

DELPHIC: Towards an Efficient Possibility-Based Epistemic Planning Framework

Alessandro Burigana¹, Paolo Felli¹ and Marco Montali¹

¹Free University of Bozen-Bolzano, Italy

Abstract

Dynamic Epistemic Logic (*DEL*) provides a very expressive framework for planning, called *epistemic planning*. In this paper, we discuss the shortcomings of the traditional *DEL* semantics, based on Kripke models, and we propose a formulation of epistemic planning based on a different kind of object, called *possibility*. We argue that the resulting framework, named *DELPHIC*, constitutes a valid alternative, since it provides a more compact representation of epistemic information and a lighter update mechanism. We conclude by discussing the advantages of *DELPHIC*, both theoretically and in terms of implementation.

1. Introduction

In recent times, there has been a growing interest in the study of Multi-Agent Systems and in their application to a wide range of settings. Autonomous agents are expected to face dynamic situations and to be able to adapt in order to reach a given goal. In these contexts, it is paramount for agents to be able to reason on the physical world that surrounds them as well as on the *knowledge* that other agents have about the world and other agents' knowledge.

Bolander and Andersen [1] introduced *epistemic planning* as a planning framework based on Dynamic Epistemic Logic (*DEL*) [2]. Although the traditional Kripke semantics of *DEL* allows for a great expressive power, only small fragments are implemented in practice [3, 4, 5, 6]. In fact, computing a solution for an epistemic planning problem requires to handle a large quantity of information, which results in poor time results. As a result, various epistemic planning systems rely on simplifying the setting of interest.

In this paper, as a first step to overcome this issue, we propose a framework for epistemic planning called *DELPHIC* (*DEL*-planning with a Possibility-based Homogeneous Information Characterization), that relies on an alternative semantics of *DEL* based on *possibilities* [7]. The choice of adopting a possibilities is motivated by the fact that they represent information in a more compact way, which has been shown to improve performance in the solver *EFP* [5]. In what follows, we introduce the main elements of *DEL* (Kripke models, event models, product update) and their corresponding ones in *DELPHIC* (*possibilities*, *eventualities*, *union update*), discussing the advantages of *DELPHIC* with respect to the Kripke semantics, both theoretically and in terms of implementation. We conclude with some future works.


OVERLAY 2022: 4th Workshop on Artificial Intelligence and Formal Verification, Logic, Automata, and Synthesis, November 28, 2022, Udine, Italy

✉ burigana@inf.unibz.it (A. Burigana); paolo.felli@unibz.it (P. Felli); montali@inf.unibz.it (M. Montali)

🆔 0000-0002-9977-6735 (A. Burigana); 0000-0001-9561-8775 (P. Felli); 0000-0002-8021-3430 (M. Montali)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

2. A Novel Semantics for Epistemic Planning


Epistemic Models. Let \mathcal{P} be a countable set of propositional atoms and $\mathcal{AG} = \{1, \dots, n\}$ be a finite set of agents. The language $\mathcal{L}_{\mathcal{P}, \mathcal{AG}}$ of *multi-agent epistemic logic* is defined by the BNF: $\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box_i\varphi$, where $p \in \mathcal{P}, i \in \mathcal{AG}$. The formula $\Box_i\varphi$ is read as “agent i knows/believes that φ ”. In the traditional *possible world* semantics of DEL, a *Kripke model* [8] represents an *epistemic model*. Epistemic models contain both factual information on multiple possible worlds and epistemic information, *i.e.*, which worlds are considered possible by agents.

Definition 1 (Kripke Model). *A Kripke model is a triple $M = (W, R, V)$ where: 1) $W \neq \emptyset$ is the set of possible worlds, 2) $R : \mathcal{AG} \rightarrow 2^{W \times W}$ assigns to each agent i an accessibility relation $R(i)$ (abbreviated as R_i), and 3) $V : \mathcal{P} \rightarrow 2^W$ assigns to each atom a set of worlds.*

A *(multi-)pointed Kripke model* is a pair (M, W_d) , where $W_d \subseteq W$ is a non-empty set of *designated worlds* denoting the real state of affairs, and it represents an *epistemic state* in DEL. Truth of formulae is defined as follows: i. $(M, w) \models p$ iff $w \in V(p)$; ii. $(M, w) \models \neg\varphi$ iff $(M, w) \not\models \varphi$; iii. $(M, w) \models \varphi \wedge \psi$ iff $(M, w) \models \varphi$ and $(M, w) \models \psi$; and iv. $(M, w) \models \Box_i\varphi$ iff $\forall v$ if wR_iv then $(M, v) \models \varphi$. Finally, $(M, W_d) \models \varphi$ iff $(M, v) \models \varphi$, for all $v \in W_d$.

Example 1. *Agents a and b are in a room where a coin is placed in a box. They both do not know whether the coin lies heads up (h) or not.*

Thus, we need two worlds to represent this uncertainty (w_1 and w_2 , respectively). As we are omniscient observers, we assume that we know that w_1 is the designated world (circled dot). In DEL, this is modeled by the epistemic state on the right.



Definition 2 (Possibility [7]). *A possibility u is a function that assigns to each atom $p \in \mathcal{P}$ a truth value $u(p) \in \{0, 1\}$ and to each agent $i \in \mathcal{AG}$ a set of possibilities $u(i)$.*

In DELPHIC, an epistemic model is represented by a *possibility*. A *possibility spectrum* is a finite set $U = \{u_1, \dots, u_k\}$ of possibilities and it represents an epistemic state in DELPHIC. Importantly, possibility spectrums represent the same information as multi-pointed Kripke models, namely factual and epistemic information, plus the designated worlds. In Kripke models, the worlds (W) and their interpretations (V) and relations (R_i) are given as separated components. On the other hand, each possibility *constitutes* a possible world and it contains both an interpretation $u(p)$ for each atom in \mathcal{P} , and a set of worlds $u(i)$ that agent i considers possible, for each agent in \mathcal{AG} . Hence, the designated worlds of the spectrum U are exactly the possibilities u_1, \dots, u_k (whereas the non-designated worlds correspond to possibilities attributed to agents by a possibility in U , but not explicitly included in U themselves). Truth of formulae in possibilities is defined as follows: i. $u \models p$ iff $u(p) = 1$; ii. $u \models \neg\varphi$ iff $u \not\models \varphi$; iii. $u \models \varphi \wedge \psi$ iff $u \models \varphi$ and $u \models \psi$; and iv. $u \models \Box_i\varphi$ iff $\forall v$ if $v \in u(i)$ then $v \models \varphi$. Finally, $U \models \varphi$ iff $v \models \varphi$, for all $v \in U$.


Example 2. *The possibility spectrum representing the situation of Example 1 is $W = \{w_1\}$, where: i. $w_1(h) = 1$ and $w_1(a) = w_1(b) = \{w_1, w_2\}$; and ii. $w_2(h) = 0$ and $w_2(a) = w_2(b) = \{w_1, w_2\}$.*

Event Models. In DEL, actions are modeled by *event models* [9, 10], that represent how new information changes the perspectives of agents. In what follows, *events* play the role of worlds and the accessibility relation Q_i describes which events are considered possible by agent i .

Definition 3 (Event Model). An event model is a quadruple $\mathcal{E} = (E, Q, pre, post)$ where: 1) $E \neq \emptyset$ is the set of events, 2) $Q : \mathcal{AG} \rightarrow 2^{E \times E}$ assigns to each agent i an accessibility relation $Q(i)$ (abbreviated as Q_i), 3) $pre : E \rightarrow \mathcal{L}_{\mathcal{P}, \mathcal{AG}}$ assigns to each event a precondition, and 4) $post : E \rightarrow (\mathcal{P} \rightarrow \mathcal{L}_{\mathcal{P}, \mathcal{AG}})$ assigns to each event a postcondition for each atom.

Informally, preconditions capture the applicability of events, whereas postconditions determine whether an atom is true after the event is applied. A (multi-)pointed event model is a pair (\mathcal{E}, E_d) , where $E_d \subseteq E$ is a non-empty set of *designated events*, and it represents an action in DEL.

Example 3. Suppose now that, in the situation of Example 1, **a** peeks into the box while **b** is distracted. Events e_1 and e_2 represent **a**'s and **b**'s perspectives, respectively (we use the notation $\langle pre, post \rangle$ for events). In this way, **a** learns that the coin lies heads up, while **b** remains ignorant about the situation. Moreover, **a** is aware of **b**'s perspective, but **b** is not aware of **a**'s point of view. In DEL, this is modeled by the action on the right.



In DELPHIC, we introduce the novel concept of *eventuality* to model actions. Let $pre \notin \mathcal{P}$ be a fresh atom and let $\mathcal{P}' = \mathcal{P} \cup \{pre\}$. In the following definition, $e(pre)$ is the precondition of e , while $e(p)$, where $p \in \mathcal{P}$, is the postcondition of p in e .

Definition 4 (Eventuality). An eventuality e is a function that assigns to each atom $p' \in \mathcal{P}'$ a formula $e(p') \in \mathcal{L}_{\mathcal{P}, \mathcal{AG}}$ and to each agent $i \in \mathcal{AG}$ a set of eventualities $e(i)$.

An *eventuality spectrum* is a finite set $E = \{e_1, \dots, e_k\}$ of eventualities and it represents an action in DELPHIC. Importantly, eventuality spectrums represent the same information as multi-pointed event models, namely pre-/postconditions of events (functions pre and $post$) and epistemic information (relations Q_i), plus the designated events. While in event models these are given as separated components, each eventuality constitutes an *event* and it contains a precondition $e(pre)$, a postcondition $e(p)$ for each $p \in \mathcal{P}$, and a set of eventualities $e(i)$ that agent i considers possible, for each $i \in \mathcal{AG}$. The designated events of E are exactly the eventualities e_1, \dots, e_k .


Example 4. The eventuality spectrum representing the action of Example 3 is $E = \{e_1\}$, where: i. $e_1(pre) = h$, $e_1(h) = h$, $e_1(a) = \{e_1\}$ and $e_1(b) = \{e_2\}$; and ii. $e_2(pre) = \top$, $e_2(h) = h$ and $e_2(a) = e_2(b) = \{e_1, e_2\}$.

Product Update. The product update formalizes how (*applicable*) actions bring about information change in epistemic states in DEL. An action (\mathcal{E}, E_d) is applicable in (M, W_d) if for each world $w \in W_d$ there exists an event $e \in E_d$ such that $(M, w) \models pre(e)$.

Definition 5 (Product Update). Let $M = (W, R, V)$ and $\mathcal{E} = (E, Q, pre, post)$. The product update of (M, W_d) with (\mathcal{E}, E_d) is $(M, W_d) \otimes (\mathcal{E}, E_d) = ((W', R', V'), W'_d)$, where:

1. $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$,
2. $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_iv \text{ and } eQ_if\}$,
3. $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$, and
4. $W'_d = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

Example 5. Applying the action of Example 3 in the epistemic state of Example 1 results in the state on the right, where $v_3 = (w_1, e_1)$, $v_1 = (w_1, e_2)$ and $v_2 = (w_2, e_2)$. Notice that w_1 (resp., w_2) and v_1 (resp., v_2) encode the same information, but they are distinct objects.



In DELPHIC, we introduce the concept of *union update* in place of the traditional product update to formalize how an (*applicable*) eventuality spectrum brings about information change in a possibility spectrum. An eventuality spectrum E is applicable in a possibility spectrum U if for each possibility $u \in U$ there exists an eventuality $e \in E$ such that $u \models e(\text{pre})$.

Definition 6 (Union Update). *The union update of a possibility u with an eventuality e is the possibility $u' = u \boxtimes e$, such that if $u \not\models e(\text{pre})$, then $u' = \emptyset$; otherwise:*

$$\begin{aligned} u'(p) &= 1 \text{ iff } u \models e(p) \\ u'(i) &= \{v \boxtimes f \mid v \in u(i), f \in e(i) \text{ and } v \models f(\text{pre})\} \end{aligned}$$

The union update of a possibility spectrum U with an eventuality spectrum E is the possibility spectrum $U \boxtimes E = \{u \boxtimes e \mid u \in U, e \in E \text{ and } u \models e(\text{pre})\}$.

Example 6. *Applying E (Example 4) to W (Example 2) results in: $W \boxtimes E = \{w_1 \boxtimes e_1\} = \{v_3\}$, where $v_3(h) = 1$, $v_3(a) = \{v_3\}$ and $v_3(b) = \{w_1 \boxtimes e_2, w_2 \boxtimes e_2\} = \{w_1, w_2\}$. Notice that the union update allows us to reuse previously calculated information (i.e., possibilities w_1 and w_2).*

3. Discussion and Future Works

The DELPHIC framework for epistemic planning has various advantages over the traditional Kripke semantics for DEL. From a conceptual standpoint, possibilities allow for a more natural representation of epistemic states: while a Kripke model represents information as *heterogeneous* components (i.e., set of worlds, accessibility relations, interpretation of atoms in each world), possibilities encode all this information within one coherent component (namely, the possibility itself). In other words, a possibility encodes both a possible *world* and the information about which *worlds* are considered possible by each agent. This is more akin to how humans think of a situation, i.e., as a set of different possibilities. The same argument holds for actions, since an eventuality is essentially a possibility that maps formulae (instead of truth values) to atoms.

From a theoretical perspective, possibilities are tightly related to Kripke models and they are proved to be a more compact representation of epistemic models [7]. This was shown to improve runtime performance in the solver EFP [5], where a fragment of DEL was implemented with both semantics. Moreover, the union update overcomes two important shortcomings of the traditional product update. First, the product update is essentially a *blind* cross-product between worlds and events, which may result in epistemic states having unreachable components. Importantly, this redundant information accumulates during the planning process and produces a significant overhead in solving time. Second, the cross-product makes it impossible to *reuse* previous information. This is a crucial aspect, since it is common to have some agents maintain their perspective after an action is carried out. Instead, in DEL we are forced to create a new world and to *copy* the previous information in it (see Example 5). Conversely, due to the semantics of union update, in DELPHIC it is possible to reuse information, as shown in Example 6.

These arguments suggest that DELPHIC allows for a lighter update mechanism, that would ultimately result in more efficient implementations. To obtain empirical evidence of such improvement, we plan to develop an encoding both of DELPHIC and the traditional DEL semantics into various solver, e.g., ASP, Prolog, SMT solvers, and comparatively evaluate the two versions.

Acknowledgments

This work is partially supported by the Italian Ministry of University and Research under the PRIN programme, grant B87G22000450001 (PINPOINT), and by the Free University of Bozen-Bolzano with the SMART-APP and ADAPTERS projects.

References

- [1] T. Bolander, M. B. Andersen, Epistemic planning for single and multi-agent systems, *J. Appl. Non Class. Logics* 21 (2011) 9–34. URL: <https://doi.org/10.3166/jancl.21.9-34>. doi:10.3166/jancl.21.9-34.
- [2] H. P. van Ditmarsch, W. van der Hoek, B. P. Kooi, *Dynamic Epistemic Logic*, volume 337, Springer Netherlands, 2007. URL: <https://doi.org/10.1007/978-1-4020-5839-4>. doi:10.1007/978-1-4020-5839-4.
- [3] C. Baral, G. Gelfond, E. Pontelli, T. C. Son, An action language for multi-agent domains: Foundations, *CoRR abs/1511.01960* (2015). URL: <http://arxiv.org/abs/1511.01960>. arXiv:1511.01960.
- [4] F. Kominis, H. Geffner, Beliefs In Multiagent Planning: From One Agent to Many, in: R. I. Brafman, C. Domshlak, P. Haslum, S. Zilberstein (Eds.), *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling, ICAPS 2015, Jerusalem, Israel, June 7-11, 2015*, AAAI Press, 2015, pp. 147–155. URL: <http://www.aaai.org/ocs/index.php/ICAPS/ICAPS15/paper/view/10617>.
- [5] F. Fabiano, A. Burigana, A. Dovier, E. Pontelli, EFP 2.0: A multi-agent epistemic solver with multiple e-state representations, in: J. C. Beck, O. Buffet, J. Hoffmann, E. Karpas, S. Sohrabi (Eds.), *Proceedings of the Thirtieth International Conference on Automated Planning and Scheduling, Nancy, France, October 26-30, 2020*, AAAI Press, 2020, pp. 101–109. URL: <https://aaai.org/ojs/index.php/ICAPS/article/view/6650>.
- [6] F. Fabiano, A. Burigana, A. Dovier, E. Pontelli, T. C. Son, Multi-agent epistemic planning with inconsistent beliefs, trust and lies, in: *PRICAI 2021: Trends in Artificial Intelligence - 18th Pacific Rim International Conference on Artificial Intelligence, PRICAI 2021, Hanoi, Vietnam, November 8-12, 2021, Proceedings, Part I*, Springer, 2021, pp. 586–597. URL: https://doi.org/10.1007/978-3-030-89188-6_44.
- [7] J. Gerbrandy, W. Groeneveld, Reasoning about information change, *J. Log. Lang. Inf.* 6 (1997) 147–169. URL: <https://doi.org/10.1023/A:1008222603071>. doi:10.1023/A:1008222603071.
- [8] S. A. Kripke, Semantical considerations on modal logic, *Acta Philosophica Fennica* 16 (1963) 83–94.
- [9] A. Baltag, L. S. Moss, S. Solecki, The logic of public announcements and common knowledge and private suspicions, in: I. Gilboa (Ed.), *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98)*, Evanston, IL, USA, July 22-24, 1998, Morgan Kaufmann, 1998, pp. 43–56.
- [10] H. van Ditmarsch, B. Kooi, *Semantic results for ontic and epistemic change*, *Texts in Logic and Games* 3, Amsterdam University Press, 2008, pp. 87–117.