

E-PDDL: A Standardized Way of Defining Epistemic Planning Problems

Francesco Fabiano
DMIF Department
University of Udine, Italy
francesco.fabiano@uniud.it

Biplav Srivastava
AI Institute
University of South Carolina, Columbia, USA
biplav.s@sc.edu

Marianna Bergamaschi Ganapini
Philosophy Department
Union College, NY, USA
bergamam@union.edu

Jonathan Lenchner and Lior Horesh and Francesca Rossi

Thomas J. Watson Research Center, Yorktown Heights, NY USA

{lenchner; lhoresh}@us.ibm.com
francesca.rossi2@ibm.com

Abstract

Epistemic Planning (EP) refers to an automated planning setting where the agent reasons in the space of knowledge states and tries to find a plan to reach a desirable state from the current state. Its general form, the *Multi-agent Epistemic Planning* (MEP) problem involves multiple agents who need to reason about both the state of the world and the information flow between agents. In a MEP problem, multiple approaches have been developed recently with varying restrictions, such as considering only the concept of *knowledge* while not allowing the idea of *belief*, or not allowing for “complex” modal operators such as those needed to handle *dynamic common knowledge*. While the diversity of approaches has led to a deeper understanding of the problem space, the lack of a standardized way to specify MEP problems independently of solution approaches has created difficulties in comparing performance of planners, identifying promising techniques, exploring new strategies like ensemble methods, and making it easy for new researchers to contribute to this research area. To address the situation, we propose a unified way of specifying EP problems - the Epistemic Planning Domain Definition Language, E-PDDL. We show that E-PDDL can be supported by leading MEP planners and provide corresponding parser code that translates EP problems specified in E-PDDL into (M)EP problems that can be handled by several planners. This work is also useful in building more general epistemic planning environments where we envision a meta-cognitive module that takes a planning problem in E-PDDL, identifies and assesses some of its features, and autonomously decides which planner is the best one to solve it.

Motivation

Multi-agent scenarios are ubiquitous in everyday life. We often need to make decisions based on information on the problem to be solved and on our knowledge, or belief, about the preferences and actions of other agents involved in, or impacted by, the decision. It is therefore an epistemic reasoning task to make decisions in this scenarios.

Often such decision involve creating a plan, that is, a sequence of actions that, when executed, will lead to a desired

goal state. The research area of *Multi-agent Epistemic Planning* (MEP) attempts to build artificial agents that can build such plans. MEP is a generalization of the concept of *Epistemic Planning* (EP), which refers to a single agent setting that reasons about knowledge and beliefs. In the MEP scenario, agents also reason about the information that other agents have of each other, and how such information flows between agents.

Several approaches to develop epistemic solvers/planners have been made across the community, see for example (Baral et al. 2015; Muise et al. 2015; Fabiano et al. 2020). While these solvers may have different restrictions, e.g., considering only the concept of *knowledge* while not allowing the idea of *belief*, or not allowing for *dynamic common knowledge*, all of them share the same goal: to plan while considering the information flows between the agents.

While the diversification of approaches in solving this problem certainly brings benefits to the entire community, with many specialized efficient planners that are suitable for specific scenarios, the lack of a standardized way to define input problems in multi-agent epistemic planning has led to the creation of ad-hoc languages that can become an obstacle when comparing, combining (i.e., ensemble), or testing different planners. We believe that a unified way of expressing MEP problems would provide researchers with a faster and less error-prone way of defining standardized benchmarks and would also relieve them from the burden of having to learn a new language each time they need to make use of a different tool. Moreover, it would allow for easy transfer of techniques and results between different planning scenarios and planner methodologies.

In this paper we therefore propose the Epistemic Planning Domain Definition Language (E-PDDL), a unified way of defining epistemic problems that, similarly to the well-known language PDDL for classical planning, could be regarded as the standard language to be used to express MEP problems. Being able to define MEP problems in E-PDDL and then to map them into specific languages for existing planners provides also the necessary environment for a real-time choice of the best planner for a given MEP problem. Indeed, this work is part of a larger effort that aims to equip artificial agents with meta-cognitive capabilities that, given internal values, resources, and criteria, can identify some of the problem features and assess what skills are required to

best solve it. In the context of this paper, this would mean to be able to select the most suitable planner for a MEP problem. This is greatly facilitated if there is a way to map the input problem specification into the language required by each specific planner.

In the rest of the paper, we first give necessary background about EP and MEP with a running example. We then introduce the E-PDDL language and its key features, and we show how it encompasses current MEP formalisms. Next, we highlight implementation details in the parser and conclude.

Background

Recently the concept of Artificial Intelligence has gained a lot of momentum both inside and outside the computer science research community. We are, in fact, witnessing an increasing number of autonomous tools that help users in tasks that range from the trivial and ordinary to the complex and very specialized. In particular, the area of automated planning, that has always been widely employed especially at the industrial level, has been constantly pushed towards improvement. Nowadays planners and schedulers are able to solve complex tasks in very short times while also accounting for factors such as incomplete knowledge, non-determinism, fault occurrences and so on (Helmert 2006; Lipovetzky and Geffner 2017; Richter and Westphal 2010).

However, the main focus of these research efforts has been about reasoning within single agent domains. This type of setting mainly requires the agent to characterize her knowledge about the world and understand in which way she can manipulate it in order to achieve a particular goal. But, while the single-agent setting can be satisfactory in many scenarios, multi-agent planning can be adopted to characterize a wider spectrum of real-world situations. It is therefore natural that, thanks to the increasing interest in the area and the more powerful hardware support, the planning problem is now also broadly studied in the context of multiple agents.

Multi-Agent Epistemic Planning

Nevertheless, the presence of numerous agents in the planning setting poses a number of challenges: among which is to study the information flows that occurs between the agents. The area of planning that reasons in an environments where these streams of “knowledge” need to be considered is known as *Epistemic Multi-Agent Planning*, or simply MEP (see (Fagin and Halpern 1994; Bolander 2017) for a more detailed introduction to the topic). MEP derives from the well-known research area of *Dynamic Epistemic Logic* which uses logic to formalize and reason and reason on both the state of the world and on the *information change* in dynamic domains. To be more concrete, as stated in (Van Ditmarsch, van Der Hoek, and Kooi 2007), “*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*.” In what follows, for brevity, we use the term “knowledge” to encapsulate *both* the notions of an agent’s knowledge and

beliefs with respect to some information. The concepts of knowledge and beliefs are captured by the same modal operator in Dynamic Epistemic Logic and their difference resides in structural properties that the epistemic states respect (see (Fagin and Halpern 1994) for more details). This means that thanks to epistemic planners we are able to reasons on domains where the actions of the agents affect, and are affected, by the state of the world and the knowledge of the agents (about both the world and the others’ knowledge).

As already mentioned, MEP is concerned with information change and inherits concepts from Dynamic Epistemic Logic (DEL). In particular, the language of well-formed DEL formulae used to express agents’ knowledge, is defined as follows:

$$\varphi ::= f \mid \neg\varphi \mid \varphi \wedge \psi \mid \mathbf{B}_i\varphi \mid \mathbf{C}_\alpha\varphi,$$

where f is a propositional atom called a *fluent*, i is an agent that belongs to the set of agents \mathcal{AG} s.t. $|\mathcal{AG}| \geq 1$, φ and ψ are belief formulae and $\emptyset \neq \alpha \subseteq \mathcal{AG}$. A *fluent formula* is a DEL formula with no occurrences of modal operators. A *belief formula* is recursively defined as follows:

- A fluent formula is a belief formula;
- If φ is a belief formula and $i \in \mathcal{AG}$, then $\mathbf{B}_i\varphi$ (“ i knows/-believes that φ ”) is a belief formula where the modal operator \mathbf{B} captures the concept of *knowledge*;
- If φ_1, φ_2 and φ_3 are belief formulae, then $\neg\varphi_3$ and $\varphi_1 \text{ op } \varphi_2$ are belief formulae, where $\text{op} \in \{\wedge, \vee, \Rightarrow\}$;
- If φ is a belief formula and $\emptyset \neq \alpha \subseteq \mathcal{AG}$ then $\mathbf{C}_\alpha\varphi$ is a belief formula where \mathbf{C}_α captures the *Common knowledge* of the set of agents α .

The formula $\mathbf{C}_\alpha\varphi$ translates intuitively into the conjunction of the following belief formulae:

- every agent in α knows φ ;
- every agent in α knows that every agent in α knows φ ;
- and so on *ad infinitum*.

The semantics of DEL formulae is traditionally expressed using *pointed Kripke structures* (Kripke 1963). We refer the interested reader to (Fagin and Halpern 1994; Bolander 2017; Gerbrandy and Groeneveld 1997) for a comprehensive introduction to how Kripke structures or similar formalisms are used to capture the idea of an epistemic state and how the concept of entailment is defined.

In this paper, we will make use of the **Coin in the Box** domain (Example 1), first introduced in (Baral et al. 2015), as a running example. In particular, we will present a three agent instance of this epistemic domain that will allow us to better illustrate some of the concepts we will later introduce.

Example 1 (Coin in the Box). Three agents, A, B, and C, are in a room where in the middle there is a box containing a coin. In the initial state it is known that:

- none of the agents know whether the coin lies heads or tails up;
- the box is locked and one needs a key to open it;
- some agents (usually only A) have the key to the box;

- in order to learn whether the coin lies heads or tails up, an agent can peek into the box, but this requires the box to be open; and
- agents can be either attentive or not towards the box.

Agents can perform the following actions, which have a straightforward meaning:

- *Opening* the box;
- *Peeking* into the box;
- *Signaling* another agent;
- *Distracting* another agent; and
- *Announcing* the coin position.

Action Languages

The field of planning has seen many representations. For example, in classical planning, there was STRIPS (Fikes and Nilsson 1971), Action Description Language (ADL) (Pednault 1994) and SAS+ (Bäckström 1995) before Planning Domain Description Language (PDDL) (McDermott et al. 1998; Fox and Long 2003) standardized the notations. Nowadays, planners routinely use PDDL for problem specification even if they may convert to other representations later for solving efficiency (Helmert 2009). PDDL envisages two files, a domain description file which specifies information independent of a problem like predicates and actions, and a problem description file which specifies the initial and goal states. PDDL provides a relaxed specification of the output plan – it is a series of time steps, each of which can have one or more instantiated actions with concurrency semantics. Specification of the correctness of plans was formalized much later when tools like VAL were created (Howey and Long 2003).

In PDDL, a planning environment is described in terms of objects in the world, predicates that describe relations that hold between these objects, and actions that bring change to the world by manipulating relations. A problem is characterized by an initial state, together with a goal state that the agent wants to transition to, both states specified as configurations of objects. When planning is used for epistemic reasoning, the objects in the problem can be physical (real world objects) as well as abstract (knowledge and beliefs).

Epistemic Action Languages We briefly present the action languages adopted by two different epistemic solvers: RP-MEP (Muise et al. 2015) and EFP (Fabiano et al. 2020). We chose these solvers as, in our opinion, they represent the state of the art epistemic planners. The former, *i.e.*, RP-MEP, adopts a conversion to classical planning in order to tackle epistemic reasoning while the latter, *i.e.*, EFP, comprehensively reasons on full-fledged epistemic states. While the two planners are able to solve the same families of domains, important differences in their solving strategies make certain problem instances more suited to one approach rather than the other. For a detailed comparison of the solvers refer to (Fabiano et al. 2020; Le et al. 2018). The main objective of this paper is to present a unified way of providing problem specification for the epistemic planning community so that the planners, starting from the two considered here, can be easily tested, confronted or even combined.

PDKB-PDDL Let us start by presenting the main characteristics of the action language PDKB-PDDL adopted by the solver RP-MEP (Muise et al. 2015; Queen’s Mu Lab 2020). This action language clearly stems from PDDL, the de-facto standard for single-agent action languages, and allows one to express formulae that contain the modal operator **B**. While this language seems to represent a version of PDDL that is suitable for epistemic reasoning, there are some limitations that render PDKB-PDDL not general enough for this purpose. In fact, while this language inherits the files schema from PDDL, also allowing for parametric action definitions, the epistemic components of the actions’ description are tailored for the RP-MEP solving process. This makes PDKB-PDDL not flexible enough to be adopted as an input specification language for other planners (*e.g.*, EFP). In particular, what makes this language not suitable for every epistemic approach is that, in order to have a valid PDKB-PDDL domain, each action must specify the entirety of its effects explicitly. While this is a common procedure in other types of reasoning in epistemic planning, the need for explicit effects specification brings some limitations. To be more concrete, each action in PDKB-PDDL needs to define not only how the action itself modifies the world but also how the knowledge of the agents are updated—that is, the belief update of each action needs to be explicit. This requirement is limiting because epistemic planners (*e.g.*, EFP) may reason through implicit knowledge update¹ and therefore “automatically” derive how the agents change their knowledge after an action execution. While having explicit updates allows one to freely characterize each action’s effects (*e.g.*, in Example 1 it would be possible to define a new action in which an agent opens the box but will believe that the coin is tails up), we feel that leaving such detail as input is partially in contrast with the idea of autonomous epistemic reasoning. Furthermore, it is easy for the user to inadvertently omit or mis-characterize “intricate” knowledge relations that should instead be handled by the autonomous solving process. For example, in writing the effects *peeking* of Example 1 the effect “the agent who peeked inside the box knows that the attentive agents know that he knows the coin position but the attentive agents themselves do not know the position of the coin” may be misrepresented. This chain of knowledge is derived by partial observability (explained in the following paragraph) and, as well as other complex knowledge structures, may be easily mis-characterized given its complex and not intuitive nature.

$m\mathcal{A}^p$ We can now introduce $m\mathcal{A}^p$, the language used to describe inputs for EFP (Fabiano et al. 2020). This language is an evolution of $m\mathcal{A}^+$ first introduced in (Baral et al. 2015). Unlike PDKB-PDDL, $m\mathcal{A}^p$ does not take inspiration from PDDL and introduces an ad-hoc syntax to define epistemic problems. This syntax, illustrated in detail in (Baral et al. 2015; Fabiano et al. 2019), describes an input problem through a single file that specifies initial and goal states (or rather the belief formulae that identify such states) along

¹The distinction between explicit and implicit knowledge update derives from the fact that RP-MEP grounds the problem while EFP reasons directly on epistemic states.

with the actions making each file a problem instance. Moreover, unlike the input language of RP-MEP, $m\mathcal{A}^p$ defines only implicit effects and by associating each action with the agents’ observability relations (which agent is aware of the action’s occurrence and/or effects) and with a “type” (*world-altering*, *sensing* and *announcement*) is able to derive the complete knowledge update. In particular, in (Baral et al. 2015), each action execution `act` associates each agent `ag` to one of the following observability relations:

- *fully observant* if `ag` knows about the execution of `act` and about its effects on the world (*e.g.*, any attentive agent that witnessed the opening of the box in Example 1);
- *partially observant* if `ag` knows about the execution of `act` but she does not know how `act` affected the world (*e.g.*, any attentive agent that sees another agent *peeking* inside the box in Example 1); and
- *oblivious* if `ag` does not know about the execution of `act` (*e.g.*, any non-attentive agent during the execution of any action in Example 1).

Moreover, as noted before, each action is associated with one of the following “types”:

- *World-altering* (or *ontic*) action, used to modify certain properties (*i.e.*, fluents) of the world, *e.g.*, the action of *opening* of Example 1.
- *Sensing* action: used by an agent to refine her beliefs about the world, *e.g.*, the action of *peeking*.
- *Announcement* action: used by an agent to affect the beliefs of other agents *e.g.*, in Example 1 the action of *announcing*.

While this language is able to characterize implicit knowledge updates (from which the explicit effects can be derived), it lacks the structure of PDDL and relies on a custom syntax that may pose an obstacle to any researcher that is trying to approach the area of epistemic planning. Moreover, another problem of $m\mathcal{A}^p$ is the lack of parametric actions that, in several cases, forces the user to repeat the definition of the same action, with small variations, over and over.

E-PDDL

In what follows we present E-PDDL making use of the Coin in the Box domain (Example 1) to better explain some of the features. Let us remark that the syntax has been chosen with the objective of minimizing the difference with standard PDDL while providing a general epistemic input language. We will start by showing the syntax of the *problem-domain*—that contains the general settings of the problem—and then we will illustrate how a *problem-instance*—that contains specific objects, initial conditions and goals—is characterized. In particular in Listings 1 we present the characterization of the Example 1 *problem-domain* and in Listings 2 we present a simple *problem-instance* where it is known that agent `a` has the key and the goal is for `a` to know the coin position. Before proceeding with the description of the E-PDDL syntax we need to introduce the meaning of the operator “[*i*]” where $i \in \mathcal{AG}$. This operator captures the modal operator \mathbf{B}_i . For example in Line 13 of Listing 1 the

formula “[*i*](has_key ?i)” reads “agent `i` knows has_key_?” where $i \in \mathcal{AG}$ and the fluent `has_key_i` encodes the fact that `i` has the key. When the operator is of the form “[α]” where $\alpha \subseteq \mathcal{AG}$ and $|\alpha| \geq 2$ then it captures the idea of common knowledge.

Problem Domain

First of all let us note that in Listings 1 the actions `signal` and `distract` are omitted to avoid clutter. In fact, these actions are world-altering and, therefore, share a similar structure with the action `open`.

Following the PDDL syntax (Fox and Long 2003) we start the *problem-domain* definition by defining the name and the requirements of the problem (Lines 1 and 2-3 of Listings 1, respectively). We included a new requirement called `:mep` to identify the need for E-PDDL. In Lines 5-8 are introduced the predicates following the PDDL standard. A small variation is the *object-type* agent that does not need to be defined and is used to define variables that capture the acting agents.

From Line 10 to Line 19 the action `open` is introduced. The action’s definition starts with its *name* (Line 10) and its *type* (Line 11). The concept of action type is inherited from $m\mathcal{A}^p$ and for now is restricted to be one among *ontic*, *sensing* or *announcement* since these are the accepted variations of an action in the MEP community. This concept allows to apply the right transition function in the context of implicit knowledge update. We plan, for the future, to include custom action types through the definition of custom event models (see (Baral et al. 2015) for a better introduction on this topic). Since such feature needs to be also implemented in the epistemic planners, we leave this contribution as a future work. Next, in Line 12 and Line 13, respectively, the action’s *parameters* and *preconditions* are defined. The parameters have the same role that they have in PDDL, that is they are used to associate the variables of the action’s definition with an *object type*. Similarly, also the field *preconditions*—identified by any belief formula—follow the standard PDDL meaning. Following the preconditions the action specifies the *effects*. The effects, contrarily from the preconditions, cannot contain arbitrary belief formulae. In fact, depending on the action type, only fluents or special fluent formulae are allowed. In particular, *ontic* actions allow for conjunctions of fluents while *epistemic* actions (*i.e.*, *sensing* and *announcements*) can only consider single fluents as effects. These limitations, inherited by $m\mathcal{A}^p$, are necessary for planners that reason on epistemic transition functions. We address the reader to (Fabiano et al. 2020) for a more detailed description of the motivations behind these limitations. Finally, the last field of the action `open` is about the *observers*. This field is used to indicate which agents are fully observant, *i.e.*, knows about the execution and the effects of the action. Knowing which agent is observant is once again useful to derive how the knowledge of the agents is updated after the action is been executed. To better characterize the set of observant agents we introduced the operator `diff` that allow to “isolate” the executor of the action (since the executor needs to be observant and should not depend on other factors). For example, the condition in Lines 15-17 reads as: “the agent `i`, *i.e.*, the executor, is fully observant” (Line 15) and “for every

agent $j \neq i$ if j is looking then j is fully observant” (Lines 16-17). The same schema is used to define partial observers, the ones that are aware of the action execution but do not know the results of such action, with the field *p_observers*; an example of partial observability is at Line 27.

```

1 (define (domain coininthebox)
2   (:requirements :strips
3     :negative-preconditions :mep)
4
5   (:predicates (opened)
6     (has_key ?i - agent)
7     (looking ?i - agent)
8     (tail))
9
10  (:action open
11    :act_type      ontic
12    :parameters   (?i - agent)
13    :precondition ([?i](has_key ?i))
14    :effect        (opened)
15    :observers    (and (?i)
16      (forall (diff (?j-agent) (?i))
17        (when(looking ?j) (?j))))
18  )
19
20  (:action peek
21    :act_type      sensing
22    :parameters   (?i - agent)
23    :precondition (and ([?i](opened))
24      ([?i](looking ?i)))
25    :effect        (tail)
26    :observers    (?i)
27    :p_observers  (forall (diff (?j-agent) (?i))
28      (when(looking ?j) (?j)))
29  )
30
31  (:action announce
32    :act_type      announcement
33    :parameters   (?i - agent)
34    :precondition ([?i](tail))
35    :effect        (tail)
36    :observers    (and (?i)
37      (forall (diff (?j-agent) (?i))
38        (when(looking ?j) (?j))))
39  )
40 )

```

Listing 1: E-PDDL Coin in the Box problem-*domain*.

The introduced fields are tailored for implicit knowledge update but in Section “From Implicit to Explicit Knowledge Update” we will explain how, with the presented syntax, it is possible to generate a valid input also for those planners that need the effects of the action to be completely explicit.

Problem Instance

In Listings 2 we present an example of an E-PDDL problem-*instance*. In Line 1 and in Line 2 the problem-*instance* name (*i.e.*, *toyinstance*) and the related problem-*domain* name are defined, respectively.

Next, in Line 3, the *object type* agent values are defined. In this particular instance we defined three agents a, b and c as in Example 1.

Following, in Line 4, the *depth* is specified. The concept of depth of a belief formula is used to identify the number of *nested* epistemic operators. For example, given two agents i and j the belief formula $B_i\varphi$ has depth 1 while $B_iC_{i,j}\varphi$ has depth 2. This field is introduced to accommodate the need for certain planners, *e.g.*, RP-MEP, to limit the depth of the belief formulae. In fact RP-MEP relies on grounding the formulae into classical planning “facts” and without bound on these formulae could be infinite. On the other hand, the limit on depth is ignored by the planners, *e.g.*, EFP, that reason directly on epistemic states.

Lines 5-12 present the belief formulae that describe the initial state. The formulae are considered to be in conjunction with each other. The *initial conditions* only require to specify when a fluent is true and consider false whichever fluent is not specified (Line 5). Moreover the initial conditions are also used to specify what is known in the initial state (Line 6-12). While the belief formulae in this field can be of any type, let us remark that in (Son et al. 2014) it is demonstrated that to create a finite number of epistemic states from a set of formulae, this set must respect a finitary **S5** logic and therefore the knowledge must be expressed in terms of common knowledge. This means that if the initial conditions do not comply with a finitary **S5** logic the planners that construct the initial epistemic state from the given specification may not work.

Finally in Line 13 the conjunction of belief formulae that represent the goals is defined. Here, to respect Example 1, only the goal that reads “a knows that the coin is tails up” is written.

```

1 (define (problem toyinstance)
2   (:domain coininthebox)
3   (:agent a b c)
4   (:depth 2)
5   (:init (tails) (has_key a) (looking a)
6     ([a b c](has_key a))
7     ([a b c](not (has_key b)))
8     ([a b c](not (has_key c))))
9     ([a b c](not (opened)))
10    ([a b c](looking a))
11    ([a b c](not (looking b)))
12    ([a b c](not (looking c))))
13  (:goal ([a](tails)))
14 )

```

Listing 2: E-PDDL Coin in the Box problem-*instance*.

From Implicit to Explicit Knowledge Update

Since we tailored E-PDDL syntax to represent actions with implicit knowledge update we need to explain how E-PDDL itself is a suitable language also for those planner, *e.g.*, RP-MEP, that need the knowledge update to be explicit. That is, we need a standard way of deriving the explicit agents’ knowledge update from an E-PDDL action description. While deriving explicit knowledge update is not a method that is embedded in the language itself in what follows we propose the strategy that we adopted in our implemented parser (presented in the following Section).

The strategy that we adopted in our parser is based on (Baral et al. 2015) transition function where the idea

of agents’ observability is used to derive consistent agents’ knowledge about the actions effects and/or execution. In what follows we present a *knowledge derivation schema* that, starting from the agents’ observability, is able to generate the explicit knowledge update related to a single action. These updates generate all the knowledge-chains of finite length ℓ where $\ell \in \{0, \dots, d\}$ and d is the value assigned to the field `:depth` in the problem-*instance* (in Listings 2 $d = 2$). Namely we will have all the following chains:

- a fully observant knows the action’s effect ($\ell = 1$);
- a fully observant knows that another fully observant knows the action’s effect ($\ell = 2$);
- a fully observant knows the chains of length 2 ($\ell = 3$);
- and so on until we have that $\ell = d$.

Moreover, when partially observant agents are defined we also need to take into consideration their perspective on the knowledge update. To do that we will need to automatically generate the following knowledge-chains (still limited by the given depth):

- a partially observant knows that any chain of fully observant agents knows the action’s effect ($\ell = 2$); and
- any chain, with $\ell \leq d - 2$, of fully/partially observant knows the chain of knowledge presented in the previous point.

In order to better integrate explicit knowledge update we decided to incorporate E-PDDL with an extra, non-mandatory field for the actions’ specification. This field, identified by `:exp_effect` can be used to identify the explicit knowledge update of the action by using an arbitrary belief formula. Let us note that when this field is defined it will completely override the automatically derived knowledge update for the planners that make use of explicit knowledge update, *e.g.*, RP-MEP. On the other hand, planners that make use of a full fledged epistemic transition function, *e.g.*, EFP, will ignore the `:exp_effect` field.

An E-PDDL Parser

In this Section we will present the implementation of an E-PDDL parser. This parser is able to take E-PDDL problem-*domain* and problem-*instance* files and automatically generate corresponding inputs for the solvers RP-MEP and EFP. That is, the main objective of our parser is to generate a PDKB-PDDL and a $m\mathcal{A}^p$ description of the given input. While this is just a first step of the integration of E-PDDL in the epistemic planning community, thanks to our parser RP-MEP and EFP can already exploit the newly presented syntax. Moreover, this parser allows to better compare and/or even combine the two approaches. Therefore, the presented approach may help in building a more general epistemic planning environment where, thanks also to a meta-cognitive reasoning module, takes a planning problem in E-PDDL, identifies and assesses some of its features, and autonomously decides which planner is the best one to solve it.

The parser source code, along with some examples and a small guide on how to use the code itself, is

available at the following url: <https://github.com/FrancescoFabiano/E-PDDL>. Our E-PDDL parser is written in Python3 (Van Rossum and Drake 2009) and inherits its basic structure from an open source PDDL parser (Magnaguagno et al. 2020).

As already said the objective of the parser is to “transform” an E-PDDL input into files that can be directly used as input for either RP-MEP or EFP. The idea is to provide the researchers with a tool that generates an output that can be directly used by the two above-mentioned planners.

While elaborating upon the details of the code will be tedious and not informative we decide to show the result of the parsing process. In Listings 3 and 4 show the problem-*domain* and the problem-*instance* (partial) input files for the planner RP-MEP resulted by calling the parser on the **Coin in the Box** instance described in Listings 1 and 2. Listing 5, on the other hand, presents (part of) the encoding of the same problem for the planner $m\mathcal{A}^p$. Let us remark that Listings 3 and 4 follow the syntax of PDKB-PDDL while Listing 5 respects the syntax of $m\mathcal{A}^p$.

```

1 (define (domain coininthebox)
2   (:agents a b c)
3   (:constants)
4   (:types)
5   (:predicates (opened) (has_key ?i - agent)
6                 (looking ?i - agent) (tails))
7
8   (action: open
9     :parameters (?i - agent)
10    :derive      (looking $agent$)
11    :precondition ([?i](has_key ?i))
12    :effects     (and (opened))
13  )
14
15  (action: peek
16    :parameters (?i - agent)
17    :derive      (never)
18    :precondition (and ([?i](opened))
19                      ([?i](looking ?i)))
20    :effects     (and [?i](tails)
21                  (forall (?j - agent)
22                    (when (looking ?j)
23                      [?j][?i](or(tails)(!tails)))
24                    )
25  )
26
27  (action: announce
28    :parameters (?i - agent)
29    :derive      (looking $agent$)
30    :precondition ([?i](tails))
31    :effects     (and (tails))
32  )

```

Listing 3: PDKB-PDDL Coin in the Box problem-*domain* generated through the parsing process of Listing 1.

In Lines 10 and 28 of Listing 3 we can see how, thanks to the PDKB-PDDL field `:derive - condition` (shorten to `:derive`), we are able to define which agents are aware of an action’s execution and its effects. This special field automatically expand the knowledge update for RP-MEP. This

conversion (that only allows belief formulae of depth up to 2 as indicated by Line 4 of Listings 2) correctly captures the idea that fully observant agents are aware of the actions' effects and also know that the other fully observants know these effects.

On the other hand, to express the complete knowledge update of the action peek we need to account also for partial observability. In Lines 20-23 the effects of the action are written so that the acting agent, *i.e.*, agent *i*, knows that the coin is tails up. The other agents, which are partially observant, only knows that *i* knows the coin position but they do not the position themselves. Given that the max depth of the formulae is set to 2 these effects suffice in capturing the behavior of peek. Let us remark that the syntax presented in Lines 20-23 of Listing 3 is not fully compatible with PDKB-PDDL but nicely captures the idea we are trying to express. Partial observability, while supported by the parser, is not yet fully supported during the conversion to PDKB-PDDL and is currently under development.

```

1 (define (problem pb1)
2   (:domain coininthebox)
3   (:depth 2)
4   (:projection )
5   (:task valid_generation)
6   (:init-type complete)
7
8   (:init
9     (tails)
10    (has_key a)
11    (looking a)
12
13    ;Length 1
14    (forall ?ag1 - agent
15      [?ag1](has_key a))
16
17    ;Length 2
18    (forall ?ag1 - agent
19      (forall ?ag2 - agent
20        [?ag2][?ag1](has_key a)))
21
22  )
23
24  (:goal
25    [a](tails)
26  )
27 )

```

Listing 4: PDKB-PDDL Coin in the Box problem-*instance* generated through the parsing process of Listing 2.

While Listing 4 has some extra fields w.r.t. to Listing 2 that are used by RP-MEP to set some internal parameters the only notable difference is the conversion of the common knowledge operator, not supported by RP-MEP, into chain of belief formulae. These chains, bounded by the depth limit (that is our case is set to two), are the chains of knowledge that are used to define the common knowledge operator. The conversion, if applied for example to Line 6 of Listing 2 “[a b c](has_key a)”, generates all the belief formulae where: every agent in {a; b; c} knows has_key a ($\ell = 1$); and every agent in {a; b; c} knows that every agent in {a; b; c} knows has_key a ($\ell = 2$).

```

1 fluent has_key_a, has_key_b, ..., tails;
2
3 action open_a, open_b, ..., announce_c;
4
5 agent a, b, c;
6
7 executable open_a if B(a,has_key_a);
8 open_a causes opened;
9 b observes open_a if look_b;
10 c observes open_a if look_c;
11 a observes open_a;
12
13 executable peek_a if B(a,opened),B(a,look_a)
14 peek_a determines tails if look_a;
15 a observes peek_a;
16 b aware_of peek_a if look_b;
17 c aware_of peek_a if look_c;
18
19 executable announce_a if B(a,tails);
20 announce_a announces tails;
21 b observes announce_a if look_b;
22 c observes announce_a if look_c;
23 a observes announce_a;
24
25 initially tails, has_key_a, look_a;
26 initially -look_b, -look_c, ...;
27
28
29 initially C([a,b,c],look_a);
30 initially C([a,b,c],has_key_a);
31 initially C([a,b,c],-opened);
32 ...
33 initially C([a,b,c],-has_key_c);
34
35 goal B(a,tails);

```

Listing 5: $m\mathcal{A}^p$ Coin in the Box problem generated through the parsing process of Listings 1 and 2.

Finally, in Listing 5, we show a partial results of the conversion of E-PDDL into $m\mathcal{A}^p$. Since $m\mathcal{A}^p$ does not allow to use variables and types to define the problem specification the parser grounds the instance of E-PDDL and use the result to generate the input for EFP. In particular, this grounded instance contains grounded fluents, *e.g.*, has_key_a, has_key_b and has_key_c, and grounded actions, *e.g.*, open_a. The grounded actions are instantiated versions of the ones declared in E-PDDL where each variable is associated with one of its possible values. Other small syntactical changes are then applied, *e.g.*, ontic actions are associated with keyword causes and partial observability is identified by the keyword aware_of.

Conclusion

In this paper, we proposed E-PDDL, an Epistemic Planning Domain Definition Language, that represents a unified way to specify *Multi-agent Epistemic Planning* (MEP) problems. We show that E-PDDL can unify and support leading MEP planners, and provide a parser code to the community that is able to generate input specification for two state of the art epistemic reasoners. Defining a unified language is the first step towards defining common benchmarks and evaluation

metrics for all the MEP community. Moreover, since the new language follows the well-known language PDDL, E-PDDL may also serve as a gateway for researchers already involved in the area of planning that want to investigate the branch of epistemic reasoning. Finally defining a unified way of specifying the input for different epistemic solvers can help in designing a general purpose epistemic solver that is automatically able to select the best solving strategy, between existing planners, starting from the problem definition that can now be “understood” by every considered approach. We leave the design of such architecture as future work.

While E-PDDL is able to capture most of the features that can be found in an epistemic domain the language only considers three types of actions: ontic, sensing and announcement. We plan to expand the language syntax so that is able to accept custom action definitions, specified with a custom event model. Nonetheless, before introducing such feature we must ensure compatibility with the considered epistemic solvers. For RP-MEP the conversion to explicit knowledge update is handled directly by the parser and, therefore, custom actions definition would only generate different chains of belief formulae. On the other hand, EFP relies on an epistemic transition function that expects the actions to be ontic, sensing or announcement and it would not be able to handle custom action types. That is why before expanding E-PDDL we need to implement in EFP the possibility of using custom action definitions.

Finally we plan to refine the parser code to better integrate the concept of partial observability and to be more flexible in the definition of the formulae.

References

- Baral, C.; Gelfond, G.; Pontelli, E.; and Son, T. C. 2015. An action language for multi-agent domains: Foundations. *CoRR* abs/1511.01960.
- Bolander, T. 2017. A gentle introduction to epistemic planning: The del approach. *arXiv preprint arXiv:1703.02192*.
- Bäckström, C. 1995. Expressive equivalence of planning formalisms. *Artificial Intelligence* 76(1):17–34. Planning and Scheduling.
- Fabiano, F.; Riouak, I.; Dovier, A.; and Pontelli, E. 2019. Non-well-founded set based multi-agent epistemic action language. In *Proceedings of the 34th Italian Conference on Computational Logic*, 242–259.
- Fabiano, F.; Burigana, A.; Dovier, A.; and Pontelli, E. 2020. EFP 2.0: A multi-agent epistemic solver with multiple e-state representations. In *Proceedings of the 30th International Conference on Automated Planning and Scheduling*, 101–109.
- Fagin, R., and Halpern, J. Y. 1994. Reasoning about knowledge and probability. *Journal of the ACM (JACM)* 41(2):340–367.
- Fikes, R. E., and Nilsson, N. J. 1971. Strips: A new approach to the application of theorem proving to problem solving. In *Proceedings of the 2nd International Joint Conference on Artificial Intelligence, IJCAI’71*, 608–620. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Fox, M., and Long, D. 2003. Pddl2.1: An extension to pddl for expressing temporal planning domains. *Journal of Artificial Intelligence Research* 20:61–124.
- Gerbrandy, J., and Groeneveld, W. 1997. Reasoning about information change. *Journal of Logic, Language and Information* 6(2):147–169.
- Helmert, M. 2006. The fast downward planning system. *Journal of Artificial Intelligence Research* 26:191–246.
- Helmert, M. 2009. Concise finite-domain representations for pddl planning tasks. *Artif. Intell.* 173(5–6):503–535.
- Howey, R., and Long, D. 2003. Val’s progress: The automatic validation tool for pddl2.1 used in the international planning competition. In *ICAPS 2003 workshop on “The Competition: Impact, Organization, Evaluation, Benchmarks”*, Trento, Italy.
- Kripke, S. A. 1963. Semantical considerations on modal logic. *Acta Philosophica Fennica* 16(1963):83–94.
- Le, T.; Fabiano, F.; Son, T. C.; and Pontelli, E. 2018. EFP and PG-EFP: Epistemic forward search planners in multi-agent domains. In *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling*, 161–170. Delft, The Netherlands: AAAI Press.
- Lipovetzky, N., and Geffner, H. 2017. Best-first width search: Exploration and exploitation in classical planning. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, 3590–3596.
- Magnaguagno, M.; Felipe, M.; Behery, M. A.; and Corfman, B. 2020. PDDL-parser. <https://github.com/pucrs-automated-planning/pddl-parser>.
- McDermott, D.; Ghallab, M.; Knoblock, C.; Wilkins, D.; Barrett, A.; Christianson, D.; Friedman, M.; Kwok, C.; Golden, K.; Penberthy, S.; Smith, D.; Sun, Y.; and Weld, D. 1998. PDDL - the planning domain definition language. Technical report, Technical Report.
- Muise, C. J.; Belle, V.; Felli, P.; McIlraith, S. A.; Miller, T.; Pearce, A. R.; and Sonenberg, L. 2015. Planning over multi-agent epistemic states: A classical planning approach. In *Proc. of AAAI*, 3327–3334.
- Pednault, E. P. D. 1994. ADL and the State-Transition Model of Action. *Journal of Logic and Computation* 4(5):467–512.
- Queen’s Mu Lab. 2020. PDKB-planning. <https://github.com/QuMuLab/pdkb-planning>.
- Richter, S., and Westphal, M. 2010. The lama planner: Guiding cost-based anytime planning with landmarks. *Journal of Artificial Intelligence Research* 39:127–177.
- Son, T. C.; Pontelli, E.; Baral, C.; and Gelfond, G. 2014. Finitary s5-theories. In *European Workshop on Logics in Artificial Intelligence*, 239–252. Springer.
- Van Ditmarsch, H.; van Der Hoek, W.; and Kooi, B. 2007. *Dynamic epistemic logic*, volume 337. Springer Science & Business Media.
- Van Rossum, G., and Drake, F. L. 2009. *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace.