ECAI 2023
K. Gal et al. (Eds.)

doi:10.3233/FAIA230286

A Semantic Approach to Decidability in Epistemic Planning

Alessandro Burigana[a;*], **Paolo Felli**[b], **Marco Montali**[a] **and Nicolas Troquard**[a]

[a]Free University of Bozen-Bolzano
[b]University of Bologna

Abstract. The use of Dynamic Epistemic Logic (DEL) in multi-agent planning has led to a widely adopted action formalism that can handle nondeterminism, partial observability and arbitrary knowledge nesting. As such expressive power comes at the cost of undecidability, several decidable fragments have been isolated, mainly based on syntactic restrictions of the action formalism. In this paper, we pursue a novel *semantic approach* to achieve decidability. Namely, rather than imposing syntactical constraints, the semantic approach focuses on the axioms of the logic for epistemic planning. Specifically, we augment the logic of knowledge $\mathrm{S5}_n$ and with an interaction axiom called *(knowledge) commutativity*, which controls the ability of agents to unboundedly reason on the knowledge of other agents. We then provide a threefold contribution. First, we show that the resulting epistemic planning problem is decidable. In doing so, we prove that our framework admits a finitary non-fixpoint characterization of common knowledge, which is of independent interest. Second, we study different generalizations of the commutativity axiom, with the goal of obtaining decidability for more expressive fragments of DEL. Finally, we show that two well-known epistemic planning systems based on action templates, when interpreted under the setting of knowledge, conform to the commutativity axiom, hence proving their decidability.

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* [5] employs the theoretical framework of Dynamic Epistemic Logic (DEL) [22] 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* (see Definition 9), that asks whether there exists a plan to achieve a goal of interest, is undecidable in general [5]. As a consequence, in the past decade, DEL has been widely studied to obtain (un)decidability and complexity results for fragments of the planning problem. A common approach (see Section 6) consists in syntactically restricting the action theory, for instance by limiting the modal depth of the preconditions and postconditions of actions to a certain bound d [8, 9, 6]. 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 such syntactic restrictions 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*. Namely, rather than imposing syntactical constraints, the semantic approach focuses on the axioms of the logic for epistemic planning. Specifically, we consider the multi-agent logic for knowledge $\mathrm{S5}_n$ (where n denotes the number of agents) and we augment it with an interaction axiom, called the *(knowledge) commutativity* axiom (where, as customary, $\Box_i \varphi$ indicates that agent i *knows* that φ holds):

$$\mathbf{C} \quad \Box_i \Box_j \varphi \to \Box_j \Box_i \varphi \quad \text{(Commutativity)}$$

This axiom imposes a principle of commutativity in the higher-order knowledge across agents. In the resulting logic, which we call C-$\mathrm{S5}_n$, while agents have their own distinct individual knowledge, higher-order levels of perspectives of agents *commute*. This assumption is well suited in cooperative planning domains [21], where it is required that agents act and communicate in an observable way, thus making knowledge of agents accessible to others.

We provide a threefold contribution. First, we show that the epistemic plan existence problem in the resulting framework becomes decidable. We do so by proving that the commutativity axiom ensures that the states in the logic C-$\mathrm{S5}_n$ are bounded in size, which entails that the search space of the plan existence problem is finite. In doing so, we show that the logic C-$\mathrm{S5}_n$ admits a *finitary non-fixpoint* characterization of common knowledge, which is often regarded as a possible solution to paradoxes involving common knowledge (see [19] for an overview).

Second, we investigate the plan existence problem with different generalized principles of commutativity. Indeed, although the commutativity axiom is better fitting for tight-knit groups of agents, it may be less suited for representing more loosely organized groups. We define suitable generalizations parametrized by fixed integer constants $b{>}1$ and $1{<}\ell{\leq}n$. The resulting axioms are the following (where π is a permutation of the sequence $\langle i_1, \dots i_\ell \rangle$ of agents, as explained in more detail in Section 4.1):

$$\begin{array}{lll} \mathbf{C}^b & (\Box_i \Box_j)^b \varphi \to (\Box_j \Box_i)^b \varphi & (b\text{-Comm.}) \\ \mathbf{wC}_\ell & \Box_{i_1} \dots \Box_{i_\ell} \varphi \to \Box_{\pi_{i_1}} \dots \Box_{\pi_{i_\ell}} \varphi & \text{(Weak comm.)} \end{array}$$

Concerning axiom $\mathbf{wC}_\ell$, we show that the plan existence problem

* Corresponding Author. Email: burigana@inf.unibz.it.

remains decidable for any $1{<}\ell{\leq}n$. Relating axiom $\mathbf{C}^b$, we show that the plan existence problem remains decidable in the presence of two agents ($n{=}2$), for any $b{>}1$. We also show that for any $n{>}2$ and any $b{>}1$ the problem becomes undecidable.

Finally, we show that the knowledge (*i.e.*, $\mathrm{S5}_n$) fragment of the well known planning system $m\mathcal{A}^*$ [3] and the system by Kominis and Geffner [16] are captured by our formalism. Thus, we prove the decidability of a fragment $m\mathcal{A}^*$, which was still an open problem, and of the action formalism in [16], confirming their previous results.

Since the axioms of the logic for epistemic planning lie at the core of the semantic approach, we consider such axioms to define *meaningful states*. In other words, when a certain principle is introduced to be an *axiom* of the logic, an epistemic state is considered to be *meaningful* if and only if such principle is satisfied. Thus, when planning under a logic L, we consider a plan to be meaningful and, in turn, valid only if all the states that it visits satisfy the axioms of L. At the same time, the semantics of the product update of DEL (see Definition 6) does not guarantee that the application of an action of a generic logic L to an epistemic state of the same logic necessarily results in an epistemic state that satisfies all axioms of L. A well-known example of this phenomenon is found in the widely studied doxastic logic $\mathrm{KD45}_n$ [10], where the *consistency axiom* **D** is not guaranteed to be preserved by the product update. Addressing the problem of preservation of axioms after action updates is not trivial. Indeed, in the literature, considerable effort has been spent in developing different techniques to handle the preservation of axiom **D** [14, 20]. Analogously to the case of **D** in $\mathrm{KD45}_n$, in our framework axioms $\mathbf{C}$, $\mathbf{C}^b$ and $\mathbf{wC}_\ell$ are not guaranteed to be preserved by the product update. As a result, as explained above, we consider a plan to be valid if it only visits meaningful epistemic states, *i.e.*, those satisfying the axioms of the considered logic. Importantly, our decidability results continue to hold even when one adopts more sophisticated revision techniques that handle the preservation problem by accepting and suitably curating non-preserving states. The development of such non-trivial techniques for our logics is independent of the analysis of decidability of the plan existence problem under the same logics, and is left as an important, follow-up work.

The paper is organised as follows. In Section 2, we recall some preliminaries and define epistemic planning tasks. In Section 3, we discuss in more detail the semantic approach, we introduce our new logic for epistemic planning and we discuss the commutativity axiom. In Section 4, we analyze decidability of epistemic planning under commutativity and its generalizations. In Section 5, we apply our decidability results to existing epistemic planning systems. Finally, in Section 6, we discuss related work.

2 Dynamic Epistemic Logic

This section is organized as follows. The syntax and semantics of DEL [22] are introduced in Section 2.1, event models and the product update in Section 2.2. In Section 2.3, we recall the axioms of the logic $\mathrm{S5}_n$. In Section 2.4 we define the plan existence problem.

2.1 Epistemic Models

Let $\mathcal{P}$ be a finite set of atomic propositions and $\mathcal{AG} = \{1, \dots, n\}$ a finite set of agents. The language $\mathcal{L}^C_{\mathcal{P},\mathcal{AG}}$ of *multi-agent epistemic logic on $\mathcal{P}$ and $\mathcal{AG}$ with common knowledge* is defined by the following BNF:

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

where $p \in \mathcal{P}$, $i \in \mathcal{AG}$, and $\varnothing \neq G \subseteq \mathcal{AG}$. 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 define $\top$, $\bot$, $\vee$, $\rightarrow$ and $\Diamond_i$ as usual.

Definition 1 (Epistemic Model and State)**.** *An* epistemic model *of $\mathcal{L}^C_{\mathcal{P},\mathcal{AG}}$ is a triple $M = (W, R, V)$ where:*
- *$W \neq \varnothing$ is a finite set of possible worlds;*
- *$R : \mathcal{AG} \rightarrow 2^{W \times W}$ assigns to each agent i an accessibility relation $R(i)$ (abbreviated as R_i);*
- *$V : \mathcal{P} \rightarrow 2^W$ assigns to each atom a set of worlds.*

An epistemic state *is a pair (M, W_d), where $W_d \subseteq W$ is a nonempty set of designated worlds.*

Intuitively, a designated world in W_d is considered the current "real" world from the perspective of an external observer (the planner) rather than of agents in $\mathcal{AG}$. Thus, $|W_d| > 1$ represents the uncertainty of the observer about the real world.

The pair (W, R) is called the *frame* of M. We use the infix notation wR_iv in place of $(w, v) \in R_i$. We also define $R_G \doteq \cup_{i \in G} R_i$, where $G \subseteq \mathcal{AG}$. The reflexive and transitive closure of R is denoted by R^*. Relations R_i capture what agents consider to be possible: wR_iv denotes the fact that, in w, agent i considers v to be possible. Throughout the paper, to support our exposition, we consider the example of the *coordinated attack problem* [12, 13]. It is a well-known problem that is often analyzed in the distributed systems literature. In what follows, we appeal to the DEL representation of this problem provided in [6].

Example 1 (The Coordinated Attack Problem)**.** *Two generals,* **a** *and* **b***, are camped with their armies on two hilltops overlooking a common valley, where the enemy is stationed. The only way for them to defeat the enemy is to attack simultaneously. They can only communicate by means of a messenger, who may be captured at any time when crossing the valley. Neither general will attack until he is sure that the other will attack as well.*

General **a** *and the messenger are initially together, and general* **a** *decides to attack at dawn. We use the atomic propositions* d *to denote that* 'general **a** will attack at dawn' *and m_i, for $i = \mathbf{a}, \mathbf{b}$, to denote that the messenger is currently at the camp of general i. In this way, $\neg m_a \wedge \neg m_b$ expresses the fact that the messenger has been captured.*

The initial situation can be described by the epistemic state s_0 shown below[1]*. Each bullet represents a world and the designated world is denoted by a circled bullet. There are two possible worlds, denoting the possibility that general* **a** *will attack at dawn (w_1), or not (w_2). Both generals know that the messenger is camped with general* **a***. The fact that general* **b** *does not know whether general* **a** *has decided to attack is represented by the indistinguishability relations between worlds w_1 and w_2. In fact, initially, general* **b** *has not enough information to know whether his ally has decided to attack.*

$$s_0 \;=\; \underset{w_1 : d, m_a}{\odot} \overset{b}{\rule[0.5ex]{4em}{0.4pt}} \underset{w_2 : m_a}{\bullet}$$

Definition 2 (Truth in epistemic states)**.** *Let $M = (W, R, V)$ be an epistemic model, $w \in W$, $i \in \mathcal{AG}$, $\varnothing \neq G \subseteq \mathcal{AG}$, $p \in \mathcal{P}$ and $\varphi, \psi \in \mathcal{L}^C_{\mathcal{P},\mathcal{AG}}$ be two formulae. Then,*

$(M, w) \models p$	*iff*	$w \in V(p)$
$(M, w) \models \neg\varphi$	*iff*	$(M, w) \not\models \varphi$
$(M, w) \models \varphi \wedge \psi$	*iff*	$(M, w) \models \varphi$ *and* $(M, w) \models \psi$
$(M, w) \models \Box_i\varphi$	*iff*	$\forall v$ *if* wR_iv *then* $(M, v) \models \varphi$
$(M, w) \models C_G\varphi$	*iff*	$\forall v$ *if* wR^*_Gv *then* $(M, v) \models \varphi$

[1] In all figures, the reflexive, symmetric and transitive closures of the relations are left implicit.

Let (M, W_d) be an epistemic state. Then,

$$(M, W_d) \models \varphi \quad \textit{iff} \quad (M, w) \models \varphi \text{ for all } w \in W_d$$

For instance, $(M, w) \models p$ means that p is true in w; $(M, w) \models \Diamond_i p$ means that in w the agent i admits the possibility of φ being true, i.e., there exists a world v that i considers possible (i.e., with wR_iv) such that $(M, v) \models \varphi$; $(M, w) \models \Box_i\varphi$ means that in w the agent i knows φ, as φ holds in all worlds that i considers possible.

We recall the notion of bisimulation for epistemic states [7].

Definition 3 (Bisimulation)**.** *Let $s = ((W, R, V), W_d)$ and $s' = ((W', R', V'), W'_d)$ be two epistemic states. We say that s and s' are* bisimilar*, denoted by $s \underline{\leftrightarrow} s'$, if there exists non-empty binary relation $Z \subseteq W \times W'$ satisfying:*

- Atoms*: if $(w, w') \in Z$, then for all $p \in \mathcal{P}$, $w \in V(p)$ iff $w' \in V'(p)$.*
- Forth*: if $(w, w') \in Z$ and wR_iv, then there exists $v' \in W'$ such that $w'R'_iv'$ and $(v, v') \in Z$.*
- Back*: if $(w, w') \in Z$ and $w'R'_iv'$, then there exists $v \in W$ such that wR_iv and $(v, v') \in Z$.*
- Designated*: if $w \in W_d$, then there exists $w' \in W'_d$ such that $(w, w') \in Z$, and vice versa.*

We say that Z is a bisimulation *between s and s'.*

Throughout the rest of the paper, we assume that *each epistemic state is minimal w.r.t. bisimulation*. We denote the fact that $(w, w') \in Z$ by $w \underline{\leftrightarrow} w'$. Finally, we introduce a notion of k-bisimulation for epistemic states. The following definition follows the one in [4] and generalizes that of [23] by considering epistemic states with (possibly) multiple designated worlds.

Definition 4 (k-bisimulation)**.** *Let $k \geq 0$ and let $s = ((W, R, V), W_d)$ and $s' = ((W', R', V'), W'_d)$ be two epistemic states. We say that s and s' are k-bisimilar, denoted by $s \underline{\leftrightarrow}_k s'$, if there exists a sequence of non-empty binary relations $Z_k \subseteq \ldots \subseteq Z_0$ (with $Z_0 \subseteq W \times W'$) satisfying (for any $i < k$):*

- Atoms*: if $(v, v') \in Z_0$, then for all $p \in \mathcal{P}$, $v \in V(p)$ iff $v' \in V'(p)$.*
- Forth*: if $(v, v') \in Z_{i+1}$ and vR_iu, then there exists $u' \in W'$ such that $v'R'_iu'$ and $(u, u') \in Z_i$.*
- Back*: if $(v, v') \in Z_{i+1}$ and $v'R'_iu'$, then there exists $u \in W$ such that vR_iu and $(u, u') \in Z_i$.*
- Designated*: if $v \in W_d$, then there exists $v' \in W'_d$ such that $(v, v') \in Z_k$, and vice versa.*

*We say that Z_k is a k-*bisimulation *between s and s'.*

2.2 Event Models and Product Update

Information change is captured by *product updates* of the current epistemic state with the *event model* of actions.

Definition 5 (Event Model and Action)**.** *An* event model *for $\mathcal{L}^C_{\mathcal{P},\mathcal{AG}}$ is a tuple $\mathcal{E} = (E, Q, pre, post)$ where:*

- *$E \neq \varnothing$ is a finite set of events;*
- *$Q : \mathcal{AG} \to 2^{E \times E}$ assigns to each agent i an accessibility relation $Q(i)$ (abbreviated as Q_i);*
- *$pre : E \to \mathcal{L}^C_{\mathcal{P},\mathcal{AG}}$ assigns to each event a* precondition*;*
- *$post : E \to (\mathcal{P} \to \mathcal{L}^C_{\mathcal{P},\mathcal{AG}})$ assigns to each event and atom a* postcondition*.*

An action *is a pair $(\mathcal{E}, E_d)$ where $E_d \subseteq E$ is a non-empty set of designated events.*

Similarly to epistemic states, the designated events in E_d represent the "real" events that are taking place from the perspective of an external observer.

The pair (E, Q) is called the *frame* of $\mathcal{E}$. We use the infix notation eQ_if in place of $(e, f) \in Q_i$. These relations are analogous to the accessibility relations of epistemic models: they are used to specify how the knowledge of each agent is affected by an action, depending on which events each agent considers possible. Intuitively, the precondition of an event e specify whether e *could happen* in a certain world w, whereas the postconditions of e describe how such event might change the factual properties of w (see Definition 6). Formally, we say that an event e is *applicable* in a world w of M if $(M, w) \models pre(e)$.

Example 2. *Imagine that general* **a** *decides to send the messenger to general* **b** *(action $send_{ab}$). While doing so, the general considers two possible outcomes: 1) the messenger safely arrives to the other side of the valley, or 2) the messenger is captured by the enemy. In the figure below, these eventualities are represented by events e^a_1 and e^a_2, respectively. The precondition of e^a_1 is $pre(e^a_1) = d \wedge m_a$, namely the message can only arrive to general* **b** *if general* **a** *has indeed decided to attack at dawn and if the messenger is at* **a***'s camp. The precondition of e^a_2 is simply $pre(e^a_2) = \top$, since the messenger could always be captured. We represent the fact that the messenger travels from one hilltop to the other[2] by having $post(e^a_1)(m_a) = \bot$ and $post(e^a_1)(m_b) = \top$. Finally, we denote the fact that the messenger is captured by having $post(e^a_2)(m_a) = \bot$ and $post(e^a_2)(m_b) = \bot$.*

$send_{ab}$ =

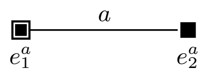

Action $send_{ba}$ is defined similarly, by having $pre(e^b_1) = d \wedge m_b$, $pre(e^b_2) = \top$, $post(e^b_1)(m_a) = \top$, $post(e^b_1)(m_b) = \bot$, $post(e^b_2)(m_a) = \bot$ and $post(e^b_2)(m_b) = \bot$.

$send_{ba}$ =

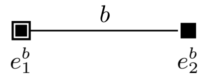

The *product update* formalizes the execution of an action $(\mathcal{E}, E_d)$ on the current epistemic state (M, W_d). Intuitively, the resulting epistemic state (M', W'_d) is computed by a cross product between the worlds in M and the events in $\mathcal{E}$. A pair (w, e) represents the world of M' that results from applying the event e on the world w. We say that $(\mathcal{E}, E_d)$ is *applicable* in (M, W_d) iff for each world $w_d \in W_d$ there exists an event $e_d \in E_d$ that is applicable in w_d.

Definition 6 (Product Update)**.** *Let $(\mathcal{E}, E_d)$ be an action applicable in an epistemic state (M, W_d), where $M = (W, R, V)$ and $\mathcal{E} = (E, Q, pre, post)$. The* product update *of (M, W_d) with $(\mathcal{E}, E_d)$ is the epistemic state $(M, W_d) \otimes (\mathcal{E}, E_d) = ((W', R', V'), W'_d)$, where:*

$$\begin{aligned}
W' &= \{(w, e) \in W \times E \mid (M, w) \models pre(e)\} \\
R'_i &= \{((w, e), (v, f)) \in W' \times W' \mid wR_iv \text{ and } eQ_if\} \\
V'(p) &= \{(w, e) \in W' \mid (M, w) \models post(e)(p)\} \\
W'_d &= \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}
\end{aligned}$$

Example 3. *Suppose that general* **a** *sends the messenger to general* **b** *(action $send_{ab}$) and that the message is successfully delivered. The situation is represented by epistemic state s_1, where $w'_1 = (w_1, e^a_1)$, $w'_2 = (w_1, e^a_2)$ and $w'_3 = (w_2, e^a_2)$ (recall that the reflexive, symmetric and transitive closures of the relations are left implicit).*

s_1 =

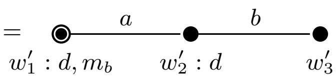

[2] For simplicity, we assume that the truth value of each atomic proposition remains unchanged unless explicitly specified.

Now general **b** *knows about the intentions of his ally* ($\Box_b d$)*, but general* **a** *does not know that general* **b** *knows. So,* **b** *decides to send the messenger back to acknowledge that the message was received (action* send_{ba}*). Here,* d *assumes the meaning of "general* **b** *has received the message". Assume again that the messenger succeeds. We obtain the epistemic state* s_2*, where* $w''_1 = (w'_1, e^b_1)$, $w''_2 = (w'_1, e^b_2)$, $w''_3 = (w'_2, e^b_2)$ *and* $w''_4 = (w'_3, e^b_2)$.

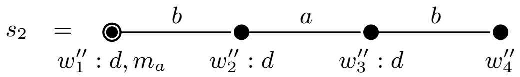

Now it holds that $\Box_a\Box_b d$*, but it does not hold that* $\Box_b\Box_a\Box_b d$*. So, general* **a** *would need to send the messenger once again to general* **b***. However, it can be intuitively seen that, regardless of how many messages the generals exchange, they will never be sure that the other will attack at dawn. This will be stated formally in Section 2.4.*

2.3 The logic S5$_n$

In the epistemic logic literature there exist many different axiomatizations of the concept of *knowledge*. In this paper, we adopt the multimodal logic S5$_n$. Its axioms are[3]:

K $\Box_i(\varphi \to \psi) \to (\Box_i\varphi \to \Box_i\psi)$ (Distribution)
T $\Box_i\varphi \to \varphi$ (Knowledge)
4 $\Box_i\varphi \to \Box_i\Box_i\varphi$ (Positive introspection)
5 $\neg\Box_i\varphi \to \Box_i\neg\Box_i\varphi$ (Negative introspection)

Axioms **T**, **4** and **5** correspond, to the following *frame properties*: reflexivity ($\forall u(uR_iu)$), transitivity ($\forall u, v, w(uR_iv \wedge vR_iw \to uR_iw)$) and Euclidicity ($\forall u, v, w(uR_iv \wedge uR_iw \to vR_iw)$). Moreover, axioms **T** and **5** together entail symmetry ($\forall u, v(uR_iv \to vR_iu)$). Thus, accessibility relations in S5$_n$ are equivalence relations. We refer to epistemic states (resp., epistemic models, actions, frames) satisfying the axioms of a logic L as L-states (resp., L-models, L-actions, L-frames). In the rest of the paper, we assume that the accessibility relations of epistemic states and actions are equivalence relations.

2.4 Plan Existence Problem

We now define our problem, adapting the formulation in [1].

Definition 7 (Planning Task)**.** *An* (epistemic) planning task *is a triple* $T = (s_0, \mathcal{A}, \varphi_g)$*, where* s_0 *is an initial epistemic state;* $\mathcal{A}$ *is a finite set of actions;* $\varphi_g \in \mathcal{L}^C_{\mathcal{P},\mathcal{AG}}$ *is a* goal formula*.*

Given a logic L, an L*-planning task* $(s_0, \mathcal{A}, \varphi_g)$ is a planning task where s_0 is an L-state and each action in $\mathcal{A}$ is an L-action. We denote the class of L-planning tasks with $\mathcal{T}_L$. We remark that, given a generic logic L, the product update of an L-state with an L-action, in general, is not necessarily an L-state. This is not a desired outcome in general, since axioms model some principles of knowledge/belief that always need to be satisfied. For instance, this is the case of the logic KD45$_n$, that captures the concept of *belief*. In the literature, there exist different approaches to guarantee the preservation of the KD45$_n$ frame properties after the product update. For instance, some techniques involve belief revision techniques [14], whereas others focus on defining some additional conditions to impose to both states and actions [20].

The case of our logic C-S5$_n$ is similar to that of KD45$_n$. In fact, the frame property corresponding to axiom **C** (see Equation 1 in Section 3) is not guaranteed to hold after the application of an action. In this paper, rather than devising some technique to guarantee the preservation of frame property 1, we instead opt for a rollback-style approach: an action is not to be applied in a state if it would lead to violate the axioms of C-S5$_n$. This leads to the following definition.

Definition 8 (Solution)**.** *A* solution *to an* L*-planning task* $(s_0, \mathcal{A}, \varphi_g)$ *is a finite sequence* $\alpha_1, \ldots, \alpha_m$ *of actions of* $\mathcal{A}$ *such that:*
1. $s_0 \otimes \alpha_1 \otimes \cdots \otimes \alpha_m \models \varphi_g$*, and*
2. For each $1 \leq k \leq m$*,* α_k *is applicable in* $s_0 \otimes \alpha_1 \otimes \cdots \otimes \alpha_{k-1}$ *and* $s_0 \otimes \alpha_1 \otimes \cdots \otimes \alpha_k$ *is an* L*-state.*

Definition 9 (Plan Existence Problem)**.** *Let* $n \geq 1$ *and* $\mathcal{T}_L$ *be a class of epistemic planning tasks for a logic* L*.* PLANEX($\mathcal{T}_L$, n) *is the following decision problem: "Given an* L*-planning task* $T = (s_0, \mathcal{A}, \varphi_g) \in \mathcal{T}_L$*, where* $|\mathcal{AG}| = n$*, does* T *have a solution?"*

Example 4. *In Example 3 we have seen that, intuitively speaking, the two generals can not coordinate a winning attack. We now state this formally. Let* $T_{\text{coord}} = (s_0, \mathcal{A}, \varphi_g)$ *be an S5*$_n$*-planning task, where* $\mathcal{A} = \{\text{send}_{ab}, \text{send}_{ba}\}$ *and* $\varphi = C_{\{a,b\}}d$*. Then,* T_{coord} *has no solution. In fact, for any number* $k \geq 0$ *of delivered messages, one can show by induction the following (where* $h \geq 0$*):*

- k=$2h$*:* $(\Box_a\Box_b)^h\Box_a d$ *holds in* s_k*, but not* $(\Box_b\Box_a)^{h+1}d$*;*
- k=$2h$+1*:* $(\Box_b\Box_a)^{h+1}d$ *holds in* s_k*, but not* $(\Box_a\Box_b)^{h+1}\Box_a d$*.*

Thus, common knowledge can not be achieved between the two generals in a finite number of steps. However, any search algorithm would never terminate, since at each step there is exactly one applicable action that, when applied, results in a new S5$_n$*-state.*

3 Semantic Approach and Commutativity

In this section, we discuss in more detail the semantic approach and we show how it can be used to obtain decidability results. Then, we introduce and analyze the commutativity axiom in the context of epistemic planning.

With the semantic approach, we aim at devising a new way to approach decidability in epistemic planning, which deviates from the common line of research in the literature focused on limiting the action theory syntactically (*e.g.*, by imposing a limit on the maximum modal depth of formulae). Our approach is motivated by the fact that to obtain decidable fragments of the general problem, one must appeal to strong syntactical constraints. To substantiate this claim, recall that the problem is still undecidable when the maximum modal depth allowed is set to d=2 (see Table 1 for more details, where $\mathcal{T}(\ell, m)$ denotes the class of epistemic planning tasks where preconditions and postconditions have modal depth at most ℓ and m, respectively). This is clearly a strong limitation of syntactic approaches, as the isolated classes are too strong in many practical cases, where reasoning about the knowledge of others is required. Instead, the semantic approach does not limit the structure of formulae of the action theory, but rather relies on devising a suitable set of axioms that guarantee desirable properties on the structure of epistemic states (*e.g.*, bounded number of possible worlds). Since, in principle, there are many ways one can obtain such desirable properties by means of modal axioms, we argue that the semantic approach constitutes a fruitful avenue of research that can be further explored in many different ways.

Towards this goal, we analyze in detail the commutativity axiom. The key insight behind the definition of such axiom is that, 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

[3] Even though **K**, **T** and **5** are sufficient to characterize S5$_n$, we include axiom **4** as it constitutes an important epistemic principle.

about the higher-order knowledge they possess about the knowledge of others. This is clear in Example 3, where at each step k we obtain an epistemic state s_k that contains a chain of worlds of the form $w_1 R_i w_2 R_j w_3 R_i \dots w_k$ (for $i, j \in \{\mathsf{a}, \mathsf{b}\}, i \neq j$), that intuitively represents i's perspective about j's perspective about i's perspective, and so forth. This idea has been exploited for building undecidability proofs of the plan existence problem in the logic S5$_n$, by showing a reduction from the halting problem of Turing machines [5] and Minsky two-counter machines [1].

To weaken this reasoning power, we introduce a principle that governs the capability of agents to reason about the knowledge of others, which is captured by the following interaction axiom (where $i \neq j$):

$$\mathbf{C} \quad \Box_i \Box_j \varphi \to \Box_j \Box_i \varphi \quad \text{(Commutativity)}$$

The commutativity axiom is well-known in many-dimensional modal logics where it is part of the axiomatisation of the product of two modal logics [11]. Here, we adopt it with a novel epistemic connotation. Namely, we can read **C** as follows: whenever an agent i knows that another agent j knows that φ, then j knows that i knows *too* that φ. Thus, intuitively, axiom **C** defines a principle of *commutativity* in the knowledge that agents have about the knowledge of others. This intuition is formalized and proved in the next section (see Lemma 2 and Theorem 2).

This axiom is instrumental in proving decidability of the plan existence problem. Moreover, it provides a useful principle for two main reasons. First, as we mentioned above, this axiom allows to govern the reasoning power of agents. As it turns out, in this way we obtain a *finitary non-fixpoint* characterization of common knowledge (see Theorem 2), which concretely shows the power of knowledge commutativity. Second, this principle constitutes a reasonable assumption in several *cooperative multi-agent planning tasks* [21] where agents are able to communicate or monitor each other. In fact, when autonomous agents cooperate to reach a shared goal, then they are expected to behave in such a way that the effects of their actions are observable by others. In other words, acting in a cooperative context results into a transparent behaviour of agents, which in turn well fits with the concept of knowledge commutativity.

We call C-S5$_n$ the logic S5$_n$ augmented with axiom **C**. As a final remark, notice that axiom **C** is a Sahlqvist formula and it corresponds to the following frame property:

$$\forall u, v, w (u R_j v \wedge v R_i w \to \exists x (u R_i x \wedge x R_j w)) \qquad (1)$$

Proposition 1. *The logic C-S5$_n$ is sound and complete with the class of reflexive, symmetric and transitive epistemic models that enjoy property (1).*

4 Epistemic Planning with Commutativity

This section is organized as follows. First, we prove the decidability of the plan existence problem under the logic C-S5$_n$. Then, in Section 4.1, we consider two generalizations of the commutativity axiom and we analyze their impact in the decidability of the plan existence problem. We first recall the following result:

Theorem 1 ([1]). *For any $n > 1$, PLANEX($\mathcal{T}_{\text{S5}}$, n) is undecidable.*

We now focus on the logic C-S5$_n$, assuming that $n > 1$. We begin by giving some preliminary results.

Lemma 1. *Let $G = \{i_1, \dots, i_m\} \subseteq \mathcal{AG}$, with $m \geq 2$ and let $\vec{v} \in G^*$ ($|\vec{v}| = \lambda \geq 2$). Let π and ρ be two permutations of elements of $\vec{v}$. Then, for any φ, in the logic C-S5$_n$ the following is a theorem:*

$$\Box_{\pi_1} \dots \Box_{\pi_\lambda} \varphi \leftrightarrow \Box_{\rho_1} \dots \Box_{\rho_\lambda} \varphi$$

For a group G of agents, the knowledge of agent i_1 about what agent i_2 knows about what agent i_3 knows, and so forth up to agent i_m, *coincides* exactly with the knowledge of any of the agents of G about what some other agent knows about some third agent, and so forth up to the m-th agent. That is, higher-order knowledge involving a group of agents is independent from the order in which we consider said agents (*i.e.*, it commutes). In other words, Lemma 1 shows us that we can *rearrange* the order of any sequence of boxes in a formula and obtain an equivalent one.

Lemma 2. *Let $G = \{i_1, \dots, i_m\} \subseteq \mathcal{AG}$, with $m \geq 2$. In the logic C-S5$_n$, for any φ and $\vec{v} \in G^*$ we have that $\Box_{i_1} \dots \Box_{i_m} \varphi \to \Box_{v_1} \cdots \Box_{v_{|\vec{v}|}} \varphi$ is a theorem.*

Lemma 2 provides the basis to show a first important result related to common knowledge in C-S5$_n$. Namely, we obtain a *finitary non-fixpoint* characterization of common knowledge.

Theorem 2. *Let $G = \{i_1, \dots, i_m\} \subseteq \mathcal{AG}$, with $m \geq 2$. In the logic C-S5$_n$, for any φ, the formula $\Box_{i_1} \dots \Box_{i_m} \varphi \leftrightarrow C_G \varphi$ is a theorem.*

Proof. ($\Leftarrow$) This follows by definition of common knowledge; ($\Rightarrow$) this immediately follows by Lemma 2. □

Corollary 1. *Let $G = \{i_1, \dots, i_m\} \subseteq \mathcal{AG}$, with $m \geq 2$. In an C-S5$_n$-model, for any $\vec{v} \in G^*$, we have that if $w R_{v_1} \circ \dots \circ R_{v_{|\vec{v}|}} w'$, then $w R_{i_1} \circ \cdots \circ R_{i_m} w'$.*

The statement above directly follows from the contrapositive of the implication in Lemma 2, under the assumption of minimality of states (w.r.t. bisimulation). Intuitively, this states that in a C-S5$_n$-model, given any subset of $m \geq 2$ agents, if a world is reachable in an arbitrary number of steps, then it is also reachable in exactly m steps. Thus, in general, any pair of worlds of a C-S5$_n$-model that are reachable from one another are connected by a path of length *at most* n. This property suggests the existence of some boundedness property on the size of C-S5$_n$-states. Indeed, we exploit Corollary 1 to prove the following lemma.

Lemma 3. *Let (M, W_d) be an C-S5$_n$-state, with $M = (W, R, V)$. For any $w, v \in W$, we have that $w \underline{\leftrightarrow}_{n+1} v \Leftrightarrow w \underline{\leftrightarrow} v$.*

Lemma 3 shows that, in the logic C-S5$_n$, to verify whether two worlds are bisimilar, we only need to check their neighborhoods up to distance $n + 1$. We exploit this intuition, under the assumption of bisimulation minimality, to prove the following lemma.

Lemma 4. *Let (M, W_d) be a bisimulation-contracted C-S5$_n$-state, with $M = (W, R, V)$. Then, $|W|$ is bounded in n and $|\mathcal{P}|$.*

Having a bound in the number of worlds of a C-S5$_n$-state immediately provides us with the following decidability result.

Theorem 3. *For any $n>1$, PLANEX($\mathcal{T}_{\text{C-S5}}$, n) is* decidable.

Proof. Let $T \in \mathcal{T}_{\text{C-S5}_n}$ be an epistemic planning task. By Lemma 4, it follows that we can perform a breadth-first search on the search space that would only visit a finite number of epistemic states (up to bisimulation contraction) to find a solution for T. Thus, we obtain the claim. □

The following example shows that in C-S5$_n$ we can effectively obtain an answer to the Coordinated Attack Problem, hence showing that common knowledge can not be achieved by the two generals.

Example 5. *As shown in Example 4, the $S5_n$-planning task T_{coord} has no solution, but any search algorithm would never terminate. We now consider the C-$S5_n$-planning task $T^{\mathcal{A}}_{\text{coord}} = (s_0, \mathcal{A}, \varphi_g)$, with s_0, $\mathcal{A}$ and φ_g defined as in Example 4 (notice that s_0 is an C-$S5_n$-state and that send$_{ab}$ and send$_{ba}$ are both C-$S5_n$-actions). We immediately note that the epistemic state s_1 of Example 3* is not *an C-$S5_n$-state. Thus, by Definition 8 and since send$_{ab}$ is the only applicable action in s_0, any search algorithm would immediately stop returning the answer* "no"*. Hence, in the logic C-$S5_n$, one can conclude that it is impossible for the two generals to coordinate an attack in a finite number of steps.*

4.1 Generalizing the Principle of Commutativity

Although the commutativity axiom is better fitting for tight-knit groups of agents, it may be less suited for representing more loosely organized groups. Thus, having established that adding axiom **C** to S5_n leads to decidability of the plan existence problem, we investigate two generalized principles of commutativity, namely *b-commutativity* and *weak commutativity*. In what follows, we consider such generalizations and we provide (un)decidability results of their corresponding plan existence problems.

b-Commutativity Let $b>1$ be a fixed constant. Then, we define the following axiom:

$$\mathbf{C}^b \quad (\Box_i\Box_j)^b\varphi \to (\Box_j\Box_i)^b\varphi \quad (b\text{-Commutativity})$$

We call C^b-S5_n the logic S5_n augmented with axiom $\mathbf{C}^b$. Axiom $\mathbf{C}^b$ generalizes commutativity by considering an arbitrary fixed amount of repetitions of box operators. Indeed, notice that $\mathbf{C}^1 = \mathbf{C}$. Moreover, since every $\Box_i$ is a monotone modality (*i.e.*, from $\varphi \to \psi$ we can infer $\Box_i\varphi \to \Box_i\psi$), it is easy to see that each axiom $\mathbf{C}^{b+1}$ leads to a weaker logic than axiom $\mathbf{C}^b$, and that every logic C^b-S5_n is weaker than C-S5_n. One could hope that the plan existence problem remains decidable when replacing axiom **C** with $\mathbf{C}^b$. But this is not true in general. In fact, we prove that it remains decidable for n=2 and any b>1 (Theorem 4) and that it becomes undecidable for any n>2 and b>1 (Theorem 5). Due to space constraints, we only provide the proof sketches (full proofs are available in the arXiv Appendix [24]).

Theorem 4. *For any b>1,* PlanEx*($\mathcal{T}_{\text{C}^b\text{-S5}}$, 2) is* decidable.

Proof. (Sketch) Analogous to the case of the logic C-S5_n. Namely, we can prove the correspondent versions of Lemma 2, Theorem 2, Corollary 1 and Lemmata 3 and 4. The claim follows by combining these results as in Theorem 3. □

Example 6. *As the Coordinated Attack Problem involves exactly two agents, for any b>1, we can define the C^b-$S5_2$-planning task $T^{\text{C}^b}_{\text{coord}} = (s_0, \mathcal{A}, \varphi_g)$, with s_0, $\mathcal{A}$ and φ_g defined as in Example 5. Then, as above, we note that the epistemic state $s_{1+2(b-1)}$ of Example 4 is* not *a C^b-$S5_2$-state. Thus, by Definition 8 and since send$_{ab}$ is the only applicable action in $s_{2(b-1)}$, a search algorithm would return the answer* "no" *in exactly $2(b-1)$ steps.*

Theorem 5. *For any n>2, b>1,* PlanEx*($\mathcal{T}_{\text{C}^b\text{-S5}}$, n) is* undecidable.

Proof. (Sketch) We adapt the proof in [1, Section 6], of the undecidability of epistemic planning in the logic $S5_n$ ($n > 1$). It is an elegant reduction from the halting problem of Minsky two-counter machines [18] to the plan existence problem.

We prove our result for the logic C^2-S5_3 (*i.e.*, having $b = 2$ and $n = 3$). Since C^2-S5_3-models are also C^b-S5_n-models for any $n > 3$ and $b > 2$, our results hold for any combination of the values of $n \geq 3$ and $b \geq 2$. Given a two-counter machine M, the procedure follows three steps:

1. We define an encoding for integers and configurations;
2. We build a finite set of actions for encoding the computation function f_M; and
3. We combine the previous steps and we encode the halting problem as an C^2-S5_3-planning task.

Finally, the claim follows by showing that the resulting planning task has a solution iff M halts. □

These results show that it is not straightforward to generalize knowledge commutativity and to maintain the decidability of the plan existence problem. However, Theorem 4 suggests that the logic C^b-S5_n could result into interesting developments in contexts where only two agents are involved (*e.g.*, epistemic games).

Weak commutativity Let $1<\ell\leq n$ be a fixed constant. Let $\langle i_1, \dots, i_\ell \rangle$ be a sequence of agents with no repetitions, and let π be any permutation of this sequence. Then, we define the following axiom (for any such π):

$$\mathbf{wC}_\ell \quad \Box_{i_1} \dots \Box_{i_\ell}\varphi \to \Box_{\pi_{i_1}} \dots \Box_{\pi_{i_\ell}}\varphi \quad \text{(Weak comm.)}$$

We call wC_ℓ-S5_n the logic S5_n augmented with axiom $\mathbf{wC}_\ell$. Axiom $\mathbf{wC}_\ell$ generalizes commutativity by extending it to more than two agents, whereas **C** corresponds to taking $\ell = 2$. Indeed, notice that $\mathbf{wC}_2 = \mathbf{C}$. Moreover, since every $\Box_i$ is a monotone modality, it is easy to see that each axiom $\mathbf{wC}_{\ell+1}$ leads to a weaker logic than axiom $\mathbf{wC}_\ell$, and that every logic wC_ℓ-S5_n is weaker than wC-S5_n.

By considering this form of generalization of axiom **C**, we are able to provide a decidability result that holds for any $1 < \ell \leq n$. The arguments adopted by the proof are similar to those of Theorem 3.

Theorem 6. *For any n>1 and $1<\ell\leq n$,* PlanEx*($\mathcal{T}_{\text{wC}_\ell\text{-S5}}$, n) is* decidable.

Proof. (Sketch) As in the proof of Theorem 3, we can prove the correspondent versions of Lemmata 1, 2, Theorem 2 and Corollary 1. From these results, we obtain that any pair of worlds of a wC_ℓ-S5_n-model that are reachable from one another are connected by a path of length *at most* n. Hence, we show that Lemmata 3 and 4 hold also in the logic wC_ℓ-S5_n (for any $\ell > 1$). Thus, to obtain the claim, we use Lemmata 3 and 4 by combining them as in Theorem 3. □

To summarize, b-commutativity and weak commutativity constitute two generalizations of axiom **C**. All of the above decidability results are outlined in Table 2.

5 Epistemic Planning Systems

In this section, we look at two well-known epistemic planning systems that adopt the DEL semantics and we show their decidability by applying our previous results. These are $m\mathcal{A}^*$ [3] and the framework by Kominis and Geffner [16]. While decidability is already known for the latter, the decidability of the former is still an open problem. In what follows, we show that both these systems are captured by our setting when knowledge is considered, namely under S5_n axioms. We begin by briefly introducing the two systems.

1. The epistemic planning framework $m\mathcal{A}^*$ [3] features three action types: *ontic*, *sensing* and *announcement* actions. In $m\mathcal{A}^*$, agents

Figure 1: Frames of $m\mathcal{A}^*$ event models for ontic actions (left) and sensing/announcement actions (right).

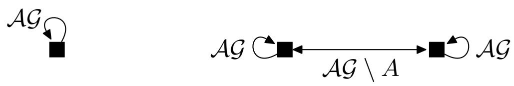

Figure 2: Frames of event models in [16] for *do* and *update* actions (left) and *sense* (right).

are partitioned in three sets: *fully observant* agents (F), that are able to observe the action corresponding to an event, *partially observant* agents (P) that only know about the execution of an action, but not the effects, and *oblivious* agents (O), that are ignorant about the fact that the action is taking place. When oblivious agents are considered, however, event models fall beyond C-S5$_n$ (they have KD45$_n$-frames), as they are not *symmetric* (see Figure 1). Consequently, we restrict ourselves to a fragment of $m\mathcal{A}^*$ that includes *public* ontic actions and *semi-private* sensing and announcement actions. This is achieved by removing from the event models of Figure 1 all events considered possible by oblivious agents. It is easy to see that the frames of the resulting event models are indeed S5$_n$-frames. We call the resulting system the *S5$_n$-fragment* of $m\mathcal{A}^*$, and we denote with $\mathcal{T}_{\text{S5}_n\text{-}m\mathcal{A}^*}$ the class of planning tasks of such system.

2. Kominis and Geffner [16] describe a system for handling beliefs in multi-agent scenarios. They describe three types of actions: *do*, *update*, and *sense*. Although their formulation is not given in terms of DEL semantics, the authors briefly describe the event models corresponding to each action type (Figure 2). We denote with $\mathcal{T}_{\text{KG}}$ the class of planning tasks of such system. Differently from $m\mathcal{A}^*$, all the described event models already have S5$_n$-frames.

We show that the two described systems fall within our logic, thus proving their decidability (see arXiv Appendix [24] for full proofs):

Lemma 5. *$\mathcal{T}_{\text{S5}_n\text{-}m\mathcal{A}^*} \subseteq \mathcal{T}_{\text{C-S5}}$ and $\mathcal{T}_{\text{KG}} \subseteq \mathcal{T}_{\text{C-S5}}$.*

Corollary 2. *For any $n > 1$,* PLANEX*($\mathcal{T}_{\text{S5}_n\text{-}m\mathcal{A}^*}$, n) and* PLANEX*($\mathcal{T}_{\textbf{KG}}$, n) are* decidable*.*

6 Related Works

The DEL semantics has been widely exploited to obtain complexity and decidability results for the plan existence problem [1, 2, 5, 6, 8, 9, 17, 23]. We identify three main lines of research. The first one restricts the class of actions that are allowed in planning tasks by means of syntactical conditions. In particular, some approaches constrain actions to be *purely epistemic* (*i.e.*, without postconditions), while others introduce limitations to the modal depth preconditions and/or postconditions. Following the notation of [6], we denote with $\mathcal{T}(\ell, m)$ the class of epistemic planning tasks in which the modal depth of the preconditions are $\leq \ell$, and that of the postconditions are $\leq m$. Similarly, planning tasks with purely epistemic actions are denoted with $\mathcal{T}(\ell, -1)$. We report the main results in Table 1.

The second line of research pivots both on syntactical limitations to formulae and on constraining the frames of event models (*e.g.*, singletons, chains, trees). In particular, PLANEX($\mathcal{T}(0, -1), n$) is NP-complete on singletons and chains and it is PSPACE-complete on trees [8], whereas PLANEX($\mathcal{T}(\ell, m), n$) (for any $\ell, m \geq 0$) on singletons is PSPACE-hard [15].

PLANEX($\mathcal{T}(0, -1), m$)	PSPACE-complete [9]
PLANEX($\mathcal{T}(1, -1), n$)	UNKNOWN [9]
PLANEX($\mathcal{T}(2, -1), n$)	UNDECIDABLE [9]
PLANEX($\mathcal{T}(0, 0), n$)	DECIDABLE [23, 2]
PLANEX($\mathcal{T}(1, 0), n$)	DECIDABLE [6]

Table 1: Decidability and complexity results of plan existence problem based on the *syntactical* approach.

Logic	Decidability
K$_n$, K$_n$, KT$_n$, K4$_n$, K45$_n$, S4$_n$, S5$_n$	UNDECIDABLE [1]
C^b-S5$_n$ (n>2)	UNDECIDABLE
C^b-S5$_2$	DECIDABLE
wC$_\ell$-S5$_n$	
C-S5$_n$	

Table 2: Decidability results of plan existence problem based on the semantic approach, compared to our results (in gray).

Finally, the third line of research revolves around the choice of the logic for epistemic models and actions [1]. This approach is more similar to the one we adopted in this paper, with the difference that the logics considered in previous works are a combination of standard and well-known axioms of epistemic logic (see Section 3). We report the main results in Table 2. For a more detailed analysis of complexity and decidability results in epistemic planning we refer the interested reader to [6].

7 Conclusions

The paper presents novel decidability results in epistemic planning. The approach adopted in this work deviates from previous ones, where syntactical conditions are imposed to actions. In particular, we pursue a novel semantic approach by introducing a principle of knowledge commutativity that is well suited for cooperative multi-agent planning contexts. In this way, 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. Starting from this key result, we studied decidability of the plan existence problem under different generalizations of the commutativity axiom, showing both positive and negative results.

Notably, our decidability results are orthogonal to the problem of preservation of commutativity during planning. In this paper we adopt the baseline strategy of rejecting action sequences that visit invalid states. A natural follow up is then to incorporate more sophisticated techniques to revise/repair such invalid states, in the style of belief revision for KD45$_n$ [14] and to single out fragments of C-S5$_n$ where preservation is guaranteed by design [20].

In the future, we plan on further investigating our semantic approach by formulating other properties to add to the logic of knowledge. In particular, we are interested into defining new generalizations of the commutativity axiom to obtain broader fragments that maintain decidability of the plan existence problem. Moreover, we are interested into verifying if and how this approach can be suitably recast in a doxastic setting, where knowledge is replaced by *belief*. This is not a trivial task, as the results of this paper do not readily apply to the logic KD45$_n$. We also intend to delve into a fine-grained analysis of the computational complexity of $\mathcal{T}_{\text{C-S5}_n}$ and to compare it with the current results in the literature.

Acknowledgements

This research has been partially supported by the Italian Ministry of University and Research (MUR) under PRIN project PINPOINT Prot. 2020FNEB27, and by the Free University of Bozen-Bolzano with the ADAPTERS project.

References

[1] Guillaume Aucher and Thomas Bolander, 'Undecidability in epistemic planning', in *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, ed., Francesca Rossi, pp. 27–33. IJCAI/AAAI, (2013).

[2] Guillaume Aucher, Bastien Maubert, and Sophie Pinchinat, 'Automata techniques for epistemic protocol synthesis', in *Proceedings 2nd International Workshop on Strategic Reasoning, SR 2014, Grenoble, France, April 5-6, 2014*, eds., Fabio Mogavero, Aniello Murano, and Moshe Y. Vardi, volume 146 of *EPTCS*, pp. 97–103, (2014).

[3] Chitta Baral, Gregory Gelfond, Enrico Pontelli, and Tran Cao Son, 'An action language for multi-agent domains: Foundations', *CoRR*, **abs/1511.01960**, (2015).

[4] Patrick Blackburn, Maarten de Rijke, and Yde Venema, *Modal Logic*, Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, 2001.

[5] Thomas Bolander and Mikkel Birkegaard Andersen, 'Epistemic planning for single and multi-agent systems', *J. Appl. Non Class. Logics*, **21**(1), 9–34, (2011).

[6] Thomas Bolander, Tristan Charrier, Sophie Pinchinat, and François Schwarzentruber, 'Del-based epistemic planning: Decidability and complexity', *Artif. Intell.*, **287**, 103304, (2020).

[7] Thomas Bolander, Lasse Dissing, and Nicolai Herrmann, 'DEL-based Epistemic Planning for Human-Robot Collaboration: Theory and Implementation', in *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning*, pp. 120–129, (11 2021).

[8] Thomas Bolander, Martin Holm Jensen, and François Schwarzentruber, '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*, eds., Qiang Yang and Michael J. Wooldridge, pp. 2791–2797. AAAI Press, (2015).

[9] Tristan Charrier, Bastien Maubert, and François Schwarzentruber, '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*, ed., Subbarao Kambhampati, pp. 1030–1036. IJCAI/AAAI Press, (2016).

[10] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi, *Reasoning About Knowledge*, MIT Press, 2004.

[11] D.M. Gabbay, *Many-dimensional Modal Logics: Theory and Applications*, Studies in logic and the foundations of mathematics, North Holland Publishing Company, 2003.

[12] Jim Gray, 'Notes on data base operating systems', in *Operating Systems, An Advanced Course*, pp. 393–481, Berlin, Heidelberg, (1978). Springer-Verlag.

[13] Joseph Y. Halpern and Yoram Moses, 'Knowledge and common knowledge in a distributed environment', *J. ACM*, **37**(3), 549–587, (1990).

[14] Andreas Herzig, Paul Sabatier, Jérôme Lang, and Pierre Marquis, 'Action progression and revision in multiagent belief structures', in *Sixth Workshop on Nonmonotonic Reasoning, Action, and Change (NRAC)*, (2005).

[15] Martin Holm Jensen, *Epistemic and Doxastic Planning*, Ph.D. dissertation, Technical University of Denmark, 2014.

[16] Filippos Kominis and Hector Geffner, 'Beliefs in multiagent planning: From one agent to many', in *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling, ICAPS 2015, Jerusalem, Israel, June 7-11, 2015*, eds., Ronen I. Brafman, Carmel Domshlak, Patrik Haslum, and Shlomo Zilberstein, pp. 147–155. AAAI Press, (2015).

[17] Benedikt Löwe, Eric Pacuit, and Andreas Witzel, 'DEL planning and some tractable cases', in *Logic, Rationality, and Interaction - Third International Workshop, LORI 2011, Guangzhou, China, October 10-13, 2011. Proceedings*, eds., Hans van Ditmarsch, Jérôme Lang, and Shier Ju, volume 6953 of *Lecture Notes in Computer Science*, pp. 179–192. Springer, (2011).

[18] Marvin L Minsky, *Computation: finite and infinite machines*, Prentice-Hall, Inc., 1967.

[19] Cédric Paternotte, 'Being realistic about common knowledge: a lewisian approach', *Synth.*, **183**(2), 249–276, (2011).

[20] Tran Cao Son, Enrico Pontelli, Chitta Baral, and Gregory Gelfond, '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*, eds., Blai Bonet and Sven Koenig, pp. 1604–1610. AAAI Press, (2015).

[21] Alejandro Torreño, Eva Onaindia, Antonín Komenda, and Michal Stolba, 'Cooperative multi-agent planning: A survey', *ACM Comput. Surv.*, **50**(6), 84:1–84:32, (2018).

[22] Hans P. van Ditmarsch, Wiebe van der Hoek, and Barteld P. Kooi, *Dynamic Epistemic Logic*, volume 337, Springer Netherlands, 2007.

[23] Quan Yu, Ximing Wen, and Yongmei Liu, 'Multi-agent epistemic explanatory diagnosis via reasoning about actions', in *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, ed., Francesca Rossi, pp. 1183–1190. IJCAI/AAAI, (2013).

[24] Alessandro Burigana, Paolo Felli, Marco Montali, and Nicolas Troquard, 'A Semantic Approach to Decidability in Epistemic Planning (Extended Version)', arXiv preprint arXiv:2307.15485, (2023).