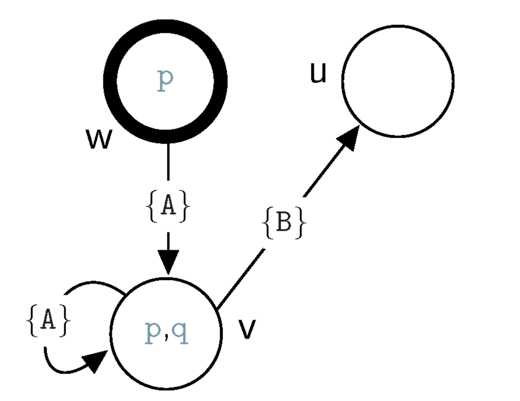

Figure 1: Transition from a possibility to a Kripke structure

Inconsistent Beliefs and Attitudes

In real-world situations, it is often the case that we learn a fact that discords with our previous beliefs. When such a discrepancy arises we talk about *inconsistent belief*. Notice that an inconsistency is not relative to the real state of the world (as in the case of *false beliefs*), but rather to the perspective of a particular agent.

Assume that agent i believes that $\neg p$ is the case in the e-state u. In our framework, there are two main sources of inconsistencies:

1. Agent i observes the real world (*sensing* action) and learns p (the opposite of what she believed).
2. Agent i learns p as a result of an *announcement* performed by another agent j.

In both scenarios, we must account for the belief of i after the action. In the former case, we make i believe that p (*i.e.*, agents trust their senses when observing the world). In the latter, we must take into account the *attitude* of the agent w.r.t. the announcer j. In fact, agent i may be skeptical or credulous, and thus she would change her belief according to her attitude towards j.

Specifically, it would be reasonable to have agent i believe the announcer j if i ***trusts*** j. We consider three attitudes for agent i:

- *Trustful* (T): i believes what the announcer has told her.
- *Mistrustful* (M): i believes the opposite of what has been announced.
- *Stubborn* (S): i does not modify her beliefs.

Sometimes it might be the case that an agent, say k, is aware that something about p is being announced (*i.e.*, $p = 0$, or $p = 1$), but she is not aware of what that *something* is. This is the case of *semi-private announcements*. In this situation, agent k is said to be *partially observant*. Agents that are also aware of what is announced about p are called *fully observant*. We consider two attitudes for partially observant agents:

- *Impassive* (I): k keeps her current beliefs.
- *Doubtful* (D): k does not believe neither what is being announced nor the opposite, regardless of her previous beliefs.

Enriched Semantics

When an agent, say j, announces something, the beliefs of the other attentive agents are affected. The key element of our semantics is that each agent will update her beliefs depending on her *attitude towards* j. The semantics of announcements is captured by a transition function $\Phi(a, u)$, that returns a new possibility u'.

Moreover, our semantics relies on a *table* $\mathcal{T}$ (contained in the input domain) containing the attitudes of all agents. Specifically, $\mathcal{T}$ is a set of tuples of the form $(i, j, \mathtt{att}, \varphi)$, meaning that agent i has attitude att towards agent j if the formula φ is true in the current e-state. We assume that, for each pair of distinct agents, at each time there is exactly one possible tuple where the condition is satisfied.

Consider the announcement action $a\langle j \rangle$, where j announces the value of p from her perspective. Confronting the tuples of $\mathcal{T}$ with the current e-state u results into a partition of the agent set $\mathcal{AG}$ that group agents depending on their attitude towards j. Such partition is called the *frame of reference* of $a\langle j \rangle$ and it is denoted by $\rho_{a\langle j \rangle} = \langle (\{j\}, \mathbf{T}_a, \mathbf{M}_a, \mathbf{S}_a), (\mathbf{I}_a, \mathbf{D}_a), \mathbf{O}_a \rangle$, where $\{j\}$ is the singleton containing the announcer, $\mathbf{O}_a$ is the set of *oblivious* (*i.e.*, non-attentive) agents, and $\mathbf{X}_a = \{i \mid \{(i, j, \mathtt{att}_{\mathbf{X}}, \varphi)\}$ and $u \models \varphi\}$. Notice that j is separated from the other fully observant agents, since she is the announcer in action a.

Applying $a\langle j \rangle$ in u results in a new possibility u'. The beliefs of agent i are updated following the transition function:

$$u'(i) = \begin{cases} u(i) & \text{if } i \in \mathbf{O}_a \\ P(a, u) & \text{if } i \in \mathbf{I}_a \cup \mathbf{D}_a \\ F(a, u, 1) & \text{if } i \in \mathbf{T}_a \\ F(a, u, 0) & \text{if } i \in \mathbf{M}_a \\ S(a, u, e(a), 1) & \text{if } i \in \mathbf{S}_a \\ S(a, u, e(a), 0) & \text{if } i = j \end{cases}$$

Intuitively, the beliefs of i are updated depending on her attitude towards j; for instance, if i is *trustful*, then the sub-function $F(a, u, 1)$ will handle the update. Notice that *mistrustful* agents are handled with the sub-function $F(a, u, 0)$, where the last parameter signals that we have to "flip" the value of p that j announced. Similarly, S handles *stubborn* agents.

Let us explore into detail the case of *partially observant* agents. In what follows, for a possibility w, the function $\chi(a, w, x)$ (resp., $\bar{\chi}(a, w, \neg x)$) is used to recursively set the value of p to x (resp., $\neg x$), in the possibilities representing the beliefs of *trustful* (resp., *mistrustful*) agents.

$$w'_p(i) = \begin{cases} w(i) & \text{if } i \in \mathbf{O}_a \\ \bigcup_{v \in w(i)} P(a, v) & \text{if } i \in \mathbf{I}_a \\ \bigcup_{v \in w(i)} \chi(a, v, 0) \cup \chi(a, v, 1) & \text{if } i \in \mathbf{D}_a \\ \bigcup_{v \in w(i)} \chi(a, v, v(p)) & \text{if } i \in \mathbf{T}_a \cup \mathbf{M}_a \cup \{j\} \\ \bigcup_{v \in w(i)} S(a, v, v(p), 1) & \text{if } i \in \mathbf{S}_a \end{cases}$$

Impassive agents do not modify their beliefs, since $P(a, v)$ does not affect the truth values of fluents. *Doubtful* agents relax their previous beliefs on p: this is achieved by including in their beliefs both a possibility with $p = 0$ and one with $p = 1$. Since partially observant agents are not aware of the value of p, they update the perceived beliefs of fully observant agents by calling the sub-functions χ and S with $v(p)$. This ensures that the beliefs of fully observers from the perspective of partial observers remain unchanged.

Functions F and S are defined in an analogous way, making sure to correctly update the nested beliefs of all agents, depending on their attitude.

We now show an illustrative example. Let $\mathcal{AG} = \{a, t, m, i, d\}$ and $\mathcal{F} = \{\mathsf{tails}\}$. In this simple example, the agents will share information about the position of a coin. On the top is displayed the possibility u where a believes that the coin lies tails up and each agent except for a is uncertain about the position of the coin (tails or $\neg$tails). The *real* state of the world is represented in boldface. On the bottom is displayed the possibility u' that results from the execution of the action instance shout⟨a⟩ where the announcer agent a announces tails. We assume the following attitudes towards a of the agents: t is *trustful*, m is *mistrustful*, i is *impassive*, and d is *doubtful*.

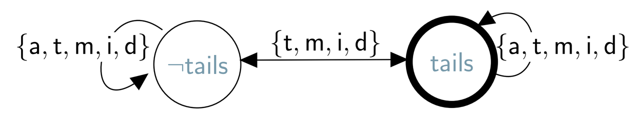

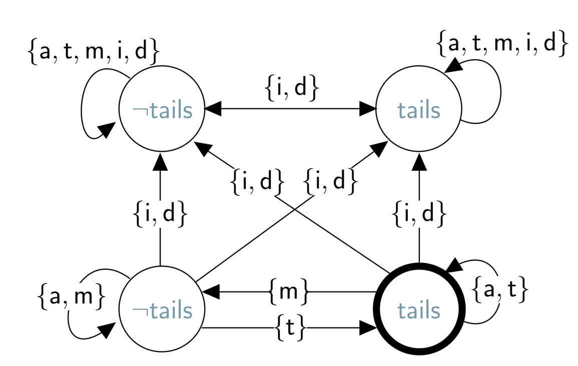

As we can see in the resulting e-state, t and m believe that the coin lies tails and heads up, respectively. Moreover, t and m believe that a shares their beliefs on the coin position. Finally, agents i and d, still do not know the coin position, but they believe that a, t and m know it.

References

[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] Francesco Fabiano, Alessandro Burigana, Agostino Dovier, and Enrico Pontelli. EFP 2.0: A multi-agent epistemic solver with multiple e-state representations. In *Proceedings of the Thirtieth International Conference on Automated Planning and Scheduling, Nancy, France, October 26-30, 2020*, pages 101–109. AAAI Press, 2020.

[3] Jelle Gerbrandy. *Bisimulations on planet Kripke*. Inst. for Logic, Language and Computation, Univ. van Amsterdam, 1999.

[4] John McCarthy. Programs with common sense. In *Proceedings of the Teddington Conference on the Mechanization of Thought Processes*, pages 75–91, London, United Kingdom, 1959.

[5] Hans Van Ditmarsch, Wiebe van Der Hoek, and Barteld Kooi. *Dynamic epistemic logic*, volume 337. Springer Science & Business Media, 2007.

Acknowledgements

This research is partially supported by Indam GNCS grants, by Uniud PRID ENCASE and by NSF grants 1914635 and 1833630.