A Dynamic Epistemic Framework for Conformant Planning

Quan Yu

Department of Computer Science, Sun Yat-sen University, China

Qiannan Normal College for Nationalities, China

Yanjun Li

Department of Philosophy, Peking University, China

Faculty of Philosophy, University of Groningen, The Netherlands

Yanjing Wang[*]

Department of Philosophy, Peking University, China

In this paper, we introduce a lightweight dynamic epistemic logical framework for automated planning under initial uncertainty. We reduce plan verification and conformant planning to model checking problems of our logic. We show that the model checking problem of the iteration-free fragment is PSPACE-complete. By using two non-standard (but equivalent) semantics, we give novel model checking algorithms to the full language and the iteration-free language.

1 Introduction

Conformant planning is the problem of finding a linear plan (a sequence of action) to achieve a goal in presence of uncertainty about the initial state (cf. [29]). For example, suppose that you are a rookie spy trapped in a foreign hotel with the following map at hand:[1]
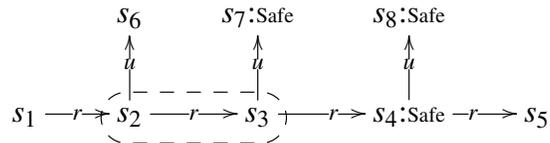
Now somebody spots you and sets up the alarm. In this case you need to move fast to one of the safe hiding places marked in the map (i.e., s_7, s_8 and s_4). However, since you were in panic, you lost your way and you are not sure whether you are at s_2 or s_3 (denoted by the circle in the above graph). Now what should you do in order to reach a safe place quickly? Clearly, merely moving r or moving u may not guarantee your safety given the uncertainty. A simple plan is to move r first and then u, since this plan will take you to a safe place, no matter where you actually are initially. This plan is *conformant* since it does not require any feedback during the execution and it should work in presence of uncertainty about the initial state. More generally, a conformant plan should also work given actions with non-deterministic effects. Such a conformant plan is crucial when there are no feedbacks/observations available during the execution of the plan.[2] Note that since no information is provided during the execution, the conformant plan is simply a finite sequence of actions without any conditional moves.

---

[*]Corresponding author

[1]It is a variant of the running example in [33].

[2]In many other cases, feedbacks may be just too 'expensive' to obtain during a plan aiming for quick actions [8].

As discussed in [9, 25], conformant planning can be reduced to classical planning, the planning problem without any initial uncertainty, over the space of *belief states*. Intuitively, a belief state is a subset of the state space, which records the uncertainty during the execution of a plan, e.g., $\{s_2, s_3\}$ is an initial belief state in the above example. In order to make sure a goal is achieved eventually, it is crucial to track the transitions of belief states during the execution of the plan, and this may traverse exponentially many belief states in the size of the original state space. As one may expect, conformant planning is computationally harder than classical planning. The complexity of checking the existence of a conformant plan is EXPSPACE-complete in the size of the variables generating the state space [19]. In the literature, people proposed compact and implicit representations of the belief spaces, such as OBDD [13, 15, 14] and CNF [31], and different heuristics are used to guide the search for a plan, e.g., [11, 12].

Besides the traditional AI approaches, we can also take an epistemic-logical perspective on planning in presence of initial uncertainties, based on dynamic epistemic logic (DEL) (cf. e.g., [16]). The central philosophy of DEL takes the meaning of an action as the change it brings to the knowledge of the agents. Intuitively, this is what we need to track the belief states during the execution of a plan[3]. Indeed, in recent years, there has been a growing interest in using DEL to handle multi-agent planning with knowledge goals (cf. e.g., [7, 24, 2, 3, 34, 26]), while the traditional AI planning focuses on the single-agent case. In particular, the event models of DEL (cf. [6]) are used to handle non-public actions that may cause different knowledge updates to different agents. In these DEL-based planning frameworks, states are epistemic models, actions are event models and the state transitions are implicitly encoded by the update product which computes a new epistemic model based on an epistemic model and an event model.

One advantage of this approach is its expressiveness in handling scenarios which require reasoning about agents' higher-order knowledge about each other in presence of partially observable actions. However, this expressiveness comes at a price, as shown in [7, 4], that multi-agent epistemic planning is undecidable in general. Many interesting decidable fragments are found in the literature [7, 24, 34, 1], which suggests that the single-agent cases and restrictions on the form of event models are the key to decidability. However, if we focus on the single-agent planning, a natural question arises: how do we compare such DEL approaches with the traditional AI planning? It seems that the DEL-based approaches are more suitable for planning with actions that change (higher-order) knowledge rather than planning with fact-changing actions, although the latter type of actions can also be handled in DEL. Moreover, the standard models of DEL are purely epistemic thus do not encode the temporal information of available actions directly. This may limit the applicability of such approaches to planning problems based on transition systems.

In this paper, we tackle the standard single-agent conformant planning problem over transition systems, by using the core idea of DEL, but not its standard formalism. Our formal framework is based on the logic proposed by Wang and Li in [33], where the model is simply a transition system with initial uncertainty as in the motivating example, and an action is interpreted in the semantics as an update on the uncertainty of the agent. Our contributions are summarized as follows:

- A lightweight dynamic epistemic framework with a simple language and a complete axiomatization.

- Non-trivial reduction of conformant planning to a model checking problem using our language with programs.

---

[3]Here the belief states are actually about knowledge in epistemic logic.

- Two novel model checking algorithms based on two alternative semantics for the proposed logic, which make the context-dependency in the original semantics explicit.

- The complexity of model checking the iteration-free fragment of our language is PSPACE-complete. The model checking problem of the full language is in EXPTIME. The model checking problem of the conformant planning is in PSPACE.

The last result may sound contradictory to the aforementioned result that the complexity of conformant planning is EXPSPACE-complete. Actually, the apparent contradiction is due to the fact that the EXPSPACE complexity result is based on the number of *state variables* which require an exponential blow up to generate an explicit transition system that we use here. We will come back to this issue at the end of Section 4.3.

Our approach has the following advantages compared to the existing planning approaches:

- The planning goals can be specified as arbitrary formulas in an epistemic language. Extra plan constraints (e.g., what actions to use) can be expressed explicitly by programs in the language. Therefore it may cover a richer class of (conformant) planning problems compared to the traditional AI approach where a goal is Boolean.[4]

- The plans can be specified as regular expressions with tests in terms of arbitrary EPDL formulas, which generalizes the knowledge-based programs in [18, 22].

- By reducing conformant planning to a model checking problem in an explicit logical language, we also see the subtleties hidden in the planning problem. In principle, there are various model checking techniques to be applied to conformant planning based on this reduction.

- Our logical language and models are very simple compared to the standard action-model based DEL approach, yet we can encode the externally given executability of the actions in the model, inspired by epistemic temporal logic (ETL) [17, 27].

- Our approach is flexible enough to provide, in the future, a unified platform to compare different planning problems under uncertainty. By studying different fragments of the logical language and model classes, we may categorize planning problems according to their complexity.

The rest of the paper is organized as follows: We introduce our basic logical framework and its axiomatization in Section 2, and extend it in Section 3 with programs to handle the conformant planning. The complexity analysis of the model checking problems is in Section 4 and we conclude in Section 5 with future directions.

2   Basic framework

2.1   Epistemic action language

To talk about the knowledge of the agent during an execution of a plan, we use the following language proposed in [33].

---

[4]The goal in the standard conformant planning is simply a set of different valuations of basic propositional variables. Our approach can even handle epistemic goals in negative forms, e.g., we want to make sure the agent knows something but does not know too much in the end.

Definition 2.1 (Epistemic Action Language (EAL)) *Given a countable set A of action symbols and a countable set P of atomic proposition letters , the language EAL_P^A is defined as follows:*[5]

$$\phi ::= \top \mid p \mid \neg\phi \mid (\phi \wedge \phi) \mid [a]\phi \mid K\phi,$$

where $p \in P$, $a \in A$. The following standard abbreviations are used: $\bot := \neg\top$, $\phi \vee \psi := \neg(\neg\phi \wedge \neg\psi)$, $\phi \rightarrow \psi := \neg\phi \vee \psi$, $\langle a \rangle \phi := \neg[a]\neg\phi$, $\hat{K}\phi := \neg K \neg\phi$.

$K\phi$ says that the agent knows that ϕ, and $[a]\phi$ expresses that if the agent can move forward by action a, then after doing a, ϕ holds. Throughout the paper, we fix some P and A, and refer to EAL_P^A by EAL.

The size of EAL-formulas (notation $|\varphi|$) is defined inductively: $|\top| = |p| = 1$; $|\neg\phi| = 1 + |\phi|$; $|\phi \wedge \psi| = 1 + |\phi| + |\psi|$; $|K\phi| = |[a]\phi| = 1 + |\phi|$. The set of subformulas of $\phi \in$ EAL, denoted as $sub(\phi)$, is defined as usual.

Definition 2.2 (Uncertainty map) *Given P and A, a (multimodal) Kripke model $\mathcal{N}$ is a tuple $\langle \mathcal{S}, \{\mathcal{R}_a \mid a \in A\}, \mathcal{V} \rangle$, where $\mathcal{S}$ is a non-empty set of states, $\mathcal{R}_a \subseteq \mathcal{S} \times \mathcal{S}$ is a binary relation labelled by a, $\mathcal{V} : \mathcal{S} \rightarrow 2^P$ is a valuation function. An uncertainty map $\mathcal{M}$ is a Kripke model $\langle \mathcal{S}, \{\mathcal{R}_a \mid a \in A\}, \mathcal{V} \rangle$ with a non-empty set $\mathcal{U} \subseteq \mathcal{S}$. Given an uncertainty map $\mathcal{M}$, we refer to its components by $\mathcal{S}_\mathcal{M}$, $\mathcal{R}_{a,\mathcal{M}}$, $\mathcal{V}_\mathcal{M}$, and $\mathcal{U}_\mathcal{M}$. A pointed uncertainty map $\mathcal{M}, s$ is an uncertainty map $\mathcal{M}$ with a designated state $s \in \mathcal{U}_\mathcal{M}$. We write $s \xrightarrow{a} t$ for $(s,t) \in \mathcal{R}_a$.*

Intuitively, a Kripke model encodes a map (transition system) and the uncertainty set $\mathcal{U}$ encodes the uncertainty that the agent has about where he is in the map. The graph mentioned at the beginning of the introduction is a typical example of an uncertainty map. Note that there may be non-deterministic transitions in the model, i.e., there may be $t_1 \neq t_2$ such that $s \xrightarrow{a} t_1$ and $s \xrightarrow{a} t_2$ for some s, t_1, t_2.

Remark 1 *It is crucial to notice that the designated state in a pointed uncertainty map must be one of the states in the uncertainty set.*

Definition 2.3 (Semantics) *Given any uncertainty map $\mathcal{M} = \langle \mathcal{S}, \{\mathcal{R}_a \mid a \in A\}, \mathcal{V}, \mathcal{U} \rangle$ and any state $s \in \mathcal{U}$, the semantics is defined as follows:*

$$
\begin{array}{lll}
\mathcal{M}, s \vDash \top & & \textit{always} \\
\mathcal{M}, s \vDash p & \Longleftrightarrow & s \in \mathcal{V}(p) \\
\mathcal{M}, s \vDash \neg\phi & \Longleftrightarrow & \mathcal{M}, s \nvDash \phi \\
\mathcal{M}, s \vDash \phi \wedge \psi & \Longleftrightarrow & \mathcal{M}, s \vDash \phi \textit{ and } \mathcal{M}, s \vDash \psi \\
\mathcal{M}, s \vDash [a]\phi & \Longleftrightarrow & \forall t \in S : s \xrightarrow{a} t \textit{ implies } \mathcal{M}|^a, t \vDash \phi \\
\mathcal{M}, s \vDash K\phi & \Longleftrightarrow & \forall u \in \mathcal{U} : \mathcal{M}, u \vDash \phi
\end{array}
$$

where $\mathcal{M}|^a = \langle \mathcal{S}, \{\mathcal{R}_a \mid a \in A\}, \mathcal{V}, \mathcal{U}|^a \rangle$ and $\mathcal{U}|^a = \{r' \mid \exists r \in \mathcal{U} \textit{ such that } r \xrightarrow{a} r'\}$. We say ϕ is valid (notation: $\vDash \phi$) if it is true on all the pointed uncertainty maps. For a action sequence $\sigma = a_1 \ldots a_n$, we write $\mathcal{U}|^\sigma$ for $(\ldots((\mathcal{U}|^{a_1})|^{a_2}) \ldots)|^{a_n}$. and write $\mathcal{M}|^\sigma$ for $(\ldots((\mathcal{M}|^{a_1})|^{a_2}) \ldots)|^{a_n}$.

Intuitively, the agent 'carries' the uncertainty set with him when moving forward and obtains a new uncertainty set $\mathcal{U}|^a$. Note that here we differ from [33] where the updated uncertainty set is further

---

[5]We do need unboundedly many action symbols to encode the desired problem in the later discussion of model checking complexity.

refined according to what the agent can observe at the new state. For conformant planning, we do not consider the observational power of the agent during the execution of a plan.

Let us call the model mentioned in the introduction $\mathscr{M}$, it is not hard to see that $\mathscr{M}|^r$ and $(\mathscr{M}|^r)|^u$ are as follows:
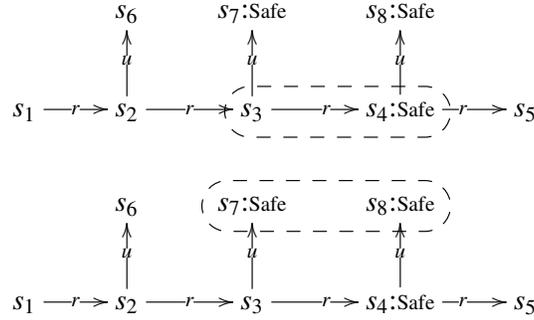
Thus we have:

- $\mathscr{M}, s_3 \vDash [r](Safe \wedge \neg KSafe)$

- $\mathscr{M}, s_3 \vDash K[r][u](Safe \wedge KSafe)$

The usual global model checking algorithm for modal logics labels the states with the subformulas that are true on the states. However, this cannot work here since the truth value of epistemic formulas on the states outside $\mathscr{U}$ is simply undefined. Moreover, the exact truth value of an epistemic formula on a state depends on 'how you get there', as the following example shows (the underlined states mark the actual states):
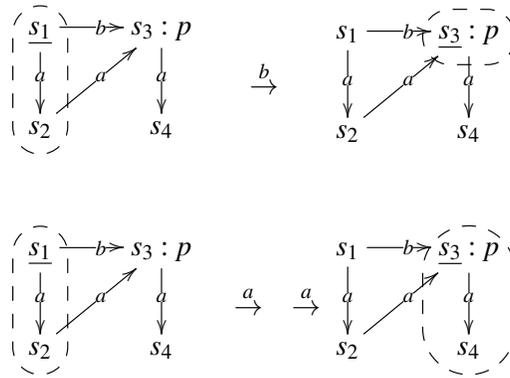
Let the left-hand-side model be $\mathscr{M}$ then it is clear that $\mathscr{M}|^b, s_3 \vDash Kp$ while $\mathscr{M}|^{aa}, s_3 \nvDash Kp$ thus $\mathscr{M}, s_1 \vDash \langle b \rangle Kp \wedge \langle a \rangle \langle a \rangle \neg Kp$. This shows that the truth value of an epistemic subformula w.r.t. a state in the model is somehow 'context-dependent', which requires new techniques in model checking. We will make this explicit in Section 4.3 when we discuss the model checking algorithm.

2.2   Axiomatization

Following the axioms proposed in [33], we give the following axiomatization for EAL w.r.t. our semantics:

System $\mathbb{SELA}$

Axioms		**Rules**	
TAUT	all axioms of propositional logic	MP	$\dfrac{\phi, \phi \rightarrow \psi}{\psi}$
DISTK	$K(p \rightarrow q) \rightarrow (Kp \rightarrow Kq)$	NECK	$\dfrac{\phi}{K\phi}$
DIST(a)	$[a](p \rightarrow q) \rightarrow ([a]p \rightarrow [a]q)$	NEC(a)	$\dfrac{\phi}{[a]\phi}$
T	$Kp \rightarrow p$	SUB	$\dfrac{\phi(p)}{\phi(\psi)}$
4	$Kp \rightarrow KKp$		
5	$\neg Kp \rightarrow K\neg Kp$		
PR(a)	$K[a]p \rightarrow [a]Kp$		
NM(a)	$\langle a \rangle Kp \rightarrow K[a]p$		

where a ranges over A, p, q range over P. PR$(\cdot)$ and NM$(\cdot)$ denote the axioms of *perfect recall* and *no miracles* respectively (cf. [32]).

Note that since we do not assume that the agent can observe the available actions, the axiom OBS(a) : $K\langle a \rangle \top \vee K\neg\langle a \rangle \top$ in [33] is abandoned. Due to the same reason, the axiom of no miracles is also simplified.

We show the completeness of $\mathbb{SELA}$ using a more direct proof strategy compared to the one used in [33].

Theorem 2.1 $\mathbb{SELA}$ *is sound and strongly complete w.r.t. EAL on uncertainty maps.*

Proof: To prove that $\mathbb{SELA}$ is sound on uncertainty maps, we need to show that all the axioms are valid and all the inference rules preserve validity. Since the uncertainty set in an UM denotes an equivalent class, axioms T, 4 and 5 are valid; due to the semantics, the validity of axioms PR$(\cdot)$ and NM$(\cdot)$ can be proved step by step; others can be proved as usual.

To prove that $\mathbb{SELA}$ is strongly complete on uncertainty maps, we only need to show that every $\mathbb{SELA}$-consistent set of formulas is satisfiable on some uncertainty map. The proof idea is that we construct an uncertainty map consisting of maximal $\mathbb{SELA}$-consistent sets (MCSs), and then with the Lindenbaum-like lemma that every $\mathbb{SELA}$-consistent set of formulas can be extended in to a MCS (we omit the proof here), we only need to prove that every formula holds on the MCS to which it belongs.

Firstly, we construct a canonical Kripke model $\mathcal{N}^c = \langle \mathcal{S}^c, \{\mathcal{R}_a^c \mid a \in A\}, \mathcal{V}^c \rangle$ as follows:

- $\mathcal{S}^c$ is the set of all MCSs;

- $s\mathcal{R}_a^c t \iff \langle a \rangle \phi \in s$ for any $\phi \in t$ (equivalently $\phi \in t$ for any $[a]\phi \in s$);

- $\mathcal{V}^c(p) = \{s \mid p \in s\}$.

Given $s \in \mathcal{S}^c$, we define $\mathcal{U}_s^c = \{u \in \mathcal{S}^c \mid K\phi \in s \text{ iff } K\phi \in u\}$, and it is obvious that $s \in \mathcal{U}_s^c$. Thus we have that for each $s \in \mathcal{S}^c$, $\mathcal{M}_s^c = \langle \mathcal{N}^c, \mathcal{U}_s^c \rangle$ is an uncertainty map, and $\mathcal{M}_s^c, s$ is a pointed uncertainty map.

Secondly, we prove the following claim.

Claim 2.1 *If* $s \xrightarrow{a} t$, *then we have* $\mathcal{U}_s^c|^a = \mathcal{U}_t^c$.

$\subseteq$: Assuming $v \in \mathscr{U}_s^c|a$, we need to show $v \in \mathscr{U}_t^c$, namely we need to show that $K\phi \in v \iff K\phi \in t$. Since $v \in \mathscr{U}_s^c|a$, we have that there is $u \in \mathscr{U}_s^c$ such that $u\mathscr{R}_a^c v$. If $K\phi \in t$, it follows by axiom 4 that $KK\phi \in t$. Thus we have $\langle a \rangle KK\phi \in s$. By axiom $\mathtt{NM}(a)$, it follows that $K[a]K\phi \in s$. By $u \in \mathscr{U}_s^c$ and axiom T, we have $[a]K\phi \in u$. It follows by $u\mathscr{R}_a^c v$ that $K\phi \in v$. If $K\phi \notin t$, we have $\neg K\phi \in t$. By axiom 5, we have $K\neg K\phi \in t$. Similarly, we have $\neg K\phi \in v$. Thus we have $K\phi \notin v$.

$\supseteq$: Assuming $v \in \mathscr{U}_t^c$, we need to show $v \in \mathscr{U}_s^c|a$, namely there is $u \in \mathscr{U}_s^c$ such that $u\mathscr{R}_a^c v$. Let u^- be $\{K\phi \mid K\phi \in s\} \cup \{\langle a \rangle \psi \mid \psi \in v\}$. Then u^- is consistent. For suppose not, we have $\vdash K\phi_1 \wedge \cdots \wedge K\phi_n \to [a]\neg\psi_1 \vee \cdots \vee [a]\neg\psi_k$ for some n and k. Since $\vdash [a]\neg\psi_1 \vee \cdots \vee [a]\neg\psi_k \to [a](\neg\psi_1 \vee \cdots \vee \neg\psi_k)$, we have $\vdash K\phi_1 \wedge \cdots \wedge K\phi_n \to [a](\neg\psi_1 \vee \cdots \vee \neg\psi_k)$. By rule NECK and axiom DISTK, we have $\vdash KK\phi_1 \wedge \cdots \wedge KK\phi_n \to K[a](\neg\psi_1 \vee \cdots \vee \neg\psi_k)$. Since $KK\phi_i \in s$ for each $1 \le i \le n$, we have $K[a](\neg\psi_1 \vee \cdots \vee \neg\psi_k) \in s$. By axiom $\mathtt{PR}(a)$, it follows that $[a]K(\neg\psi_1 \vee \cdots \vee \neg\psi_k) \in s$. It follows by $s\mathscr{R}_a^c t$ that $K(\neg\psi_1 \vee \cdots \vee \neg\psi_k) \in t$. Since $v \in \mathscr{U}_t^c$, by axiom T, we have $\neg\psi_1 \vee \cdots \vee \neg\psi_k \in v$. This is contrary with $\psi_i \in v$ for each $1 \le i \le k$. Thus u^- is consistent. By Lindenbaum-like Lemma, there exists a MCS u extending u^-. It follows by $u^- \subseteq u$ that $u \in \mathscr{U}_s^c$ and $u\mathscr{R}_a^c v$. We conclude that $v \in \mathscr{U}_s^c|a$.

Finally, we will show that $\mathscr{M}_s^c, s \vDash \phi$ iff $\phi \in s$. we prove it by induction on ϕ. Please note that the 'existence lemmas' (that $\neg[a]\phi \in s$ implies $\neg\phi \in t$ for some t such that $s \xrightarrow{a} t$ and that $\neg K\phi \in s$ implies $\neg\phi \in s'$ for some $s' \in \mathscr{U}_s^c$) also hold in the model $\mathscr{N}^c$. We only focus on the case of $[a]\phi$. With Claim 2.1, it follows that $\mathscr{M}_t^c = \mathscr{M}_s^c|a$ if $s \xrightarrow{a} t$. Then by the induction hypothesis and the existence lemmas, it is easy to show that $\mathscr{M}_s^c, s \vDash [a]\phi$ iff $[a]\phi \in s$.   ■

3 An extension of EAL for conformant planning

3.1 Epistemic PDL over uncertainty maps

In this section we extend the language of EAL with programs in propositional dynamic logic and use this extended language to express the existence of a conformant plan.

Definition 3.1 (Epistemic PDL) *The Epistemic PDL Language (EPDL) is defined as follows:*

$$\phi ::= \top \mid p \mid \neg\phi \mid (\phi \wedge \phi) \mid [\pi]\phi \mid K\phi$$
$$\pi ::= a \mid ?\phi \mid (\pi;\pi) \mid (\pi+\pi) \mid \pi^*$$

where $p \in P$, $a \in A$. We use $(\!|\pi|\!)\phi$ to denote $[\pi]\phi \wedge \langle\pi\rangle\phi$, which is logically equivalent to $[\pi]\phi \wedge \langle\pi\rangle\top$. Given a finite $B \subseteq A$, we write B^* for $(\Sigma_{a \in B}a)^*$, i.e., the iteration over the 'sum' of all the action symbols in B. The size of EPDL formulas/programs is given by: $|[\pi]\phi| = |\pi| + |\phi|$, $|a| = 1$, $|\pi_1;\pi_2| = 1 + |\pi_1| + |\pi_2|$, $|?\phi| = |\pi^*| = 1 + |\phi|$, and $|\pi_1 + \pi_2| = 1 + |\pi_1| + |\pi_2|$.

Given any uncertainty map $\mathscr{M} = \langle \mathscr{S}, \{\mathscr{R}_a \mid a \in A\}, \mathscr{V}, \mathscr{U}\rangle$, any state $s \in \mathscr{U}$, the semantics is given by a mutual induction on ϕ and π (we only show the case about $[\pi]\phi$, other cases are as in EAL):

$$
\begin{aligned}
\mathscr{M}, s \vDash [\pi]\phi &\Leftrightarrow \text{for all } \mathscr{M}', s' : (\mathscr{M}, s)[\![\pi]\!](\mathscr{M}', s') \\
&\qquad \text{implies } \mathscr{M}', s' \vDash \phi \\
(\mathscr{M}, s)[\![a]\!](\mathscr{M}', s') &\Leftrightarrow \mathscr{M}' = \mathscr{M}|a \text{ and } s \xrightarrow{a} s' \\
(\mathscr{M}, s)[\![?\psi]\!](\mathscr{M}', s') &\Leftrightarrow (\mathscr{M}', s') = (\mathscr{M}, s) \text{ and } \mathscr{M}, s \vDash \psi \\
(\mathscr{M}, s)[\![\pi_1;\pi_2]\!](\mathscr{M}', s') &\Leftrightarrow (\mathscr{M}, s)[\![\pi_1]\!] \circ [\![\pi_2]\!](\mathscr{M}', s') \\
(\mathscr{M}, s)[\![\pi_1 + \pi_2]\!](\mathscr{M}', s') &\Leftrightarrow (\mathscr{M}, s)[\![\pi_1]\!] \cup [\![\pi_2]\!](\mathscr{M}', s') \\
(\mathscr{M}, s)[\![\pi^*]\!](\mathscr{M}', s') &\Leftrightarrow (\mathscr{M}, s)[\![\pi]\!]^\star(\mathscr{M}', s')
\end{aligned}
$$

where $\circ, \cup, {}^\star$ at the right-hand side denote the usual composition, union and reflexive transitive closure of binary relations respectively. Clearly this semantics coincides with the semantics of EAL on EAL formulas.

Note that each program π can be viewed as a set of computation sequences, which are sequences of actions in A and tests with $\phi \in$ EPDL:

$$\mathcal{L}(a) = \{a\}$$
$$\mathcal{L}(?\phi) = \{?\phi\}$$
$$\mathcal{L}(\pi;\pi') = \{\sigma\eta \mid \sigma \in \mathcal{L}(\pi) \text{ and } \eta \in \mathcal{L}(\pi')\}$$
$$\mathcal{L}(\pi+\pi') = \mathcal{L}(\pi) \cup \mathcal{L}(\pi')$$
$$\mathcal{L}(\pi^*) = \{\varepsilon\} \cup \bigcup_{n>0}(\mathcal{L}(\underbrace{\pi\cdots\pi}_{n})) \text{ where } \varepsilon \text{ is the empty sequence}$$

Here are some valid formulas which are useful in our latter discussion:

$$\langle\pi;\pi'\rangle\phi \leftrightarrow \langle\pi\rangle\langle\pi'\rangle\phi$$
$$[\pi+\pi']\phi \leftrightarrow [\pi]\phi \wedge [\pi']\phi$$
$$[?\psi]\phi \leftrightarrow (\psi \rightarrow \phi)$$

We leave the complete axiomatization of EPDL on uncertainty maps to future work.

3.2   Conformant planning via model checking EPDL

Definition 3.2 (Conformant planning) *Given an uncertainty map $\mathcal{M}$, a goal formula $\phi \in$ EPDL, and a set $B \subseteq A$, the conformant planning problem is to find a finite (possibly empty) sequence $\sigma = a_1 a_2 \cdots a_n \in \mathcal{L}(B^*)$ such that for each $u \in \mathcal{U}_{\mathcal{M}}$ we have $\mathcal{M}, u \vDash (\!|a_1|\!)(\!|a_2|\!) \cdots (\!|a_n|\!)\phi$. The existence problem of conformant planning is to test whether such a sequence exists.*

Recall that $(\!|\pi|\!)\phi$ is the shorthand of $[\pi]\phi \wedge \langle\pi\rangle\phi$. Intuitively, we want a plan which is both executable and safe w.r.t. non-deterministic actions and initial uncertainty of the agent. It is crucial to observe the difference between $(\!|a_1|\!)(\!|a_2|\!) \cdots (\!|a_n|\!)\phi$ and $(\!|a_1;a_2;\cdots;a_n|\!)\phi$ by the following example:

Example 1 *Given uncertainty map $\mathcal{M}$ depicted as follows, we have $\mathcal{M}, s_1 \vDash (\!|a;b|\!)p$ but $\mathcal{M}, s_1 \nvDash (\!|a|\!)(\!|b|\!)p$.*
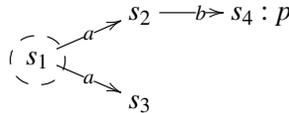
Given $\mathcal{M}$ and ϕ, to verify whether $\sigma \in \mathcal{L}(\pi)$ is a conformant plan can be formulated as the model checking problem: $\mathcal{M}, u \vDash K(\!|a_1|\!)(\!|a_2|\!) \cdots (\!|a_n|\!)\phi$. On the other hand, the existence problem of a conformant plan is more complicated to formulate: it asks whether there *exists* a $\sigma \in \mathcal{L}(B^*)$ such that it can be verified as a conformant plan. The simple-minded attempt would be to check whether $\mathcal{M}, u \vDash K\langle B^*\rangle\phi$ holds. Despite the $\langle\cdot\rangle$-vs.-$(\!|\cdot|\!)$ distinction, $K\langle B^*\rangle\phi$ may hold on a model where the sequences to guarantee ϕ on different states in $\mathcal{U}_M$ are different, as the following example shows:

Example 2 *Given uncertainty map $\mathcal{M}$ depicted as follows, let the goal formula be p and $B = \{a,b\}$. We have $\mathcal{M}, s_1 \vDash K\langle B^*\rangle p$, but there is no solution to this conformant planning problem.*
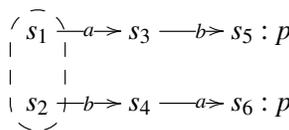
The right formula to check for the existence of a conformant plan w.r.t. $B \subseteq A$ and $\phi \in \text{EPDL}$ is:

$$\theta_{B,\phi} = \langle (\Sigma_{a \in B}(?K\langle a \rangle \top; a))^* \rangle K\phi.$$

For example, if $B = \{a_1, a_2\}$ then $\theta_{B,\phi} = \langle ((?K\langle a_1 \rangle \top; a_1) + (?K\langle a_2 \rangle \top; a_2))^* \rangle K\phi$. Intuitively, the confrmant plan consists of actions that are always executable given the uncertainty of the agent (guaranteed by the guard $K\langle a \rangle \top$). In the end the plan should also make sure that ϕ must hold given the uncertainty of the agent (guaranteed by $K\phi$). In the following, we will prove that this formula is indeed correct.

First, we observe that the rule of substitution of equivalents is valid ($\phi(\psi/\chi)$ is obtained by replacing any occurrence of χ by ψ, similar for $[\![\pi(\psi/\chi)]\!]$):

Proposition 3.1 *If* $\vDash \psi \leftrightarrow \chi$, *then:*

(1) $\vDash \phi \leftrightarrow \phi(\psi/\chi)$;

(2) $[\![\pi]\!] = [\![\pi(\psi/\chi)]\!]$.

Proposition 3.2 $\vDash K(\!|a|\!)\phi \leftrightarrow \langle ?K\langle a \rangle \top; a \rangle K\phi$

Proof: Since $\vDash K(\!|a|\!)\phi \leftrightarrow (K[a]\phi \wedge K\langle a \rangle \phi)$ and $\vDash (K\langle a \rangle \top \wedge \langle a \rangle K\phi) \leftrightarrow \langle ?K\langle a \rangle \top; a \rangle K\phi$, we only need to show that $\vDash (K[a]\phi \wedge K\langle a \rangle \phi) \leftrightarrow (K\langle a \rangle \top \wedge \langle a \rangle K\phi)$.

Left to right:

(L1) $\vDash K[a]\phi \rightarrow [a]K\phi$, by validity of Axiom $\text{PR}(a)$

(L2) $\vDash K\langle a \rangle \phi \rightarrow \langle a \rangle \top \wedge K\langle a \rangle \top$, by semantics

(L3) $\vDash \langle a \rangle \top \wedge [a]K\phi \rightarrow \langle a \rangle K\phi$, by semantics

(L4) $\vDash K[a]\phi \wedge K\langle a \rangle \phi \rightarrow K\langle a \rangle \top \wedge \langle a \rangle K\phi$, by (L1)-(L3)

Right to left:

(R1) $\vDash \langle a \rangle K\phi \rightarrow K[a]\phi$, by validity of Axiom $\text{NM}(a)$

(R2) $\vDash K[a]\phi \wedge K\langle a \rangle \top \rightarrow K\langle a \rangle \phi$, by semantics

(R3) $\vDash K\langle a \rangle \top \wedge \langle a \rangle K\phi \rightarrow K[a]\phi \wedge K\langle a \rangle \phi$, by R(1)-R(2) ∎

Lemma 3.1 *For any* $a_1 a_2 \cdots a_n \in \mathscr{L}(A^*)$:

$$\vDash K(\!|a_1|\!)(\!|a_2|\!) \cdots (\!|a_n|\!)\phi \leftrightarrow \langle ?K\langle a_1 \rangle \top; a_1; \ldots; ?K\langle a_n \rangle \top; a_n \rangle K\phi$$

Proof: It is trivial when $n = 0$ (i.e., the sequence is ε), since the claim then boils down to $K\phi \leftrightarrow K\phi$. We prove the non-trivial cases by induction on $n \geq 1$. When $n = 1$, it follows from Proposition 3.2. Now, as the induction hypothesis, we assume that:

$$\vDash K(\!|a_1|\!)(\!|a_2|\!) \cdots (\!|a_k|\!)\phi \leftrightarrow \langle ?K\langle a_1 \rangle \top; a_1; \ldots; ?K\langle a_k \rangle \top; a_k \rangle K\phi.$$

We need to show:

$$\vDash K(\!|a_1|\!)(\!|a_2|\!) \cdots (\!|a_{k+1}|\!)\phi \leftrightarrow \langle ?K\langle a_1 \rangle \top; a_1; \ldots; ?K\langle a_{k+1} \rangle \top; a_{k+1} \rangle K\phi.$$

By IH,

$$\vDash K(\!|a_1|\!)(\!|a_2|\!) \cdots (\!|a_{k+1}|\!)\phi \leftrightarrow \langle ?K\langle a_1 \rangle \top; a_1; \ldots; ?K\langle a_k \rangle \top; a_k \rangle K(\!|a_{k+1}|\!)\phi. \qquad (1)$$

Due to Propositions 3.1 and 3.2, we have:

$$\vDash \langle ?K\langle a_1 \rangle \top; a_1; \ldots; ?K\langle a_k \rangle \top; a_k \rangle K(\!|a_{k+1}|\!)\phi \leftrightarrow \langle ?K\langle a_1 \rangle \top; a_1; \ldots; ?K\langle a_n \rangle \top; a_k \rangle \langle ?K\langle a_{k+1} \rangle \top; a_{k+1} \rangle K\phi. (2)$$

The conclusion is immediate by combining (1) and (2). ∎

The following theorem follows from the above lemma.

Theorem 3.1 *Given a pointed uncertainty map $\mathscr{M},s$, an EPDL formula ϕ and a set $B \subseteq A$, the following two are equivalent:*

(1) *There is a $\sigma = a_1 \ldots a_n \in \mathscr{L}(B^*)$ such that $\mathscr{M},s \vDash K(\!|a_1|\!)(\!|a_2|\!) \cdots (\!|a_n|\!)\phi$;*

(2) *$\mathscr{M},s \vDash \langle (\Sigma_{a \in B}(?K\langle a\rangle\top;a))^* \rangle K\phi$.*

We would like to emphasise that the K operator right before ϕ in the definition of $\theta_{B,\phi}$ cannot be omitted, as demonstrated by the following example:

Example 3 *Given uncertainty map $\mathscr{M}$ depicted as follows, let the goal formula be p. As we can see, there is no solution to this conformant planning problem. Indeed $\mathscr{M},s_1 \nvDash \langle (\Sigma_{a \in B}(?K\langle a\rangle\top;a))^* \rangle Kp$ with $B = \{a,b\}$, but we could have $\mathscr{M},s_1 \vDash \langle (\Sigma_{a \in B}(?K\langle a\rangle\top;a))^* \rangle p$.*
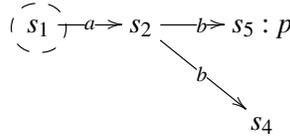
We close this section with an example about planning with both positive and negative epistemic goals (the agent should know something, but not too much).

Example 4 *Given uncertainty map $\mathscr{M}$ depicted as follows, let the goal be Kp then both a and b are conformant plans. If the goal is $Kp \wedge \neg Kq$, only a is a good plan.*
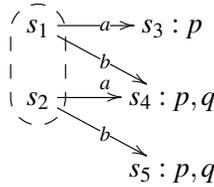
4 Model checking EPDL: complexity and algorithms

In this section, we first focus on the model checking problem of the following star-free fragment of EPDL (call it EPDL$^-$):

$$\phi ::= \top \mid p \mid \neg\phi \mid (\phi \wedge \phi) \mid [\pi]\phi \mid K\phi$$
$$\pi ::= a \mid ?\phi \mid (\pi;\pi) \mid (\pi + \pi)$$

We will show that model checking EPDL$^-$ is PSPACE-complete. In particular, the upper bound is shown by making use of an alternative context-dependent semantics. Then we give an EXPTIME algorithm for the model checking problem of the full EPDL inspired by another alternative semantics based on 2-dimensional models. Finally we give a PSPACE algorithm for the conformant planning problem in EPDL. Note that throughout this section, we focus on uncertainty maps with finitely many states and assume $\mathscr{R}_a = \emptyset$ for *co-finitely* many $a \in A$.

4.1 Complexity of model checking EPDL$^-$

4.1.1 Lower Bound

To show the PSPACE lower bound, we provide a polynomial reduction of QBF (*quantified Boolean formula*) truth testing to the model checking problem of EPDL$^-$. Note that to determine whether a given

QBF (even in prenex normal form based on a conjunctive normal form) is true or not is known to be PSPACE-complete [30]. Our method is inspired by [28] which discusses the complexity of model checking temporal logics with past operators. Surprisingly, we can use the uncertainty sets to encode the 'past' and use the dual of the knowledge operator to 'go back' to the past. This intuitive idea will become more clear in the proof.

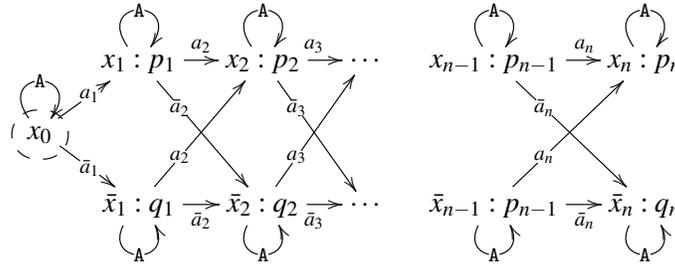QBF formulas are $Q_1 x_1 Q_2 x_2 \ldots Q_n x_n \phi(x_1, \ldots, x_n)$ where:

- For $1 \leq n \leq n, Q_i$ is $\exists$ if i is odd, and Q_i is $\forall$ if i is even.

- ϕ is a propositional formula in CNF based on variables $x_1, \ldots, x_n$,

For each such QBF α with n variables, we need to find a pointed model $\mathcal{M}_n, x_0$ and a formula θ_α such that α is true iff $\mathcal{M}_n, x_0 \models \theta_\alpha$. The model $\mathcal{M}_n$ is defined below.

Definition 4.1 *Let $A = \{a_i, \bar{a}_i \mid i \geq 1\}$ and $P = \{p_k, q_k \mid k \geq 1\}$, the uncertainty map $\mathcal{M}_n = \langle \mathcal{S}, \{\mathcal{R}_a \mid a \in A\}, \mathcal{V}, \mathcal{U} \rangle$ is defined as:*

- $\mathcal{S} = \{x_0\} \cup \{x_i \mid 1 \leq i \leq n\} \cup \{\bar{x}_i \mid 1 \leq i \leq n\}$

- $\mathcal{V}(x_0) = \emptyset$, *and* $\mathcal{V}(x_i) = \{p_i\}, \mathcal{V}(\bar{x}_i) = \{q_i\}$ *for* $1 \leq i \leq n$.

- $\overset{a_i}{\rightarrow} = \{(s,s) \mid s \in \mathcal{S}\} \cup \{(x_{i-1}, x_i), (\bar{x}_{i-1}, x_i)\}$

- $\overset{\bar{a}_i}{\rightarrow} = \{(s,s) \mid s \in \mathcal{S}\} \cup \{(x_{i-1}, \bar{x}_i), (\bar{x}_{i-1}, \bar{x}_i)\}$

- $\mathcal{U} = \{x_0\}$

$|\mathcal{M}_n|$ is linear in n and can be depicted as the following:

Given $\alpha = Q_1 x_1 Q_2 x_2 \ldots Q_n x_n \phi(x_1, \ldots, x_n)$, the formula θ_α is defined as
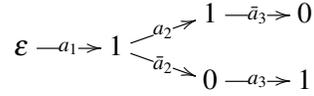
$$QT_1 \cdots QT_n \psi(\hat{K}p_1, \cdots, \hat{K}p_n, \hat{K}q_1, \cdots, \hat{K}q_n)$$

where QT_i is $\langle (a_i + \bar{a}_i); ?(p_i \vee q_i) \rangle$ if i is odd and QT_i is $[(a_i + \bar{a}_i); ?(p_i \vee q_i)]$ if i is even, and ψ is obtained from $\phi(x_1, \ldots, x_n)$ by replacing each x_i with $\hat{K}p_i$ and $\neg x_i$ with $\hat{K}q_i$.

To ease the latter proof, we first define the valuation tree below.

Definition 4.2 (V-tree) *A V-tree τ is a rooted tree such that 1) each node is 0 or 1 (except the root ε); 2) each internal node in an even level has only one successor; 3) each internal node in an odd level has two successors: one is 0 and the other one is 1; 4) each edge to node 0 of level i is labelled $\bar{a}_i$; 5) each edge to node 1 of level i is labelled a_i. Given a V-tree with depth n, a path σ is a sequence of $A_1 \ldots A_n$ where $A_i = a_i$ or $A_i = \bar{a}_i$. A path σ can also be seen as a valuation assignment for $x_1, \ldots, x_n$ with the convention that $\sigma(x_i) = 1$ if a_i occurs in σ and $\sigma(x_i) = 0$ if $\bar{a}_i$ occurs in σ. Let path(τ) be the set of all paths of τ.*

As an example, a V-tree τ can be depicted as below:

$$\varepsilon \longrightarrow a_1 \rightarrow 1 \begin{array}{c} \overset{a_2}{\nearrow} 1 \longrightarrow \bar{a}_3 \rightarrow 0 \\ \underset{\bar{a}_2}{\searrow} 0 \longrightarrow a_3 \rightarrow 1 \end{array}$$

It is not hard to see the following:

Proposition 4.1 *For each $1 \le i \le n$, we have: $\alpha = Q_1 x_1 \dots Q_i x_i Q_{i+1} x_{i+1} \dots Q_n x_n \phi$ is true iff there exists a V-tree τ with depth i such that for each $\sigma \in path(\tau)$ $\sigma(Q_{i+1} x_{i+1} \dots Q_n x_n \phi) = 1$ (σ as a valuation).*

Now let us see the update result of running a path $\sigma \in path(\tau)$ on $\mathcal{M}_n$. Due to the lack of space, we omit the proofs of the following two propositions.

Proposition 4.2 *Given $\mathcal{M}_n$, let $\sigma = A_1 \dots A_i$ $(1 \le i \le n)$ be a sequence of actions such that $A_k = a_k$ or $A_k = \bar{a}_k$ for each $1 \le k \le i$, then we have $\mathcal{U}|^\sigma = \{x_0, X_1, \dots, X_i\}$ where $X_k = x_k$ if $A_k = a_k$ else $X_k = \bar{x}_k$ for each $1 \le k \le i$.*

Given $\sigma = A_1 \dots A_n$ where A_i is a_i or $\bar{a}_i$ for each $1 \le i \le n$, let $g(\sigma) = x_n$ if $A_n = a_n$ and $g(\sigma) = \bar{x}_n$ if $A_n = \bar{a}_n$. By Proposition 4.2, we always have $g(\sigma) \in \mathcal{U}_{\mathcal{M}_k}|^\sigma$ with $k > n$. Thus given $\mathcal{M}_k$ and $\sigma = A_1 \dots A_n$ and $k > n$, $\mathcal{M}_k|^\sigma, g(\sigma)$ is a pointed uncertainty map.

Proposition 4.3 *For each $1 \le i \le n$, we have $\mathcal{M}_k, x_0 \vDash QT_1 \dots QT_i QT_{i+1} \dots QT_n \psi$ iff there exists a V-tree τ with depth i such that $\mathcal{M}_k|^\sigma, g(\sigma) \vDash QT_{i+1} \dots QT_n \psi$ for each $\sigma \in path(\tau)$, where $k > n$ and $g(\sigma)$ is the state corresponds to the last edge of σ, e.g., $g(a_1 \bar{a}_2) = \bar{x}_2$.*

Theorem 4.1 *The following two are equivalent:*

- $\alpha = Q_1 x_1 Q_2 x_2 \dots Q_n x_n \phi(x_1, \dots, x_n)$ *is true*
- $\mathcal{M}_n, x_0 \vDash QT_1 \cdots QT_n \psi(\hat{K} p_1 \cdots \hat{K} p_n, \hat{K} q_1 \cdots \hat{K} q_n)$ *in which ψ is obtained from ϕ by replacing each x_i with $\hat{K} p_i$ and $\neg x_i$ with $\hat{K} q_i$.*

Proof: By Propositions 4.1 and 4.3, we only need to show that given V-tree τ with depth n, $\sigma(\phi) = 1$ if and only if $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \psi$ for each $\sigma \in path(\tau)$. Since ϕ is in CNF, ψ is also in CNF-like form obtained by replacing each x_i with $\hat{K} p_i$ and each $\neg x_i$ with $\hat{K} q_i$ for $1 \le i \le n$. Thus we only need to show that $\sigma(x_i) = 1$ iff $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \hat{K} p_i$ and $\sigma(\neg x_i) = 1$ iff $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \hat{K} q_i$. Since $\sigma(x_i) = 1$ iff $\sigma(\neg x_i) = 0$, we only need to show that $\sigma(x_i) = 1$ iff $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \hat{K} p_i$ and $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \hat{K} p_i$ iff $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \neg \hat{K} q_i$. By the definition of τ, we know that $\sigma = A_1 \dots A_n$ where A_i is a_i or $\bar{a}_i$ for each $1 \le i \le n$.

Firstly, we will show that $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \hat{K} p_i$ if and only if $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \neg \hat{K} q_i$. To verify the right-to-left direction, if $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \hat{K} p_i$, it follows by the definition of $\mathcal{M}_n$ that $x_i \in \mathcal{U}|^\sigma$. Then it must be the case that a_i occurs in σ. Suppose not, $\bar{a}_i$ occurs in σ. It follows by Proposition 4.2, $\mathcal{U}|^\sigma = \{x_0, X_1, \dots, X_{i-1}, \bar{x}_i, X_{i+1}, \dots, X_n\}$. This is contrary with $x_i \in \mathcal{U}|^\sigma$. Thus it must be that a_i occurs in σ. It follows by Proposition 4.2 that $\mathcal{U}|^\sigma = \{x_0, X_1, \dots, X_{i-1}, x_i, X_{i+1}, \dots, X_n\}$. Thus $\bar{x}_i \notin \mathcal{U}|^\sigma$. By the definition of $\mathcal{M}_n$ and the semantics, we have $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \neg \hat{K} q_i$. To verify the left-to-right direction, $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \neg \hat{K} q_i$ implies that $\bar{x}_i \notin \mathcal{U}|^\sigma$. For the similar reason as above, it must be the case that $\bar{a}_i$ does not occur in σ. Thus we have that a_i occurs in σ. It follows by Proposition 4.2 that $x_i \in \mathcal{U}|^\sigma$. Thus we have $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \hat{K} p_i$.

Next we will show that $\sigma(x_i) = 1$ iff $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \hat{K} p_i$. To verify the right-to-left direction, $\sigma(x_i) = 1$ implies that $A_i = a_i$. It follows by Proposition 4.2 that $x_i \in \mathcal{U}|^\sigma$. Thus we have $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \hat{K} p_i$. To verify the left-to-right direction, we will show that $\sigma(x_i) = 0$ implies $\mathcal{M}_n|^\sigma, g(\sigma) \vDash \hat{K} q_i$. It follows by the definition of $\sigma(x_i) = 0$ that $A_i = \bar{a}_i$. It follows by Proposition 4.2 that $\bar{x}_i \in \mathcal{U}|^\sigma$. Thus we have

$\mathcal{M}_n|^{\sigma}, g(\sigma) \vDash \hat{K} q_i.$ ∎

This gives us the desired lower bound:

Theorem 4.2 *The model checking problem for* EPDL⁻ *is* PSPACE-*hard.*

4.1.2   Upper Bound

In this section we give a non-trivial model checking algorithm for EPDL⁻ inspired by an equivalent semantics.

As we mentioned earlier, the semantics of EPDL is 'context-dependent': reaching the same state through different paths may affect the truth value of an epistemic subformula. This means that the usual global model checking algorithm for modal logics may not work here. In order to establish the upper bound, we first give the following equivalent semantics to EPDL⁻ which makes the context dependency explicit in order to facilitate a local model checking algorithm. The idea is to keep the model intact but record the scope of action modalities in order to compute the right uncertainty set for epistemic subformulas. Similar idea appeared in [32] to give an alternative semantics of public announcement logic.

Definition 4.3 *Given an uncertainty map* $\mathcal{M} = \langle \mathcal{S}, \{\mathcal{R}_a \mid a \in A\}, \mathcal{V}, \mathcal{U} \rangle$ *and any state* $s \in \mathcal{S}$, *the satisfaction relation* ⊩ *is defined using the auxiliary satisfaction relation* ⊩$_\sigma$ *and auxiliary relation* $\overset{\omega_\sigma}{\to}$, *where* σ *is a finite (possibly empty) sequence of actions in* A:

$$
\begin{array}{ll}
\mathcal{M}, s \Vdash \phi & \Leftrightarrow \mathcal{M}, s \Vdash_\varepsilon \phi \\
\mathcal{M}, s \Vdash_\sigma \top & \Leftrightarrow \text{always} \\
\mathcal{M}, s \Vdash_\sigma p & \Leftrightarrow p \in \mathcal{V}(s) \\
\mathcal{M}, s \Vdash_\sigma \neg\phi & \Leftrightarrow \mathcal{M}, s \nVdash_\sigma \phi \\
\mathcal{M}, s \Vdash_\sigma \phi \wedge \psi & \Leftrightarrow \mathcal{M}, s \Vdash_\sigma \phi \text{ and } \mathcal{M}, s \Vdash_\sigma \psi \\
\mathcal{M}, s \Vdash_\sigma K\phi & \Leftrightarrow \text{for all } v \in \mathcal{U}|^\sigma : \mathcal{M}, v \Vdash_\sigma \phi \\
\mathcal{M}, s \Vdash_\sigma \langle \pi \rangle \phi & \Leftrightarrow \text{there exists } \omega \in \mathcal{L}(\pi) \text{ and } t \in \mathcal{S} \\
 & \qquad \text{such that } s \overset{\omega_\sigma}{\to} t \text{ and } \mathcal{M}, t \Vdash_{\sigma r(\omega)} \phi \\
s \overset{\varepsilon_\sigma}{\to} t & \Leftrightarrow s = t \\
s \overset{(a\omega')_\sigma}{\to} t & \Leftrightarrow \text{there exists } s' \text{ such that } s \overset{a}{\to} s' \text{ and } s' \overset{\omega'_{(\sigma a)}}{\to} t \\
s \overset{(?\phi\omega')_\sigma}{\to} t & \Leftrightarrow \mathcal{M}, s \Vdash_\sigma \phi \text{ and } s \overset{\omega'_\sigma}{\to} t
\end{array}
$$

where $r(\omega)$ *is the sequence of actions obtained by eliminating all the tests in* ω.

Note that ω in the above definition is a computation sequence, i.e., a finite sequence of actions and EPDL⁻-tests, while σ is a test-free sequence of actions.

The following can be proved by induction on η:

Proposition 4.4 *Given an uncertainty map* $\mathcal{M}$ *and sequences of actions and tests* η, ω, ω' *such that* $\eta = \omega\omega'$, *we have* $(s,t) \in \overset{\eta_\sigma}{\to}$ *iff* $(s,t) \in \overset{\omega_\sigma}{\to} \circ \overset{\omega'_{\sigma r(\omega)}}{\to}$ *for any sequence of actions* σ.

Proof: We prove it by induction on $|\eta|$. If $|\eta| \leq 2$, it is obvious by the definition. If $|\eta| > 2$, there are two cases, that is, $\eta = a\eta'$ or $\eta = ?\phi\eta'$.

Case $\eta = a\eta'$: We have $\omega = a\omega''$ for some initial segment ω'' of η', and $(s,t) \in \overset{(a\eta')_\sigma}{\to}$ iff there exists s' such that $s \overset{a}{\to} s'$ and $(s',t) \in \overset{\eta'_{\sigma a}}{\to}$. By IH, we have $\overset{\eta'_{\sigma a}}{\to} = \overset{\omega''_{\sigma a}}{\to} \circ \overset{\omega'_{\sigma a r(\omega'')}}{\to}$. Thus we have $(s',t) \in \overset{\eta'_{\sigma a}}{\to}$ iff there

exists t' such that $(s',t') \in \overset{\omega''_{\sigma a}}{\rightarrow}$ and $(t',t) \in \overset{\omega'_{\sigma ar(\omega'')}}{\rightarrow}$. By definition, we have that $s \overset{a}{\rightarrow} s'$ and $(s',t') \in \overset{\omega''_{\sigma a}}{\rightarrow}$ iff $(s,t') \in \overset{a\omega''_{\sigma}}{\rightarrow}$. Thus we have $(s,t) \in \overset{a\omega''_{\sigma}}{\rightarrow} \circ \overset{\omega'_{\sigma ar(\omega'')}}{\rightarrow}$, namely $(s,t) \in \overset{\omega_{\sigma}}{\rightarrow} \circ \overset{\omega'_{\sigma r(\omega)}}{\rightarrow}$.

Case $\eta = ?\phi\eta'$: We have $\omega = ?\phi\omega''$ for some initial segment ω'' of η', and $(s,t) \in \overset{(?\phi\eta')_{\sigma}}{\rightarrow}$ iff $\mathcal{M}, s \Vdash_{\sigma} \phi$ and $s \overset{\eta'_{\sigma}}{\rightarrow} t$. By IH, we have $s \overset{\eta'_{\sigma}}{\rightarrow} t$ iff $(s,t) \in \overset{\omega''_{\sigma}}{\rightarrow} \circ \overset{\omega'_{\sigma r(\omega'')}}{\rightarrow}$. Thus we have there exists s' such that $(s,s') \in \overset{\omega''_{\sigma}}{\rightarrow}$ and $(s',t) \in \overset{\omega'_{\sigma r(\omega'')}}{\rightarrow}$. This follows that $(s,s') \in \overset{(?\phi\omega'')_{\sigma}}{\rightarrow}$, and $(s,t) \in \overset{(?\phi\omega'')_{\sigma}}{\rightarrow} \circ \overset{\omega'_{\sigma r(?\phi\omega'')}}{\rightarrow}$, namely $(s,t) \in \overset{\omega_{\sigma}}{\rightarrow} \circ \overset{\omega'_{\sigma r(\omega)}}{\rightarrow}$. ∎

In the following we show that $\Vdash$ coincides with $\vDash$.

Theorem 4.3 *Given an uncertainty map $\mathcal{M}$ and an action sequence σ, if $\mathcal{U}|^{\sigma} \neq \emptyset$, we have that for each $s \in \mathcal{U}|^{\sigma}$,*

(i) $\mathcal{M}|^{\sigma}, s[\![\pi]\!]\mathcal{M}', s'$ iff there exists $\omega \in \mathcal{L}(\pi)$ such that $\mathcal{M}' = \mathcal{M}|^{\sigma r(\omega)}$ and $s \overset{\omega_{\sigma}}{\rightarrow} s'$,

(ii) $\mathcal{M}|^{\sigma}, s \vDash \phi$ iff $\mathcal{M}, s \Vdash_{\sigma} \phi$.

Proof: The proof is by simultaneous induction on π and ϕ (due to the test actions). For (i), we will only focus on the case of $\pi_1; \pi_2$; the other cases are straightforward.

Case $\pi_1; \pi_2$: We only show the direction from left to right; the other direction is similar. It follows by assumption that there is pointed uncertainty map $\mathcal{N}, t$ such that $\mathcal{M}|^{\sigma}, s[\![\pi_1]\!]\mathcal{N}, t$ and $\mathcal{N}, t[\![\pi_2]\!]\mathcal{M}', s'$. By IH, we have that there exists $\omega \in \mathcal{L}(\pi_1)$ such that $\mathcal{N} = \mathcal{M}|^{\sigma r(\omega)}$ and $s \overset{\omega_{\sigma}}{\rightarrow} t$. Since $\mathcal{N}, t$ is a pointed uncertainty map and $\mathcal{N} = \mathcal{M}|^{\sigma r(\omega)}$, we have $t \in \mathcal{U}|^{\sigma r(\omega)}$. By IH and $\mathcal{M}|^{\sigma r(\omega)}, t[\![\pi_2]\!]\mathcal{M}', s'$, we have that there exists $\omega' \in \mathcal{L}(\pi_2)$ such that $\mathcal{M}|^{\sigma r(\omega)r(\omega')} = \mathcal{M}|^{\sigma r(\omega\omega')} = \mathcal{M}'$ and $t \overset{\omega'_{\sigma r(\omega)}}{\rightarrow} s'$. By Proposition 4.4, it follows that $\omega\omega' \in \mathcal{L}(\pi_1; \pi_2)$ and $s \overset{(\omega\omega')_{\sigma}}{\rightarrow} s'$.

For (ii), we will focus on the case of $\langle\pi\rangle\phi$; the other cases are straightforward.

Case $\langle\pi\rangle\phi$: We have $\mathcal{M}|^{\sigma}, s \vDash \langle\pi\rangle\phi$ if and only if there is pointed uncertainty map $\mathcal{M}', s'$ such that $\mathcal{M}|^{\sigma}, s[\![\pi]\!]\mathcal{M}', s'$ and $\mathcal{M}', s' \vDash \phi$. By (i), it follows that $\mathcal{M}|^{\sigma}, s[\![\pi]\!]\mathcal{M}', s'$ iff there exists $\omega \in \mathcal{L}(\pi)$ such that $\mathcal{M}' = \mathcal{M}|^{\sigma r(\omega)}$ and $s \overset{\omega_{\sigma}}{\rightarrow} s'$. By IH, it follows that $\mathcal{M}|^{\sigma r(\omega)}, s' \vDash \phi$ iff $\mathcal{M}, s' \Vdash_{\sigma r(\omega)} \phi$. Thus we have $\mathcal{M}, s \Vdash \langle\pi\rangle\phi$. ∎

Let σ be ε, we have the equivalence of $\Vdash$ and $\vDash$.

Corollary 4.1 *Given pointed uncertainty map $\mathcal{M}, s$, we have $\mathcal{M}, s \vDash \phi$ iff $\mathcal{M}, s \Vdash \phi$ for each $\phi \in$ EPDL⁻.*

This alternative semantics induces a natural algorithm to compute the truth value of an EPDL⁻ formula w.r.t. to a pointed uncertainty map. The idea is to recursively call a function $MC(\mathcal{M}, s, \sigma, \phi)$ which returns the truth value of a subformula ϕ on state s given the context of σ while keeping $\mathcal{M}$ intact. Note that, we do not need to compute all the $MC(\mathcal{M}, s, \sigma, \phi)$ for each σ and each subformula ϕ. The only tricky part comes when evaluating $\langle\pi\rangle\phi$ formulas since it is too space consuming to compute the whole set of $\mathcal{L}(\pi)$ in the search of the right ω. Instead, we can generate one by one in some lexicographical order all the possible sequences up to a bound based on the atomic actions and tests occurring in the formula, and then test whether it belongs to the program π. Note that in this way, we can use the space repeatedly, and the membership testing of $\mathcal{L}(\pi)$ is not expensive (NLOGSPACE-complete according to [20]).

In the appendix we present three algorithms based on matrix representation of the model: Algorithm 1 computes the uncertainty set $\mathcal{U}|^{\sigma}$; Algorithm 2 computes $\overset{w_{\sigma}}{\rightarrow}$ and Algorithm 3 is the main model

checking algorithm. Note that Algorithms 2 and 3 involve mutual recursion of each other due to the tests in programs. However, the depth of the recursion is bounded by the length of the formula, and for each call polynomial space suffices. The detailed algorithms and complexity analysis can be found in the appendix. It is not hard to show the following (based on Theorem 4.2)

Theorem 4.4 (Upper bound) *The model checking problem of EPDL$^-$ is in* PSPACE. *Thus it is* PSPACE-*complete.*

4.2   Upper Bounds for model checking EPDL

In this section, we give an EXPTIME model checking method for the full EPDL via model checking EPDL over two-dimensional models with both epistemic and action relations. Let us first define such models.

Definition 4.4 (Epistemic Temporal Structure) *An Epistemic Temporal Structure (ETS) is a Kripke model with both epistemic and action relations. Formally, an ETS model $\mathfrak{M}$ is a tuple $\langle \mathscr{S}, \{\mathscr{R}_a \mid a \in A\}, \sim, \mathscr{V} \rangle$, where $\mathscr{R}_a$ is a binary relation on $\mathscr{S}$, $\sim$ is an equivalence relation on $\mathscr{S}$ and $\mathscr{V} : \mathscr{S} \to 2^P$ is a valuation function.*

Now we define an alternative semantics of EPDL over ETSs.[6]

Definition 4.5 (ETS Semantics) *Given any ETS model $\mathfrak{M} = \langle \mathscr{S}, \{\mathscr{R}_a \mid a \in A\}, \sim, \mathscr{V} \rangle$ and any state $s \in \mathscr{S}$, the satisfaction relation for EPDL formulas is defined as follows (the Boolean cases are as in the standard modal logic):*

$$
\begin{array}{ll}
\mathfrak{M}, s \Vdash K\phi & \Leftrightarrow \forall u \in \mathscr{S} : s \sim u \text{ implies } \mathfrak{M}, u \Vdash \phi \\
\mathfrak{M}, s \Vdash [\pi]\phi & \Leftrightarrow \forall t \in S : s \xrightarrow{\pi} t \text{ implies } \mathfrak{M}, t \Vdash \phi \\
\xrightarrow{a} & = \mathscr{R}_a \\
\xrightarrow{?\phi} & = \{(s,s) \mid \mathfrak{M}, s \Vdash \phi\} \\
\xrightarrow{\pi_1 ; \pi_2} & = \xrightarrow{\pi_1} \circ \xrightarrow{\pi_2} \\
\xrightarrow{\pi_1 + \pi_2} & = \xrightarrow{\pi_1} \cup \xrightarrow{\pi_2} \\
\xrightarrow{\pi^*} & = (\xrightarrow{\pi})^\star
\end{array}
$$

where $\circ, \cup, \star$ at right-hand side denote the usual composition, union and reflexive transitive closure of binary relations respectively.

We can turn a Kripke model without the epistemic relation into an ETS model by essentially considering all the possible uncertainty sets.

Definition 4.6 *Given any Kripke model $\mathscr{M} = \langle \mathscr{S}, \{\mathscr{R}_a \mid a \in A\}, \mathscr{V} \rangle$, we define the ETS model $\mathscr{M}^\bullet$ as follows:*

$$
\begin{array}{ll}
\mathscr{S}^\bullet & = \{s_\Gamma \mid s \in \mathscr{S}, \Gamma \in 2^{\mathscr{S}}, s \in \Gamma\} \\
\mathscr{R}_a^\bullet & = \{(s_\Gamma, t_\Delta) \mid s \xrightarrow{a} t, \Delta = \Gamma|^a\} \\
\sim^\bullet & = \{(s_\Gamma, t_\Delta) \mid \Gamma = \Delta\} \\
\mathscr{V}^\bullet(s_\Gamma) & = \mathscr{V}(s)
\end{array}
$$

where $\Gamma|^a = \{t \in \mathscr{S} \mid \exists s \in \Gamma \text{ such that } s \xrightarrow{a} t\}$. For any Kripke model $\mathscr{M}$ and any $\Gamma \in 2^{\mathscr{S}} \setminus \{\emptyset\}$, let $\mathscr{M}^\Gamma$ be the uncertainty map $\langle \mathscr{M}, \Gamma \rangle$.

---

[6]Here we abuse the notation $\Vdash$ to denote the new semantics. Note that it is different from the alternative semantics in the previous section.

Note that each s_Γ can be viewed as an uncertainty set (Γ) with a designated state (s), and the definition of $\mathscr{R}_a$ captures the update in the $\models$ semantics of EPDL, and $\mathscr{M}^\bullet$ unravels all the updates in a whole picture. Note that the size of $\mathscr{M}^\bullet$ is $|\mathscr{S}| \cdot 2^{|\mathscr{S}|-1}$ where $\mathscr{S}$ is the set of states of $\mathscr{M}$.

Now we can show that $\models$ and $\Vdash$ coincide w.r.t. uncertainty map $\mathscr{M}^\Gamma$ and ETS model $\mathscr{M}^\bullet$ (the proofs are omitted due to the lack of space).

Proposition 4.5 *Given any map $\mathscr{M}$, we have*

  (i) $\mathscr{M}^\Gamma, s[\![\pi]\!]\mathscr{M}^\Delta, t$ *iff* $s_\Gamma \xrightarrow{\pi} t_\Delta$ *in* $\mathscr{M}^\bullet$;[7]

  (ii) $\mathscr{M}^\Gamma, s \models \phi$ *iff* $\mathscr{M}^\bullet, s_\Gamma \Vdash \phi$.

Corollary 4.2 *Given an uncertainty map $\mathscr{M} = \langle \mathscr{N}, \mathscr{U} \rangle$ and $s \in \mathscr{U}$, we have $\mathscr{M}, s \models \phi$ iff $\mathscr{N}^\bullet, s_\mathscr{U} \Vdash \phi$.*

Based on the above corollary we can have a model checking method via model checking EPDL over ETS models.

Proposition 4.6 *The model checking problem of EPDL on uncertainty maps is in* EXPTIME.

Proof: Given an uncertainty map $\mathscr{M} = \langle \mathscr{N}, \mathscr{U} \rangle$, the construction of ETS $\mathscr{N}^\bullet$ can be done in exponential time in the size of $\mathscr{N}$ due to the fact that there are at most $|\mathscr{N}|$ a-successors t_Δ of each s_Γ since $\Delta = \Gamma|^a$. By modifying the algorithm for PDL in [23], we can get an algorithm to check EPDL formula ϕ on $\mathscr{N}^\bullet$ w.r.t. $\Vdash$, and its time complexity is $O(|\phi|^2 \cdot |\mathscr{N}^\bullet|^3)$. Thus, the time complexity of model checking ϕ on $\mathscr{M}$ is bounded by $O(|\phi|^2 \cdot |\mathscr{S}_\mathscr{N}|^3 \cdot 2^{3|\mathscr{S}_\mathscr{N}|-3})$. ∎

We conjecture that the model checking problem of full EPDL is EXPTIME-complete, and leave the lower bound to the extended version of this paper.

4.3 Complexity of conformant planning

In the rest of this section, let us look at the complexity of conformant planning in terms of EPDL model checking. Although the model checking problem of full EPDL is likely to be EXPTIME-complete, the complexity of model checking the EPDL formula which encodes the conformant planning problem (cf. Theorem 3.1) is in PSPACE if the goal formula is program-free. More precisely, we can show the following:

Theorem 4.5 *The problem of model checking EPDL formulas in the shape of $\langle (\Sigma_{a \in B}(?K\langle a \rangle \top; a))^* \rangle K\phi$, where ϕ is an epistemic formula (i.e. program-free) and $B \subseteq A$, is in* PSPACE.

Proof: (Sketch) Note that $(\Sigma_{a \in B}(?K\langle a \rangle \top; a))^*$ is a special program which has only simple epistemic tests depending on the structure of the underlying Kripke model. Now given a Kripke model $\mathscr{N}$ and a set $B \subseteq A$ we can define an ETS model $\mathscr{N}^\circ$ similar to $\mathscr{N}^\bullet$ but with a different definition for the action relations:

$$\mathscr{R}_a^\circ = \{(s_\Gamma, t_\Delta) \mid s \xrightarrow{a} t, \Delta = \Gamma|^a, \forall u \in \Gamma \exists v \text{ st. } u \xrightarrow{a} v.\}$$

Note that the extra condition guarantees that the action a is always executable w.r.t. the whole Γ, thus fulfilling the test $?K\langle a \rangle \top$. Now we can have an analog of Corollary 4.2, and reduce the problem of checking $\langle \mathscr{N}, \mathscr{U} \rangle, s \models (\Sigma_{a \in B}(?K\langle a \rangle \top; a))^* K\phi$ to the reachability problem in $\mathscr{N}^\circ$: whether there is a path from $s_\mathscr{U}$ in $\mathscr{N}^\circ$ such that it can reach a state $t_{\mathscr{U}'}$ where $K\phi$ holds. Since ϕ is $[\pi]$-free, we can check it easily given $\mathscr{U}'$ using polynomial space, thus the main task is to find the reachable $t_{\mathscr{U}'}$. Note

---

[7]Cf. the definition of $\xrightarrow{\pi}$ in Def. 4.5.

that, in the size of $\mathcal{N}$, there are exponentially many such $t_{\mathcal{U}'}$ and the maximal length of the plan is also exponential. However, we do not need to build the whole $\mathcal{N}^\circ$ and the bisection-like algorithm behind the proof of Savitch's Theorem will do the job.[8] More precisely, we first pick up a $t_{\mathcal{U}'}$, and then run the recursive bisection method to see whether $t_{\mathcal{U}'}$ is reachable from $s_{\mathcal{U}}$ within $2^{|\mathcal{N}|}$ steps. The depth of the recursion is bounded by $log_2(2^{|\mathcal{N}|}) = |\mathcal{N}|$ and at each recursion we need to record the choice of the state which can be encoded by a $(0,1)$-vector using $log_2(2^{|\mathcal{N}|}) = |\mathcal{N}|$ space (plus one bit to record the result). Moreover, at the bottom of the recursion we only need to verify one step reachability, i.e., whether two states in $\mathcal{N}^\circ$ are linked by $\mathscr{R}_a^\circ$, without building the whole $\mathcal{N}^\circ$. Thus the whole procedure of model checking can be done using polynomial space.   ∎

As we mentioned in the introduction, the conformant planning problems in the AI literature are usually given by using state variables and actions with preconditions and (conditional) effects, rather than explicit transition systems. The corresponding explicit transition system can be generated by taking all the possible valuations of the state variables as the state space (an exponential blow up), and computing the transitions among the valuations according to the preconditions and the postconditions of the actions. In terms of the size of explicit transition systems, our above result is consistent with the EXPSPACE complexity result in the AI literature for conformant planning with Boolean and modal goals [21, 8]. Actually, the complexity result of Theorem 4.5 can be strengthened to PSPACE-complete based on the corresponding complexity result in the AI literature.

However, not all the transition systems can be generated in this way since the preconditions and postconditions are (usually) purely propositional and thus two states that share the same valuation must have the same executable actions. In an arbitrary transition system, multiple states with the same valuation may have different available actions due to some underlying protocol or other (external) factors not modelled by basic propositions.

5   Conclusions and future work

In this work we first introduce the logical language EAL over uncertainty maps and axiomatize it completely. EAL is then extended to EPDL with programs to specify conformant and conditional plans. We show that the conformant planning problems can be reduced to model checking problems of EPDL. Finally we showed that model checking star-free EPDL over uncertainty maps is PSPACE-complete and model checking the full fragment is in EXPTIME. On the other hand, model checking the conformant planning problem is in PSPACE.

Note that our EPDL is a powerful language which can already express conditional plans, for example, $(?p; a+?\neg p; b); c$. This suggests that we can use the very EPDL language (EPDL$^-$ is enough) to *verify* plans in contingent planning w.r.t. a variant of the semantics which can handle feedbacks during the execution. In fact, observational power about the availability of the actions has been already incorporated in [33], which can be extended to general feedbacks discussed in the literature of contingent planning (cf. e.g., [10]). On the other hand, to check the *existence* of a conditional plan, we are not sure whether EPDL is expressive enough, as subtleties may arise as in the case of conformant planning. We leave the contingent planning to future work.

Another natural extension is to go probabilistic, and reduce the probabilistic planning over MDP to some model checking problem of the probabilistic version of our EPDL. Our ultimate goal is to cast all the standard AI planning problems into one unified logical framework in order to facilitate careful

---

[8]A similar algorithm was used to pinpoint complexity of the conformant planning in AI, cf.[21].

comparison and categorization. We will then see clearly how the form of the goal formula, the constructor of the plan, and the observational ability matter in the theoretical and practical complexity of planning, in line with the research pioneered in [5].

References

[1] Mikkel Birkegaard Andersen, Thomas Bolander & Martin Holm Jensen: *Don't plan for the unexpected: Planning based on plausibility models. Logique et Analyse.* To appear.

[2] Mikkel Birkegaard Andersen, Thomas Bolander & Martin Holm Jensen (2012): *Conditional epistemic planning.* In: *Logics in Artificial Intelligence*, Springer, pp. 94–106, doi:10.1007/978-3-642-33353-8_8.

[3] Guillaume Aucher (2012): *DEL-sequents for regression and epistemic planning.* Journal of Applied Non-Classical Logics 22(4), pp. 337–367, doi:10.1080/11663081.2012.736703.

[4] Guillaume Aucher & Thomas Bolander (2013): *Undecidability in Epistemic Planning.* In: *IJCAI*, pp. 27–33. Available at `https://hal.inria.fr/hal-00824653`.

[5] Christer Bäckström & Peter Jonsson (2011): *All PSPACE-Complete Planning Problems Are Equal but Some Are More Equal than Others.* In: *SOCS 2011.* Available at `http://www.aaai.org/ocs/index.php/SOCS/SOCS11/paper/view/4009`.

[6] Alexandru. Baltag & Larry Moss (2004): *Logics for epistemic programs.* Synthese 139, pp. 165–224, doi:10.1023/B:SYNT.0000024912.56773.5e.

[7] Thomas Bolander & M. Birkegaard Andersen (2011): *Epistemic planning for single and multi-agent systems.* Journal of Applied Non-Classical Logics 21(1), pp. 9–34, doi:10.3166/jancl.21.9-34.

[8] Blai Bonet (2010): *Conformant plans and beyond: Principles and complexity.* Artificial Intelligence. 174(3-4), pp. 245–269, doi:10.1016/j.artint.2009.11.001.

[9] Blai Bonet & Hector Geffner (2000): *Planning with Incomplete Information as Heuristic Search in Belief Space.* In: *ICAPS 2000*, pp. 52–61, doi:10.1.1.38.8535. Available at `http://www.aaai.org/Library/AIPS/2000/aips00-006.php`.

[10] Blai Bonet & Hector Geffner (2012): *Width and Complexity of Belief Tracking in Non-Deterministic Conformant and Contingent Planning.* In: *AAAI 2012.* Available at `http://www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/view/5022`.

[11] Ronen I. Brafman & Jörg Hoffmann (2004): *Conformant Planning via Heuristic Forward Search: A New Approach.* In: *ICAPS 2004*, pp. 355–364, doi:10.1016/j.artint.2006.01.003. Available at `http://www.aaai.org/Library/ICAPS/2004/icaps04-042.php`.

[12] Daniel Bryce, Subbarao Kambhampati & David E. Smith (2006): *Planning Graph Heuristics for Belief Space Search.* Journal of Artificial Intelligence Research 26, pp. 35–99, doi:10.1613/jair.1869.

[13] Alessandro Cimatti & Marco Roveri (2000): *Conformant Planning via Symbolic Model Checking.* Journal of Artificial Intelligence Research 13, pp. 305–338, doi:10.1613/jair.774.

[14] Alessandro Cimatti & Marco Roveri (2011): *Conformant Planning via Symbolic Model Checking.* CoRR abs/1106.0252. Available at `http://arxiv.org/abs/1106.0252`.

[15] Alessandro Cimatti, Marco Roveri & Piergiorgio Bertoli (2004): *Conformant planning via symbolic model checking and heuristic search.* Artificial Intelligence 159(1-2), pp. 127–206, doi:10.1016/j.artint.2004.05.003.

[16] Hans van Ditmarsch, Wiebe van der Hoek & Barteld Kooi (2007): *Dynamic epistemic logic*. Springer, doi:10.1007/978-1-4020-5839-4.

[17] R. Fagin, J. Halpern, Y. Moses & M. Vardi (1995): *Reasoning about knowledge*. MIT Press, Cambridge, MA, USA.

[18] R. Fagin, J. Y. Halpern, Y. Moses & M. Y. Vardi (1997): *Knowledge-based programs. Distributed Computing* 10(4), pp. 199–225, doi:10.1007/s004460050038.

[19] Patrik Haslum & Peter Jonsson (1999): *Some Results on the Complexity of Planning with Incomplete Information*. In: *ECP 1999*, pp. 308–318, doi:10.1007/10720246_24.

[20] Tao Jiang & B. Ravikumar (1991): *A note on the space complexity of some decision problems for finite automata. Information Processing Letters* 40, pp. 25–31, doi:10.1016/S0020-0190(05)80006-7.

[21] Jon Kleinberg & Éva Tardos (2005): *Algorithm Design*. Addison-Wesley.

[22] Jérôme Lang & Bruno Zanuttini (2012): *Knowledge-Based Programs as Plans - The Complexity of Plan Verification*. In: *ECAI 2012*, pp. 504–509, doi:10.3233/978-1-61499-098-7-504.

[23] Martin Lange (2006): *Model checking propositional dynamic logic with all extras. Journal of Applied Logic* 4, pp. 39–49, doi:10.1016/j.jal.2005.08.002.

[24] Benedikt Löwe, Eric Pacuit & Andreas Witzel (2011): *DEL planning and some tractable cases*. In: *LORI 2011*, Springer, pp. 179–192, doi:10.1007/978-3-642-24130-7_13.

[25] Héctor Palacios & Hector Geffner (2006): *Compiling Uncertainty Away: Solving Conformant Planning Problems using a Classical Planner (Sometimes)*. In: *AAAI 2006*, pp. 900–905. Available at `http://www.aaai.org/Library/AAAI/2006/aaai06-142.php`.

[26] Pere Pardo & Mehrnoosh Sadrzadeh (2013): *Strong Planning in the Logics of Communication and Change*. In: *Declarative Agent Languages and Technologies X*, Springer, pp. 37–56, doi:10.1007/978-3-642-37890-4_3.

[27] Rohit. Parikh & R. Ramanujam (1985): *Distributed Processes and the Logic of Knowledge*. In: *Proceedings of Conference on Logic of Programs*, Springer-Verlag, London, UK, pp. 256–268, doi:10.1007/3-540-15648-8_21.

[28] Philippe Schnoebelen (2003): *The Complexity of Temporal Logic Model Checking*. In Philippe Balbiani, Nobu-Yuki Suzuki, Frank Wolter & Michael Zakharyaschev, editors: *AiML 2002*, King's College Publication, Toulouse, France, pp. 393–436. Available at `http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/Sch-aiml02.pdf`. Invited paper.

[29] David E. Smith & Daniel S. Weld (1998): *Conformant Graphplan*. In: *AAAI 1998*, pp. 889–896. Available at `http://www.aaai.org/Library/AAAI/1998/aaai98-126.php`.

[30] Larry J Stockmeyer & Albert R Meyer (1973): *Word problems requiring exponential time (Preliminary Report)*. In: *STOC 1973*, ACM, pp. 1–9, doi:10.1145/800125.804029.

[31] Son Thanh To, Tran Cao Son & Enrico Pontelli (2010): *A New Approach to Conformant Planning Using CNF**. In: *ICAPS 2010*, pp. 169–176. Available at `http://www.aaai.org/ocs/index.php/ICAPS/ICAPS10/paper/view/1461`.

[32] Yanjing Wang & Qinxiang Cao (2013): *On axiomatizations of public announcement logic. Synthese* 190, pp. 103–134, doi:10.1007/s11229-012-0233-5.

[33] Yanjing Wang & Yanjun Li (2012): *Not All Those Who Wander Are Lost: Dynamic Epistemic Reasoning in Navigation*. In: *AiML 2012*, pp. 559–580. Available at `http://www.aiml.net/volumes/volume9/Wang-Li.pdf`.

[34] Quan Yu, Ximing Wen & Yongmei Liu (2013): *Multi-Agent Epistemic Explanatory Diagnosis via Reasoning about Actions*. In: *IJCAI 2013*, pp. 1183–1190. Available at `http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6631`.

A Algorithms for EPDL$^-$

Definition A.1 (Matrix representation) *Let $B_{n \times m}$ denote a (0,1)-matrix of size $n \times m$. A matrix $B_{n \times 1}$, or B_n for short, is called a vector. Given finite uncertainty map $\mathcal{M}$, its domain $\mathcal{S}$ can be linearly ordered as $\{s_1, \cdots, s_n\}$. Thus $\mathcal{M}$ can be represented by a set $\{B_{n \times n}^a \mid a \in A\}$ of adjacency matrices for accessibility relation, a vector $B_n^{\mathcal{U}}$ for $\mathcal{U}$ and a set $\{B_n^p \mid p \in P\}$ of vectors for atomic propositions.*

Definition A.2 *Given (0,1)-matrices $B'_{n \times k}, B_{k \times m}$, their product $B''_{n \times m}$ is defined as: $B''_{n \times m}[i, j] = 1$ iff there exists $r \leq n$ such that $B'_{n \times k}[i, r] = B_{k \times m}[r, j] = 1$ for all $1 \leq i \leq n, 1 \leq j \leq m$.*

The following algorithms are to check whether ϕ holds on a pointed uncertainty map $\mathcal{M}, s$ by Definition 4.3. The main algorithm (Algorithm 3) recursively calls itself for each non-trivial subformula of ϕ. The complex cases are for the subformulas in the form of $\langle \pi \rangle \phi$ and $K\phi$. By Definition 4.3, to check $\mathcal{M}, s \Vdash_\sigma \langle \pi \rangle \phi$, we need to make sure that there exists a sequence $\omega \in \mathcal{L}(\pi)$ and a state $t \in \mathcal{S}$ such that $s \overset{\omega_\sigma}{\to} t$ and $\mathcal{M}, t \Vdash_{\sigma r(\omega)} \phi$. Since π is star-free, $|\omega| \leq |\pi|$ for each $\omega \in \mathcal{L}(\pi)$. It is clear that we cannot compute and store the whole set of $\mathcal{L}(\pi)$ within polynomial space. Instead, *one by one* we generate all the possible sequences that are shorter than $|\pi|$ and are formed from the alphabet of π (cf. line 14), and check whether they are in $\mathcal{L}(\pi)$. We can order the possible sequences lexicographically according to an ordering of the basic actions and tests in *Sig*, and compute the next sequence merely from the current one using function *next*. *memb_chec(ω, π)* checks whether it is the case $\omega \in \mathcal{L}(\pi)$. If $\omega \in \mathcal{L}(\pi)$, we need to check whether there exists $s_j \in \mathcal{S}_\mathcal{M}$ such that $s \overset{\omega_\sigma}{\to} s_j$ (Algorithm 2) and $\mathcal{M}, s_j \Vdash_{\sigma r(\omega)} \phi$, where $r(\omega)$ is the test-free subsequence of ω which is easy to compute. For the case of $K\phi$, we need to calculate the state set $\mathcal{U}|^\sigma$ (Algorithm 1).

B Complexity analysis

We suppose $|\mathcal{S}_\mathcal{M}| = n$ and $|\phi| = k$. Algorithm 1 uses one variable A to record the uncertainty set which requires $O(n)$ space. Note that there is a mutual recursion in Algorithm 2 and 3, but the depth of the overall recursion is bounded by k. In Algorithm 2, the variable consuming the most of the space is the matrix $B_{n \times n}$ recording the (intermediate) relation. Since σ and ω are also variables in the main algorithm and $|\omega| + |\sigma| \leq k$ due to the construction in Algorithm 3, the space usage of Algorithm 2 before the recursive calls of *PW* and *MC* is bounded by $O(k + n^2)$. For Algorithm 3, the most space-demanding part is the $\langle \pi \rangle \phi$ case, where we need to store π, *Sig*, and keep track one ω and one state s in the loop, which are bounded by either k or s. Moreover, according to [20], the complexity of *memb_chec* is NLOGSPACE-complete in the size of *Sig*, i.e., the alphabet of π which is bounded again by k. Thus before calling *MC* and *PW* again in the $\langle \pi \rangle \phi$ case, the space requirement is at most linear in both k and n, which is less demanding than *PW* for each recursion. Recall that the overall recursion depth of *MC* (and *PW*) is bounded by k thus the space usage of the whole algorithm is bounded by $O(k(k + n^2)) = O(k^2 + kn^2)$.

Algorithm 1: Function CNU($\mathscr{U}, \sigma$): Calculate the the new uncertainty set $\mathscr{U}|^{\sigma}$

---

    input : $\mathscr{U}, \sigma$
    output: $B_n^{\mathscr{U}|^{\sigma}}$

1   $A \leftarrow B_n^{\mathscr{U}}$;
2   **for** $i \leftarrow 1$ **to** $|\sigma|$ **do**
3     $\lfloor \quad A \leftarrow A \times B_{n \times n}^{\sigma[i]}$;
4   **return** A;

Algorithm 2: Function $PW(\omega, \sigma)$: Calculate the binary relation $\xrightarrow{\omega_\sigma}$

---

    input : computation sequence ω, action sequence σ
    output: $B_{n \times n}$

1   **switch** ω_σ **do**
2   **case** ε_σ **return** $Matrix(\{(s,s) \mid s \in \mathscr{S}\})$ /* $Matrix(R)$ is the $(0,1)$-matrix
            representation of the binary relation R */;
3   **case** $(?\phi\omega')_\sigma$ **return** $Matrix(\{(s,s) \mid MC(\mathscr{M}, s, \sigma, \phi) = true \}) \times PW(\omega', \sigma)$;
4   **case** $(a\omega')_\sigma$ **return** $B_{n \times n}^a \times PW(\omega', \sigma a)$;

Algorithm 3: Function $MC(\mathscr{M}, s, \sigma, \phi)$: Model checking algorithm for EPDL$^-$ (Boolean cases omitted)

---

    input : The pointed uncertainty map $(\mathscr{M}, s)$, sequence of actions σ, $\phi \in$ EPDL$^-$.
    output: true if $\mathscr{M}, s \Vdash_\sigma \phi$.

1   **switch** ϕ **do**
2   **case** $\langle \pi \rangle \varphi$
3   Let Sig be the array consisting of atomic programs and formulas in π ordered according to their first appearances;
4        $\omega \leftarrow Sig[1]$ /* ω is the candidate sequence we want to test */;
5   **while** $|\omega| \leq |\pi|$ **do**
6   **if** *memb_chec(ω, π)* **then**
7   **for** *i = 1* **to** $\mathscr{S}_{\mathscr{M}}$ **do**
8   **if** $(s, s_i) \in PW(\omega, \sigma)$ **then**
9   **if** $MC(\mathscr{M}, s_j, \sigma r(\omega), \varphi)$ **then return** true;
10          $\omega \leftarrow next(\omega, Sig)$ /* calculate the next sequence lexicographically according to the order Sig */;
11   **return** false;
12   **case** $K\varphi$
13        $B_n^{\mathscr{U}|^{\sigma}} = CNU(\mathscr{U}, \sigma)$ /* calculate the vector representation of $\mathscr{U}|^{\sigma}$ */
14   **for** *m = 1* **to** $|\mathscr{S}_{\mathscr{M}}|$ **do**
15   **if** $(B_n^{\mathscr{U}|^{\sigma}})_m = 1$ **and** $MC(\mathscr{M}, s_m, \sigma, \varphi) = false$ **then return** false;
16   **return** true;