

# Design of a Solver for Multi-Agent Epistemic Planning

Francesco Fabiano

Dipartimento DMIF, Università di Udine, Udine, Italy

`francesco.fabiano@uniud.it`

As the interest in Artificial Intelligence continues to grow it is becoming more and more important to investigate formalization and tools that allow us to exploit logic to reason about the world. In particular, given the increasing number of multi-agents systems that could benefit from techniques of automated reasoning, exploring new ways to define not only the world's status but also the agents' information is constantly growing in importance. This type of reasoning, *i.e.*, about agents' perception of the world and also about agents' knowledge of her and others' knowledge, is referred to as *epistemic reasoning*.

In our work we will try to formalize this concept, expressed through *epistemic logic*, for dynamic domains. In particular we will attempt to define a new action-based language for *multi-agent epistemic planning* and to implement an epistemic planner based on it. This solver should provide a tool flexible enough to be able to reason on different domains, *e.g.*, economy, security, justice and politics, where reasoning about others' beliefs could lead to winning strategies or help in changing a group of agents' view of the world.

**Keywords:** Epistemic reasoning · Planning · Multi-agent · Action languages · Heuristics  
Information change · Knowledge representation

## 1 Introduction

The proliferation of agent-based and IoT technologies has enabled the development of novel applications involving hundreds of agents. Considering that self-driving cars and other autonomous devices that can control several aspects of our daily life are going to be available en mass in just a few years it will not be long until massive systems of autonomous agents, each acting upon its own knowledge and beliefs to achieve its own (or group) goals, become available and widely deployed.

To maximize the potentials of such autonomous systems, *multi-agent planning* and scheduling research [1, 8–10, 24, 28] will need to keep pace. Moreover creating a plan for multiple agents to achieve a goal will need to take into consideration agents' knowledge and beliefs, to account for aspects like trust, dishonesty, deception, and incomplete knowledge. The planning problem in this new setting is referred to as *epistemic planning* in the literature; that is epistemic planners are not only interested in the state of the world but also in the knowledge or beliefs of the agents.

Nevertheless, reasoning about knowledge and beliefs is not as direct as reasoning on the “physical” state of the world. That is because expressing, for example, belief relations between a group of agents often implies to consider *nested* and *group* beliefs that are not easily extracted from the state description by a human reader. For this reasons it is necessary to develop a complete and accessible action language to model multi-agent epistemic domains [2] and to advance also in the study of epistemic solvers [4, 19, 23, 26, 34].

B. Bogaerts, E. Erdem, P. Fodor, A. Formisano,  
G. Ianni, D. Incezan, G. Vidal, A. Villanueva,  
M. De Vos, F. Yang (Eds.): International  
Conference on Logic Programming 2019 (ICLP'19).  
EPTCS 306, 2019, pp. 403–412, doi:10.4204/EPTCS.306.54

In our research we are exploring the epistemic planning problem with particular focus on formalizing an updated version of the action language proposed in [2] and on developing a multi-agent epistemic solver flexible enough to be used in several scenarios such as: economy, security, justice and politics. During this study we also intent to confront problems related to *epistemic dynamic logic* such as: i) the complexity of the plan existence problem (as explored in [3]); and ii) the use of alternative or implicit data structures that would help in representing epistemic states more efficiently.

## 2 Background

### 2.1 Dynamic Epistemic Logic

Logicians have always been interested in describing *the state of the world* through formalism that would allow to reason on the world with logic itself. This interest has lead, among other things, to the formalization of the well-known planning problem [29] and to the introduction of several *modal logics* [6, 30, 33] used to describe different types of domains. The differences between these logics are not merely syntactical but they carry implications in both expressiveness and complexity. Let us take for example, without going into details, the boolean *propositional logic* and the *linear temporal logic* (LTL). The first one, being one of the simplest logics, is mostly used to encode the world as a series of facts that can be true or not and, therefore, allows to “reduce” properties of the domain to *boolean formulae*. The latter instead, even if it is based on propositional logic, introduces new *modal operators* that allow to reason also about time (with a little abuse of the term). The absence of these operators in the first one makes propositional logic, adopted to represent problems such as *n-queens* or *circuit-SAT*, not expressive enough to encode problems, *e.g.*, *timeline-based planning problem* [16], that LTL can deal with. So, in general, we have that different logics have diverse operators and therefore are suitable for different type of reasoning on the world.

Nevertheless, even if different, both the two logics introduced above are limited in reasoning only on the state of world, *i.e.*, on its “physical” properties and on their changes. *Dynamic Epistemic Logic* (DEL), on the other hand, is used to reason not only on the state of the world but also on *information change*. As said in [33] *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*. The idea behind DEL is, therefore, to have a formalization that allows to reason on dynamic domains where, not only the state world is taken into consideration, but also the knowledge/beliefs that the agents have about the world and about the knowledge/beliefs of each other are considered.

#### 2.1.1 Epistemic Logic

Dynamic Epistemic Logic is, clearly, connected to *epistemic logic*: that is the logic that allows to reason on the knowledge/belief of agents in static domains. This logic is based on two main concepts: i) *Kripke structures*, a data structure that is widely use in literature [2, 13, 33] to model its semantics; and ii) *belief formulae*, a type of formula that takes into consideration epistemic operators and is used to represent the knowledge/beliefs of the agents. As it is beyond the scope of this work to give an exhaustive introduction on epistemic logic we will provide only the fundamental definitions and intuitions; the reader who has interest in a more detailed description can refer to [13].

Let  $\mathcal{A}$  be a set of agents and let  $\mathcal{F}$  be a set of propositional variables, called *fluents*. We have that each *world* is described by a subset of elements of  $\mathcal{F}$  (intuitively, those that are “true” in the world).

Moreover in epistemic logic each agent  $ag \in \mathcal{AG}$  is associated with an epistemic modal operator  $\mathbf{B}_{ag}$  that intuitively represents the knowledge/belief of  $ag$ . Finally, epistemic *group operators*  $\mathbf{E}_\alpha$  and  $\mathbf{C}_\alpha$  are also introduced in epistemic logic. Intuitively  $\mathbf{E}_\alpha$  and  $\mathbf{C}_\alpha$  represent the knowledge/belief of a group of agents  $\alpha$  and the *common knowledge/belief* of  $\alpha$  respectively. To be more precise, as in [2], we have that:

**Definition 2.1.** A *fluent formula* is a propositional formula built using the propositional variables in  $\mathcal{F}$  and the traditional propositional operators  $\wedge, \vee, \Rightarrow, \neg$ . We will use  $\top$  and  $\perp$  to indicate *True* and *False*, respectively. A *fluent atom* is a formula composed by just an element  $f \in \mathcal{F}$ , instead a *fluent literal* is either a fluent atom  $f \in \mathcal{F}$  or its negation  $\neg f$ . During this work we will refer to fluent literals simply as *fluents*.

**Definition 2.2.** A *belief formula* is defined as follow:

- A fluent formula is a belief formula;
- let  $\varphi$  be belief formula and  $ag \in \mathcal{AG}$ , then  $\mathbf{B}_{ag}\varphi$  is a belief formula;
- let  $\varphi_1, \varphi_2$  and  $\varphi_3$  be belief formulae, then  $\neg\varphi_3$  and  $\varphi_1 \text{ op } \varphi_2$  are belief formulae, where  $\text{op} \in \{\wedge, \vee, \Rightarrow\}$ ;
- all the formulae of the form  $\mathbf{E}_\alpha\varphi$  or  $\mathbf{C}_\alpha\varphi$  are belief formulae, where  $\varphi$  is itself a belief formula and  $\emptyset \neq \alpha \subseteq \mathcal{AG}$ .

From now on we will denote with  $\mathcal{L}_{\mathcal{AG}}^{\mathbf{C}}$  the language of the belief formulae over the sets  $\mathcal{F}$  and  $\mathcal{AG}$  and with  $\mathcal{L}_{\mathcal{AG}}$  as the language over beliefs formulae that does not allow the use of  $\mathbf{C}$ .

**Example 2.1.** Let us consider the formula  $\mathbf{B}_{ag_1}\mathbf{B}_{ag_2}\varphi$ . This formula expresses that the agent  $ag_1$  believes that the agent  $ag_2$  believes that  $\varphi$  is true, instead,  $\mathbf{B}_{ag_1}\neg\varphi$  expresses that the agent  $ag_1$  believes that  $\varphi$  is false.

As mentioned above the classical way of providing a semantics for the language of epistemic logic is in terms of *pointed Kripke structure* [22]. More formally:

**Definition 2.3.** A *Kripke structure* is a tuple  $\langle S, \pi, \mathcal{B}_1, \dots, \mathcal{B}_n \rangle$ , such that:

- $S$  is a set of worlds;
- $\pi : S \mapsto 2^{\mathcal{F}}$  is a function that associates an interpretation of  $\mathcal{F}$  to each element of  $S$ ;
- for  $1 \leq i \leq n$ ,  $\mathcal{B}_i \subseteq S \times S$  is a binary relation over  $S$ .

**Definition 2.4.** A *pointed Kripke structure* is a pair  $(M, s)$  where  $M$  is a Kripke structure as defined above, and  $s \in S$ , where  $s$  represents the real world.

Following the notation of [2], we will indicate with  $M[S], M[\pi]$ , and  $M[i]$  the components  $S, \pi$ , and  $\mathcal{B}_i$  of  $M$ , respectively. Intuitively  $M[S]$  captures all the worlds that the agents believe to be possible and  $M[i]$  encodes the beliefs of each agent. More formally the semantics on pointed Kripke structure is as follows:

**Definition 2.5.** Given the belief formulae  $\varphi, \varphi_1, \varphi_2$ , an agent  $ag_i$ , a group of agents  $\alpha$ , a pointed Kripke structure  $(M = \langle S, \pi, \mathcal{B}_1, \dots, \mathcal{B}_n \rangle, s)$ :

- (i)  $(M, s) \models \varphi$  if  $\varphi$  is a fluent formula and  $\pi(s) \models \varphi$ ;
- (ii)  $(M, s) \models \mathbf{B}_{ag_i}\varphi$  if for each  $t$  such that  $(s, t) \in \mathcal{B}_i$  it holds that  $(M, t) \models \varphi$ ;
- (iii)  $(M, s) \models \mathbf{E}_\alpha\varphi$  if  $(M, s) \models \mathbf{B}_{ag_i}\varphi$  for all  $ag_i \in \alpha$ ;
- (iv)  $(M, s) \models \mathbf{C}_\alpha\varphi$  if  $(M, s) \models \mathbf{E}_\alpha^k\varphi$  for every  $k \geq 0$ , where  $\mathbf{E}_\alpha^0\varphi = \varphi$  and  $\mathbf{E}_\alpha^{k+1}\varphi = \mathbf{E}_\alpha(\mathbf{E}_\alpha^k\varphi)$ ;
- (v) the semantics of the traditional propositional operators is as usual.

### 2.1.2 Knowledge or Belief

As pointed out in the previous paragraph the modal operator  $\mathbf{B}_{ag}$  represents the worlds' relation in a Kripke structure and, as expected, different relations' properties imply different meaning for  $\mathbf{B}_{ag}$ . In particular in this work we are interested in representing the knowledge or the beliefs of the agents. The problem of formalizing these two concepts has been studied in depth bringing to an accepted formalization for both knowledge and beliefs. In fact we have that when a relation<sup>1</sup> respects all the axioms presented in Table 1 is called an **S5** relation and it encodes the concept of knowledge while when it encodes all the axioms but **T** it characterizes the concept of belief. Following these characterization we will refer to knowledge and belief as **S5** and **KD45** logic respectively.

Property of $\mathcal{B}$	Axiom
$\mathcal{B}_i \varphi \Rightarrow \varphi$	<b>T</b>
$\mathcal{B}_i \varphi \Rightarrow \mathcal{B}_i \mathcal{B}_i \varphi$	<b>4</b>
$\neg \mathcal{B}_i \varphi \Rightarrow \mathcal{B}_i \neg \mathcal{B}_i \varphi$	<b>5</b>
$\neg \mathcal{B}_i \perp$	<b>D</b>
$(\mathcal{B}_i \varphi \wedge \mathcal{B}_i (\varphi \Rightarrow \psi)) \Rightarrow \mathcal{B}_i \psi$	<b>K</b>

Table 1: Knowledge and beliefs axioms. [13].

Intuitively the difference between the two logics is that an agent cannot *know* something that is not true in **S5** but she can *believe* it in **KD45**. As this introduction is not supposed to explore in depth this topic we will not go into further detail and we address the interested reader to [13].

### 2.1.3 Complexity Overview

Finally, as last note on DEL we will present two short, but informative, tables that summarize the complexity results in the epistemic logic and in the epistemic planning fields (Table 2a and Table 2b respectively). The information presented in these tables serve to provide the reader with a general idea on “how hard” the problem of reasoning on information change is.

SAT Complexity	Epistemic logic
NP-complete	<b>S5</b> <sub>1</sub> , <b>KD45</b> <sub>1</sub>
PSPACE-complete	<b>S5</b> <sub>n</sub> , <b>KD45</b> <sub>n</sub> with $n \geq 2$
EXPTIME-complete	<b>S5</b> <sup>C</sup> <sub>n</sub> , <b>KD45</b> <sup>C</sup> <sub>n</sub> with $n \geq 2$

(a) Complexity of the satisfiability problem w.r.t. knowledge and beliefs logics [13].

Action type	Plan Existence Complexity
Non factual with propositional preconditions	EXPSpace
Factual with propositional preconditions	NON-ELEMENTARY
Factual with epistemic preconditions	UNDECIDABLE

(b) Complexity of the *plan existence problem* [3].

Table 2: Complexity results in dynamic epistemic logic.

<sup>1</sup>In our case the relation between the world in a Kripke structure.

## 2.2 Multi-Agent Epistemic Planning

*Epistemic planning* [4] refers to the generation of plans for multiple agents to achieve goals which can refer to the state of the world, the beliefs of agents, and/or the knowledge of agents. It has recently attracted the attention of researchers from various communities such as planning, dynamic epistemic logic, and knowledge representation.

With the introduction of the classical planning problem, in the early days of artificial intelligence, several action languages (e.g.,  $\mathcal{A}$ ,  $\mathcal{B}$ , and  $\mathcal{C}$ ) have been developed [14] and have also provided the foundations for several successful approaches to automated planning. However, the main focus of these research efforts has been about reasoning within single agent domains. In single agent domains reasoning about information change mainly involves reasoning about what the agent knows about the world and how she can manipulate it to reach particular states. In multi-agent domains, on the other hand, an agent's action may change the world, other agents' knowledge about it and their knowledge about other agents' knowledge about the world. Similarly, goals of an agent in a multi-agent domain may involve manipulating the knowledge of other agents—in particular, this may concern not just their knowledge about the world, but also their knowledge about other agents' knowledge about the world.

As said before, epistemic planning is not only interested in the state of the world but also in the beliefs and the knowledge of the agents. Although there is a large body of research on multi-agent planning [1, 8, 9, 17, 18], very few efforts address the above mentioned aspects of multi-agent domains, which pose a number of new research challenges in representing and reasoning about actions and change.

Due to its complexity, the majority of search based epistemic planners (e.g., [7, 11, 19–21, 26, 34]) impose certain restrictions, such as the finiteness of the levels of nested beliefs. Such restrictions permit, for instance, in [19, 26, 34] to solve the problem by translating it into classical planning.

On the other hand, to the best of our knowledge, only few systems [23, 25] can reason about epistemic knowledge in multi-agent domains without these limitations. In particular [23] and the language  $m\mathcal{A}^*$ , that is the language that [23] implements, are the starting points of this research work. As we will explain in the following sections the main objectives of this thesis are: i) to formalize a complete and flexible epistemic action-based language; ii) to study alternative representations for epistemic models; iii) to implement a competitive planner to reason on information change; and iv) to explore the concept of heuristics in epistemic planning. Given the amount of information to properly introduce  $m\mathcal{A}^*$  and the concept of *event update semantics* we redirect the reader who is not familiar with these two topics to [2] for a clear and complete explanation.

## 3 Goals of the research

As already pointed in the previous sections our work tries to explore in depth the multi-agent epistemic planning problem. In particular this research aims to tackle the problem from several points of view, with the ultimate goal to provide a design (and an implementation) of an efficient *epistemic planner* that can reason on the full extent of  $\mathbf{S5}_n^C$  and  $\mathbf{KD45}_n^C$  with  $n \geq 1$ .

To reach this main objective we are investigating multi-agent epistemic planning from different levels. Following the the major sub-goals of our work will be presented.

- To formalize an **action-based language** that would be used to express epistemic domains. This new language will use [2] as base. Our design of a language will be mainly interested in dealing with: i) “*false beliefs*”, i.e., the situation when an agent has false beliefs about the world and tries to modify the world itself. Solutions to this problem may be found in literature, e.g. [23], but are somewhat

unsatisfactory. ii) We will also be interested in providing a formalization for the concepts of *trust*, *lies* and *deception*. iii) Finally, we would also like to increase the expressiveness of  $m\mathcal{A}^*$ , e.g., by adding the possibility to perform announcements of belief formulae.

- To study **alternative data structures** for epistemic states. This can be done on two different levels:
  - Introducing completely new data structures exploiting, for example, concepts from the widely studied knowledge representation [5] and graph theory [27] literature.
  - Another interesting approach, following what is done in the *model checking* community, is to represent the states through symbolic data structures. This will also imply the introduction of a transition function that can deal with symbolic state representation.
- As last, and maybe most substantial, contribution of the thesis we would like to design and implement a flexible and complete **epistemic solver**. Whilst the formalization of the language and the introduction of new data structures will surely help in reducing the search time, the inherent complexity of DEL will always be a strong limitation for epistemic planners. That is why our work will be focused on designing a solver that makes a strong use of heuristics. In fact, as shown in [23], reasoning on the full extent of  $\mathbf{KD45}_n^C$  at the moment is only possible when the length of the problems is really short (w.r.t. to the classical planning ones). That is why we believe that to provide a significant upgrade in performance it is necessary to introduce heuristics in this setting. As final note on the planner it is important to clarify that our solver could also be useful when exploited without heuristics. For example it could be used to reason on small domain where the information about the world is so intricate that would be impossible for humans to reason about it. Providing examples of the various scenarios where the planner could be used, even with different configurations, will also be taken into consideration during our research.

## 4 Status of the research

In this section we quickly present the work done on the goals introduced previously. Given the early stage of the research the results are often not completely demonstrated or fully tested; nevertheless we are positive to accomplish that in the next future. We will now provide, without tedious technical details, the status of the research of each one of the three sub-goals presented in Section 3; the reader who is interested in a more complete introduction to these is mainly addressed to [12, 23].

- **Language formalization:** we are currently investigating different lines of work that consider the description of a new epistemic action-based language from diverse points of view. First of all let us stress once again that this new language will be strongly based on  $m\mathcal{A}^*$  [2].
  - In respect to the “false belief” problem we sketched a solution that involves to apply some small changes to the transition function of  $m\mathcal{A}^*$ . Intuitively we think that by allowing epistemic actions (i.e., *sensing* and *announcement*) to modify not only the edges of the state but also its nodes could provide a clean solution to this problem.
  - On the other hand we introduced in  $m\mathcal{A}^*$  the concept of *trust* as a globally visible relation between agents. From this concept we formalized, thanks to the introduction of two new *events* in the announcement transition function, the *(un)trustworthy announcement* that takes into consideration the trust that the listener has w.r.t. the announcer.

- Starting from the concept of (un)trustworthy announcements we were also able to expand<sup>2</sup> the language from being able to characterize only fluent formulae announcement to being able to announce belief formulae as well.
- **Alternative data structures:** as already said Kripke structure is the classical way to provide semantics for epistemic logic but this does not mean that is the only one. In particular in [12] we introduced an extension of  $m\mathcal{A}^*$ , called  $m\mathcal{A}^P$ , that bases the state representation on *possibilities*, a data structure introduced in [15] and derived from *non-well-founded sets* theory. As shown in the paper, using these structures allows to have a more accurate concept of epistemic state equality. Moreover, being possibilities based on sets, we think that some aspects of epistemic planning, *e.g.*, the group operators, could be reduced to set operations providing a cleaner semantics and a faster implementation.
- **Epistemic Solver:** as main objective of our research we hope to provide a multi-agent epistemic planner flexible enough to be used in real world scenarios. As starting point we took the solver presented in [23] and we reformatted it. At this point of our study we have a modular C++ solver that:
  - encodes the language  $m\mathcal{A}^*$  and is as powerful as the ones presented in [23];
  - exploits the dynamic programming paradigm to encode more efficiently the search-space;
  - has a more clean transition function that prunes, when necessary, useless nodes of the Kripke structure to reduce the memory overhead;
  - allows parametric state representation. Thanks to this it is now only necessary to describe a new state representation, *e.g.*, possibilities, in the programming language and the solver will be able to reason on it. This feature will permit to have a planner with several options that could be used differently depending on the domain.
- **Heuristics:** with [23] we started to explore the concept of heuristics in epistemic planning. At this point of the our study we did not implement new heuristics yet but we are constantly exploring new options. Few informal ideas that we have had are:
  - to provide a normal form for belief formulae to introduce *mutex* in the *Epistemic Planning-Graph*;
  - to solve epistemic problems as an instance of classical planning (with limited nested knowledge) to score the e-states;
  - to use techniques of local search where, as parameters to optimize, we take into consideration the knowledge or the ignorance of a group of agents;
  - and finally, to exploit more classical approaches, *e.g.*, the Monte Carlo strategy, to try to reduce the search-space.

In this brief section we defined what we think are the main aspects of our research. Nevertheless there are other directions that our work can follows as epistemic planning is a relatively new field of study and needs to be explored from several points of view. In the following a short and summarizing list of alternative research directions is given.

- In [31] the computation of the *initial state* in **S5** is addressed but it is still an open question how to allow the same type of construction in **KD45**. We formalized a possible solution for this problem but is yet to be verified thoroughly.
- Due to their high computational power, ASP solvers could be exploited to solve the epistemic planning problem or at least some of its components (as in [32]).

---

<sup>2</sup>Correctness still to be proven.

- On the same line of the previous point there is the possibility to exploit well known techniques from the model checking field to deal with the common knowledge entailment. This operator is, in fact, strongly connected to the notion of *fixed point*.
- Finally, an important feature of classical planning problems that has not been formalized yet for the epistemic ones are *static laws*. Static laws could play a really important role from both the expressiveness and flexibility points of view.

## 5 Open issues

Multi-agent planning and DEL are widely studied areas but their combination is relatively new and less explored. In addition to that multi-agent epistemic planning, given its heavy formalism and inherent complexity, it is not a really “accessible” field of study and therefore not easy to be tackled. In the following we list the main problems that we are facing in working on this project.

- First of all a small non-technical issue: we noticed a substantial lack of real-world examples where epistemic planning can be adopted and that could help in responding to the question “why using epistemic planners when the classical ones are extremely faster?” We, ourselves, are trying to provide real-world scenarios that could exploit multi-agent epistemic planning but we recognize that experts from other fields, *e.g.*, law or economics, could be very helpful in formalizing more attractive examples.
- Another problem is the one relative to heuristics. As we are not dealing only with world’s properties we should include in our heuristics also the concepts of belief (or knowledge). The main problem that we are facing when considering heuristics is that is not always known what an epistemic planner seeks but, potentially, different combinations of ignorance and knowledge for each agent.
- From the action-language point of view the problems reside in the complexity of its formalism. While understanding every single detail of an epistemic planning language can be really challenging, implementing or changing it is even harder. This is reflected into a general difficulty of the topic that often discourages new researchers in investigating this field. A cleaner formalization of the epistemic planning problem could help the community to grow.
- One last problem is the subjectivity of some epistemic-related concepts. For example what is a lie? Or, again, when is an agent being deceptive? Formalizing these concepts requires rules that may change from person to person or from situation to situation.

## 6 Conclusions

In this paper we presented the main concepts regarding our ideas on how to develop a solver for multi-agent epistemic planning problems. We identified three major sub-goals that, in our opinion, are the ones that are the most significant for the epistemic planning community and we provided some research directions that we intend to follow in the next future. Moreover, we identified the major issues we are encountering while developing our ideas that could be used as starting points for related works that could bring benefits to the entire community.



## 7 Acknowledgements

The work described in this paper has been partially supported by the Universit di Udine PRID ENCASE project, and by GNCS-INdAM 2017 and 2019 projects.

## References

- [1] Martin Allen & Shlomo Zilberstein (2009): *Complexity of decentralized control: Special cases*. In: *Advances in Neural Information Processing Systems*, pp. 19–27.
- [2] Chitta Baral, Gregory Gelfond, Enrico Pontelli & Tran Cao Son (2015): *An Action Language for Multi-Agent Domains: Foundations*. CoRR abs/1511.01960. Available at <http://arxiv.org/abs/1511.01960>.
- [3] T. Bolander, M.H. Jensen & F. Schwarzentruher (2015): *Complexity results in epistemic planning*. In: *IJCAI International Joint Conference on Artificial Intelligence*, 2015-January, pp. 2791–2797.
- [4] Thomas Bolander & Mikkel Birkegaard Andersen (2011): *Epistemic planning for single-and multi-agent systems*. *Journal of Applied Non-Classical Logics* 21(1), pp. 9–34, doi:10.1016/0010-0277(83)90004-5.
- [5] Ronald J. Brachman & Hector J. Levesque, editors (1985): *Readings in Knowledge Representation*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- [6] Alexander Chagrov (1997): *Modal Logic*. Oxford University Press.
- [7] Matthew Crosby, Anders Jonsson & Michael Rovatsos (2014): *A single-agent approach to multiagent planning*. In: *Proceedings of the Twenty-first European Conference on Artificial Intelligence*, IOS Press, pp. 237–242.
- [8] Mathijs De Weerd, André Bos, Hans Tonino & Cees Witteveen (2003): *A resource logic for multi-agent plan merging*. *Annals of Mathematics and Artificial Intelligence* 37(1-2), pp. 93–130, doi:10.1023/A:1020236119243.
- [9] Mathijs De Weerd & Brad Clement (2009): *Introduction to planning in multiagent systems*. *Multiagent and Grid Systems* 5(4), pp. 345–355, doi:10.3233/MGS-2009-0133.
- [10] Agostino Dovier, Andrea Formisano & Enrico Pontelli (2013): *Autonomous agents coordination: Action languages meet CLP() and Linda*. *Theory and Practice of Logic Programming* 13(2), pp. 149–173, doi:10.1016/S0004-3702(00)00031-X.
- [11] Thorsten Engesser, Thomas Bolander, Robert Mattmüller & Bernhard Nebel (2017): *Cooperative epistemic multi-agent planning for implicit coordination*. *arXiv preprint arXiv:1703.02196*.
- [12] Francesco Fabiano, Idriss Riouak, Agostino Dovier & Enrico Pontelli (2019): *Non-well-founded set based multi-agent epistemic action language*. In: *Proceedings of the 34rd Italian Conference on Computational Logic, CEUR Workshop Proceedings 2396*, Trieste, Italy, pp. 242–259. Available at <http://ceur-ws.org/Vol-2396/paper38.pdf>.
- [13] Ronald Fagin & Joseph Y Halpern (1994): *Reasoning about knowledge and probability*. *Journal of the ACM (JACM)* 41(2), pp. 340–367, doi:10.1145/174652.174658.
- [14] Michael Gelfond & Vladimir Lifschitz (1998): *Action Languages*. *Electron. Trans. Artif. Intell.* 2, pp. 193–210. Available at <http://www.ep.liu.se/ej/etai/1998/007/>.
- [15] J. Gerbrandy & W. Groeneveld (1997): *Reasoning about information change*. *Journal of Logic, Language and Information* 6(2), pp. 147–169, doi:10.1023/A:1008222603071.
- [16] Nicola Gigante (2019): *Timeline-based planning: Expressiveness and Complexity*. CoRR abs/1902.06123.
- [17] Claudia V Goldman & Shlomo Zilberstein (2004): *Decentralized control of cooperative systems: Categorization and complexity analysis*. *J. Artif. Intell. Res. (JAIR)* 22, pp. 143–174, doi:10.1613/jair.1427.
- [18] Carlos Guestrin, Daphne Koller & Ronald Parr (2002): *Multiagent planning with factored MDPs*. In: *Advances in neural information processing systems*, pp. 1523–1530.

- [19] X. Huang, B. Fang, H. Wan & Y. Liu (2017): *A general multi-agent epistemic planner based on higher-order belief change*. In: *IJCAI International Joint Conference on Artificial Intelligence*, pp. 1093–1101.
- [20] Filippos Kominis & Hector Geffner (2015): *Beliefs In Multiagent Planning: From One Agent to Many*. In: *Proceedings of the International Conference on Automated Planning and Scheduling, ICAPS*, pp. 147–155.
- [21] Filippos Kominis & Hector Geffner (2017): *Multiagent Online Planning with Nested Beliefs and Dialogue*. In: *Proceedings of the International Conference on Automated Planning and Scheduling, ICAPS*, Pittsburgh, Pennsylvania, USA, pp. 186–194. Available at <https://aaai.org/ocs/index.php/ICAPS/ICAPS17/paper/view/15748>.
- [22] Saul A. Kripke (1963): *Semantical Considerations on Modal Logic*. *Acta Philosophica Fennica* 16(1963), pp. 83–94.
- [23] Tiep Le, Francesco Fabiano, Tran Cao Son & Enrico Pontelli (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*, AAAI Press, Delft, The Netherlands, pp. 161–170. Available at <https://aaai.org/ocs/index.php/ICAPS/ICAPS18/paper/view/17733>.
- [24] Nir Lipovetzky & Hector Geffner (2017): *Best-First Width Search: Exploration and Exploitation in Classical Planning*. In: *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, San Francisco, California, USA, pp. 3590–3596. Available at <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14862>.
- [25] Qiang Liu & Yongmei Liu (2018): *Multi-agent Epistemic Planning with Common Knowledge*. In: *IJCAI International Joint Conference on Artificial Intelligence*, pp. 1912–1920.
- [26] Christian J. Muise, Vaishak Belle, Paolo Felli, Sheila A. McIlraith, Tim Miller, Adrian R. Pearce & Liz Sonenberg (2015): *Planning Over Multi-Agent Epistemic States: A Classical Planning Approach*. In: *Proc. of AAAI*, pp. 3327–3334.
- [27] E.G. Omodeo, A. Policriti & A.I. Tomescu (2017): *On sets and graphs: Perspectives on logic and combinatorics*. Springer, doi:10.1007/978-3-319-54981-1.
- [28] Silvia Richter & Matthias Westphal (2010): *The LAMA planner: Guiding cost-based anytime planning with landmarks*. *Journal of Artificial Intelligence Research* 39, pp. 127–177, doi:10.1613/jair.2972.
- [29] Stuart Russell & Peter Norvig (2009): *Artificial Intelligence: A Modern Approach*, 3rd edition. Prentice Hall Press, Upper Saddle River, NJ, USA.
- [30] Raymond R Smullyan (2012): *First-order logic*. 43, Springer Science & Business Media.
- [31] Tran Cao Son, Enrico Pontelli, Chitta Baral & Gregory Gelfond (2014): *Finitary S5-theories*. In: *European Workshop on Logics in Artificial Intelligence*, Springer, pp. 239–252, doi:10.1093/jigpal/jzm059.
- [32] Tran Cao Son, Enrico Pontelli, Michael Gelfond & Marcello Balduccini (2016): *An Answer Set Programming Framework for Reasoning About Truthfulness of Statements by Agents*. In: *Technical Communications of the 32nd International Conference on Logic Programming, ICLP 2016 TCS, October 16-21, 2016, New York City, USA*, pp. 8:1–8:4, doi:10.4230/OASlcs.ICLP.2016.8.
- [33] Hans Van Ditmarsch, Wiebe van Der Hoek & Barteld Kooi (2007): *Dynamic epistemic logic*. 337, Springer Science & Business Media.
- [34] Hai Wan, Rui Yang, Liangda Fang, Yongmei Liu & Huada Xu (2015): *A Complete Epistemic Planner without the Epistemic Closed World Assumption*. In: *IJCAI International Joint Conference on Artificial Intelligence*, Buenos Aires, Argentina, pp. 3257–3263. Available at <http://ijcai.org/Abstract/15/459>.