Reasoning about Non-Deterministic Observability and Hypothetical Action Occurrences in Multi-Agent Domains

Loc Pham
New Mexico State University
Las Cruces, New Mexico
locpham@nmsu.edu

Tran Cao Son
New Mexico State University
Las Cruces, New Mexico
tson@cs.nmsu.edu

Enrico Pontelli
New Mexico State University
Las Cruces, New Mexico
epontelli@cs.nmsu.edu

ABSTRACT

This paper proposes a method based on high-level action description languages for reasoning about hypothetical action occurrences in a multi-agent environments. In order to accommodate this type of reasoning, one needs to consider *non-deterministic observability of action occurrences* of agents. The paper presents an extension of the language m$\mathcal{A}^*$, called m$\mathcal{A}^*_e$, to allow for non-deterministic observability and hypothetical actions. The paper defines the semantics of the new language using edge-conditioned update models and discusses properties of the new language, including differences from approaches that use Dynamic Epistemic Logic. The new definitions are illustrated using the well-known story of two stockbrokers from the literature.

KEYWORDS

Epistemic Reasoning, Multi-Agent Domains, Action Language, Non-Deterministic Observability, Hypothetical Action Occurrences.

ACM Reference Format:
Loc Pham, Tran Cao Son, and Enrico Pontelli. 2023. Reasoning about Non-Deterministic Observability and Hypothetical Action Occurrences in Multi-Agent Domains. In *Proceedings of ACM SAC Conference (SAC'23).* ACM, New York, NY, USA, Article 4, 8 pages. https://doi.org/10.1145/3555776.3577612

1 INTRODUCTION

Reasoning about occurrences of hypothetical actions has been an important part of Reasoning about Actions and Change (RAC). In single-agent environments, this enables an agent to predict the state of the world after the execution of an action sequence, in turn allowing the agent to perform planning. In multi-agent environments, an action occurrence is often associated with the agent(s) executing such action and affects both the state of the world and the beliefs of agents. For this reason, when an agent conceives that an action might occur, they need to take into consideration the actor who executes the action and the potential changes in beliefs of every agent in the domain. Furthermore, once an agent thinks that the action might have occurred but they cannot witness the action occurrence, they will not be able to return to the previous state and continue with their reasoning, as their beliefs might have changed and thus cannot assume that all agents' beliefs are the same as before. This is highlighted in the following example, discussed in [14].

Example 1.1. Two stockbrokers Anne and Bill are sitting at a table in a Wall Street bar. A messenger comes in and delivers a letter to Anne that contains information about the status of United Agents. Suppose that both Anne and Bill did not have any information about the status of United Agents before their meeting. Suppose that United Agents is doing well. Consider two scenarios:

(1) **(mayread)** Bill leaves the table and orders a drink at the bar so that Anne may have read the letter while he was away (Anne does not read the letter).
(2) **(bothmayread)** Bill orders a drink at the bar while Anne goes to the bathroom. Each may have read the letter while the other was away from the table (Both read the letter).

For simplicity of the presentation, we focus on the first scenario. In this scenario, **mayread** is a hypothetical action occurrence of the action *"Anne reads the letter"* and Bill who cannot observe whether the action has actually occurred. If Bill hypothesizes that Anne reads the letter then he will need to update the belief of Anne about United Agents (i.e., Anne knows that the company is doing well); Bill also updates his belief about Anne's belief about United Agents. On the other hand, if he believes that Anne does not read the letter then he will not update anything. However, after conceiving that the action might occur and *not being able to observe whether the action occurs,* Bill's belief changes and should not revert to the state of his belief at the start of the meeting with Anne, because of his uncertainty about the action occurrence. We say that Bill has a *non-deterministic observability of the occurrence of* **mayread**.

In this paper, we investigate the problem of reasoning with non-deterministic observability and hypothetical action occurrences. We propose an approach to deal with these features in m$\mathcal{A}^*$ [3], by devising m$\mathcal{A}^*_e$, an extension of m$\mathcal{A}^*$ which supports these features. We use m$\mathcal{A}^*$ because this is among the first action languages that utilize the concept of update models, a well-known notion used in RAC in multi-agent domains, originally introduced in [1, 2] (under the name of *action model*) and later extended to *update model* [11, 15]. The novelty of this language lies in that it provides a method to automatically generate the update model of an action occurrence given the domain description and the state of the world and beliefs (captured by pointed Kripke structures). A similar approach has recently been developed by [10]. This construction is simple, in that it only uses update models with at most three events. Furthermore, early versions of the languages m$\mathcal{A}^*$ have proved to be useful in the development of epistemic planning systems, such as those proposed in [6, 8].

SAC'23, March 27 – March 31, 2023, Tallinn, Estonia

ACM ISBN 978-1-4503-9517-5/23/03…$15.00
https://doi.org/10.1145/3555776.3577612

In m$\mathcal{A}^*$, the observability of an action occurrence by agents is classified into three different classes: *full observers*, *partial observers* and *oblivious*. It has been pointed out in [3] that the simplicity of m$\mathcal{A}^*$ prevents it from being able to handle the hypothetical action occurrences of **mayread** as discussed in Example 1.1.

We begin with a short review of the language m$\mathcal{A}^*$ and some necessary notions, such as belief formula, Kripke structure, pointed Kripke structure, edge-conditioned update models, and update of Kripke structures by edge-conditioned update models. We then define the semantics of m$\mathcal{A}^*_e$ using edge-conditioned update models and show how m$\mathcal{A}^*_e$ can be used in situations mentioned in the previous example. Finally, we prove relevant properties of the new semantics for m$\mathcal{A}^*_e$.

2 BACKGROUND

Belief Formulae. A *multi-agent* domain $\langle \mathcal{AG}, \mathcal{F} \rangle$ includes a finite and non-empty set of agents $\mathcal{AG}$ and a set of fluents (atomic propositions) $\mathcal{F}$, used to encode the properties of the world. *Belief formulae* over $\langle \mathcal{AG}, \mathcal{F} \rangle$ are defined by the BNF: "$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid \mathbf{B}_i\varphi$" where $p \in \mathcal{F}$ is a fluent and $i \in \mathcal{AG}$. We refer to a belief formula which does not contain any occurrence of $\mathbf{B}_i$ as *a fluent formula*. In addition, for a formula φ and a non-empty set $\alpha \subseteq \mathcal{AG}$, $\mathbf{B}_\alpha\varphi$ and $\mathbf{C}_\alpha\varphi$ denote $\bigwedge_{i\in\alpha} \mathbf{B}_i\varphi$ and $\bigwedge_{k=1}^{\infty} \mathbf{B}_\alpha^k\varphi$, where $\mathbf{B}_\alpha^1\varphi$=$\mathbf{B}_\alpha\varphi$ and $\mathbf{B}_\alpha^k\varphi$=$\mathbf{B}_\alpha^{k-1}\mathbf{B}_\alpha\varphi$ for $k > 1$, respectively. $\mathcal{L}_{\mathcal{AG}}$ denotes the set of belief formulae over $\langle \mathcal{AG}, \mathcal{F} \rangle$.

Satisfaction of belief formulae is defined over *pointed Kripke structures* [7]. A Kripke structure M is a tuple $\langle S, \pi, \{\mathcal{B}_i\}_{i\in\mathcal{AG}} \rangle$, where S is a set of worlds (denoted by $M[S]$), $\pi : S \mapsto 2^{\mathcal{F}}$ is a function that associates an interpretation of $\mathcal{F}$ to each element of S (denoted by $M[\pi]$), and for $i \in \mathcal{AG}$, $\mathcal{B}_i \subseteq S \times S$ is a binary relation over S (denoted by $M[i]$). For convenience, we will often draw a Kripke structure M as a directed labeled graph, whose set of labeled nodes represents S and whose set of labeled edges contains $s \xrightarrow{i} t$ iff $(s, t) \in \mathcal{B}_i$; the label of each node is the name of the world and its interpretation is displayed as a text box next to it. For $u \in S$ and a fluent formula φ, $M[\pi](u)$ and $M[\pi](u)(\varphi)$ denote the interpretation associated to u via π and the truth value of φ with respect to $M[\pi](u)$. For a world $s \in M[S]$, (M, s) is a *pointed Kripke structure*, hereafter called a *state*.

The satisfaction relation $\models$ between belief formulae and a state (M, s) is defined as follows:

(1) $(M, s) \models p$ if p is a fluent and $M[\pi](s)(p)$ is true;
(2) $(M, s) \models \neg\varphi$ if $(M, s) \not\models \varphi$;
(3) $(M, s) \models \varphi_1 \wedge \varphi_2$ if $(M, s) \models \varphi_1$ and $(M, s) \models \varphi_2$;
(4) $(M, s) \models \varphi_1 \vee \varphi_2$ if $(M, s) \models \varphi_1$ or $(M, s) \models \varphi_2$;
(5) $(M, s) \models \mathbf{B}_i\varphi$ if $\forall t.[(s, t) \in \mathcal{B}_i \Rightarrow (M, t) \models \varphi]$.

Edge-Conditioned Update Models. The formalism of *update models* has been used to describe transformations of states according to a predetermined transformation pattern (see, e.g., [1, 11]). This formalism makes use of the notion of $\mathcal{L}_{\mathcal{AG}}$-substitution, which is a set $\{p_1 \to \varphi_1, \ldots, p_k \to \varphi_k\}$, where each p_i is a distinct fluent in $\mathcal{F}$ and each $\varphi_i \in \mathcal{L}_{\mathcal{AG}}$. $SUB_{\mathcal{L}_{\mathcal{AG}}}$ denotes the set of all $\mathcal{L}_{\mathcal{AG}}$-substitutions. To handle the non-deteministic observability in m$\mathcal{A}^*$, we will utilize the *edge-conditioned event update models* as proposed in [4]. In edge-conditioned event update models, the assumption that full and partial observers know observability of all agents is not required. An edge-conditioned event update model Σ is a tuple $\langle \Sigma, \{R_i\}_{i\in\mathcal{AG}}, pre, sub \rangle$ where Σ is a set of events, $R_i \subseteq \Sigma \times \mathcal{L}_{\mathcal{AG}} \times \Sigma$ is the accessibility relation of agent i between events, $pre : \Sigma \to \mathcal{L}_{\mathcal{AG}}$ is a function mapping each event $e \in \Sigma$ to a formula in $\mathcal{L}_{\mathcal{AG}}$, $sub : \Sigma \to SUB_{\mathcal{L}_{\mathcal{AG}}}$ is a function mapping each event $e \in \Sigma$ to a substitution in $SUB_{\mathcal{L}_{\mathcal{AG}}}$. Elements of R_i are in the form (e_1, γ, e_2) where γ is a belief formula. In the graph representation, such an accessibility relation is shown by a directed edge from e_1 to e_2 with the label $i : \gamma$. We will omit γ and write simply i as label of the edge when $\gamma = \top$. Given an edge-conditioned update model Σ, an *update instance* ω is a pair (Σ, e) where e is an event in Σ, referred to as a *designated event* (or *true event*). For simplicity of the presentation, we often draw an update instance as a graph whose events are rectangles and whose links represent the accessibility relations between events with double square representing the *designated event*.

Given a Kripke structure M and an edge-conditioned update model $\Sigma = \langle \Sigma, \{R_i\}_{i\in\mathcal{AG}}, pre, sub \rangle$, the *update* of M induced by Σ results in a new Kripke structure M', denoted by $M' = M \otimes \Sigma$, defined by:

(i) $M'[S] = \{(s, \tau) \mid \tau \in \Sigma, s \in M[S], (M, s) \models pre(\tau)\}$;
(ii) $((s, \tau), (s', \tau')) \in M'[i]$ iff $(s, \tau), (s', \tau') \in M'[S]$, $(s, s') \in M[i]$, $(\tau, \gamma, \tau') \in R_i$ and $(M, s) \models \gamma$;
(iii) For all $(s, \tau) \in M'[S]$ and $f \in \mathcal{F}$, $M'[\pi]((s, \tau)) \models f$ if and only if $f \to \varphi \in sub(\tau)$ and $(M, s) \models \varphi$.

For simplicity of the presentation, we will use update model instead of edge-conditioned update model from now on. An *update template* is a pair (Σ, Γ), where Σ is an update model with the set of events Σ and $\Gamma \subseteq \Sigma$. The update of a pointed Kripke structure (M, s) given an update template (Σ, Γ) is a set of pointed Kripke structures, denoted by $(M, s) \otimes (\Sigma, \Gamma)$, where $(M, s) \otimes (\Sigma, \Gamma) = \{(M \otimes \Sigma, (s, \tau)) \mid \tau \in \Gamma, (M, s) \models pre(\tau)\}$.

Syntax of m$\mathcal{A}^*$. An action theory in the language m$\mathcal{A}^*$ over $\langle \mathcal{AG}, \mathcal{F} \rangle$ consists of a set of action instances $\mathcal{AI}$ of the form $a\langle\alpha\rangle$, representing that a set of agents α performs action a, and a collection of statements of the following forms:

a	**executable_if**	ψ	(1)
a	**causes** ℓ **if**	ψ	(2)
a	**determines**	φ	(3)
a	**announces**	φ	(4)
z	**observes** a **if**	δ_z	(5)
z	**aware_of** a **if**	θ_z	(6)
	initially	ψ	(7)

where ℓ is a fluent literal (a fluent $f \in \mathcal{F}$ or its negation $\neg f$), ψ is a belief formula, φ, δ_z and θ_z are fluent formulae, a $\in \mathcal{AI}$, and $z \in \mathcal{AG}$. Statement (1) encodes the executability condition of a and ψ is referred as the *precondition of* a. Statement (2) describes the effect of the ontic action a, i.e., if ψ is true then ℓ will be true after the execution of a. Statement (3) enables the agents who execute a to learn the value of the formula φ. Statement (4) encodes an *announcement* action, whose owner announces that φ is true. Statements of the forms (5)–(6) encode the observability of agents given an occurrence of a. Statement (5) indicates that agent z is a full observer of a if δ_z holds. Statement (6) states that agent z is a partial observer of a if θ_z holds. z, a, and δ_z (resp. θ_z) are referred to as the observed agent, the action instance, and the condition of (5) (resp.

(6)). It is assumed that the sets of ontic actions, sensing actions, and announcement actions are pairwise disjoint. Furthermore, for every pair of a and z, if z and a occur in a statement of the form (5) then they do not occur in any statement of the form (6) and vice versa. Also, we assume that z **observes** a **if** *False* (or z **aware_of** a **if** *False*) is in the action theory if the information about δ_z (or θ_z) is not given. Statements of the form (7) indicate that ψ is true in the initial state. An action domain is a collection of statements (1)–(6). An action theory is a pair of an action domain and a set of statements of the form (7). By this definition, action domains are deterministic in that each ontic action, when executed in a world, results in a unique world.

The semantics of $m\mathcal{A}^*$ is defined by a transition function (denoted by Φ) which maps pairs of action instances and *states* (pointed Kripke structures) into sets of states. It starts with the definition of the update template[1] for the execution of a, denoted with $\omega(\text{a})$. The result of executing a in (M, s) is then defined by $(M', s') \otimes \omega(\text{a})$ where (M', s') is the result of a belief revision, necessary to allow agents to correct some false beliefs prior to the execution of a. The belief revision process, however, is not needed for the cases of non-deterministic observability considered in this paper.

3 NON-DETERMINISTIC OBSERVABILITY AND HYPOTHETICAL ACTION OCCURRENCES

In this section, we will extend $m\mathcal{A}^*$ to deal with non-deterministic observability and hypothetical action occurrences. We will use Example 1.1 as a running example to illustrate the proposed definitions.

Representation. To encode non-deterministic observability and hypothetical action occurrences, we introduce statements of the following forms:

$$z \ \ \textbf{may\_observe} \ \ \text{a} \ \ \textbf{if} \ \ \gamma_z \tag{8}$$

$$z \ \ \textbf{may\_aware\_of} \ \ \text{a} \ \ \textbf{if} \ \ \vartheta_z \tag{9}$$

where γ_z and ϑ_z are fluent formula, $\text{a} \in \mathcal{AI}$, and $z \in \mathcal{AG}$. Statements of the form (8) and (9) encode two types of non-deterministic observability. Statement (8) indicates that agent z is not sure whether action a actually happens, but if it does then she should know all the effect of a (she must be a full observer of a). Statement (9) indicates that agent z is not sure whether action a actually happens, but if it does then she must be a partial observer of a. We assume δ_z, θ_z, γ_z and ϑ_z are mutually exclusive. Again, we assume that z **may_observe** a **if** *False* and z **may_aware_of** a **if** *False* are in the action theory if the information about γ_z and ϑ_z are not given. For reference, we will denote with $m\mathcal{A}^*_e$ the language $m\mathcal{A}^*$ with (8) and (9).

Let us denote the multi-agent domain described in Example 1.1 by D_{letter}. For this domain, we have that $\mathcal{AG} = \{A, B\}$ where A and B represent Anne and Bill, respectively. The set of fluents $\mathcal{F}$ for this domain consists of *doing_well* (United Agents is doing well) and at_table_x (agent x is at the table where $x \in \{A, B\}$). D_{letter} has three actions: *read*, *leave* and *return*. The main focus of this work is about situations where agents are uncertain about the occurrences of some actions; for this example, we focus on the action *read*. We will provide information about the other two actions and their effects only when needed. The *read* action can be represented by the following $m\mathcal{A}^*_e$ statements:

$$read\langle x\rangle \ \ \textbf{determines} \ \ doing\_well \tag{10}$$

$$x \ \ \textbf{observes} \ \ read\langle x\rangle \tag{11}$$

$$y \ \ \textbf{aware\_of} \ \ read\langle x\rangle \ \ \textbf{if} \ \ at\_table_x \tag{12}$$

$$y \ \ \textbf{may\_aware\_of} \ \ read\langle x\rangle \ \ \textbf{if} \ \ \neg at\_table_x \tag{13}$$

where $x, y \in \{A, B\}$ and $x \neq y$. Initially, both Anne and Bill do not know whether United Agents is doing well and both are at the table. So the initial state is given in the following figure where at_table_A and at_table_B are true in both s_0 and s_1:

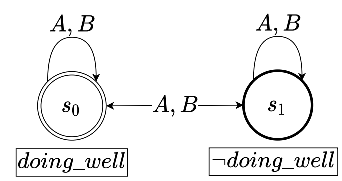

Figure 1: The initial state (M_0, s_0)

For space limitation, we only mention the fluent *doing_well* in Figure 1 and other figures in the paper. The values of other fluents (at_table_A and at_table_B) will be specified whenever it is needed.

We will next define the update template $\omega(\text{a})$ for an action occurrence a when agents have non-deterministic observability. Similar to the definition of the semantics of $m\mathcal{A}^*$ [3], we will define $\omega(\text{a})$ for different types of actions assuming that the domain contains some statements of the form (8) or (9). This is because $\omega(\text{a})$ has already been define when no statement of these types is present.

3.1 Ontic Actions with Non-Deterministic Observability

We assume that an action domain D is given. As in $m\mathcal{A}^*$, there is no partial observer for ontic actions, i.e., we assume that there is no statement in the form (6) and (9) for a if a is an ontic action instance.

Definition 3.1. Let a be an ontic action instance with the precondition ψ. The update model for a, denoted by $\omega(\text{a})$, is defined by $\langle \Sigma, \{R_i\}_{i \in \mathcal{AG}}, pre, sub\rangle$ where

- $\Sigma = \{\sigma, \epsilon\}$;
- $R_i = \{(\sigma, \delta_i \vee \gamma_i, \sigma), (\sigma, \neg\delta_i, \epsilon), (\epsilon, \neg\delta_i \wedge \gamma_i, \sigma), (\epsilon, \top, \epsilon)\}$ where[2] "i **observes** a **if** δ_i" and "i **may_observe** a **if** γ_i" belong to D;
- $pre(\sigma) = \psi$ and $pre(\epsilon) = \top$; and
- $sub(\epsilon) = \emptyset$ and $sub(\sigma) = \{p \rightarrow \Psi^+(p, \text{a}) \vee (p \wedge \neg\Psi^-(p, \text{a})) \mid p \in \mathcal{F}\}$, where
  $\Psi^+(p, \text{a}) = \bigvee\{\varphi \mid [\text{a} \ \textbf{causes} \ p \ \textbf{if} \ \varphi] \in D\}$ and
  $\Psi^-(p, \text{a}) = \bigvee\{\varphi \mid [\text{a} \ \textbf{causes} \ \neg p \ \textbf{if} \ \varphi] \in D\}$.

[1] The original semantics of $m\mathcal{A}^*$ does not use edge-conditioned update model.

[2] Recall that we assume that statements of these forms are in the action theory for every agent i.

When an ontic action occurs, an agent may or may not observe its occurrence or be unsure about its occurrence. As such, $\omega(\text{a})$ has two events. σ is the designated event representing the true occurrence of the action whereas ϵ denotes the null event representing that the action does not occur. σ is the event full observers believe occurring and ϵ is the event seen by oblivious agents. For the agents who think a could happen (but are not sure about that), they should consider both σ and ϵ are possible.

It is instructive to point out that the accessibility relation for each agent i is well-defined as δ_i and γ_i are mutually exclusive. The proposed update model differs from a recently proposed edge-conditioned update model for ontic actions in [9] in that it includes γ_i in the condition for the relations $(\sigma, \delta_i \vee \gamma_i, \sigma)$ and $(\epsilon, \neg\delta_i \wedge \gamma_i, \sigma)$. The first one says that i recognizes that σ is the true event (a occurs) if i is a full observer (δ_i is true) or non-deterministically observes that a occurs. The second relation says that when γ_i is true then i also believes that ϵ is the true event.

The update model from Definition 3.1 allows us to reason with non-deterministic observability. It also enables us to reason with hypothetical action occurrences. The key difference lies in the difference designated events for different types of reasoning. Specifically,

- When a actually occurs, the update template is $(\omega(\text{a}), \{\sigma\})$;
- For hypothetical reasoning
  - if a indeed occurs, the update template is still $(\omega(\text{a}), \{\sigma\})$.
  - if a does not occur, the update template in this case is $(\omega(\text{a}), \{\epsilon\})$.

It is worth noticing that the update models (templates) are constructed from the perspective of an external observer. We assume that the external observer has no uncertainly and thus consider only two situations when the action does (or does not) actually occur. The situation when the external observer has uncertainty about action occurrences is an interesting one and deserves an in-depth exploration. It is left as a topic of interest for our future work.

Figure 2 shows the edge-conditioned update template of a 'real' ontic action a while Figure 3 illustrates the edge-conditioned update template for a hypothetical ontic action occurrence. In the figure, we use $i \in X : \delta_i$ as a shorthand for the set of links with labels $\{i : \delta_i \mid i \in X\}$.

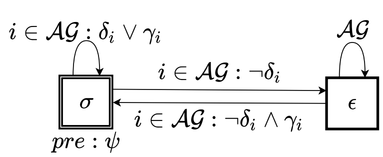

Figure 2: Edge-conditioned update template for an ontic action occurrence with non-deterministic observability.

3.2 Sensing/Announcement Actions with Non-Deterministic Observability

For sensing and announcement actions, an agent can fully observe, partially observe, or be oblivious of its occurrence. As such, we define the update template for their occurrences as follows.

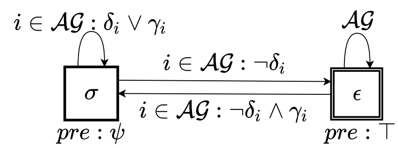

Figure 3: Edge-conditioned update template for a hypothetical ontic action occurrence when the action does not actually occur.

Definition 3.2. Let a be a sensing action instance that senses φ or an announcement action instance that announces φ with the precondition ψ. The update model for a, denoted by $\omega(\text{a})$, is defined by $\langle \Sigma, \{R_i\}_{i \in \mathcal{AG}}, pre, sub \rangle$ where

- $\Sigma = \{\sigma, \tau, \epsilon\}$;
- $R_i = \{(\sigma, \delta_i \vee \theta_i \vee \gamma_i \vee \vartheta_i, \sigma), (\tau, \delta_i \vee \theta_i \vee \gamma_i \vee \vartheta_i, \tau), (\sigma, \neg(\delta_i \vee \gamma_i) \wedge (\theta_i \vee \vartheta_i), \tau), (\tau, \neg(\delta_i \vee \gamma_i) \wedge (\theta_i \vee \vartheta_i), \sigma), (\sigma, \neg(\delta_i \vee \theta_i), \epsilon), (\tau, \neg(\delta_i \vee \theta_i), \epsilon), (\epsilon, \neg(\delta_i \vee \theta_i) \wedge (\gamma_i \wedge \vartheta_i), \sigma), (\epsilon, \neg(\delta_i \vee \theta_i) \wedge (\gamma_i \wedge \vartheta_i), \tau), (\epsilon, \top, \epsilon)\}$ where "i **observes** a **if** δ_i", "i **aware_of** a **if** θ_i", "i **may_observe** a **if** γ_i", and "i **may_aware_of** a **if** ϑ_i" belong to D;
- $pre(\sigma) = \psi \wedge \varphi$, $pre(\tau) = \psi \wedge \neg\varphi$ and $pre(\epsilon) = \top$;
- $sub(x) = \emptyset$ for each $x \in \Sigma$.

Similar to ontic action, the update model from Definition 3.2 allow us to reason with non-deterministic observability for sensing and announcement actions. It also enables us to reason with hypothetical action occurrences. The key difference lies in the difference designated events for different types of reasoning. Specifically,

- When a actually occurs, the update template is $(\omega(\text{a}), \{\sigma, \tau\})$ for sensing action, while for announcement action is $(\omega(\text{a}), \{\sigma\})$;
- For hypothetical reasoning
  - if a indeed occurs, the update template is still $(\omega(\text{a}), \{\sigma, \tau\})$ for sensing action and $(\omega(\text{a}), \{\sigma\})$ for announcement action.
  - if a does not occur, the update template in this case is $(\omega(\text{a}), \{\epsilon\})$.

Observe that an update model of a sensing or announcement action has three events. However, when the action actually happens in real life, the true event can be σ or τ for sensing action whereas for announcement actions, the true event is σ. As in ontic actions, ϵ is the null event representing that the action does not occur. Therefore, in the case of *hypothetical action occurrence*, ϵ would become the true event for all types of actions. Sensing and announcement actions do not alter the state of the world and thus *sub* is empty (or is the identity mapping) for every event. Full observers learn the value of the formula while partial observers only know that the action has taken place without learning the actual outcome. For the agents that have non-deterministic observability, a similar argument as in ontic actions can be applied here: if agent i thinks she could be a full observer of a, but she is not sure about that, then she can not distinguish between σ and ϵ (τ and ϵ also). If i thinks action a could happen but she does not know the effect of a, then she would consider all three events σ, τ, and ϵ possible. Figure 4 and Figure 5 illustrate the new edge-conditioned update model for

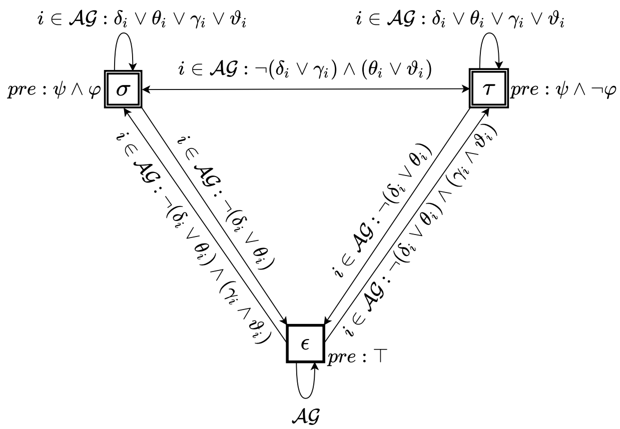

Figure 4: Edge-conditioned update template for a sensing action occurrence

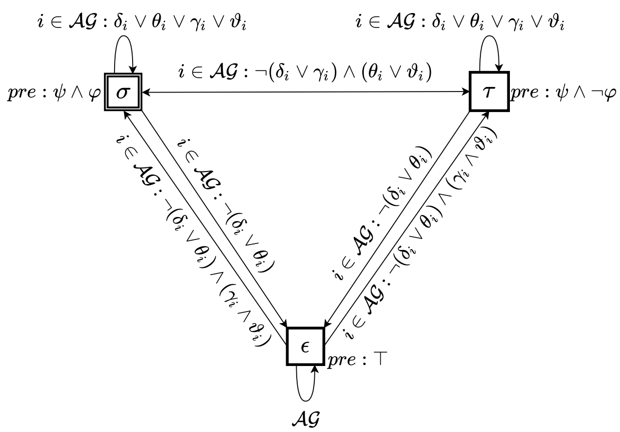

Figure 5: Edge-conditioned update template for an announcement action occurrence

a sensing action a that determines φ and an announcement a that truthfully announces φ respectively; while Figure 6 encodes the update template for a hypothetical sensing/announcement action occurrence.

3.3 Transition Function for m$\mathcal{A}_e^*$

Similar to m$\mathcal{A}^*$, the semantic of m$\mathcal{A}_e^*$ is also defined by a transition function Φ_D that maps pairs of action occurrences (represented by update models) and states (represented by pointed Kripke models) into a set of states. Intuitively, $\Phi_D(\text{a}, (M, s))$ encodes the state that is the result of the execution of a in (M, s). The definition of Φ_D in m$\mathcal{A}^*$ takes into consideration false/incorrect beliefs of agents in some situations. However, this belief revision process is not needed for the cases of non-deterministic observability and hypothetical action occurrences considered in this paper. For simplicity of the representation, we will assume that in the following, (M, s) is the state that is obtained after the belief update process in m$\mathcal{A}^*$.

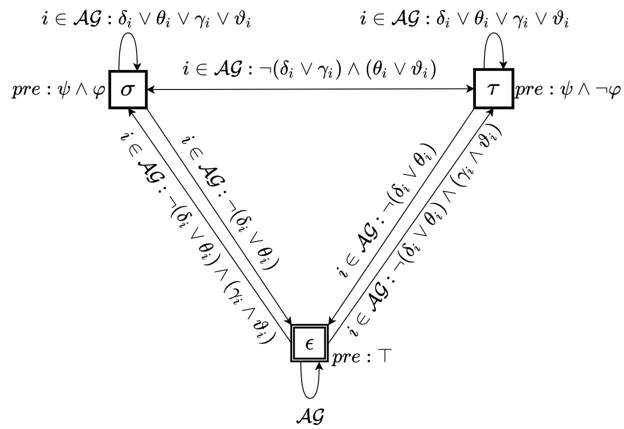

Figure 6: Edge-conditioned update template for a hypothetical sensing/announcement action occurrence when the action does not occury

Consider an action instance a $\in \mathcal{AI}$ whose precondition is ψ, a state (M, s), and a set of agents α. We say a is executable in (M, s) if $(M, s) \models \psi$. The result of executing a in (M, s) is a set of states, denoted by $\Phi_D(\text{a}, (M, s))$ and defined as follows:

- If a is not executable in (M, s) then $\Phi_D(\text{a}, (M, s)) = \emptyset$;
- If a is executable in (M, s) and $(\mathcal{E}, E_d)$ is the update template of action instance a in (M, s) then $\Phi_D(\text{a}, (M, s)) = (M, s) \otimes (\mathcal{E}, E_d)$.

Finally, for a set of states $\mathcal{M}$,

- If a is not executable in some $(M, s) \in \mathcal{M}$ then $\Phi_D(\text{a}, \mathcal{M}) = \emptyset$;
- If a is executable in every $(M, s) \in \mathcal{M}$ then
$$\Phi_D(\text{a}, \mathcal{M}) = \bigcup_{(M,s) \in \mathcal{M}} \Phi_D(\text{a}, (M, s)).$$

3.4 Reasoning in D_{letter}

We will now illustrate the use of m$\mathcal{A}_e^*$ for reasoning about hypothetical action occurrences discussed in Example 1.1. For the first scenario, the hypothetical sensing action occurrence of **mayread** is represented by (13) with $x = A$ and $y = B$. In other words, Bill thinks that Anne could read the letter ($read\langle A \rangle$) while he is away ($\neg at\_table_B$), but in fact Anne did not read. After this hypothetical reasoning, Bill thinks that Anne could know the status of United Agents or if Anne had not read the letter then she could not know. The update template and the result state for this hypothetical action occurrence[3] are described in Firgure 7 and Figure 8, respectively. Note that before Bill leaves the table to order a drink, he knows that Anne does not know about the status of United Agents ($\mathbf{B}_B(\neg \mathbf{B}_A doing\_well \wedge \neg \mathbf{B}_A \neg doing\_well)$ is true in (M_0, s_0)). However, when he comes back to the table with his drink, he can no longer derive the same conclusion about Anne's belief

[3]Strictly speaking, s_0 and s_1 in this figure and subsequent figures must be changed as at_table_x for $x \in \{A, B\}$ changes its value in each situation. We omit this change for brevity.

since $\mathbf{B}_B(\neg\mathbf{B}_A doing\_well \wedge \neg\mathbf{B}_A\neg doing\_well)$ is no longer true in the state resulting from his hypothetical reasoning (Figure 8).

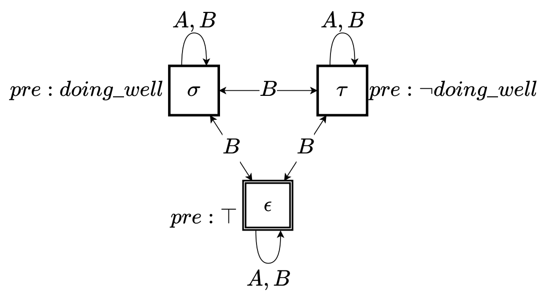

Figure 7: Update template for hypothetical action occurrence $read\langle A\rangle$ **when the action does not occur.**

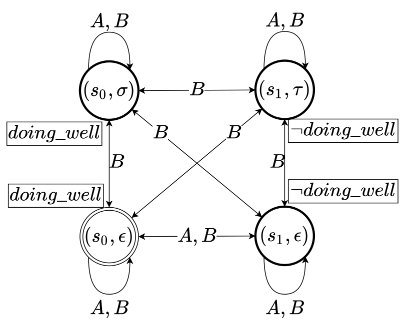

Figure 8: Result state after hypothetical action occurrence $read\langle A\rangle$ **and the action does not occur.**

The second scenario is more complex than the first one; this is a sequence of two action occurrences $read\langle A\rangle$ and $read\langle B\rangle$ where Anne is the full observe in $read\langle A\rangle$ and she is not sure if she is a partial observer in $read\langle B\rangle$ or not (vice versa for Bill). The update template and epistemic state after $read\langle A\rangle$ are showed in Figure 9 and Figure 10, respectively. The update template and epistemic state after $read\langle B\rangle$ are showed in Figure 12 and Figure 13, respectively.

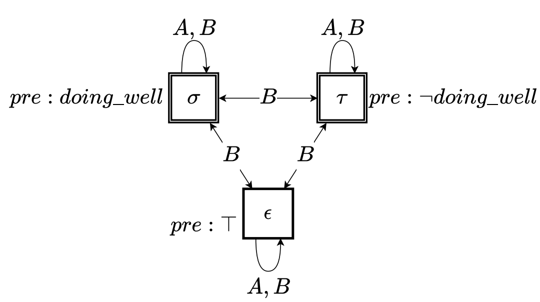

Figure 9: Update template for hypothetical action occurrence of $read\langle A\rangle$ **when the action actually happens.**

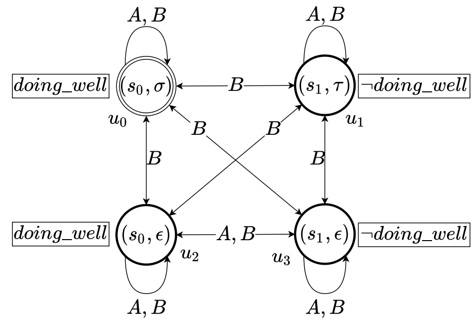

Figure 10: Result state after action occurrence $read\langle A\rangle$ **from Figure 9.**

Observe from Figure 10, Bill does not know if Anne knows about the status of United Agents or not; but still, Bill considers it is possible that Anne knows $doing\_well$, or knows $\neg doing\_well$. Comparing with the case where Bill has no idea about this action occurrence (Bill is an oblivious agent, the state represents that situation is detailed in Figure 8): as we can observe from Figure 10 and Figure 11, the accessibility relations of Bill are a bit different. There do not have any B's loops and links between u'_0 and u'_1, and every B's links from u'_0, u'_1 to u'_2, u'_3 are one-way only since Bill has no doubt about Anne read the letter, so Bill does not consider the exist of the two worlds u'_0 and u'_1. These differences highlight the fact that when Bill is totally oblivious about the action occurrence, his belief about everything must stays the same as before; but when Bill hypothesize that $read\langle A\rangle$ may happen (but not sure about that), he would change his belief according to his assumption.

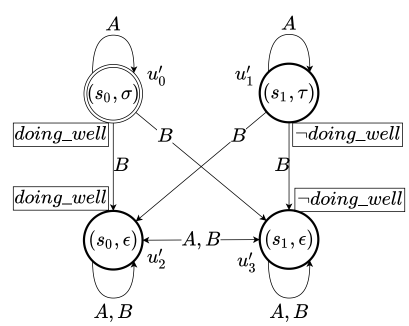

Figure 11: Result state after action occurrence $read\langle A\rangle$ **when** B **is an oblivious agent.**

After these two actions ($read\langle A\rangle$ and $read\langle B\rangle$), both Anne and Bill know the status of United Agents, but both also are not sure about others' beliefs as well. This result show that our proposed work is also cable of handle hypothetical reasoning in multi-agent domains for sequence of actions.

3.5 Properties of Edge-Conditioned Update Models for Non-Deterministic Observability and Hypothetical Action Occurrences

In this section, we will prove some important properties of the transition functions of $m\mathcal{A}^*_e$. For simplicity of the presentation, the

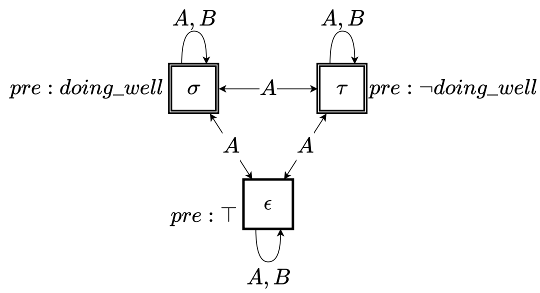

Figure 12: Update template for the hypothetical action occurrence of $read\langle B\rangle$ and the action occurs.

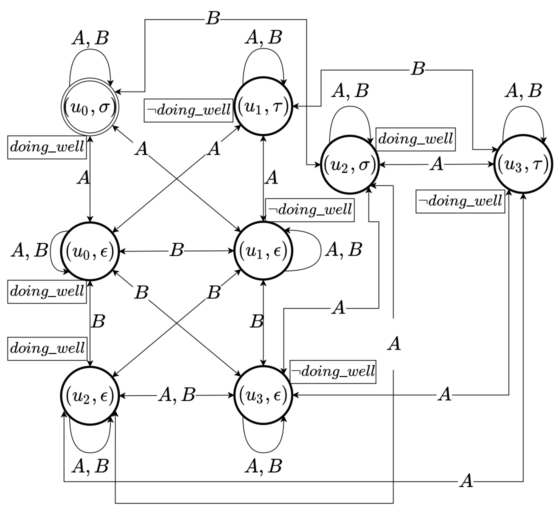

Figure 13: Result state after the hypothetical action occurrence of $read\langle B\rangle$ and the action actually happens.

properties are considered at the update template models because the transition function definition relies on the update template of the action occurrence.

Proposition 3.3. *Let (M, s) be a state and* a *be an ontic action instance that is executable in (M, s) and $\omega(\text{a})$ be given in Definition 3.1. Consider an agent $i \in \mathcal{AG}$ and a belief formula η. Assume that* $[x$ **may_observe** a **if** $\gamma_x]$ *and* $[\text{a}$ **causes** ℓ **if** $\varphi]$ *belong to D. If $(M, s) \models \gamma_x$, $(M, s) \models \mathbf{B}_x\varphi$, $(M, s) \models \mathbf{B}_x\eta$ and $(M', s') = (M, s) \otimes (\omega(\text{a}), \sigma)$ then:*

(1) Exist a world $w \in M'[S]$ such that $(M', w) \models \ell$ and $(s', w) \in M'[x]$;
(2) Exist a world $w' \in M'[S]$ such that $(M', w') \models \eta$ and $(s', w') \in M'[x]$; and
(3) $(w, w') \in M'[x]$ and $(w', w) \in M'[x]$ where w, w' are from item 1 and item 2.

Proof. We have that $s' = (s, \sigma)$. Assume that the fluent in ℓ is p, i.e., $\ell = p$ or $\ell = \neg p$. Let $\Psi^+(p, \text{a}) = \bigvee\{\varphi \mid [\text{a}$ **causes** p **if** $\varphi] \in D\}$ and $\Psi^-(p, \text{a}) = \bigvee\{\varphi \mid [\text{a}$ **causes** $\neg p$ **if** $\varphi] \in D\}$ and $\gamma = \Psi^+(p, \text{a}) \vee \Psi^-(p, \text{a})$. By Definition 3.1, $p \rightarrow \gamma \in sub(\sigma)$. Also, for every $u' \in M'[S]$ such that $(s', u') \in M'[x]$, it holds that $u' = (u, \sigma)$ or $u' = (u, \epsilon)$ for some $u \in M[S]$ and $(s, u) \in M[x]$.

(1) *Proof of the first item:* We have that for a world $u \in M[S]$ and $(s, u) \in M[x]$, then $(M, u) \models \varphi$ since $(M, s) \models \mathbf{B}_x\varphi$. By Definition 3.1 and Definition of $\otimes$, we have that there is $u' = (u, \sigma) \in M'[S]$ and $(M', u') \models \ell$. We also have $(s', u') \in M[x]$ since $(M, s) \models \gamma_x$. This implies the first item of the proposition.
(2) *Proof of the second item:* By the construction of M', we have the following observations:
   - For every $u \in M[S]$ iff $(u, \epsilon) \in M'[S]$;
   - For every $z \in \mathcal{AG}$, $(u, v) \in M[z]$ iff $((u, \epsilon), (v, \epsilon)) \in M'[z]$;
   - For every $u \in M[S]$ and $p \in \mathcal{F}$, $M'[\pi]((u, \epsilon)) \models p$ iff $(M', (u, \epsilon)) \models$ because $sub(\epsilon) = \emptyset$.

   These observations allow us to conclude for every formula η, $(M, u) \models \eta$ iff $(M', (u, \epsilon)) \models \eta$. Consider $u \in M[S]$ such that $(s, u) \in M[x]$, by Definition 3.1 and Definition of $\otimes$, we have that $u' = (u, \epsilon) \in M'[S]$ and $(s', u') \in M[x]$. Also, because $(M, s) \models \mathbf{B}_x\eta$, then $(M, u) \models \eta$. This implies $(M', u') \models \eta$. That concludes the second item of the proposition.
(3) *Proof of the third item:* Consider $w, w' \in M'[S]$ such that w has the same properties as in item 1 and w' as the same properties as in item 2. This mean $u, v \in M[S]$ such that $w = (u, \sigma)$, $w' = (v, \epsilon)$ and $(s, u), (s, v) \in M[x]$. Since $(M, s) \models \gamma_x$, by Definition 3.1 and Definition of $\otimes$, we have that for every pair $(w, w'), (w', w) \in M'[x]$. This concludes the third item of the proposition. □

Proposition 3.3 shows that if an agent is not certain about an ontic action occurrence, then in the next state, they would consider some world in the next state where the action could happened (item 1). But they also keep their (old) belief in some other world in the next state (item 2). Furthermore, they cannot distinguish those worlds with each others (item 3). Similar propositions can be established for the sensing and announcement actions and are omitted for brevity. These propositions highlight that $m\mathcal{A}^*_e$ can properly deal with non-deterministic observability.

Proposition 3.4. *Let (M, s) be a state and* a *be an ontic action instance that is executable in (M, s) and $\omega(\text{a})$ be given in Definition 3.1. Consider a hypothetical action occurrence of* a *that does not actually happen. Consider an agent $i \in \mathcal{AG}$ and a belief formula η. Assume that* $[x$ **observes** a **if** $\delta_x]$ *and* $[\text{a}$ **causes** ℓ **if** $\varphi]$ *belong to D. If $(M, s) \models \delta_x$, $(M, s) \models \mathbf{B}_x\varphi$, $(M, s) \models \mathbf{B}_x\eta$ and $(M', s') = (M, s) \otimes (\omega(\text{a}), \sigma)$ then $(M', s') \models \mathbf{B}_x\eta$.*

Proof. The observations from the *proof of the second item* of Proposition 3.3 allow us to conclude for every formula η, $(M, u) \models \eta$ iff $(M', (u, \epsilon)) \models \eta$. Since this is a hypothetical action occurrence that does not actually occur, we have that $s' = (s, \epsilon)$. Therefor for every formula η such that $(M, s) \models \mathbf{B}_x\eta$ and $(M, s) \models \delta_x$, we have that $(M', s') \models \mathbf{B}_x\eta$. □

Proposition 3.4 illustrates the fact that full observers of a hypothetical action occurrence that does not happen would not change their belief at all. For non-deterministic observers–agents who believe that this hypothetical action occurrence could happen–similar properties as in Proposition 3.3 can be established and are omitted for brevity.

It is worth to mention that the use of edge-conditioned update models allows $m\mathcal{A}^*_e$ to overcome some current issues of $m\mathcal{A}^*$ (see, [9] for a good discussion). Observe that if we remove statements of the form (8) and (9) from an $m\mathcal{A}^*_e$ action domain then it is a 'normal' $m\mathcal{A}^*$ whose edge-conditioned models are studied in [9]; therefore, the transition function of $m\mathcal{A}^*_e$ also satisfies the properties about second-order belief discussed in [9].

4 DISCUSSION AND RELATED WORK

Our work is inspired by the line of research to formalize actions in multi-agent systems that emphasizes the use of logic, such as Dynamic Epistemic Logic (DEL) [12, 15, 16]. The language that DEL uses to represent and reason about non-deterministic actions over a set of agent $\mathcal{AG}$ and a set of atomic propositions $\mathcal{F}$ is $\mathcal{L}_{!}(\mathcal{AG},\mathcal{F})$. It is the union of the *formulas* $\mathcal{L}^{stat}_{!}(\mathcal{AG},\mathcal{F})$ and the *actions* $\mathcal{L}^{act}_{!}(\mathcal{AG},\mathcal{F})$, defined by:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi \mid C_B\varphi \mid [\alpha]\psi$$
$$\alpha ::= ?\varphi \mid L_B\beta \mid (\alpha \,!\, \alpha) \mid (\alpha \,¡\, \alpha) \mid (\alpha \,;\, \beta') \mid (\alpha \cup \alpha')$$

where $p \in \mathcal{F}$, $a \in \mathcal{AG}$, $B \subseteq \mathcal{AG}$ and $\psi \in \mathcal{L}^{stat}_{!}(gr(\alpha),\mathcal{F})$, $\beta \in \mathcal{L}^{act}_{!}(B,\mathcal{F})$ and $\beta' \in \mathcal{L}^{act}_{!}(gr(\alpha),\mathcal{F})$. The *group* $gr(\alpha)$ of an action α is defined as: $gr(?\varphi) = \emptyset$, $gr(L_B\alpha) = B$, $gr(\alpha \,!\, \alpha') = gr(\alpha)$, $gr(\alpha \,¡\, \alpha') = gr(\alpha')$, $gr(\alpha \,;\, \alpha') = gr(\alpha')$ and $gr(\alpha \cup \alpha') = gr(\alpha) \cap gr(\alpha')$. Action $?\varphi$ is a *test*. Operator L_B is called *learning* and the construct $L_B?\varphi$ is pronounced as 'group B learn that φ'. Action $(\alpha \,!\, \alpha')$ and $(\alpha \,¡\, \alpha')$ are *(left/right) local choices*. Action $(\alpha \,;\, \alpha')$ is *sequential execution* - 'first α, then α'' and action $(\alpha \cup \alpha')$ is *non-deterministic choice* between α and α'. Instead of $(\alpha \,!\, \alpha')$ (or $(\alpha \,¡\, \alpha')$), it is often written as $(!\alpha \cup \alpha')$ (or $(\alpha \cup \,!\alpha')$. The meaning of *local choice* is that in $L_B(\alpha \,!\, \alpha')$, everybody in B but not in learning operator occurring in α, α', in unaware of the choice between α and α'. DEL with assignments, which allows for the representation of ontic actions, has also been investigated (e.g., [16]). The semantics of DEL formulae is defined over the **S5** logic over $\mathcal{AG}$ and $\mathcal{F}$. We omit the full definition for brevity.

One significant difference between our work and DEL lies in that we do not assume **S5** while DEL assumes **S5**. This means that $m\mathcal{A}^*_e$ could accommodate agents with false beliefs, and perhaps, lying agents—if update templates for this type of action can be developed—in the context of epistemic planning. On the other hand, new logic needs to be developed to consider false announcements (e.g., [13]).

In recent work, [5] proposed a new semantic for the action language $m\mathcal{A}^*$ that maintains two tiers of information: knowledge and belief. However, this work does not consider situations where the observability of agents can be non-deterministic. Therefore it cannot handle such actions as in Example 1.1.

Another work that also uses event update models for reasoning in multi-agent domains is [10] within the language DER. In this language, the observability of agents is encoded by an observations set O and no distinction between ontic, sensing, and announcement actions are made. Comparing with the update models used in [10], we can see that the update models used in the present paper have a fixed number of events, given the type of the action: two events for ontic actions and three events for sensing/announcement actions. On the other hand, the number of events in DER can vary given the number of statements specifying its effects and observations. We believe that this feature might bring some advantages if update models are used for planning, where efficient construction of update models is critical (in $m\mathcal{A}^*_e$, the model needs to be constructed only once).

5 CONCLUSIONS AND FUTURE WORK

In this paper, we define $m\mathcal{A}^*_e$, an extension of the language $m\mathcal{A}^*$, to allow reasoning with non-deterministic observability of action occurrences of agents in multi-agent domains. The new extension also allows an external observer to reason about hypothetical action occurrences. We rely on the notion of edge-conditioned models in defining the semantics of $m\mathcal{A}^*_e$ and prove properties of the new transition function that highlights that $m\mathcal{A}^*_e$ can properly deal with non-deterministic observability. We illustrate the usefulness of $m\mathcal{A}^*_e$ on an example that is frequently used in DEL literature but cannot be formalized in $m\mathcal{A}^*$ and some other action languages for multi-agent domains. We also discuss the difference between this extension with DEL or DER, a recent language for RAC in multi-agent domains. As a future work, we plan to investigate the situations when external observers are uncertain about action occurrences.

Acknowledgments. The authors have been partially supported by NSF grants 2151254, 1914635 and 1757207. Tran Cao Son was also partially supported by NSF grant 1812628.

REFERENCES

[1] A. Baltag and L. Moss. 2004. Logics for epistemic programs. *Synthese* (2004).
[2] A. Baltag, L. Moss, and S. Solecki. 1998. The logic of public announcements, common knowledge, and private suspicions. In *7th TARK*. 43–56.
[3] Chitta Baral, Gregory Gelfond, Enrico Pontelli, and Tran Cao Son. 2022. An action language for multi-agent domains. *Artificial Intelligence* 302 (2022), 103601.
[4] Thomas Bolander. 2018. *Seeing Is Believing: Formalising False-Belief Tasks in Dynamic Epistemic Logic*. Springer International Publishing, Cham, 207–236.
[5] David Buckingham, Daniel Kasenberg, and Matthias Scheutz. 2020. Simultaneous Representation of Knowledge and Belief for Epistemic Planning with Belief Revision. (9 2020), 172–181.
[6] Francesco Fabiano, Alessandro Burigana, Agostino Dovier, and Enrico Pontelli. 2020. EFP 2.0: A Multi-Agent Epistemic Solver with Multiple E-State Representations. *Proceedings of the International Conference on Automated Planning and Scheduling* 30, 1 (Jun. 2020), 101–109. https://ojs.aaai.org/index.php/ICAPS/article/view/6650
[7] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. 1995. *Reasoning about Knowledge*. MIT press.
[8] Tiep Le, Francesco Fabiano, Tran Cao Son, and Enrico Pontelli. 2018. EFP and PG-EFP: Epistemic Forward Search Planners in Multi-Agent Domains. In *International Conference on Automated Planning and Scheduling (ICAPS)*. AAAI Press.
[9] Loc Pham, Yusuf Izmirlioglu, Tran Cao Son, and Enrico Pontelli. 2022. A New Semantics for Action Language mA*. In *Proceedings of the 24th International Conference on Principles and Practice of Multi-Agent Systems (PRIMA 22)*. Springer, 553–562. https://doi.org/10.1007/978-3-031-21203-1_33
[10] David Rajaratnam and Michael Thielscher. 2021. Representing and Reasoning with Event Models for Epistemic Planning. In *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning*. 519–528. https://doi.org/10.24963/kr.2021/49
[11] Johan van Benthem, Jan van Eijck, and Barteld P. Kooi. 2006. Logics of communication and change. *Inf. Comput.* 204, 11 (2006), 1620–1662.
[12] H. van Ditmarsch. 2005. Prolegomena to dynamic logic for belief revision. *Synthese (Knowledge, Rationality & Action)* 147 (2005), 229–275.
[13] Hans van Ditmarsch, Petra Hendriks, and Rineke Verbrugge. 2020. Editors' Review and Introduction: Lying in Logic, Language, and Cognition. *Top. Cogn. Sci.* 12, 2 (2020), 466–484. https://doi.org/10.1111/tops.12492
[14] H. van Ditmarsch, W. van der Hoek, and B. Kooi. 2007. *Dynamic Epistemic Logic*. Springer.
[15] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. 2007. *Dynamic Epistemic Logic* (1st ed.). Springer Publishing Company, Incorporated.
[16] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld P. Kooi. 2005. Dynamic epistemic logic with assignment. In *4th International Joint Conference on Autonomous Agents and Multiagent Systems*. ACM, 141–148.