

Tractable reasoning about Agent Programming in Dynamic Preference Logic

Marlo Souza
Institute of Mathematics and Statistics
UFBA
Salvador, Brazil
Email: marlo@dcc.ufba.br

Álvaro Moreira
Institute of Informatics
UFRGS
Porto Alegre, Brazil
Email: afmoreira@inf.ufrgs.br

Renata Vieira
Faculty of Informatics
PUCRS
Porto Alegre, Brazil
Email: renata.vieira@pucrs.br

Abstract—While several BDI logics have been proposed in the area of Agent Programming, it is not clear how these logics are connected to the agent programs they are supposed to specify. More yet, the reasoning problems in these logics, being based on modal logic, are not tractable in general, limiting their usage to tackle real-world problems. In this work, we use of Dynamic Preference Logic to provide a semantic foundation to BDI agent programming languages and investigate tractable expressive fragments of this logic to reason about agent programs. With that, we aim to provide a way of implementing semantically grounded agent programming languages with tractable reasoning cycles.

Index Terms—Dynamic Epistemic Logic; Agent Programming; Formal Semantics; BDI Logics;

I. INTRODUCTION

In the study of rational action and agency, several different logics and formal theories for practical reasoning have been proposed. Particularly, the Belief, Desire Intention framework [1] has become a popular approach to practical reasoning in the areas of Artificial Intelligence and Autonomous Agents, giving rise to the construction of several programming languages and computer systems.

Having a formal definition of its semantics is essential for proving properties about a programming language and also for providing a formal framework to specify and to verify properties about a system's behaviour. For agent programming languages, a formal semantics also clarifies the notion of agency it carries.

Recently, it has been proposed that Dynamic Preference Logic (DPL) can be used to reason about BDI Agent Programming with declarative mental attitudes [2], providing a computable two-way translation between specifications in the logic and agent programs. Reasoning in DPL is, however, not tractable in general. As such, the use of DPL for reasoning about agent programming - while theoretically relevant for the analysis of a programming language semantics - is of very limited practical use.

In this work, we investigate expressive fragments of the language of DPL that yield tractable reasoning problems. The reasoning problems discussed in this paper are concerned with knowing whether an agent knows (believes, desires or intends) a certain propositional property φ and how to compute the resulting mental state of an agent after performing a

belief/desire revision or contraction. With this, we aim to provide a tractable fragment that may be used to implement an actual agent programming language with declarative mental attitudes having a well-defined logical semantics based on Kripke frames.

This work is structured as follows: in Section II we present the logic of agency proposed here, based on Dynamic Preference Logic; in Section III we discuss the connection between the logic proposed and Agent Programming. In Section IV, we discuss a tractable expressive subset of the language which can be used to implement agent programming languages. In section V, we present the related work and compare their contributions to ours. Finally, in Section VI, we present our final considerations.

II. A DYNAMIC LOGIC FOR BDI PROGRAMMING

In this section, we present a dynamic propositional modal logic of agency which will be used to specify an agent's mental state. Throughout this work, we will assume the existence of a set P of propositional symbols and we will denote by \mathcal{L}_0 the language of propositional logic constructed over the symbols in P .

We assume a BDI agent has a library of plans describing which actions she can perform on the environment. For the sake of simplicity, we will assume that the plans are deterministic, completely specified and STRIPS-like. Aware of these restrictions, we introduce the notion of plan library.

Definition 1. We call $\mathcal{A} = \langle \Pi, pre, pos \rangle$ a plan library, iff Π is a finite set of plans symbols, $pre, pos : \Pi \rightarrow \mathcal{L}_0$ are functions that map each plan to its preconditions and post-conditions, respectively. We further require that the post-conditions of any plan is a consistent conjunction of propositional literals. We say $\alpha \in \mathcal{A}$ for any plan symbol $\alpha \in \Pi$.

With this definition in mind, we can establish the language we will use as a base for our constructions.

Definition 2. Let $\mathcal{A} = \langle \Pi, pre, pos \rangle$ be a plan library. We define the language $\mathcal{L}_{\leq P, \leq D}(\mathcal{A})$ by the following grammar (where $p \in P$ and $\alpha \in \mathcal{A}$):

$$\begin{aligned} \varphi ::= & p \mid \neg\varphi \mid \varphi \wedge \varphi \mid A\varphi \mid [\leq_P]\varphi \mid [<_P]\varphi \\ & \mid [\leq_D]\varphi \mid [<_D]\varphi \mid [\alpha]\varphi \mid I\alpha \end{aligned}$$

The formula $A\varphi$ means that it is universally true that φ holds, while the formula $[\leq_D]\varphi$ ($[\leq_P]\varphi$) means that in all worlds equally or more desirable (plausible) than the current one, φ holds and $[\lt_D]\varphi$ ($[\lt_P]\varphi$) meaning that in all worlds strictly more desirable (plausible) than the current one, φ holds. Finally, the formulas $[\alpha]\varphi$ and $I\alpha$ mean that after carrying out the plan α , the property φ holds, and that the agent intends to execute a plan α , respectively. As usual, we will define $E\varphi \equiv \neg A\neg\varphi$ and $\langle \leq \square \rangle \varphi \equiv \neg[\leq \square]\neg\varphi$ with $\square \in \{P, D\}$.

To interpret these formulas, we introduce the notion of *agent model* - which will be used to represent an agent's mental state.

Definition 3. Let $\mathcal{A} = \langle \Pi, pre, pos \rangle$ be a plan library. An agent model is a tuple $M = \langle W, \leq_P, \leq_D, I, v \rangle$ where W is a set of possible worlds, and both \leq_P and \leq_D are pre-orders over W with well-founded strict parts \lt_P and \lt_D , $I \subseteq \Pi$ is a set of adopted plans (or intentions), and v is a valuation function.

To model the effect of executing a plan $\alpha \in \mathcal{A}$ given an agent model M , we will define the notion of model update, as commonly used in the area of Dynamic Epistemic Logic.

Definition 4. Let $\mathcal{A} = \langle \Pi, pre, pos \rangle$ be a plan library, $\alpha \in \mathcal{A}$ a plan, and $M = \langle W, \leq_P, \leq_D, I, v \rangle$ an agent model. The update of model M by execution of plan α is defined as the model $M \otimes [\mathcal{A}, \alpha] = \langle W', \leq'_P, \leq'_D, I', v' \rangle$ where

$$\begin{aligned} W' &= \{w \in W \mid M, w \models pre(\alpha)\} \\ \leq'_\square &= \leq_\square \cap W' \times W' \\ I' &= I \\ v'(p) &= \begin{cases} W' & \text{if } pos(\alpha) \models p \\ \emptyset & \text{if } pos(\alpha) \models \neg p \\ v(p) \cap W' & \text{otherwise} \end{cases} \end{aligned}$$

The interpretation of the formulas is defined as usual, with each modality corresponding to an accessibility relation.

$$\begin{aligned} M, w \models [\leq_P]\varphi &\text{ iff } \forall w' \in W : w' \leq_P w \Rightarrow M, w' \models \varphi \\ M, w \models [\lt_P]\varphi &\text{ iff } \forall w' \in W : w' \lt_P w \Rightarrow M, w' \models \varphi \\ M, w \models [\leq_D]\varphi &\text{ iff } \forall w' \in W : w' \leq_D w \Rightarrow M, w' \models \varphi \\ M, w \models [\lt_D]\varphi &\text{ iff } \forall w' \in W : w' \lt_D w \Rightarrow M, w' \models \varphi \\ M, w \models [\alpha]\varphi &\text{ if } M, w \models pre(\alpha) \text{ then } M \otimes [\mathcal{A}, \alpha], w \models \varphi \\ M, w \models I\alpha &\text{ iff } \alpha \in I \end{aligned}$$

In this work, we denote the most plausible (similarly, desirable) worlds in the set S by $Min_{\leq_P} S$. In other words,

$$Min_{\leq_P} S = \{w \in S \mid \forall w' \in S : w' \leq_P w \Rightarrow w \leq_P w'\}$$

Let $\varphi \in \mathcal{L}_{\leq_P, \leq_D}(\mathcal{A})$, we also define the formula $\mu_P\varphi \equiv (\varphi \wedge \neg[\lt_P]\varphi)$ (similarly, $\mu_D\varphi$), which is satisfied exactly by the minimal worlds according to the order \leq_P (similarly for \leq_D) which satisfy the formula φ , i.e., $\llbracket \mu_P\varphi \rrbracket = Min_{\leq_P} \llbracket \varphi \rrbracket$. These formulas will be useful to encode mental attitudes in this logic.

A. Encoding mental attitudes

Following Souza et al [3] [2], we introduce a codification of mental attitudes in the language $\mathcal{L}_{\leq_P, \leq_D}(\mathcal{A})$. In this work, we interpret the notion of 'possible world' as epistemically possible, not metaphysically possible. As such, the universal modality can be used to encode the knowledge held by an agent.

$$K\varphi \equiv A\varphi$$

We encode the (KD45) notion of belief as what is true in the worlds that the agent believes to be the most plausible ones. As such, our notion $B\varphi$ means that 'it is most plausible that φ '.

$$B\varphi \equiv A(\mu_P\top \rightarrow \varphi)$$

Similar to belief, we propose a codification of (consistent) desires as everything that is satisfied in all most desirable worlds.

$$G(\varphi) \equiv A(\mu_D\top \rightarrow \varphi)$$

Our language possesses the notion of procedural intentions by the formula $I\alpha$. To encode Bratman's [1] notion of prospective intention, we will define a formula $Int\varphi$. First, however, we must encode the restrictions imposed by Bratman for consistency of an intention by means of a formula $AdmInt(\varphi)$ meaning that 'it is admissible for the agent to intend that φ '.

$$AdmInt(\varphi) \equiv G(\varphi) \wedge E(\varphi) \wedge \neg B(\varphi)$$

With this notion, we can define the notion of having an 'intention that φ '.

$$Int(\varphi) \equiv AdmInt(\varphi) \wedge \bigvee_{\alpha \in \mathcal{A}} (I\alpha \wedge B(pre(\alpha) \wedge [\alpha]\varphi))$$

Notice that, while we imposed several restrictions in our codification for an agent to rationally hold some prospective intention, none of these restrictions have been required for an agent to hold a procedural intention, i.e. an *intention to do* - here represented by the set of adopted plans I in the model. To model the kind of agent that satisfies Bratman's restrictions, we define the notion of a coherent agent model, i.e. an agent having a coherent mental state.

Definition 5. Let \mathcal{A} be a plan library and $M = \langle W, \leq_P, \leq_D, I, v \rangle$ be an agent model. We say a set $I \subset \Pi$ of plans is \mathcal{A} -coherent in M if for all $\alpha \in I$, $M \models B(pre(\alpha))$ and $M \models AdmInt(pos(\alpha))$. If I is \mathcal{A} -coherent in M , we say M is a coherent agent model.

B. Dynamic operations on agent's mental states

Once established the basic language and the encodings of mental attitudes, we define some well-behaved mental operations, which will be used to implement agents' practical reasoning. Here we explore three dynamic operations on agent models, each a representative of the three basic mental operations as studied by the Belief Revision Theory [4]: expansion, revision and contraction.

The first operation we introduce is that of public announcement. This operation corresponds, in a sense, to the operation of expansion from Belief Revision Theory. Based on the codifications we provided in the previous section, this operation can be interpreted as the mental operation of knowledge acquisition.

Definition 6. Let $M = \langle W, \leq_P, \leq_D, I, v \rangle$ be a coherent agent model and φ a formula of \mathcal{L}_0 . We say the model $M_{! \varphi} = \langle W_{! \varphi}, \leq_{P! \varphi}, \leq_{D! \varphi}, I_{! \varphi}, v_{! \varphi} \rangle$ is the result of public announcement of φ in M , where:

$$\begin{aligned} W_{! \varphi} &= \{w \in W \mid M, w \models \varphi\} \\ \leq_{\square! \varphi} &= \leq_{\square} \cap (W_{! \varphi} \times W_{! \varphi}) \\ I_{! \varphi} &\text{ is the maximal subset of } I \text{ that is } \mathcal{A} \text{ - coherent} \\ v_{! \varphi}(p) &= v(p) \cap W_{! \varphi} \end{aligned}$$

The radical upgrade of an agents beliefs by an information φ results in a model such that all worlds satisfying φ are deemed more plausible than those not satisfying it. This operation corresponds to a operation of belief revision from belief Revision Theory.

Definition 7. Let $M = \langle W, \leq_P, \leq_D, I, v \rangle$ be a coherent agent model and φ a formula of \mathcal{L}_0 . We say the model $M_{\uparrow \varphi} = \langle W, \leq_{P\uparrow \varphi}, \leq_D, I_{\uparrow P \varphi}, v \rangle$ is the result of the radical upgrade on the plausibility of M by φ , where

$$\begin{aligned} \leq_{P\uparrow \varphi} &= (\leq_P \setminus \{\langle w, w' \rangle \mid M, w \not\models \varphi \text{ and } M, w' \models \varphi\}) \cup \\ &\quad \{\langle w, w' \rangle \mid M, w \models \varphi \text{ and } M, w' \not\models \varphi\} \\ I_{\uparrow P \varphi} &\text{ is the maximal subset of } I \text{ that is } \mathcal{A} \text{ - coherent} \end{aligned}$$

We can similarly define the radical upgrade of the agents desires by the operation $\uparrow_D \varphi$, which updates the desirability relation, instead of the plausibility relation.

Lastly, we introduce the operation of lexicographic contraction [5]. This operation corresponds to performing a re-ordering of the worlds in a way that both φ and $\neg\varphi$ are considered equally plausible (or equally desirable) to the agent. To define this operation in a more elegant way, we will define the notion of the plausibility degree of a world.

Definition 8. Let $M = \langle W, \leq_P, \leq_D, I, v \rangle$ be a coherent agent model, $w \in W$ a possible world and $\varphi \in \mathcal{L}_{\leq_P, \leq_D}(\mathcal{A})$ a formula. We say that w has plausibility degree $n \in \mathbb{N}$ for ϕ , denoted $n = d_{P\varphi}(w)$, if $M, w \models \varphi$ and there is a maximal chain $w_0 <_P w_1 <_P w_2 <_P \dots <_P w_n$ s.t. $\forall i : w_i \in \llbracket \varphi \rrbracket$, $w_0 \in \text{Min}_{\leq_P} \llbracket \varphi \rrbracket$ and $w_n = w$.

We can define the desirability degree of a world w for φ , denoted by $d_{D\varphi}(w)$, the same way. With that notion, we define the lexicographic contraction as below.

Definition 9. Let $M = \langle W, \leq_P, \leq_D, I, v \rangle$ be a coherent agent model and φ a formula of \mathcal{L}_0 . We say the model $M_{\Downarrow P \varphi} = \langle W, \leq_{P\Downarrow \varphi}, \leq_D, I_{\Downarrow P \varphi}, v \rangle$ is the lexicographic contraction of the plausibility of M by φ , where:

$$w \leq_{P\Downarrow \varphi} w' \text{ iff } \begin{cases} w \leq_P w' & \text{if } w, w' \in \llbracket \varphi \rrbracket \\ w \leq_P w' & \text{if } w, w' \notin \llbracket \varphi \rrbracket \\ d_{P\varphi}(w) < d_{P\neg\varphi}(w') & \text{if } w \in \llbracket \varphi \rrbracket \text{ and } \\ & w' \notin \llbracket \varphi \rrbracket \\ d_{P\neg\varphi}(w) < d_{P\varphi}(w') & \text{if } w \notin \llbracket \varphi \rrbracket \text{ and } \\ & w' \in \llbracket \varphi \rrbracket \end{cases}$$

$I_{\Downarrow P \varphi}$ is the maximal subset of I that is \mathcal{A} - coherent

As before, we can similarly define the lexicographic contraction on the desirability of M by φ (denoted by $M_{\Downarrow D \varphi}$). These operations correspond to the contraction (or abandonment) of a belief/desire by the agent.

For each operation \star defined before, we introduce a new modality $[\star\varphi]\psi$ in our language, meaning “after the operation of \star by φ , ψ holds”. which can be interpreted as

$$M, w \models [\star\varphi]\psi \quad \text{iff} \quad M_{\star\varphi}, w \models \psi$$

An important result about the dynamified logic is that, if we consider some special kind of model, which includes the models we will use in Section III to reason about Agent Programming, it has the same expressibility as the static logic presented before [2]. In fact, the formulas $[\uparrow\varphi]\psi$, $[\uparrow\Box\varphi]\psi$ and $[\Downarrow\Box\varphi]\psi$, with $\Box \in \{P, D\}$, are definable in the language $\mathcal{L}_{\leq_P, \leq_D}(\mathcal{A})$ by a set of reduction axioms.

III. REASONING ABOUT BDI AGENTS USING DPL

An interesting property of Preference Logic - the logic used as a foundation to construct $\mathcal{L}_{\leq_P, \leq_D}(\mathcal{A})$ - is that preference models can be encoded by means of some structures known as priority graphs [6]. Exploring this connection, we will show how we can use agent programs with stratified mental attitudes, e.g. beliefs annotated with their credence/plausibility level, to obtain a model for the agent’s mental state.

In most BDI agent programming languages, an agent program is defined by means of a tuple $ag = \langle K, B, D, I \rangle$, where K , B and D are sets of (ranked) propositional formulas representing the agent’s knowledge, beliefs and desires, respectively, and I is a set of plans adopted by the agent.

A ranked formula $\langle \varphi, n \rangle$ expresses that the agent has a certain degree of uncertainty in the information φ . As such, if the agent has $\langle \varphi, n \rangle$ in her belief base, it means that the agent believes that φ is true with an uncertainty degree of n .

Definition 10. Let P be a set of propositional variables and \mathcal{A} a plan library. We call an agent program over \mathcal{A} , a tuple of finite sets $ag = \langle K, B, D, I \rangle$ where:

- $K \subset \mathcal{L}_0$ is the knowledge base;
- $B \subset \mathcal{L}_0 \times \mathbb{N}$ is called the stratified belief base.
- $D \subset \mathcal{L}_0 \times \mathbb{N}$ is called a stratified goal base.
- $I \subset \mathcal{A}$ is the (procedural) intention base.

When the plan library \mathcal{A} is clear, we will often call the tuple $ag = \langle K, B, D, I \rangle$ an agent program.

Given the definition above, we define the mental attitudes of an agent, i.e. what she knows, believes, etc. by means of the components of an agent program. Notice that our belief and desire bases are stratified, in the sense that the beliefs (similarly desires) of an agent are ranked according to their

plausibility. Since some of these beliefs may be contradictory with each other, we must compute the maximal set of beliefs (desires) that is consistent - respecting the stratification of the base.

Definition 11. Let $\Gamma \subset \mathcal{L}_0 \times \mathbb{N}$ be a finite set of pairs $\langle \varphi, i \rangle$ and let $\Gamma_i = \{\varphi \mid \langle \varphi, i \rangle \in \Gamma\}$. We define the maximal consistent subset of Γ , the set $\Gamma^{Max} \subset \mathcal{L}_0$, s.t.

- $\Gamma^{Max} \subseteq \bigcup \Gamma_i$ and if $\langle \varphi, i \rangle \in \Gamma$ and $\varphi \in \Gamma^{Max}$ then $\Gamma_i \subseteq \Gamma^{Max}$;
- $\forall \Gamma' \subseteq \bigcup_{i=1}^{\infty} \Gamma_i : (\exists \Gamma_i \subseteq \Gamma' \wedge \Gamma_i \not\subseteq \Gamma^{Max} \Rightarrow \Gamma' \models \perp$ or $\exists \Gamma_j \subseteq \Gamma^{Max} \wedge \Gamma_j \not\subseteq \Gamma'$ and $j < i$)

With this in mind, we can provide interpretations of the notions of belief and desire by means of such bases.

Definition 12. Let $ag = \langle K, B, D, I \rangle$ be an agent program and $\varphi \in \mathcal{L}_0$. We define what an agent believes, desires and intends as:

- $ag \models K\varphi$, iff $K \models \varphi$
- $ag \models B\varphi$, iff $B^{Max} \models \varphi$
- $ag \models G\varphi$, iff $D^{Max} \models \varphi$
- $ag \models I\varphi$, iff $ag \models G\varphi$ and $\exists \alpha \in I$, s.t. $pos(\alpha) \models \varphi$

While we placed no condition on agent programs, since in this work we adhere to Bratman's [1] notion of intention, our declarative mental attitudes must satisfy some constraints in order for the agent to be considered rational.

Definition 13. Let $ag = \langle K, B, D, I \rangle$ be an agent program. We say ag is coherent iff all of the conditions below hold.

- 1) Knowledge consistency: $K \not\models \perp$
- 2) Belief-Knowledge consistency: $\varphi \in K$ iff $\langle \varphi, 0 \rangle \in B$
- 3) Desire-Knowledge consistency: $\varphi \in K$ iff $\langle \varphi, 0 \rangle \in D$
- 4) Intention-Desire consistency: for all $\alpha \in I$ there is a φ in D^{Max} s.t. $pos(\alpha) \models_{\mathcal{L}_0} \varphi$;
- 5) Pursuable plan: $\forall \alpha \in I : ag \models B(pre(\alpha))$;
- 6) Intention consistency: $\{pos(\alpha) \mid \alpha \in I\} \not\models \perp$;
- 7) Plans are relevant: $\forall \alpha \in I : ag \not\models B(pos(\alpha))$.

Liu [6] shows that preference relations can be equivalently represented by means of syntactical constructs, known as priority graphs. A priority graph, however, is nothing more than a partial order over propositional formulas, much similar to the stratified bases we introduced here. As such, we can use the same reasoning to compute the plausibility and desirability orders of an agent model by means of belief and desire bases.

Definition 14. Let $\Gamma \subset \mathcal{L}_0 \times \mathbb{N}$ be a stratified base, W a set of possible worlds and $v : \mathcal{L}_0 \rightarrow W$ a valuation function. Considering $\Gamma_i = \{\varphi \mid \langle \varphi, i \rangle \in \Gamma\}$ and $w \models X$ to stand for $\forall \varphi \in X : (w \in v(\varphi))$, then we define the pre-order $\leq_{\Gamma} \subseteq W \times W$ s.t.

$$w \leq_{\Gamma} w' \text{ iff } \forall i \in \mathbb{N} : (w' \models \Gamma_i \Rightarrow w \models \Gamma_i) \vee (\exists j < i \text{ s.t. } (w \models \Gamma_j \text{ and } w' \not\models \Gamma_j))$$

Using this construction, we are able to construct an agent model from an agent program.

Definition 15. Let $ag = \langle K, B, D, I \rangle$ be an agent program, we define the model induced by ag as $M_{ag} = \langle \llbracket K \rrbracket, \leq_B, \leq_D, I, v \rangle$ where $\llbracket K \rrbracket \subset 2^P$ are all the propositional valuations that satisfy the set K , $\leq_B \subset \llbracket K \rrbracket \times \llbracket K \rrbracket$ and $\leq_D \subset \llbracket K \rrbracket \times \llbracket K \rrbracket$ are the preference relations induced by the bases B and D , as described in Definition 14, and $w \in v(p)$ iff $p \in w$.

Finally, since preference relations can always be encoded as priority graphs [6], we can always compute agent programs describing mental models.

Proposition 16. Let $M = \langle W, \leq_P \leq_D, I, v \rangle$ be an agent model, with $W \subset 2^P$, then there is an agent program $ag = \langle K, B, D, I \rangle$ s.t. $M = M_{ag}$. More yet, M is a coherent agent model iff ag is a coherent agent program.

From this result and the encodings of mental attitudes in both the logic and in agent programs, it is not difficult to see that the mental notions coincide.

Corollary 17. Let $ag = \langle K, B, D, I \rangle$ be a coherent agent program and $\varphi \in \mathcal{L}_0$ be a propositional formula, then

$$\begin{array}{ll} M_{ag} \models K(\varphi) & \text{iff } ag \models K\varphi \\ M_{ag} \models B(\varphi) & \text{iff } ag \models B\varphi \\ M_{ag} \models G(\varphi) & \text{iff } ag \models G\varphi \\ M_{ag} \models Int(\varphi) & \text{iff } ag \models I\varphi \\ M_{ag} \models I\alpha & \text{iff } \alpha \in I \end{array}$$

We have two considerations to make about the codification presented here. First, regarding the complexity of reasoning about agent programs' attitudes, to compute an agents beliefs (or goals), it requires a linear number of propositional satisfiability checks on the depth of the base. Second, regarding the notion of mental attitudes encoded here, notice that we adopted a notion of goal as a maximal set of consistent desires - consistent with other works in BDI programming. It is not difficult, however, to treat other notions, as that of Van Riemsdij et al [7], in our framework.

IV. TRACTABLE FRAGMENTS OF DPL

We have seen so far that we can use the logic $\mathcal{L}_{\leq_P, \leq_D}(\mathcal{A})$ to reason about agent programs. The computational complexity of reasoning about agents, however, is far too great to be useful for real-world problems. In this section, we investigate some restrictions on the kinds of agent programs and agent models that guarantee that the reasoning problems in the resulting logic are tractable.

Since agent programs are defined over propositional formulas - and reasoning about propositional satisfiability is a well-known NP-complete problem - we define a restriction of agent programs for which reasoning will be proved to be tractable.

Definition 18. Let $\Gamma \subset \mathcal{L}_0 \times \mathbb{N}$ be a stratified base, we say Γ is conjunctive iff for all $\langle \varphi, i \rangle \in \Gamma$, φ is a conjunction of literals, i.e. $\varphi = \bigwedge l_k$, with $l_k = p_k$ or $l_k = \neg p_k$.

A conjunctive agent program is, thus, an agent program in which all of its bases are conjunctive.

Definition 19. Let $ag = \langle K, B, D, I \rangle$ be an agent program, we say ag is a conjunctive agent program iff K is a set of conjunctive formulas, B and D are conjunctive stratified bases and for any plan $\alpha \in I$, $pre(\alpha)$ and $pos(\alpha)$ are conjunctive formulas.

First, we must show that we can compute the maximal consistent subsets, such as B^{Max} and D^{Max} , in polynomial time. To do so, we provide the Algorithm $Max(\Gamma)$, depicted in Figure 1.

Algorithm $Max(\Gamma)$
Input : a conjunctive stratified base Γ
Output : a set of literals Γ^{Max} corresponding to the maximal consistent subset of the base Γ

```

[1]  $\Gamma^{Max} := \{\}$ 
[2]  $n :=$  maximal depth of  $\Gamma$ 
[3] for  $i := 1$  to  $n$ 
[4]    $\Gamma_i := \{l \mid \langle \varphi, i \rangle \in \Gamma \text{ and } l \text{ appears in } \varphi\}$ 
[5]   if  $\neg l \in \Gamma_i$  and  $l \in \Gamma_i$  for some  $l$  then
[6]     continue
[7]   else
[8]     if  $l \in \Gamma_i$  and  $\neg l \in \Gamma^{Max}$  for some  $l$  then
[9]       continue
[10]    else
[11]       $\Gamma^{Max} := \Gamma^{Max} \cup \Gamma_i$ 
[12] return  $\Gamma^{Max}$ 

```

Fig. 1: Algorithm for the maximal consistent subset of Γ .

Proposition 20. Let $\Gamma \subset \mathcal{L}_0 \times \mathbb{N}$ be a conjunctive stratified base, then the algorithm Max presented in Figure 1 is correct, i.e. computes Γ^{Max} in $O(n^3m^2)$ time, where n is the size of Γ and m is the size of the biggest formula in Γ .

As a consequence, we can always decide whether a conjunctive agent program knows (beliefs, desires or intends) a certain formula φ in polynomial time.

Corollary 21. Let $ag = \langle K, B, D, I \rangle$ be a conjunctive agent program and $\varphi \in \mathcal{L}_0$ a propositional formula in disjunctive normal form. We can compute whether $ag \models K(\varphi)$ ($ag \models B(\varphi)$ or $ag \models D(\varphi)$) in polynomial time in the size of K (B or D) and φ .

Corollary 21 guarantees that we can reason about the agent's mental state at any point in time in the program execution in polynomial time. The execution of an agent program, however, is usually determined by its reasoning cycle, i.e. the execution of certain mental changing operations that describe the agent's reasoning. These mental operations are usually described by means of changes in the agent's mental state. As such, to provide a truly tractable semantic framework to reason about agent programming, we must ensure that these mental changing operations can be computed in polynomial time.

We now dedicate our attention to this problem. We aim to provide tractable operations on agent programs to compute the dynamic operations discussed in Section II.

First, based on the work of Girard [8] and of Liu [6], let's show how we can compute knowledge acquisition - interpreted here as a public announcement - using agent programs.

Proposition 22. Let $ag = \langle K, B, D, I \rangle$ be an agent program, $\varphi \in \mathcal{L}_0$, and $ag' = \langle K \cup \{\varphi\}, B', D', I' \rangle$, where

$$\begin{aligned}
B' &= (B \cup \{\langle \varphi, 0 \rangle\}) \\
D' &= (D \cup \{\langle \varphi, 0 \rangle\}) \\
I' &= \{\alpha \in I \mid (B')^{Max} \models pre(\alpha) \text{ and } (B')^{Max} \not\models pos(\alpha) \\
&\quad \text{and } \exists \varphi \in D' : pos(\alpha) \models \varphi\}
\end{aligned}$$

be the agent program resulting of agent ag obtaining a knowledge φ . Then $M_{ag'} = M_{ag|\varphi}$. We denote ag' by $ag|\varphi$.

As a result of this encoding, we can compute knowledge acquisition/public announcement in polynomial time.

Corollary 23. Let $ag = \langle K, B, D, I \rangle$ be a conjunctive agent program and $\varphi, \psi \in \mathcal{L}_0$ conjunctive propositional formulas. We can compute whether $ag|\varphi \models K(\psi)$ ($ag|\varphi \models B(\psi)$ or $ag|\varphi \models D(\psi)$) in polynomial time in the size of K (B or D), φ and ψ .

As Radical Upgrade can also be computed by means of transformation on the agent programs, we can represent the mental operation of belief revision in our framework.

Proposition 24. Let $ag = \langle K, B, D, I \rangle$ be a coherent agent program and $\varphi \in \mathcal{L}_0$, let yet $ag' = \langle K, B', D, I' \rangle$, where

$$\begin{aligned}
B' &= \{\langle \psi, 0 \rangle \mid \psi \in K\} \cup \{\langle \psi, i+2 \rangle \mid \langle \psi, i \rangle \in B\} \\
&\quad \cup \{\langle \varphi, 1 \rangle\} \\
I' &= \{\alpha \in I \mid (B')^{Max} \models pre(\alpha) \text{ and } (B')^{Max} \not\models pos(\alpha)\}
\end{aligned}$$

be the agent program resulting of agent ag revising her beliefs by information φ . Then $M_{ag'} = M_{ag \uparrow_P \varphi}$. We denote ag' by $ag \uparrow_P \varphi$.

As a corollary, reasoning about the resulting mental state of the agent after belief revision is a tractable problem.

Corollary 25. Let $ag = \langle K, B, D, I \rangle$ be a conjunctive agent program and $\varphi, \psi \in \mathcal{L}_0$ conjunctive propositional formulas. We can compute whether $ag \uparrow_P \varphi \models K(\psi)$ ($ag \uparrow_P \varphi \models B(\psi)$ or $ag \uparrow_P \varphi \models D(\psi)$) in polynomial time in the size of K (B or D), φ and ψ .

A similar result can be stated for the radical upgrade of the agents desires, instead of beliefs. This operation represents the adoption of a given goal.

To implement lexicographic contraction, we use the algorithm depicted in Figure 2. We represent by $\varphi[\psi|l]$ the substitution of literal l appearing in φ by the formula ψ .

Proposition 26. Let $ag = \langle K, B, D, I \rangle$ be a coherent agent program and $\varphi \in \mathcal{L}_0$, let yet $ag' = \langle K, B', D, I' \rangle$, where

$$\begin{aligned}
B' &= Cont(B, \varphi) \\
I' &= \{\alpha \in I \mid (B')^{Max} \models pre(\alpha) \text{ and } (B')^{Max} \not\models pos(\alpha)\}
\end{aligned}$$

be the agent program resulting of agent ag contracting her beliefs by information φ . Then $M_{ag'} = M_{ag \downarrow_P \varphi}$, i.e. the

Algorithm $Cont(\Gamma, \varphi)$
Input : a conjunctive stratified base Γ
a negated conjunctive formula $\neg\varphi$
Output : $\Gamma_{\Downarrow\varphi}$ the lexicographic contraction of Γ by φ
[1] $\Gamma_{\Downarrow\varphi} := \{\}$
[1] **for each** $\langle\psi, i\rangle \in \Gamma$
[2] $\psi' := \psi$
[2] **for each** propositional symbol p in φ
[2] $\psi' := \psi[\top|p][\top|\neg p]$
[3] $\Gamma_{\Downarrow\varphi} := \Gamma_{\Downarrow\varphi} \cup \{\langle\psi', i\rangle\}$
[12] **return** $\Gamma_{\Downarrow\varphi}$

Fig. 2: Algorithm for the contraction of a base Γ by a formula φ .

algorithm $Cont(\Gamma, \varphi)$ depicted in Figure 2 is correct. We denote ag' by $ag_{\Downarrow_P\varphi}$.

As before, we can reason about the changes in the mental state of the agent after the contraction of a belief, or similarly the withdraw of a goal, in polynomial time to the size of the agent program and the formulas.

Corollary 27. *Let $ag = \langle K, B, D, I \rangle$ be a conjunctive agent program and $\varphi \in \mathcal{L}_0$ a disjunctive propositional formula and $\psi \in \mathcal{L}_0$ conjunctive propositional formula, we can compute whether $ag_{\Downarrow_P\varphi} \models K(\psi)$ ($ag_{\Downarrow_P\varphi} \models B(\psi)$ or $ag_{\Downarrow_P\varphi} \models D(\psi)$) in polynomial time in the size of K (B or D), φ and ψ .*

With these results, we've provided a restriction of the logic which with which reasoning about agents' mental states is tractable and provided a way to translate from agent models to agent programs.

V. RELATED WORK

From the Agent Programming perspective, the two most important works on modelling BDI mental attitudes are, in our opinion, the seminal work of Cohen and Levesque [9] and the work of Rao and Georgeff [10] describing the logic BDI-CTL. While their contribution to the area is undeniable, much criticism has been drawn to both approaches. Particularly, they are difficult to connect with agent programming languages, by the use of a possible-world model semantics.

Works as that of Wobcke [11] and of Hindriks and Van der Hoek [12] propose ways to connect the semantics of a given programming language to some appropriate logic to reason about agent's mental attitudes. These logics, however, cannot represent the mental actions that characterize the practical reasoning process of the agent, i.e. the agent program execution. As such, it is not clear how to establish connection between the logic and agent programs.

Perhaps the work most related to ours in spirit is that of Hindriks and Meyer [13]. They propose a dynamic logic for agents and show that this logic has an equivalent semantics based on the operational semantics of the programming language. The main difference between from our approach is that they choose to work in the framework of Situation Calculus and, as such, mental actions are only implicitly defined in their framework, while the inclusion of such actions in the language is exactly the main advantage advocated by us. In some sense,

our work can be seen as a generalisation of their work, since by employing Dynamic Preference Logic the equivalence they seek between operational semantics and declarative semantics can be automatically achieved by the results of Liu [6].

VI. FINAL CONSIDERATIONS

Our work has investigated the use of a Dynamic Preference Logic to encode BDI mental attitudes and its connections to Agent Programming. We provided an expressive fragment of the logic for which reasoning about agents' mental states is tractable and how this can be computed by means of agent programs. With this, we believe we provided a roadmap to use Dynamic Preference Logic as a semantic framework to specify and also implement the formal semantics of BDI agent programming languages with declarative mental attitudes.

While we provide a fairly simple encoding of the mental attitudes in this work, the logic discussed here is expressive enough to encode different notions of desires, goals and intentions. For example, we can represent the semantics of goals as proposed by Van Riemsdijk et al [7] in our framework.

As a future work, we aim to implement a simple fragment of an agent programming language implementing declarative mental attitudes in this language by means of the codifications proposed in this work. We believe such an implementation can be used to understand the notions of mental attitudes imbued in the language.

REFERENCES

- [1] M. E. Bratman, *Intention, plans, and practical reason*. Cambridge, US: Harvard University Press, 1999.
- [2] M. Souza, Á. Moreira, R. Vieira, and J.-J. C. Meyer, "A dynamic preference logic for reasoning about agent programming," in *Intelligent Systems (BRACIS), 2017 Brazilian Conference on*. IEEE, 2017, pp. 133–138.
- [3] M. Souza, "Choices that make you change your mind: a dynamic epistemic logic approach to the semantics of bdi agent programming languages," Ph.D. dissertation, Universidade Federal do Rio Grande do Sul, 2016.
- [4] C. E. Alchourrón, P. Gärdenfors, and D. Makinson, "On the logic of theory change: Partial meet contraction and revision functions," *Journal of Symbolic Logic*, vol. 50, no. 2, pp. 510–530, 1985.
- [5] R. Ramachandran, A. C. Nayak, and M. A. Orgun, "Three approaches to iterated belief contraction," *Journal of philosophical logic*, vol. 41, no. 1, pp. 115–142, 2012.
- [6] F. Liu, *Reasoning about preference dynamics*. New York, US: Springer, 2011, vol. 354.
- [7] M. B. Van Riemsdijk, M. Dastani, and J.-J. C. Meyer, "Goals in conflict: semantic foundations of goals in agent programming," *Autonomous Agents and Multi-Agent Systems*, vol. 18, no. 3, pp. 471–500, 2009.
- [8] P. Girard, "Modal logic for belief and preference change," Ph.D. dissertation, Stanford University, 2008.
- [9] P. R. Cohen and H. J. Levesque, "Intention is choice with commitment," *Artificial Intelligence*, vol. 42, no. 2-3, pp. 213–261, 1990.
- [10] A. S. Rao and M. P. Georgeff, "Decision procedures for BDI logics," *Journal of Logic and Computation*, vol. 8, no. 3, pp. 293–343, 1998.
- [11] W. Wobcke, "Model theory for PRS-like agents: Modelling belief update and action attempts," in *Proceedings of the 8th Pacific Rim International Conference on Artificial Intelligence*. Berlin, DE: Springer-Verlag, 2004, pp. 595–604.
- [12] K. Hindriks and W. Van der Hoek, "Goal agents instantiate intention logic," in *Logics in Artificial Intelligence*. New York, US: Springer, 2008, pp. 232–244.
- [13] K. V. Hindriks and J.-J. C. Meyer, "Toward a programming theory for rational agents," *Autonomous Agents and Multi-Agent Systems*, vol. 19, no. 1, pp. 4–29, 2009.