Update Models for Lying and Misleading Announcements

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

Lying and misleading announcements are a significant part of multi-agent communications. Providing a formal account of such announcements is important for representing and reasoning about effects of actions and epistemic planning in multi-agent domains. This paper defines update models for lying and misleading announcements in the high-level action language m$\mathcal{A}^*$. Applying our update models on a pointed Kripke structure—representing the current state of knowledge and beliefs of agents—results in a pointed Kripke structure in which lying (or misleading) announcements achieve the intended effects of deceiving other agents who are uncertain about the announced formula to believe in the wrong information.

KEYWORDS

Lying Announcement, Misleading Announcement, Action Language, Update Model

ACM Reference Format:

Loc Pham, Tran Cao Son, and Enrico Pontelli. 2022. Update Models for Lying and Misleading Announcements. In *The 37th ACM/SIGAPP Symposium on Applied Computing (SAC '22), April 25–29, 2022, Virtual Event,* . ACM, New York, NY, USA, Article 4, 3 pages. https://doi.org/10.1145/3477314.3507197

1 INTRODUCTION AND MOTIVATION

Reasoning about actions and change (RAC) forms the foundation for temporal reasoning, planning, explanation, and diagnosis in dynamic environments. In recent years, RAC in multi-agent domains has attracted a lot of attention. Significant challenges for approaches to RAC in multi-agent domains lie in that they need to deal with actions, that do not appear in single-agent domains such as *announcement*, and that agents might not be aware of actions occurrences, and therefore, have false beliefs. In addition, agents might intentionally act so that other agents believe in some false information.

One of the most well-known notions for RAC in multi-agent domains is the notion of *action model*, introduced in [1, 2] and later extended to *update model* [6, 11]. Intuitively, an update model represents different views of agents about an action occurrence, which ultimately affects the beliefs of agents about the state of the world and the state of beliefs of other agents after the action occurrence. Formally, given a pointed Kripke structure (M, s) representing the current state of the world and of the knowledge/beliefs of agents, the result of the execution of an action a in (M, s) can be computed by a cross product of (M, s) and the update model $U(\mathrm{a})$— corresponding to the occurrence of a in (M, s).

High-level action languages such as m$\mathcal{A}$ and its subsequent developments m$\mathcal{A}^*$ have proposed the context that help to *automatically* construct $U(\mathrm{a})$ given (M, s) and the description of a's effects and preconditions [3]. These languages have proved to be useful in the development of some epistemic planning systems such as those proposed in [7, 9]. However, the current m$\mathcal{A}^*$ only deals with ontic actions, sensing actions, and public and private *truthful* announcements. It does not deal with false or misleading announcements. On the other hand, the ability to reason with untruthful announcements is needed for multi-agent planners to generate plans which include deception of opposing agents, e.g., using announcements to lead them to believe false facts. Such type of plans is desirable in strategy games, in military operations, or even in economy. They also represent a crucial component for building negotiation agents that can reason with dishonest information, as shown, for example, in [10].

In this paper, we explore how reasoning about lying/misleading announcements could be done by autonomous agents. Being able to do so is an important step in both detecting and countering dishonest behaviors of autonomous agents. It is also well-known that, even when agents are fully cooperative, a white lie might be needed (e.g., [13]). Our goal is to propose a way for dealing with false and misleading announcements in m$\mathcal{A}^*$, thereby providing the foundation for multi-agent planners which can include lying and misleading announcements in their plans.

2 PRELIMINARY: ACTION LANGUAGE M$\mathcal{A}^*$

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 $\mathcal{F}$ encoding 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 ψ and a non-empty set $\alpha \subseteq \mathcal{AG}$, $\mathbf{B}_\alpha\psi$ denote $\bigwedge_{i \in \alpha} \mathbf{B}_i\psi$. $\mathcal{L}_{\mathcal{AG}}$ denotes the set of belief formulae over $\langle \mathcal{AG}, \mathcal{F} \rangle$.

Satisfaction of belief formulae is defined over *pointed Kripke structures* [8]. 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 $u \in S$ 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: **(i)** $(M, s) \models p$ if p is a fluent and $M[\pi](s)(p)$

SAC '22, April 25–29, 2022, Virtual Event,

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.
https://doi.org/10.1145/3477314.3507197

is true; **(ii)** $(M, s) \models \neg\varphi$ if $(M, s) \not\models \varphi$; **(iii)** $(M, s) \models \varphi_1 \wedge \varphi_2$ if $(M, s) \models \varphi_1$ and $(M, s) \models \varphi_2$; **(iv)** $(M, s) \models \varphi_1 \vee \varphi_2$ if $(M, s) \models \varphi_1$ or $(M, s) \models \varphi_2$; **(v)** $(M, s) \models \mathbf{B}_i\varphi$ if $\forall t.[(s, t) \in \mathcal{B}_i \Rightarrow (M, t) \models \varphi]$.

In order to handle lying and misleading announcements, we utilize the edge-conditioned update models as proposed in [5]. We present a simplified version of edge-conditioned update models, as needed for announcement actions. An edge-conditioned update model Σ is a tuple $\langle \Sigma, R_1, \ldots, R_n, pre \rangle$ 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, $pre : \Sigma \rightarrow \mathcal{L}_{\mathcal{AG}}$ is a function mapping each event $e \in \Sigma$ to a formula in $\mathcal{L}_{\mathcal{AG}}$. Elements of R_i are in the form (e_1, γ, e_2) where γ is a belief formula.

Given a Kripke structure M and an edge-conditioned update model $\Sigma = \langle \Sigma, \{R_i\}_{i \in \mathcal{AG}}, pre \rangle$, the *update* induced by Σ, $M'{=}M{\otimes}\Sigma$, is 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 $(M, s) \models \gamma$; and **(iii)** $\forall (s, \tau) \in M'[S], f \in \mathcal{F}$, $[M'[\pi]((s, \tau)){\models}f$ iff $M[\pi](s){\models}f]$.

An *update template* is a pair (Σ, Γ), where Σ is an update model with the set of events Σ and $\Gamma \subseteq \Sigma$ (Γ is the set of *designated events* (or *true events*). The update of 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 $(M, s) \otimes (\Sigma, \Gamma) = \{(M \otimes \Sigma, (s, \tau)) \mid \tau \in \Gamma, (M, s) \models pre(\tau)\}$.

m$\mathcal{A}^*$ is an action theory over $\langle \mathcal{AG}, \mathcal{F} \rangle$ consists of a set of action instances $\mathcal{AI}$, of the form $a\langle\alpha\rangle$, indicating that the set of agents α jointly performs action a. In this paper, we focus on one class of actions, *announcement actions*—a more complete description of m$\mathcal{A}^*$ can be found in [4]. We assume that an announcement can always be executed.

$$\text{a } \mathbf{announces}\ \varphi \quad (1) \qquad z\ \mathbf{observes}\ \text{a}\ \mathbf{if}\ \delta_z \quad (2) \qquad z\ \mathbf{aware\_of}\ \text{a}\ \mathbf{if}\ \theta_z \quad (3)$$

where φ, δ_z and θ_z are fluent formulae, a $\in \mathcal{AI}$, and $z \in \mathcal{AG}$. (1) encodes an *announcement* action, whose executors announce that φ is true. (2) indicates that agent z is a full observer of a if δ_z holds. (3) states that agent z is a partial observer of a if θ_z holds. In other words, an agent i can be classified as *full observer*, or *partial observer*, or *oblivious* (neither δ_i nor θ_i holds) with an action's occurrence. We will refer to φ in (1) as the *announced formula* of a. We also assume that the announcer of an announcement is always a full observer of its occurrence.

Given an action theory D, a state (M, s), and an occurrence of a, the *frame of reference* for the execution of a in (M, s) is a tuple $(F_D(\text{a}, M, s), P_D(\text{a}, M, s), O_D(\text{a}, M, s))$ where: $F_D(\text{a}, M, s) = \{x \in \mathcal{AG} \mid [x\,\mathbf{observes}\,\text{a}\,\mathbf{if}\,\delta_x] \in D$ such that $(M, s) \models \delta_x\}$; $P_D(\text{a}, M, s) = \{x \in \mathcal{AG} \mid [x\ \mathbf{aware\_of}\ \text{a}\ \mathbf{if}\ \theta_x] \in D$ such that $(M, s) \models \theta_x\}$; and $O_D(\text{a}, M, s) = \mathcal{AG} \setminus (F_D(\text{a}, M, s) \cup P_D(\text{a}, M, s))$.

Intuitively, $F_D(\text{a}, M, s)$ (resp. $P_D(\text{a}, M, s)$ and $O_D(\text{a}, M, s)$) are the agents that are fully observant (resp. partially observant and oblivious/other) of the execution of a in the state (M, s)[1]. m$\mathcal{A}^*$ assumes that the sets $F_D(\text{a}, M, s)$, $P_D(\text{a}, M, s)$, and $O_D(\text{a}, M, s)$ are pairwise disjoint. In the following, whenever we say that an agent i is a fully or partially observer, we mean that there exists a statement of the form (2) or (3) such that $(M, s) \models \delta_i$ or $(M, s) \models \theta_i$, respectively.

[1]We will use F (resp. P and O) instead of $F_D(\text{a}, M, s)$ (resp. $P_D(\text{a}, M, s)$ and $O_D(\text{a}, M, s)$ later on for brevity.

3 LYING/MISLEADING ANNOUNCEMENTS

Our focus in this paper is to study announcement actions that have not been considered in m$\mathcal{A}^*$. Specifically, our aim is to define update models for untrustworthy announcements. Since an announcement is not a forcible action, it is up to the agents who are fully aware of the action occurrence to believe in the announced formula; we make the following assumptions:

> (**A1**) if an agent is certain about the truth of a formula, even if it is incorrect in the actual world, or if she realizes that the announcement is untruthful (i.e. she knows that the announcers make a false statement) then she will not change her belief about the formula, regardless of what the announcers say;
> (**A2**) if an agent is uncertain about the truth of a formula and cannot reason that the announcers are untruthful then she will believe what the announcers say.

The assumption (**A1**) indicates that agents are "opinionated" about their own beliefs, while (**A2**) suggests that agents are eager to remove uncertainty in their beliefs. There can be finer or different classifications of agents' attitudes with respect to an announcement, e.g., agents' attitudes can depend on who makes the announcement and what type of information is announced. For example, in [12], agents can be credulous, skeptical, or revising reasoners, and they will update their beliefs based on their attitude. This is an interesting subject but outside the scope of this paper.

3.1 Update Model for Lying Announcement

Let us consider an occurrence of a lying announcement about φ, a $= a\langle\alpha\rangle$, in a state (M, s). Formally, we say that the occurrence of a in (M, s) is a *lying announcement* if $(M, s) \models \mathbf{B}_\alpha\neg\varphi$. Observe that we do not require that φ is true in the actual state. As such, it is possible that a lying announcement might result in some agents believe in a formula that is true in the actual world, i.e., a lie indirectly becomes a true announcement for some agents. From our assumptions (**A1**) and (**A2**), if i is an agent in $\mathcal{AG}$, we have the following situations:

- i is a *fully observer* of a, we have the following sub-cases:
  - $(M, s) \models \mathbf{B}_i\mathbf{B}_\alpha\neg\varphi$ or $(M, s) \models \mathbf{B}_i\neg\varphi$ or $(M, s) \models \mathbf{B}_i\varphi$: by (**A1**), the belief of i should not be changed;
  - $(M, s) \models \neg(\mathbf{B}_i\varphi \vee \mathbf{B}_i\neg\varphi) \wedge \neg\mathbf{B}_i\mathbf{B}_\alpha\neg\varphi$: by (**A2**), the belief of i about φ should be changed and $\mathbf{B}_i\varphi$ is true after the lying announcement.
- i is a *partially observer* of action a. As i is unaware of what is announced about the formula her belief about φ does not change. However, she would assume that people who are fully observant know the truth value of the formula after the action occurrence.
- i is not aware of the execution of action a: nothing changes for i.

Based on the above discussion, the update model for lying announcement can be defined as below.

Definition 3.1 (Update Model for Lying Announcement). Let a $= a\langle\alpha\rangle$ be an announcement of formula φ and (M, s) be a state such that $(M, s) \models \mathbf{B}_\alpha\neg\varphi$. The update model for a in (M, s), $\omega(\text{a}, (M, s))$, is defined by $\langle \Sigma, \{R_i\}_{i \in \mathcal{AG}}, pre \rangle$ as follow:

- $\Sigma = \{\sigma, \zeta, \tau, \chi, \mu, \epsilon\}$;

- $R_i = \{\ (\sigma, i \in F : i \in \alpha \vee \mathbf{B}_i\mathbf{B}_\alpha\neg\varphi, \sigma), (\sigma, i \in F : \mathbf{B}_i\neg\varphi \vee \mathbf{B}_i\mathbf{B}_\alpha\neg\varphi, \zeta), (\sigma, i \in F : \neg\mathbf{B}_i\neg\varphi, \tau), (\sigma, i \in P, \chi), (\sigma, i \in P, \mu), (\sigma, i \in O, \epsilon), (\zeta, i \in F : \mathbf{B}_i\neg\varphi \vee \mathbf{B}_i\mathbf{B}_\alpha\neg\varphi, \zeta), (\zeta, i \in F : \neg\mathbf{B}_i\neg\varphi, \tau), (\zeta, i \in P, \chi), (\zeta, i \in P, \mu), (\zeta, i \in O, \epsilon), (\tau, i \in F : \mathbf{B}_i\neg\varphi \vee \mathbf{B}_i\mathbf{B}_\alpha\neg\varphi, \zeta), (\tau, i \in F, \tau), (\tau, i \in P, \chi), (\tau, i \in P, \mu), (\tau, i \in O, \epsilon), (\chi, i \in F \cup P, \chi), (\chi, i \in P \vee i \in F : \mathbf{B}_i\varphi, \mu), (\chi, i \in O, \epsilon), (\mu, i \in F \cup P, \mu), (\mu, i \in P \vee i \in F : \mathbf{B}_i\neg\varphi, \chi), (\mu, i \in O, \epsilon), (\epsilon, i \in \mathcal{AG}, \epsilon)\ \}$;
- The preconditions are: $pre(\sigma) = \mathbf{B}_\alpha\neg\varphi$; $pre(\zeta) = \neg\varphi$; $pre(\tau) = \varphi$; $pre(\chi) = \neg\varphi$; $pre(\mu) = \varphi$; and $pre(\epsilon) = \top$.

Observe that there are six events in this update model. The designated event σ represents the real event created by α, the executors of the announcement. ζ encodes the event recorded by agents who have the correct belief about the announced formula (φ is false); they will maintain their belief. τ encodes the event under the perspective of those agents who have false belief or do not know whether φ. χ and μ are events under the perspective of partial observers, who are aware of the occurrence of the announcement but do not receive its content. Finally, ϵ is the event under the perspective of oblivious agents, for whom nothing happens.

3.2 Update Model for Misleading Announcements

Consider an occurrence of a misleading announcement of φ, a $= a\langle\alpha\rangle$, in a state (M, s). Formally, we say that the occurrence of a in (M, s) is a misleading announcement if for every $i \in \alpha$, $(M, s) \models \neg(\mathbf{B}_i\varphi \vee \mathbf{B}_i\neg\varphi)$. Given the assumptions (**A1**) and (**A2**), the update of beliefs of the agents will be as follows:

- i is a *fully observer* of a, we have the following sub-cases:
  - $(M, s) \models \mathbf{B}_i\neg(\mathbf{B}_\alpha\neg\varphi \vee \mathbf{B}_\alpha\varphi)$ or $(M, s) \models \mathbf{B}_i\neg\varphi$ or $(M, s) \models \mathbf{B}_i\varphi$: by (**A1**), the belief of i should not be changed;
  - $(M, s) \models \neg(\mathbf{B}_i\varphi \vee \mathbf{B}_i\neg\varphi)$ then:
    - $i \in \alpha$: she is one of the agents who executes this action and knows this is a misleading announcement, her belief should not be changed;
    - $i \notin \alpha$: by (**A2**), the belief of i about φ should be changed, i.e., $\mathbf{B}_i\varphi$ should be true for i.
- i is a *partially observer* of a. Since i is unaware of what is announced about the formula, her belief about φ does not change. But how about her belief about others? Again, we have two sub-cases here:
  - $(M, s) \models \mathbf{B}_i\neg(\mathbf{B}_\alpha\neg\varphi \vee \mathbf{B}_\alpha\varphi)$: i realizes that this is a misleading announcement since α is talking about something they do not know, therefor the belief of i would not be changed at all;
  - $(M, s) \models \mathbf{B}_i\mathbf{B}_\alpha\neg\varphi \vee \mathbf{B}_i\mathbf{B}_\alpha\varphi$: i would recognize that fully observers have the knowledge of the formula after the occurrence.
- i is not aware of the execution of action a. Since i is oblivious of the action occurrence, nothing changes for i.

Definition 3.2 (Update Model for Misleading Announcement). Let a $= a\langle\alpha\rangle$ be an announcement with the announced formula φ and (M, s) be a state such that $(M, s) \models \neg(\mathbf{B}_i\varphi \vee \mathbf{B}_i\neg\varphi)$. The update model for a in (M, s), $\omega(\text{a}, (M, s))$, is defined by $\langle\Sigma, \{R_i\}_{i\in\mathcal{AG}}, pre\rangle$ as follow:

- $\Sigma = \{\sigma, \zeta, \tau, \chi, \mu, \epsilon\}$;
- $R_i = \{\ (\sigma, i \in F \cup P : i \in \alpha \vee \mathbf{B}_i\neg(\mathbf{B}_\alpha\neg\varphi \vee \mathbf{B}_\alpha\varphi), \sigma), (\sigma, i \in F : \mathbf{B}_i\neg\varphi \vee \mathbf{B}_i\neg(\mathbf{B}_\alpha\neg\varphi \vee \mathbf{B}_\alpha\varphi), \zeta), (\sigma, i \in F : \neg\mathbf{B}_i\neg\varphi, \tau), (\sigma, i \in P, \chi), (\sigma, i \in P, \mu), (\sigma, i \in O, \epsilon), (\zeta, i \in F : \mathbf{B}_i\neg\varphi \vee \mathbf{B}_i\neg(\mathbf{B}_\alpha\neg\varphi \vee \mathbf{B}_\alpha\varphi), \zeta), (\zeta, i \in F : \neg\mathbf{B}_i\neg\varphi, \tau), (\zeta, i \in P, \chi), (\zeta, i \in P, \mu), (\zeta, i \in O, \epsilon), (\tau, i \in F : i \notin \alpha \wedge (\mathbf{B}_i\neg\varphi \vee \mathbf{B}_i\neg(\mathbf{B}_\alpha\neg\varphi \vee \mathbf{B}_\alpha\varphi)), \zeta), (\tau, i \in F, \tau), (\tau, i \in P, \chi), (\tau, i \in P, \mu), (\tau, i \in O, \epsilon), (\chi, i \in F \cup P, \chi), (\chi, i \in P \vee i \in F : \mathbf{B}_i\varphi, \mu), (\chi, i \in O, \epsilon), (\mu, i \in F \cup P, \mu), (\mu, i \in P \vee i \in F : \mathbf{B}_i\neg\varphi, \chi), (\mu, i \in O, \epsilon), (\epsilon, i \in \mathcal{AG}, \epsilon)\ \}$;
- The preconditions are: $pre(\sigma) = \neg(\mathbf{B}_\alpha\neg\varphi \vee \mathbf{B}_\alpha\varphi)$; $pre(\zeta) = \neg\varphi$; $pre(\tau) = \varphi$; $pre(\chi) = \neg\varphi$; $pre(\mu) = \varphi$; and $pre(\epsilon) = \top$.

We observe that the update model of a misleading announcement has the same set of events as that of a lying announcement. The key difference are the precondition of the designated event σ; the addition condition for agents in P in σ; the requirement for an agent i who know about α's belief is $\mathbf{B}_i\neg(\mathbf{B}_\alpha\neg\varphi \vee \mathbf{B}_\alpha\varphi)$ instead of $\mathbf{B}_i\mathbf{B}_\alpha\neg\varphi$ as in a lying announcement; and the addition $i \notin \alpha$ for relation from τ to ζ which make sure anyone who believe in the misleading announcement would also believe α knows φ too.

4 CONCLUSION

In this paper, we presented a formal account for representing and reasoning about public and private untruthful announcements. The presented work could be viewed as an continuation of the language m$\mathcal{A}^*$. We distinguished two relevant forms of untruthful announcements, provided the modeling of such announcements using update models with edge-conditions. The proposed investigation represents the first step towards modeling and reasoning in presence of untruthful agents.

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. 2012. An Action Language for Reasoning about Beliefs in Multi-Agent Domains. In *Proceedings of NMR*.
[4] Chitta Baral, Gregory Gelfond, Enrico Pontelli, and Tran Cao Son. 2020. An Action Language for Multi-Agent Domains: Foundations. *arXiv.org* (2020), https://arxiv.org/abs/1511.01960v3.
[5] Thomas Bolander. 2018. *Seeing Is Believing: Formalising False-Belief Tasks in Dynamic Epistemic Logic*. Springer International Publishing, Cham, 207–236. https://doi.org/10.1007/978-3-319-62864-6_8
[6] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. 2007. *Dynamic Epistemic Logic* (1st ed.). Springer Publishing Company, Incorporated.
[7] 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
[8] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. 1995. *Reasoning about Knowledge*. MIT press.
[9] 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.
[10] Tran Cao Son, Enrico Pontelli, Ngoc-Hieu Nguyen, and Chiaki Sakama. 2014. Formalizing Negotiations Using Logic Programming. *ACM Trans. Comput. Log.* 15, 2 (2014), 12.
[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] Hans van Ditmarsch. 2014. Dynamics of lying. *Synth.* 191, 5 (2014), 745–777. https://doi.org/10.1007/s11229-013-0275-3
[13] Gilad Zlotkin and Jeffrey S. Rosenschein. 1991. Incomplete information and deception in multi-agent negotiation. In *IJCAI'91: Proceedings of the 12th international joint conference on Artificial intelligence*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 225–231.