

A Semantic Approach to Decidability in Epistemic Planning

Extended Abstract

Alessandro Burigana

Free University of Bozen-Bolzano
Bolzano, Italy
burigana@inf.unibz.it

Marco Montali

Free University of Bozen-Bolzano
Bolzano, Italy
montali@inf.unibz.it

Paolo Felli

University of Bologna
Bologna, Italy
paolo.felli@unibo.it

Nicolas Troquard

Free University of Bozen-Bolzano
Bolzano, Italy
nicolas.troquard@unibz.it

ABSTRACT

Dynamic Epistemic Logic (DEL) provides a very rich planning formalism that can handle nondeterminism, partial observability and arbitrary knowledge nesting. The general framework is notoriously undecidable. In this paper, we pursue a novel *semantic approach* to achieve decidability, by focussing on the logic for epistemic planning, rather than to limit the syntax of the accepted modal formulae. Specifically, we augment the logic $S5_n$ by introducing a new interaction axiom that we call *knowledge alignment*, in order to control the ability of agents to unboundedly reason on the knowledge of other agents. We show that the resulting epistemic planning problem is decidable. In doing so, we prove that this framework admits a finitary non-fixpoint characterization of common knowledge, which is of independent interest.

KEYWORDS

Epistemic Planning; Dynamic Epistemic Logic; Decidability

ACM Reference Format:

Alessandro Burigana, Paolo Felli, Marco Montali, and Nicolas Troquard. 2023. A Semantic Approach to Decidability in Epistemic Planning: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 3 pages.

1 INTRODUCTION

Multi-agent systems find applications in a wide range of settings where the agents need to be able to reason about both the physical world and the *knowledge* that other agents possess—that is, their *epistemic state*. *Epistemic planning* [2] employs the theoretical framework of Dynamic Epistemic Logic (DEL) [12] in the context of automated planning. The resulting formalism is able to represent nondeterminism, partial observability and arbitrary knowledge nesting. That is, agents have the power to reason about higher-order knowledge of other agents with no limitations.

Due to the high expressive power of the DEL framework, the *plan existence problem* [1], that asks whether there exists a plan to achieve a goal of interest, is undecidable in general [2]. One of the critical aspects that leads to undecidability is the reasoning power

of agents. In fact, in the logic $S5_n$ there is no rule or principle that describes how the knowledge of one agent should interact with the knowledge of another agent. Hence, there is no restriction on the ability of agents to reason about the higher-order knowledge they possess about the knowledge of others.

To tackle this problem, a common approach consists in syntactically restricting the action theory, for instance by limiting the modal depth of the pre- and postconditions of actions to a certain bound d [3–5]. Nonetheless, the problem remains undecidable even with $d=2$ when only *purely epistemic actions* are allowed, and with $d=1$ when factual change is involved. This suggests that syntactic restrictions of the action theory are too strong in many practical cases, where reasoning about the knowledge of others is required.

For this reason, in this paper we pursue a different strategy that we call *semantic approach*. Specifically, we consider the multi-agent logic for knowledge $S5_n$ and we augment it with a novel interaction axiom, called the (*knowledge*) *alignment* axiom. Such axiom ensures that the epistemic states in such augmented logic are bounded in size. This, in turn, guarantees that the size of the state search space is finite and that the plan existence problem is decidable.

In what follows, we introduce and discuss the knowledge alignment axiom and we summarize our main (un)decidability results.

2 EPISTEMIC PLANNING AND ALIGNMENT

Before introducing our axiom, we briefly recall the syntax of epistemic logic (see [12] for a complete introduction). We fix a finite set of agents $\mathcal{AG} = \{1, \dots, n\}$ and a finite set of atomic propositions \mathcal{P} . We consider the following language of epistemic logic:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid \Box_i\varphi \mid C_G\varphi,$$

where $p \in \mathcal{P}$, $i \in \mathcal{AG}$, and $\emptyset \neq G \subseteq \mathcal{AG}$. Formulae $\Box_i\varphi$ and $C_G\varphi$ are respectively read as “agent i knows that φ ” and “group G has common knowledge that φ ”.

We are now ready to introduce our axiom:

$$\mathbf{A} \quad \Box_i\Box_j\varphi \rightarrow \Box_i\Box_j\Box_i\varphi \quad (\text{Alignment})$$

We augment the logic $S5_n$ with axiom **A** and we call the resulting logic $A\text{-}S5_n$. We can read **A** as follows: whenever an agent i knows that some other agent j knows that φ , then i also knows that j knows that i too knows that φ . Thus, intuitively, axiom **A** defines a principle of *alignment* in the higher-order knowledge across agents. In other words, the *misalignment* between their knowledge is possible only up to a certain modal depth.

When we introduce a certain principle to be an *axiom* of our logic, we are also implicitly considering an epistemic state to be meaningful if and only if such principle is satisfied. Thus, when planning under our logic, we consider a plan to be valid only if all the states it visits satisfy not only the axioms of $S5_n$, but also axiom **A**. At the same time, it is not guaranteed that the application of an action in a $A-S5_n$ -state actually results in an epistemic state that satisfies **A**. This is not unusual in epistemic planning and in fact applies to the well-studied doxastic logic $KD45_n$ [6], where axiom **D** is not guaranteed to be preserved. In the literature, different techniques for the preservation of axiom **D** are studied [7, 10]. The development of such techniques for a logic is independent of the analysis of decidability of the plan existence problem under that logic and, for this reason, it is left as future work. In the paper we adopt a rollback-style approach to reject invalid plans, thus visiting only meaningful epistemic states. Our decidability results continue to hold even when one adopts more sophisticated revision approaches that accept and suitably curate inconsistent states.

In $A-S5_n$, while agents have their own distinct individual knowledge, higher-order levels of perspectives of agents are aligned. This assumption is well suited in cooperative planning domains [11], in which agents are able to maintain some alignment in their knowledge, for instance when modelling homogenous agents that received common training (e.g., firefighters or rescue teams).

To illustrate, consider the following example, where i, j and k are agents and, as customary, we use $\Box_i\varphi$ to indicate that agent i knows that φ holds. Suppose that agent i knows whether a certain proposition p holds, i.e., $Kw(i, p) \equiv \Box_i p \vee \Box_i \neg p$, and that agent j does not know whether p holds, i.e., $\neg Kw(j, p) \equiv \neg \Box_j p \wedge \neg \Box_j \neg p$. That is, the individual knowledge of agents i and j is *not* aligned. This is allowed by the alignment axiom. Nonetheless, if j knows that k knows that i knows whether p (i.e., $\Box_j \Box_k Kw(i, p)$), then, by virtue of the alignment axiom, it can not be the case that j knows that k knows that j knows that i does not know whether p (i.e., $\Box_j \Box_k \Box_j \neg Kw(i, p)$). Therefore, the higher-order knowledge across j and k (about i 's perspective) is aligned. Hence, agents are no longer able to consider unboundedly nested perspectives on the knowledge of others.

Informally, the limited reasoning power of agents affects the size of states in the logic $A-S5_n$. Indeed, we are able to prove that $A-S5_n$ -states are bounded in size, which, in turn, ensures that the search space of the plan existence problem is finite.

Consequently, we obtain the following contribution:

THEOREM 2.1. *For any $n > 1$, the plan existence problem in the logic $A-S5_n$ is DECIDABLE.*

Additionally, the alignment axiom has important consequences on properties of common knowledge. Indeed, we show that the logic $A-S5_n$ admits a *finitary non-fixpoint* characterization of common knowledge, which is often regarded as a possible solution to paradoxes involving common knowledge (see [9] for an overview). Specifically, we prove the following result:

THEOREM 2.2. *Let $G = \{i_1, \dots, i_m\} \subseteq \mathcal{AG}$, with $m \geq 2$. In $A-S5_n$, for any φ , the formula $\Box_{i_1} \dots \Box_{i_m} \varphi \leftrightarrow C_G \varphi$ is a theorem.*

Although the alignment axiom is better fitting for tight-knit groups of agents, it may be less suited for representing more loosely organized groups. Thus, we investigate the plan existence problem

Logic	Decidability
$K_n, KT_n, K4_n, K45_n, S4_n, S5_n$	UNDECIDABLE [1]
$A^b-S5_n (n > 2)$	UNDECIDABLE
A^b-S5_2	DECIDABLE
$A-S5_n$	DECIDABLE

Table 1: Decidability results of plan existence problem based on the semantic approach, compared to our results (in gray).

with a weaker principle of alignment, which is parametrized by an integer $b > 1$. The resulting axiom is the following:

$$A^b \quad (\Box_i \Box_j)^b \varphi \rightarrow (\Box_i \Box_j)^b \Box_i \varphi \quad (\text{Weak alignment})$$

One could hope that the plan existence problem remains decidable when replacing axiom **A** with A^b . But this is not true in general. Namely, we show that the plan existence problem remains decidable in the presence of two agents. We also show that for $n > 2$ the problem becomes undecidable, thus establishing the frontier of decidability using the semantic approach. Formally:

THEOREM 2.3. *For any $b > 1$, the plan existence problem in the logic $A-S5_2^b$ is DECIDABLE.*

THEOREM 2.4. *For any $n > 2$ and $b > 1$, the plan existence problem in the logic $A-S5_n^b$ is UNDECIDABLE.*

To prove Theorem 2.3, we follow the same idea as in Theorem 2.1, namely we show that $A-S5_2^b$ -states are bounded in size. Finally, to prove Theorem 2.4, we appeal to a reduction to the halting problem for two-counter machines [8], similar to the one developed by Aucher and Bolander [1]. We summarize our results and we compare them with the literature in Table 1.

3 DISCUSSION

The paper presents a novel decidability result in epistemic planning. The approach adopted in this work deviates from previous ones, where syntactical conditions are imposed on actions. In particular, we pursue a novel semantic approach and we introduce a principle of knowledge alignment that is well suited for cooperative multi-agent planning contexts, by which we govern the extent to which agents can reason about the knowledge of their peers. This results in a boundedness property of the size of epistemic states, which in turn guarantees that the search space is finite. Additionally, we show that the alignment axiom leads to a finitary non-fixpoint characterization of common knowledge.

In the future, we plan to further exploit the axiom-based approach by formulating other properties to add to the logic of knowledge. Moreover, we are interested into moving from the domain of knowledge to the one of *belief*. This is not a trivial task, as the results of this paper do not readily apply in the logic $KD45_n$. Further, we plan on determining the computational complexity of the plan existence problem in $A-S5_n$ and to compare it with the current results in the literature. Finally, we plan on determining sufficient conditions for preserving axiom **A** during the planning process. This is similar to what is done in [10] for the preservation of axioms of the logic $KD45_n$.

REFERENCES

- [1] Guillaume Aucher and Thomas Bolander. 2013. Undecidability in Epistemic Planning. In *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, Francesca Rossi (Ed.), IJCAI/AAAI, 27–33.
- [2] Thomas Bolander and Mikkel Birkegaard Andersen. 2011. Epistemic planning for single and multi-agent systems. *J. Appl. Non Class. Logics* 21, 1 (2011), 9–34.
- [3] Thomas Bolander, Tristan Charrier, Sophie Pinchinat, and François Schwarzentruber. 2020. DEL-based epistemic planning: Decidability and complexity. *Artif. Intell.* 287 (2020), 103304.
- [4] Thomas Bolander, Martin Holm Jensen, and François Schwarzentruber. 2015. Complexity Results in Epistemic Planning. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, Qiang Yang and Michael J. Wooldridge (Eds.). AAAI Press, 2791–2797.
- [5] Tristan Charrier, Bastien Maubert, and François Schwarzentruber. 2016. On the Impact of Modal Depth in Epistemic Planning. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, Subbarao Kambhampati (Ed.). IJCAI/AAAI Press, 1030–1036.
- [6] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. 2004. *Reasoning About Knowledge*. MIT Press.
- [7] Andreas Herzig, Jérôme Lang, and Pierre Marquis. 2005. Action progression and revision in multiagent belief structures. In *Sixth Workshop on Nonmonotonic Reasoning, Action, and Change (NRAC)*.
- [8] Marvin L Minsky. 1967. *Computation: finite and infinite machines*. Prentice-Hall, Inc.
- [9] Cédric Paternotte. 2011. Being realistic about common knowledge: a Lewisian approach. *Synth.* 183, 2 (2011), 249–276.
- [10] Tran Cao Son, Enrico Pontelli, Chitta Baral, and Gregory Gelfond. 2015. Exploring the KD45 Property of a Kripke Model After the Execution of an Action Sequence. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, Blai Bonet and Sven Koenig (Eds.). AAAI Press, 1604–1610.
- [11] Alejandro Torreño, Eva Onaindia, Antonín Komenda, and Michal Stolba. 2018. Cooperative Multi-Agent Planning: A Survey. *ACM Comput. Surv.* 50, 6 (2018), 84:1–84:32.
- [12] Hans P. van Ditmarsch, Wiebe van der Hoek, and Barteld P. Kooi. 2007. *Dynamic Epistemic Logic*. Vol. 337. Springer Netherlands.