DEL planning and some tractable cases

Benedikt Löwe[1,2], Eric Pacuit[3], Andreas Witzel[4]

[1] Institute for Logic, Language and Computation, Universiteit van Amsterdam, Postbus 94242, 1090 GE Amsterdam, The Netherlands; b.loewe@uva.nl
[2] Department Mathematik, Universität Hamburg, Bundesstrasse 55, 20146 Hamburg, Germany
[3] Tilburg Center for Logic and Philosophy of Science, Tilburg Universiteit van Tilburg, Postbus 90153, 5000 LE Tilburg, The Netherlands; E.J.Pacuit@uvt.nl
[4] Bioinformatics Group, Courant Institute of Mathematical Sciences, New York University, 715 Broadway, New York, NY 10003, USA; awitzel@nyu.edu

Abstract. We describe the planning problem within the framework of dynamic epistemic logic (DEL), considering the tree of sequences of events as the underlying structure. In general, the DEL planning problem is computationally difficult to solve. On the other hand, a great deal of fruitful technical advances have led to deep insights into the way DEL works, and these can be exploited in special cases. We present a few properties that will lead to considerable simplifications of the DEL planning problem and apply them in a toy example.

1 Introduction

Dynamic epistemic logic (DEL) is one of the standard conceptual models for epistemic situations and change. Semantically, it is based on an operation called *product update* that allows to apply an *event model* to an *epistemic model* in order to describe epistemic change. Possible future epistemic states correspond to sequences of event models applied successively to an initial state. DEL is conceptually very clear; this makes it a promising framework for epistemic planning.

In this paper, we shall propose a general DEL planning framework. In general, the DEL planning problem is intractable; however, we argue that in some cases, our understanding of the DEL semantics will allow us to restrict the attention to a tractable situation. We shall give a number of conditions that will allow us to reduce the complexity of the DEL planning problem drastically, and apply these insights in a toy example.

We do not consider DEL planning as a merely theoretical framework, but aim at real applications in the setting of computer games (and fittingly, our toy example is also about a concrete computer game). Making planning into a crucial element of narrative design for computer games has been proposed by Riedl and Young [20, 21]. The first and second author have proposed to use formalisms based on epistemic logic for the formalization of narratives [15, 16], and the third author (in collaboration with Kennerly and Zvesper; [29, 13]) has formulated a simple knowledge-based action situation from a computer game.

We consider this paper as a first step towards combining these logical aspects of computer gaming.

Related Work.

In the logic community, the potential to use dynamic epistemic logic for concrete implementations of reasoning processes permeates the literature; many papers mention concrete applications as motivation for studying dynamic epistemic logic.

Renardel de Lavalette and van Ditmarsch [19] discuss updating and maintaining a minimal epistemic model and identifying subclasses of DEL for which that is possible. They provide model minimization for so-called simple actions in order to allow efficient model checking. Our toy application in § 5 has non-simple (though propositional) actions, and uses non-S5 models.

Related to our proposal is the work by van der Hoek and Wooldridge [25] on planning with epistemic goals, based on the idea of Giunchilia and Traverso to use model checking as a planning algorithm [11]. Their planning algorithm is based on S5 ATEL using a similar model to the one used in [13]; we argue here that DEL provides a more flexible framework. Also related is the work by Ågotnes and van Ditmarsch on public announcement logic, examining which public announcement to make in a strategic setting with goals (assuming truthfulness) [1].

The aim of [26] is very similar to ours: van Ditmarsch, Herzig and Lima combine DEL and planning; their work is complementary to ours in that the authors start from the situation calculus fragment from [14] and offer a subset of DEL (public announcement and public "questions") as equivalent to it, while we start from the rich framework of DEL and try to show how planning may be incorporated.

The closest cognate of our paper is [5]: Bolander and Andersen consider a setting very similar to the one described in §§ 2 and 3 and prove that the planning problem for the single agent case is decidable [5, Theorem 17] and that the multi-agent case is undecidable [5, Theorem 20]. The results of Bolander and Andersen were independently discovered (however, their paper cites a preprint version of this paper in their section on related work). Despite the conceptual agreement with Bolander and Andersen, our aim in this paper is quite different: while they are interested in decidability and undecidability results, we aim at finding subclasses for which planning becomes tractable.

Outside of the logic literature, there are many approaches to planning under uncertainty and knowledge representation in the large and mature field of *AI planning*. The authors of [18] point out that a syntactic (what they call "knowledge-based" or "knowledge level") representation can represent uncertainty more succinctly. While we focus on DEL model theory, we do not want to dictate that an implementation has to be model-based and preclude a more proof-theoretic realization. There is a trade-off of representational complexity versus computational complexity, and these are just two sides of the same coin. Petrick and Bacchus consider a restricted language to stay computationally tractable, while we look at restricted classes of models to stay spatially tractable. In certain cases, one viewpoint may allow more natural or more effective restrictions, and we think that ultimately joining both kinds of restrictions will be necessary.

The fact that a semantic approach has its advantages was pointed out by Lakemeyer and Levesque who provide an elegant semantical view on the situation calculus [14].

While these approaches are more expressive than ours in some sense by using (fragments of) first- or second-order logic, they deal with only a single agent, and only with S5 knowledge. This simplifies the semantics to just an information set for that player (although Lakemeyer and Levesque in [14] do add some structure in terms of their sensing agreement relation $\simeq$ among possible worlds).

On the semantical level, the conceptual extension to multiple agents and their (possibly sub-S5) beliefs about each other corresponds to adding more structure in the form of agents' accessibility relations. This makes the semantic objects and algorithms more complex, so for driving this extension we simplify matters by focusing on propositional modal logic.[5]

Outline of this paper

In § 2 we shall give a standard introduction to dynamic epistemic logic, following roughly the textbook [27]. In § 3, we define several versions of *DEL planning problems*, and argue that we need to find efficient techniques for solving it. In § 4, we discuss some examples of techniques to that end, and we apply them in § 5 to the setting of the doxastically enriched computer game *Thief* from [29, 13]. We close the paper with a few pointers to future work in § 6.

2 Dynamic Epistemic Logic

In this section, we give an overview of product updates due to Baltag, Moss and Solecki [3] for the non-expert reader, following closely the textbook [27] (where the reader can find more details).

Let $\mathcal{A}$ be a finite set of agents and At a set of atomic propositions. An **epistemic model** $\mathcal{M}$ is a tuple $\langle W, \{R_i\}_{i\in\mathcal{A}}, V\rangle$ where $W =: \mathrm{D}(\mathcal{M})$ is a non-empty set called the **domain of** $\mathcal{M}$, for each $i \in \mathcal{A}$, $R_i \subseteq W \times W$ is a binary relation on W (typically an equivalence relation) and $V : \mathsf{At} \to 2^W$ is a valuation function. If $w \in \mathrm{D}(\mathcal{M})$, we call $(\mathcal{M}, w)$ a **pointed epistemic model**. If w is

[5] Note that we are here using dynamic *epistemic* logic for a situation in which we describe *beliefs*. The main difference between knowledge and belief is that beliefs can be false whereas knowledge—in standard formalizations—cannot. If i discovers that j *does* know, the product update will produce a model in which i considers *no* state of the world possible. For a more graceful handling and revising of inconsistent beliefs, we could use a doxastic version of DEL [4]. Since DEL works for the simple examples in this paper, we ignore these issues for the present discussion.

clear from the context, we may omit it from the notation. The elements of W constitute "states of the world" and the relations R_i are **accessibility relations**, i.e., for states $w, v \in W$, wR_iv means "in state w, agent i would consider state v possible."

The set of multiagent epistemic formulas, denoted $\mathcal{L}_\mathcal{A}$, is the smallest set of formulas generated by the following grammar:

$$p \mid \neg\varphi \mid \varphi \wedge \psi \mid \Box_i\varphi$$

where $p \in \mathsf{At}$ and $i \in \mathcal{A}$. We use the usual abbreviations for the other propositional connectives ($\vee, \rightarrow, \leftrightarrow$), and we use $\mathcal{L}_{\mathrm{Prop}}$ to denote the propositional sub-language (i.e., not containing $\Box_i$). Truth of formulas $\varphi \in \mathcal{L}_\mathcal{A}$ is defined as usual in Kripke models.

An **event (or action) model** $\mathcal{E}$ is a tuple $\langle S, \{\rightarrow_i\}_{i\in\mathcal{A}}, \mathsf{pre}\rangle$, where S is a nonempty set, for each $i \in \mathcal{A}$, $\rightarrow_i \subseteq S \times S$ is i's **accessibility relation**, and $\mathsf{pre} : S \rightarrow \mathcal{L}_\mathcal{A}$ is the **pre-condition function**. The set S is called the domain of $\mathcal{E}$, denoted $\mathrm{D}(\mathcal{E})$. We call $\mathcal{E}$ **propositional** if pre goes into $\mathcal{L}_{\mathrm{Prop}}$, i.e., all preconditions are propositional. The **product update** operation updates an epistemic model $\mathcal{M} = \langle W, \{R_i\}_{i\in\mathcal{A}}, V\rangle$ with an event model

$$\mathcal{E} = \langle S, \{\rightarrow_i\}_{i\in\mathcal{A}}, \mathsf{pre}\rangle$$

and is defined as $\mathcal{M} \otimes \mathcal{E} = \langle W', \{R'_i\}_{i\in\mathcal{A}}, V'\rangle$ with

(i) $W' = \{(w, e) \mid w \in W, e \in S \text{ and } \mathcal{M}, w \models \mathsf{pre}(e)\}$,
(ii) $(w_1, e_1)R'_i(w_2, e_2)$ iff $w_1R_iw_2$ in $\mathcal{M}$ and $e_1 \rightarrow_i e_2$ in $\mathcal{E}$, and
(iii) $V'(p) = \{(w, e) \in W' \mid w \in V(p)\}$.

For pointed models, the point of the product is the pair of the factors' points.

If $\mathcal{E}$ and $\mathcal{E}'$ are two propositional event models, we can analogously define the **product event model** $\mathcal{E} \otimes \mathcal{E}'$, simply using the conjunction $\mathsf{pre}(e_1) \wedge \mathsf{pre}(e_2)$ for $\mathsf{pre}((e_1, e_2))$. It satisfies a type of associative law $(\mathcal{M} \otimes \mathcal{E}) \otimes \mathcal{E}' \simeq \mathcal{M} \otimes (\mathcal{E} \otimes \mathcal{E}')$, where $\simeq$ denotes isomorphism of epistemic models.

The usual notion of equivalence used in modal logic is the weaker notion of **bisimulation** [27, Definition 2.14], denoted $\leftrightarrows$. Bisimilar pointed models are equivalent in the sense that they satisfy exactly the same formulas. For any model $\mathcal{M}$ and events $\mathcal{E}$, $\mathcal{E}'$, if $\mathcal{E} \leftrightarrows \mathcal{E}'$ then $\mathcal{M} \otimes \mathcal{E} \leftrightarrows \mathcal{M} \otimes \mathcal{E}'$ (note that the converse does not hold, see [28, Observation 13]). For each model $\mathcal{M}$, the union of all bisimulations of $\mathcal{M}$ with itself is again a bisimulation. The quotient structure of this bisimulation gives us the most compact model satisfying the same formulas, called the **bisimulation contraction**. By $\lfloor\mathcal{M}\rfloor$, we denote the cardinality of the domain of the bisimulation contraction of $\mathcal{M}$.

Now, adding to $\mathcal{L}_\mathcal{A}$ a modal operator $\langle\mathcal{E}, e\rangle$ for each pointed event model $(\mathcal{E}, e)$, we obtain the language $\mathcal{L}_{\mathrm{DEL}}$. Truth for these modalities is defined as

$$\mathcal{M}, w \models \langle\mathcal{E}, e\rangle\varphi \text{ iff } \mathcal{M}, w \models \mathsf{pre}(e) \text{ and } \mathcal{M} \otimes \mathcal{E}, (w, e) \models \varphi.$$

Given a pointed model $(\mathcal{M}, w)$ and a formula $\varphi \in \mathcal{L}_{\mathcal{A}}$, checking $\mathcal{M}, w \models \varphi$ can be done in time polynomial in the size of $\mathcal{M}$ and φ [24, § 3.8]. Some care must be taken with respect to the product update, as at first glance it can potentially lead to exponentially growing models (cf. § 4).

Planning will happen in the tree of temporal sequences of events: fix a pointed epistemic model $(\mathcal{M}, w)$ and a finite set of (pointed) event models $\mathfrak{E}$. For prefixes of finite sequences $\sigma := (\mathcal{E}_0, \ldots, \mathcal{E}_n)$ of models in $\mathfrak{E}$, we have a natural notion of **immediate successor**, viz. extension by one additional model. Let

$$\mathcal{M} \otimes \sigma := \mathcal{M} \otimes \mathcal{E}_0 \otimes \ldots \otimes \mathcal{E}_n;$$

then the collection of these epistemic models forms a tree structure with successor structure derived from the finite sequences and $\mathcal{M} = \mathcal{M}_{\varnothing}$ at the root. It is this tree structure that we consider to be the natural temporal setting for DEL planning (cf. § 3).

Slightly more precisely, for a pointed epistemic model $(\mathcal{M}, w)$ and $(\mathcal{E}, e) \in \mathfrak{E}$, we say that $(\mathcal{E}, e)$ is **possible at** $(\mathcal{M}, w)$ if $\mathcal{M}, w \models \langle \mathcal{E}, e \rangle \top$. We say that a sequence σ is **legal** if it is empty or its (uniquely determined) immediate predecessor σ^* is legal and $(\mathcal{E}^*, e^*)$ is possible at $\mathcal{M} \otimes \sigma^*$, with $(\mathcal{E}^*, e^*)$ being the last element of σ. The set of legal sequences, denoted LS, contains exactly those sequences that can be performed in the given order, since the preconditions of each event are met at the appropriate time. They form a subtree of our tree structure. If we want to impose further external restrictions on the possible courses of action, we can consider a subtree $T \subseteq \mathrm{LS}$.

To conclude our description of DEL, we should note that we do not consider events that change actual facts (i.e., the valuation function). This is a serious restriction but doesn't affect the example in § 5. The definition of the product update can be extended to deal with factual change [24], but for the sake of simplicity, we restricted ourselves to purely epistemic events (cf. § 6).

3 DEL planning

The classical *planning problem* consists of a description of the world, the agent's goal and a description of the possible actions in some appropriate formal language. A planning algorithm consists of sequences of possible actions which when executed will achieve the goal.

To attempt an initial definition of a DEL planning problem, we fix a pointed epistemic model $(\mathcal{M}, w)$ and a finite set $\mathfrak{E}$ of pointed event models. As mentioned at the end of § 2, we consider the tree LS of legal sequences σ from $\mathfrak{E}$ as our space of possible plans, and we allow to impose additional rules on when events can occur in the form of specifying a subtree $T \subseteq \mathrm{LS}$.

Definition 1 (Absolute DEL planning problem) *Given $(\mathcal{M}, w)$, $\mathfrak{E}$, a subtree $T \subseteq \mathrm{LS}$, and a formula $\varphi \in \mathcal{L}_{\mathrm{DEL}}$, produce a sequence $\sigma \in T$ such that $\mathcal{M} \otimes \sigma \models \varphi$.*

The absolute DEL planning problem does not talk about agents and whether the plan is realizable by a single agent. In its standard formulation, the DEL formalism does not assign events to particular agents, so we need to supplement the formalism with such an assignment: consider a function power : $\mathcal{A} \to \wp(\mathfrak{E})$ that tells us which events an agent can bring about. Here, if $i \in \mathcal{A}$, we interpret $\mathcal{E} \in \mathrm{power}(i)$ as "agent i can perform action $\mathcal{E}$". If σ is a sequence of actions, we write $\sigma \in \mathrm{power}(i)$ if for each $\mathcal{E} \in \mathrm{Set}(\sigma)$, we have $\mathcal{E} \in \mathrm{power}(i)$. Similarly, for a partial sequence, i.e., a partial function $\widehat{\sigma} : \{0, \ldots, N\} \to \mathfrak{E}$ (for some N), we write $\widehat{\sigma} \in \mathrm{power}(i)$, if for all $n \in \mathrm{dom}(\widehat{\sigma})$, we have $\widehat{\sigma}(n) \in \mathrm{power}(i)$.[6]

Definition 2 (Single-agent DEL planning problem) *Given* $(\mathcal{M}, w)$, $\mathfrak{E}$, *a subtree* $T \subseteq \mathrm{LS}$, *a function* power, *an agent* $i \in \mathcal{A}$, *and a formula* $\varphi \in \mathcal{L}_{\mathrm{DEL}}$, *produce a sequence* $\sigma \in \mathrm{LS}$ *such that* $\mathcal{M} \otimes \sigma \models \varphi$ *and* $\sigma \in \mathrm{power}(i)$.

A special case is the situation where $\mathcal{A} = \{i\}$. This case is called SINGLE-AGENT EPISTEMIC PLANNING in [5]. More interesting and much harder is the multi-agent version that essentially asks for the existence of a *winning strategy*. Bolander and Andersen called it MULTI-AGENT EPISTEMIC PLANNING in [5]:

Definition 3 (Multi-agent DEL planning problem) *Given* $(\mathcal{M}, w)$, $\mathfrak{E}$, *a subtree* $T \subseteq \mathrm{LS}$, *a function* power, *an agent* $i \in \mathcal{A}$, *and a formula* $\varphi \in \mathcal{L}_{\mathrm{DEL}}$, *produce a partial function* $\widehat{\sigma} : \{0, \ldots, N\} \to \mathfrak{E}$ *such that*

1. $\widehat{\sigma} \in \mathrm{power}(i)$,
2. *there is a sequence* $\sigma \in T$ *such that* $\widehat{\sigma} \subseteq \sigma$, *and*
3. *for all sequences* $\sigma \in T$ *with* $\widehat{\sigma} \subseteq \sigma$, *we have* $\mathcal{M} \otimes \sigma \models \varphi$.

Bolander and Andersen prove that MULTI-AGENT EPISTEMIC PLANNING in this sense is undecidable [5, Theorem 20].[7]

Defining the planning problem is only the very first step towards solving it. In general, planning problems are intractable (even if they are decidable), and a large part of the work in AI planning is spent on finding compact representations of large spaces of situations (or plans), and time-efficient ways of traversing them in the search of given goals (cf. [10]). Our DEL planning problems tend to be more complex than classical planning problems: not only can the space of reachable situations explode with the number of events at the planner's disposal, but even the situations themselves, being product models of previous situations and events, may prima facie grow combinatorially.

[6] These definitions are easily extended to sets of agents if we want to include actions that can only jointly be performed by a group of agents.

[7] Note that this formulation assumes a worst-case scenario. Ultimately, we should like to study a general **strategic DEL planning problem** where agents take (interfering or cooperative) actions of third agents into account, and may to some extent try to anticipate them. However, while there exists work in the context of DEL on defining the information content of events in strategic settings [9] and modeling goals and preferences [23], things are far less clear in such a case, and we therefore leave it to future work.

Consider, for example, the classical STRIPS planning formalism [8], which is by now somewhat dated but still widely used as point of reference. In the most commonly used variant of that formalism, situations are valuations over a certain universe of atomic facts, and the effects of an event (or action) e are specified by a set e^- of facts that it removes (makes false) and a set e^+ of facts that it adds (makes true) in any situation that it is applied to.[8] While the number of possible situations does grow exponentially with the size of the universe of facts, once that universe is fixed the size of situations is determined and cannot grow out of control. In contrast, which formulas hold after a DEL update depends very much on the situation before the update, so it is not straightforwardly possible to give the effects of an event just by sets e^+ and e^- (cf. Footnote 8). An embedding of DEL planning into STRIPS thus would necessarily have to encode this information in some complicated way. In our setting, we decided to represent DEL models directly in order to be able to exploit certain structural properties that would not be perspicuous in an embedding into a classical formalism such as STRIPS.

Considering planning as a search task in the tree of sequences of event, we should like to highlight and discuss an interesting phenomenon with respect to two standard approaches from the classical planning literature: *progression* and *regression*.[9] In *progression*, plan search starts from an initial state and successively expands applicable events until the goal state is reached. In *regression*, the search starts from the goal state and makes backward steps to find situations from which the goal state can be reached, until the initial state is hit.

Typically, progression is typically more straightforward to formalize and implement. This also holds for our model-theoretic view of DEL planning: Given a situation and an event model, it is straightforward to compute the successor situation resulting from application of the event, but it is an open problem how possible predecessor situations can be found, to which the event can be applied yielding the situation in question.

However, in a syntactic view of DEL planning, the situation reverses: the central piece of the completeness proof for axiomatizations of DEL is a term rewriting system consisting of axioms that are sometimes called *reduction axioms* [3, § 4]. These reduction axioms directly yield the required preconditions of an event, given its postconditions, making regression very simple. On the other hand, it is not known whether there is a simply definable way of associating to each event model $\mathcal{E}$ a map $F_\mathcal{E}$ mapping sets of formulas to sets of formulas such that for all $\mathcal{M}$, if $\mathcal{M} \models \Phi$, then $\mathcal{M} \otimes \mathcal{E} \models F_\mathcal{E}(\Phi)$. Such a function would correspond to progression in the syntactic view.

---

[8] So, if S is the set of true facts before the event e takes place, then $(S \backslash e^-) \cup e^+$ is the set of facts that is true after e takes place. Note that e^+ and e^- are independent of S.

[9] We should like to thank Hans van Ditmarsch for comments on the presentation of an earlier version of this paper in Delhi (January 2011) and the e-mail discussion developing from these comments; these formed the basis of the discussion of progression and regression.

4 Some properties to ensure tractability

We first introduce some notions. Let $\mathcal{E}$ and $\mathcal{E}'$ be event models. We say that

- $\mathcal{E}$ is **self-absorbing** if for all models $\mathcal{M}$ we have $\mathcal{M} \otimes \mathcal{E} \otimes \mathcal{E} \leftrightarrows \mathcal{M} \otimes \mathcal{E}$;
- $\mathcal{E}$ and $\mathcal{E}'$ **commute** if for all models $\mathcal{M}$ we have $\mathcal{M} \otimes \mathcal{E} \otimes \mathcal{E}' \leftrightarrows \mathcal{M} \otimes \mathcal{E}' \otimes \mathcal{E}$; and that
- $\mathcal{E}$ is **almost-mutex** ("almost-mutually-exclusive") if there is at most one atomic event $e_\top \in \mathrm{D}(\mathcal{E})$ with $\mathsf{pre}(e_\top) = \top$ and $e_\top \to_i e_\top$ for all $i \in \mathcal{A}$, and the formulas $\mathsf{pre}(e)$ with $e \neq e_\top$ are pairwise inconsistent.

Lemma 4 *Propositional event models commute.*

Proof. Let $\mathcal{E}, \mathcal{E}'$ be propositional event models. We prove that $\mathcal{E} \otimes \mathcal{E}' \leftrightarrows \mathcal{E}' \otimes \mathcal{E}$. To see that this holds, consider the smallest relation ρ with $(s, s')\rho(s', s)$ for all $(s, s') \in \mathrm{D}(\mathcal{E}) \times \mathrm{D}(\mathcal{E}')$. This relation is a bisimulation due to commutativity of logical conjunction. □

Lemma 5 *Almost-mutex event models with transitive accessibility relations are self-absorbing.*

Proof. Let $\mathcal{E} = \langle S, \{\to_i\}_{i \in \mathcal{A}}, \mathsf{pre} \rangle$ be an almost-mutex event model with point $s \in S$. We again prove that $\mathcal{E} \otimes \mathcal{E} \leftrightarrows \mathcal{E}$. Consider the smallest relation $\rho \subseteq (\mathrm{D}(\mathcal{E}) \times \mathrm{D}(\mathcal{E})) \times \mathrm{D}(\mathcal{E})$ such that $(e, e)\rho e$, $(e, e_\top)\rho e$, and $(e_\top, e)\rho e$ hold for all $e \in \mathrm{D}(\mathcal{E})$. We show that this is a bisimulation on the submodels of $\mathcal{E} \otimes \mathcal{E}$ and $\mathcal{E}$ generated by (s, s) and s. To see this, note first that $(s, s)\rho s$.

Next, assume that $(e_1, e_2)\rho e$ and $(e_1, e_2) \to_i (e'_1, e'_2)$. We have to show that there is e' with $e \to_i e'$ and $(e'_1, e'_2)\rho e'$. By definition of ρ, we are in one of three cases:

- $e_1 = e_2 = e$. From $(e, e) \to_i (e'_1, e'_2)$ it follows that $e \to_i e'_1$ and $e \to_i e'_2$. If $e'_1 = e'_2$ then $(e'_1, e'_2)\rho e'_1$ by definition of ρ and we are done. Otherwise $e'_1 \neq e'_2$. Since $(e'_1, e'_2) \in \mathrm{D}(\mathcal{E} \otimes \mathcal{E})$, $\mathsf{pre}(e'_1)$ and $\mathsf{pre}(e'_2)$ cannot be inconsistent, and with $\mathcal{E}$ being almost-mutex it follows that one of the two events is $e_\top$. If $e'_1 = e_\top$ then $(e'_1, e'_2)\rho e'_2$ by definition of ρ, and analogously if $e'_2 = e_\top$.
- $e_1 = e$ and $e_2 = e_\top$. From $(e, e_\top) \to_i (e'_1, e'_2)$ it follows that $e \to_i e'_1$ and $e_\top \to_i e'_2$. If $e'_1 = e'_2$ then $(e'_1, e'_2)\rho e'_1$ by definition of ρ. Otherwise $e'_1 \neq e'_2$. Since $(e'_1, e'_2) \in \mathrm{D}(\mathcal{E} \otimes \mathcal{E})$, $\mathsf{pre}(e'_1)$ and $\mathsf{pre}(e'_2)$ cannot be inconsistent, and with $\mathcal{E}$ being almost-mutex it follows that one of the two events is $e_\top$. If $e'_2 = e_\top$ then $(e'_1, e'_2)\rho e'_1$ by definition of ρ and we are done since $e \to_i e'_1$. Otherwise $e'_1 = e_\top$, and from $e \to_i e'_1 = e_\top \to_i e'_2$, by transitivity we get $e \to_i e'_2$. By definition of ρ, $(e_\top, e'_2)\rho e'_2$.
- $e_1 = e_\top$ and $e_2 = e$. Analogous to the previous case.

Finally, assume that $(e_1, e_2)\rho e$ and $e \to_i e'$. We have to show that there is (e'_1, e'_2) with $(e_1, e_2) \to_i (e'_1, e'_2)$ and $(e'_1, e'_2)\rho e'$. This is easy to see with a similar case distinction as above, noting that $e_\top \to_i e_\top$ by assumption. □

For a sequence $\sigma = \mathcal{E}_1 \ldots \mathcal{E}_k$ of event models, let $\mathrm{Set}(\sigma) = \{\mathcal{E}_1, \ldots, \mathcal{E}_k\}$, and let $\mathcal{M} \otimes \sigma = \mathcal{M} \otimes \mathcal{E}_1 \otimes \cdots \otimes \mathcal{E}_k$ for a model $\mathcal{M}$.

Proposition 6 *For any model $\mathcal{M}$ and any sequences σ_1, σ_2 of propositional, almost-mutex events with transitive accessibility relations, if* $\mathrm{Set}(\sigma_1) = \mathrm{Set}(\sigma_2)$ *then $\mathcal{M} \otimes \sigma_1 \Leftrightarrow \mathcal{M} \otimes \sigma_2$.*

Proof. Follows immediately from Lemma 4 and 5. □

5 A toy application

The video game *Thief: The Dark Project*™ by Eidos Interactive (1998) is themed as a game of stealth, in which the player (the thief) avoids being detected by computer-simulated guards. The player exploits the guard's—possibly mistaken—beliefs about the thief's presence. The following is an epistemically enhanced version of a scene from *Thief* as presented in [13].

We assume that the scene starts with thief and guard present, each uncertain of the other's presence, and that agents cannot enter or leave. In our formalization, we consider the following kinds of events:

- $n_{\mathrm{t}}, n_{\mathrm{g}}$: The thief (the guard) makes some noise.
- $b_{\mathrm{t}}, b_{\mathrm{g}}$: The thief (the guard) sees the other one from behind.
- f: Thief and guard see each other face to face.

The intuitive epistemic effects of these events are as follows:

- n_{t}: The guard learns that a thief is present; the thief learns that, if a guard is present, the guard learns that the thief is present.
- b_{t}: The thief learns that a guard is present; the guard believes nothing has happened (he is not paranoid enough to constantly suspect being seen from behind).
- f: Thief and guard commonly learn that both are present.

The effects of n_{g} and b_{g} are analogous. To model this situation in DEL, we use the set of atomic propositions $\mathsf{At} = \{p_{\mathrm{t}}, p_{\mathrm{g}}\}$, with the reading that the thief, respectively the guard, is present. We formalize the initial situation by the pointed model $\mathcal{I}$ and the events described above by the set of pointed event models $E = \{\mathcal{N}_{\mathrm{t}}, \mathcal{N}_{\mathrm{g}}, \mathcal{B}_{\mathrm{t}}, \mathcal{B}_{\mathrm{g}}, \mathcal{F}\}$, as depicted in Figure 1. From now on we omit the qualifier "pointed".

Proposition 7 *Let σ by any sequence of events from $E = \{\mathcal{N}_{\mathrm{t}}, \mathcal{N}_{\mathrm{g}}, \mathcal{B}_{\mathrm{t}}, \mathcal{B}_{\mathrm{g}}, \mathcal{F}\}$ Then $\lfloor \mathcal{I} \otimes \sigma \rfloor \leq 6$.*

Proof (sketch). Using Proposition 6. First note that $\mathcal{F} \otimes \mathcal{E} \Leftrightarrow \mathcal{F}$ for any $\mathcal{E} \in E$, so $\lfloor \mathcal{I} \otimes \sigma \rfloor = 1$ for any σ containing $\mathcal{F}$. Also, $\mathcal{N}_{\mathrm{t}} \otimes \mathcal{N}_{\mathrm{g}} \Leftrightarrow \mathcal{F}$, so the same holds for any sequence containing these two events. Due to symmetry we are left with 6 cases to check: $\sigma \in \{\mathcal{N}_{\mathrm{t}}, \mathcal{B}_{\mathrm{t}}, \mathcal{N}_{\mathrm{t}}\mathcal{B}_{\mathrm{t}}, \mathcal{N}_{\mathrm{t}}\mathcal{B}_{\mathrm{g}}, \mathcal{B}_{\mathrm{t}}\mathcal{B}_{\mathrm{g}}, \mathcal{N}_{\mathrm{t}}\mathcal{B}_{\mathrm{t}}\mathcal{B}_{\mathrm{g}}\}$. □

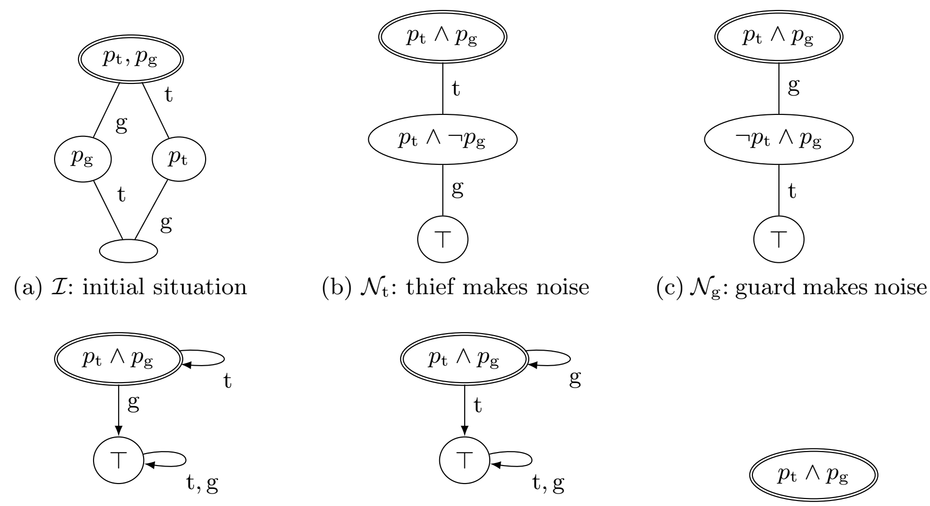

(a) $\mathcal{I}$: initial situation (b) $\mathcal{N}_t$: thief makes noise (c) $\mathcal{N}_g$: guard makes noise

(d) $\mathcal{B}_t$: thief sees guard from behind (e) $\mathcal{B}_g$: guard sees thief from behind (f) $\mathcal{F}$: face to face

Fig. 1: Models for initial situation and events. Undirected edges represent bidirectional accessibilities. In models without directed edges, reflexive accessibilities are omitted.

Together with the fact that the bisimulation contraction can be computed in linear time [7], this shows that our toy model indeed stays a toy model. As mentioned above, this may not say much about more realistic models, but guarantees may be found there with similar techniques.

We can also show the fact stated above, saying that we do not need to consider belief-revision mechanisms in our simple scenario, since the agents never reach inconsistent belief states (although their beliefs may be mistaken).

Proposition 8 *For any sequence σ of events from E and any agent $i \in \mathcal{A}$, $\mathcal{I} \otimes \sigma \not\models \Box_i \bot$.*

Proof (sketch). Since $\mathcal{F} \otimes \mathcal{E} \leftrightarrows \mathcal{F}$ for any $\mathcal{E} \in E$ and $\mathcal{I} \otimes \mathcal{F} \not\models \Box_i \bot$, with Proposition 6 we get that $\mathcal{I} \otimes \sigma \not\models \Box_i \bot$ for any σ containing $\mathcal{F}$. Assume there is some σ with $\mathcal{I} \otimes \sigma \models \Box_i \bot$, then there must be no state that i considers possible at the point of $\mathcal{I} \otimes \sigma$. By definition of $\otimes$, the same would then hold for the point of $\mathcal{I} \otimes \sigma \otimes \mathcal{F}$, which is a contradiction. □

6 Conclusions

DEL is about knowledge, but as our toy example showed, the applications we have in mind, are mostly about belief rather than knowledge. So a next step would be to phrase the planning problems of this paper in a *doxastic* version of DEL [4] instead. There are also a number of simplifications we made in our set-up that could be removed: We did not consider the possibility of events that change the valuation function (cf. [24]) nor the possibility that an agent's ability to perform an action may change over time. Generalizing our framework to include these possibilities would be a natural next step.

While DEL has a well-developed proof theory, its model theory is particularly appealing due to its clarity and intuitiveness. It will be a topic of future research to determine in what cases it can be directly implemented and in what cases a syntactic representation is better suited (possibly to be used with a generic theorem prover), or a "hybrid" approach such as binary decision diagrams (BDDs). For the model-based approach, two crucial topics for further research are the following:

– Examine the long-term expansion of models under iterated updates (related to [22], but we are foremost interested in finite models) and identify natural and general classes of actions that allow arguments such as our crucial Proposition 7 in § 5.
– Find compact representations of models (cf. techniques from model checking [6]) on which the product operation can operate directly.

Another interesting question to investigate is whether DEL planning problems can be compiled down to a simpler and more established formalism, as in the work by [17] where a certain class of planning problems with uncertainty is translated to problems that can be solved using the classical FF planner ("Fast Forward") by [12]. Classical planning has developed many effective and proven optimization strategies and other techniques, which may be exploited in this way.

We envision a use of DEL as a general engine (called "knowledge module" by Kennerly, Witzel, and Zvesper in [29, 13]) which allows for flexible specification of situations and events and then maintains the agents' mental models throughout the progress of the game or scene, much like a physics engine maintains a model of the physical state of the world. We shall illustrate this in a slightly richer scenario taken from an actual computer game. Note that, while we focus on maintaining one central epistemic model as part of the simulation engine, such a model can also be distributed and maintained by the individual agents [2].

References

1. T. Ågotnes and H. van Ditmarsch. What will they say?—Public announcement games, 2009. Paper presented at *Logic, Game Theory and Social Choice 6* in Tsukuba, Japan.
2. G. Aucher. *Perspectives on Belief and Change*. PhD thesis, Université de Toulouse and University of Otago, 2008.
3. A. Baltag, L. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicions. In I. Gilboa, editor, *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98)*, pages 43–56. Morgan Kaufmann, 1998.

4. A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, volume 3 of *Texts in Logic and Games*, pages 9–58. Amsterdam University Press, 2008.
5. T. Bolander and M. B. Andersen. Epistemic planning for single- and multi-agent systems. *Journal of Applied Non-Classical Logics*, to appear.
6. E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. MIT Press, 1999.
7. A. Dovier, C. Piazza, and A. Policriti. A fast bisimulation algorithm. In G. Berry, H. Comon, and A. Finkel, editors, *Computer Aided Verification, 13th International Conference, CAV 2001, Paris, France, July 18-22, 2001, Proceedings*, volume 2102 of *Lecture Notes in Computer Science*, pages 79–90. Springer, 2001.
8. R. Fikes and N. Nilsson. Strips: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.
9. J. Gerbrandy. Communication strategies in games. *Journal of Applied Non-Classical Logics*, 17(2):197–211, 2007.
10. M. Ghallab, D. Nau, and P. Traverso. *Automated Planning: Theory & Practice*. Morgan Kaufmann, May 2004.
11. F. Giunchiglia and P. Traverso. Planning as model checking. In S. Biundo and M. Fox, editors, *Recent Advances in AI Planning, 5th European Conference on Planning, ECP'99, Durham, UK, September 8-10, 1999, Proceedings*, volume 1809 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2000.
12. J. Hoffmann and B. Nebel. The FF planning system: fast plan generation through heuristic search. *Journal of Artificial Intelligence Research*, 14(1):253–302, 2001.
13. E. Kennerly, A. Witzel, and J. A. Zvesper. Thief belief. In B. Löwe, editor, *LSIR-2: Logic and the Simulation of Interaction and Reasoning Workshop at IJCAI-09*, pages 47–51, 2009.
14. G. Lakemeyer and H. J. Levesque. Semantics for a useful fragment of the situation calculus. In L. Pack Kaelbling and A. Saffiotti, editors, *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30-August 5, 2005*, pages 490–496. Morgan Kaufmann Publishers Inc., 2005.
15. B. Löwe and E. Pacuit. An abstract approach to reasoning about games with mistaken and changing beliefs. *Australasian Journal of Logic*, 6:162–181, 2008.
16. B. Löwe, E. Pacuit, and S. Saraf. Identifying the structure of a narrative via an agent-based logic of preferences and beliefs: Formalizations of episodes from CSI: Crime Scene Investigation™. In M. Duvigneau and D. Moldt, editors, *Proceedings of the Fifth International Workshop on Modelling of Objects, Components and Agents (MOCA'09)*, FBI-HH-B-290/09, pages 45–63, 2009.
17. H. Palacios and H. Geffner. From conformant into classical planning: Efficient translations that may be complete too. In M. Boddy, M. Fox, and S. Thiébaux, editors, *Proceedings of the Seventeenth International Conference on Automated Planning and Scheduling (ICAPS 2007)*, 2007.
18. R. Petrick and F. Bacchus. A Knowledge-Based approach to planning with incomplete information and sensing. In M. Ghallab, J. Hertzberg, and P. Traverso, editors, *Proceedings of the Sixth International Conference on Artificial Intelligence Planning and Scheduling (AIPS'02)*, pages 212–221, 2002.
19. G. R. Renardel de Lavalette and H. van Ditmarsch. Epistemic actions and minimal models, 2002. Available on the authors' websites.
20. M. O. Riedl and R. M. Young. Character-focused narrative planning. Unpublished manuscript, 2003.

21. M. O. Riedl and R. M. Young. From linear story generation to branching story graphs. *IEEE Computer Graphics and Applications*, 26(3):23–31, 2006.
22. T. Sadzik. Exploring the iterated update universe, 2006. ILLC Publications PP-2006-26.
23. J. van Benthem and F. Liu. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logic*, 17(2):157–182, 2007.
24. J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.
25. W. van der Hoek and M. Wooldridge. Tractable multiagent planning for epistemic goals. In C. Castelfranchi, editor, *The First International Joint Conference on Autonomous Agents & Multiagent Systems, AAMAS 2002, July 15-19, 2002, Bologna, Italy, Proceedings*, pages 1167–1174. ACM, 2002.
26. H. van Ditmarsch, A. Herzig, and T. de Lima. From situation calculus to dynamic epistemic logic. Unpublished manuscript, 2010.
27. H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2008.
28. J. van Eijck, J. Ruan, and T. Sadzik. Action emulation. Draft paper, 2008.
29. A. Witzel, J. Zvesper, and E. Kennerly. Explicit knowledge programming for computer games. In C. Darken and M. Mateas, editors, *Proceedings of the Fourth Artificial Intelligence and Interactive Digital Entertainment Conference, October 22-24, 2008, Stanford, California, USA*. AAAI Press, 2008.