

Modeling and Shadowing Paraconsistent BDI Agents

Barbara Dunin-Kępcicz¹ and Andrzej Szalas^{1,2}

¹ Institute of Informatics, University of Warsaw, Banacha 2, Warsaw, Poland

² Department of Computer and Information Science, Linköping University, Sweden
{kepcicz@mimuw.edu.pl, andrzej.szalas@mimuw.edu.pl, liu.se}

Abstract. For over three decades researchers have been studying the BDI model of agency. Many robust multiagent systems have been developed, and a number of BDI logics have been studied. Following this intensive development phase, the importance of integrating BDI models with inconsistency handling and revision theory have been emphasized. There is also a demand for a tighter connection between BDI-based implementations and BDI logics. In this paper, we address these postulates by introducing a novel, paraconsistent logical BDI model close to implementation, with building blocks that can be represented as SQL/rule-based databases. Importantly, tractability is achieved by reasoning as querying. This stands in a sharp contrast to the high complexity of BDI logics. We also extend belief shadowing, a shallow and lightweight alternative to deep and computationally demanding belief revision, to encompass agents' motivational attitudes.

Keywords: Beliefs-Desires-Intentions models · paraconsistent reasoning · doxastic reasoning · shadowing · reasoning by querying

1 Introduction

The BDI model of agency has been intensively discussed over years, starting from the seminal research of Bratman [8], and creatively developed by many others, just to mention some of them [7,9,11,14,15,23,19,24,28,32,33,37,42]. Taking into account the experience both from the implemented systems as well as extensive formal studies, it is now time to pay closer attention to a more computationally friendly approach to modeling and reasoning about BDI agents. Tackling these agents within a logical framework, while employing a database perspective, has been presented in [38]. Beliefs and intentions are organized there in two temporal databases providing means for further reasoning. Within the realm of classical truth values \mathfrak{t} (true) and \mathfrak{f} (false), beliefs and intentions are required there to be internally consistent, at the same time preserving the consistency of intentions with beliefs.

A similar in spirit shift in perspective we find in [17]. Instead of reasoning about beliefs in modal logics or other high-complexity formalisms, a tractable approach based on querying belief bases has been advocated for. That is, rather than being concerned with general properties that apply to all models, agents are seen to be primarily concerned with how to act in the specific world (model) they are embedded in. For such purposes reasoning-by-querying is more appropriate. Any realistic world model that is potentially acquired from multiple information sources, typically includes inconsistencies and/or gaps of knowledge. To ensure the required expressiveness and modeling

convenience, two additional truth values are adopted: \perp (inconsistent), and \perp (unknown). This adds the dimensions of paraconsistent and paracomplete reasoning, that are central to our approach.

Real-world AI applications often need agents to switch between different roles and/or groups. This calls for a well-controlled and computationally-friendly adjustment of relevant beliefs, particularly when the structures and organizations evolve dynamically. In such volatile circumstances, a full-fledged reasoning about beliefs is too demanding. Instead, a shallow belief change may be realized by *belief shadowing* [5]. Since beliefs influence agents' motivational attitudes, shadowing needs to be extended to include goals and intentions. Our formalism covers this case.

In the context of BDI systems, in [24] the authors postulated to provide:

- a logical semantics underlying the implementations;
- a link between logical approach and revision theory;
- a paraconsistent reasoning about desires,

while stressing the importance of tractability: “one reason explaining the gap between theory and practice is the too high complexity of existing logics.” We address these postulates and ensure the desideratum of tractability by reporting original developments that include:

- a paraconsistent and paracomplete approach to beliefs, desires and intention.
- shadowing mechanism: a lightweight form of transient revision of BDI attitudes;
- an intuitive and tractable formalism used in querying and shadowing BDI bases.

In order to distinguish it from traditional BDI models and emphasize its four-valued nature, the introduced formalism is referred to as 4BDI.

The rest of the paper is structured as follows. In Section 2, we discuss our approach to modeling BDI agents. In Section 3 we present the intended architecture of 4BDI systems. Section 4 introduces the language for reasoning about beliefs, desires and intentions. Section 5 is devoted to semantical structures underlying 4BDI. In Section 6 belief shadowing is extended to desires and intentions. Complexity of the approach is addressed in Section 7, showing tractability of 4BDI querying and shadowing wrt size of the databases involved. Section 8 summarizes and discusses the design choices we have made. In Sections 9 and 10 we discuss related work and conclude the paper.

2 Modeling BDI Agents

BDI model of rational agency combines agents informational and motivational attitudes: beliefs, intentions and desires. Intuitively, beliefs represent information an agent has about the world, including environment, other agents and itself. If choosing a potentially elaborated and nuanced logical representation of beliefs, typically in dedicated multimodal logics, we know pretty well how to reason about them. But this comes at a price of high complexity. On the other hand, desires stand for those states of affairs the agent wishes to bring about. Finally, intentions represent the chosen desires that it has committed to achieve, potentially by realizing an appropriate plan consisting of actions. The entire process leading from an initial setting of beliefs and desires to the relevant

plan is a subject of practical reasoning [8]: the process of deciding, step by step, which action to perform next. To put it simply, it consists of the deliberation phase of choosing the proper intention to focus on for the time being, and the planning phase, resulting in a specific plan, that is, most of the time, chosen from the precompiled library of plan skeletons. Thus, during the practical reasoning phase, after recapitulating what to do, the agent needs to reflect how to achieve it.

Following [42]: “[...] a fundamental problem in developing such a theory of rational agency is to give an account of the relationships that exist between an agents’ mental states. In particular, a complete agent theory would explain how an agent’s mental states lead it to select and perform rational actions, thus realizing the mapping from perception to action.” Moving towards this goal, instead of playing with a family of multimodal logics, we decided to follow the database perspective: modeling beliefs, desires, and intentions requires appropriate semantical structures that are simple to use and implement. The deliberation processes are encapsulated in bridging functions of the 4BDI formalism, reflecting the appropriate transitions between 4BDI components.

Let us first discuss the components informally before providing rigorous definitions. The basic building blocks of 4BDI are *ground literals* (possibly negated atomic formulas without variables) used to construct:

- *world aspects*: finite sets of ground literals;
- *clusters*: finite sets of world aspects.

To illustrate concepts defined and discussed throughout the paper we will use a running example that starts below.

Example 1 (Running example). Eve is a medicine student who has to complete an anatomy course. To pass the course she needs to pass a test, and give a presentation or write an essay. She enjoys traveling but, to pass the course, she cannot afford to waste time on travels. The related part of her belief base contains the following world aspects:

$$\{prepare(pres), \neg prepare(essay), pass(test), \neg travel\}, \quad (1)$$

$$\{\neg prepare(pres), prepare(essay), pass(test), \neg travel\}. \quad (2)$$

Each world aspect, (1) and (2), is internally consistent. However, at the level of cluster consisting of (1) and (2) the literals $prepare(pres)$ and $prepare(essay)$ are inconsistent until Eve makes her choice and, in effect, one of the aspects is removed. \square

In 4BDI structures, inconsistencies can occur at two structural levels:

- inside a single world aspect, when the aspect contains both a literal and its negation;
- inside a single cluster, when different aspects contain contradictory literals.

Furthermore, information gathered is frequently incomplete, what calls for introducing means for paraconsistent and paracomplete reasoning.

Remark 1. World aspects represent perspectives on the world as perceived by heterogeneous information sources, like sensors, classifiers, and so on. When compared to the worlds in Kripke structures, aspects represent complementary views rather than alternative worlds. They can be implemented as SQL-like databases. The aspects may

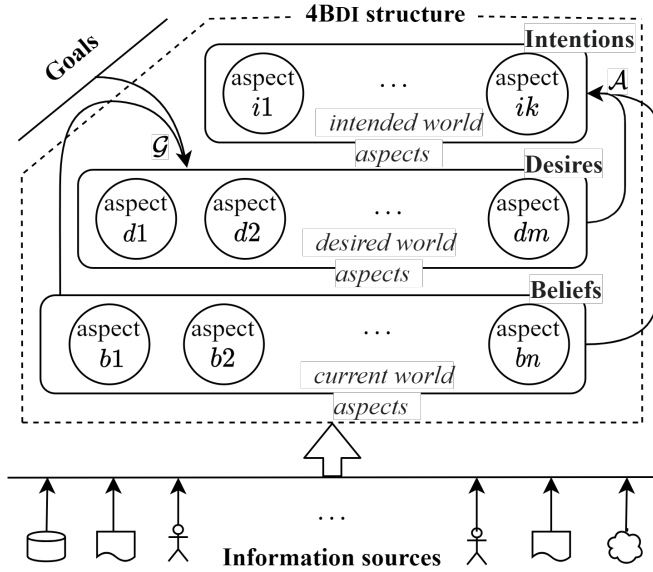


Fig. 1. High level architecture of 4BDI structures. Rounded arrows represent goal-desires (\mathcal{G}) and desires-intentions (\mathcal{A}) bridging functions.

also be computed as models of rule-based languages, such as variants of DATALOG [1] (two-valued models), ASP [21,22] (three-valued models with the third truth value \perp), paraconsistent 4QL [30] (four-valued models, additionally with \imath), etc. \square

Clusters provide an abstraction for databases used in 4BDI frames and structures. As shown in Figure 1, 4BDI frames and structures include the following components:

- *belief base* containing aspects of the current world, possibly obtained from multiple information sources;
- *desire base* containing agent’s desired world aspects;
- *intention base* containing world aspects the agent has chosen to accomplish;
- *goal-desires* bridging function (\mathcal{G}) used to determine desires;
- *desires-intentions* bridging function (\mathcal{A}) used to determine intentions;
- *goals* delivered from outside and specified by formulas.

3 The Architecture of 4BDI Systems

The intended architecture of 4BDI systems is shown in Figure 2. We assume that a system designer provides bridging functions \mathcal{G} and \mathcal{A} tailored to a specific application domain. When a goal and a belief base are provided, these functions return a desire base, followed by an intention base. This way, a frame enables the creation of a 4BDI structure containing belief, desire, and intention bases. Both functions encapsulate a number

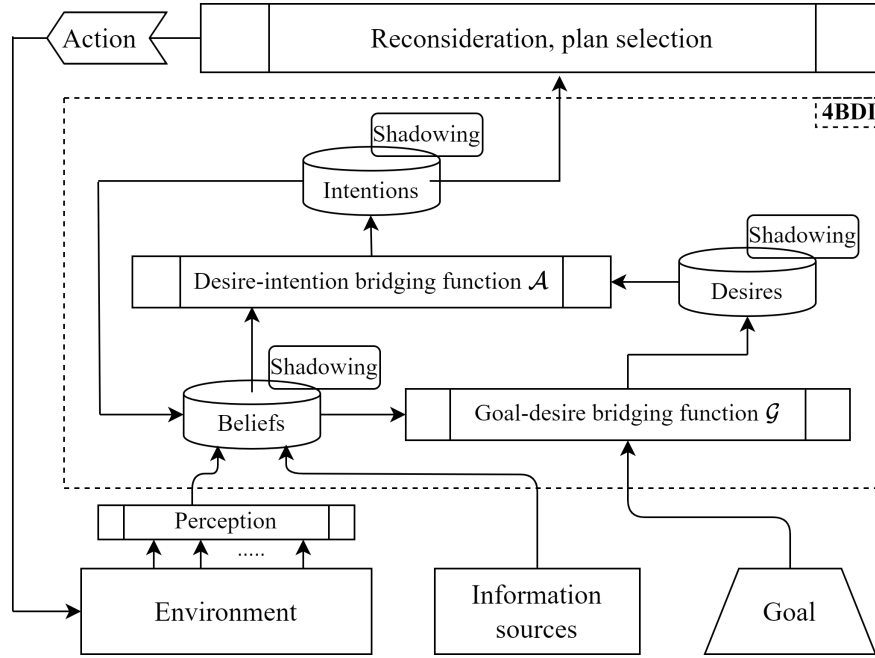


Fig. 2. The architecture of 4BDI.

of database related strategies concerning desires and, especially, intentions. On the other hand, other strategies ensuring persistence of achieving intentions, and reconsideration methods, are realized in the main body of an agent. From the agents' perspective, the 4BDI framework provides 4BDI structure construction with query answering services.

In 4BDI belief, desire and intention bases are relatively independent of other components of 4BDI agents: only query and update interfaces serve the purpose of interaction between them. This independence assumption is similar to the modular architecture of [33]. The capabilities component of [33] which “serves as a technical shortcut for encapsulating implementation of agent’s capabilities wrt its environment”, in our architecture is encapsulated in desires-intentions and possibly goal-desires bridging function.

In the 4BDI agent’s loop, first a goal for the agent is set. Next, the agent uses a pre-defined frame to determine desires and intentions based on its beliefs and goal. The loop restarts whenever the goal or relevant beliefs are updated.

4 The 4BDI Language

Let Ag be a finite set of *agent names*, uniquely identifying agents. We will use the syntax of classical first-order logic extended by operators $\mathbb{B}_a(A)$ (“ A is in a ’s beliefs”), $\mathbb{D}_a(A)$ (“ A is desired by a ”) and $\mathbb{I}_a(A)$ (“ A is intended by a ”).

Let \mathcal{P} and \mathcal{C} be finite sets of relation symbols and constants.³ The tuple $\Sigma = \langle \mathcal{P}, \mathcal{C} \rangle$ is called the language *signature*, the set of first-order variables is denoted by \mathcal{V} .

Definition 1 (Formulas, literals, ground literals). *The following BNF grammar defines syntax of formulas, where $\langle \mathcal{T} \rangle$ and $\langle \mathcal{F} \rangle$ are nonterminal symbols respectively denoting truth values and formulas, $p \in \mathcal{P}$ is a k -argument relation symbol, $\overline{arg} \in (\mathcal{C} \cup \mathcal{V})^k$ is a tuple of constants and/or variables, $x \in \mathcal{V}$ is a variable, and $a \in Ag$ is an agent name:*

$$\langle \mathcal{F} \rangle ::= \langle \mathcal{T} \rangle \mid p(\overline{arg}) \mid \neg \langle \mathcal{F} \rangle \mid \langle \mathcal{F} \rangle \wedge \langle \mathcal{F} \rangle \mid \langle \mathcal{F} \rangle \vee \langle \mathcal{F} \rangle \mid \quad (3)$$

$$\forall x(\langle \mathcal{F} \rangle) \mid \exists x(\langle \mathcal{F} \rangle) \mid \quad (4)$$

$$\mathbb{B}_a(\langle \mathcal{F} \rangle) \mid \mathbb{D}_a(\langle \mathcal{F} \rangle) \mid \mathbb{I}_a(\langle \mathcal{F} \rangle) \quad (5)$$

The set of all formulas is denoted by \mathcal{F} . Formulas of the form $p(\overline{arg})$, $\neg p(\overline{arg})$ are called literals (respectively, positive and negative ones). When \overline{arg} contains only constants, the literals are called ground. \square

Definition 2 (Plain formulas). *Formulas defined by the grammar specified in lines (3)–(4) are called plain. The set of plain formulas is denoted by \mathcal{F}^I .*

By a free occurrence of a variable x in a formula A we mean an occurrence of x outside of the scope of any quantifier, $\forall x, \exists x$, binding x . A formula without free variables is called closed. \square

Definition 3 (4BDI formulas). *A 4BDI formula is any formula of the form $\mathbb{O}_a(A)$, where $\mathbb{O} \in \{\mathbb{B}, \mathbb{D}, \mathbb{I}\}$, $a \in Ag$ and $A \in \mathcal{F}$. The set of 4BDI formulas is denoted by \mathcal{F}^{bdi} .* \square

All formulas under consideration are assumed to be closed. Since we deal with finite domains, generality is not lost in this manner. Rather than computing queries as sets of tuples satisfying the formula, we verify the truth value of formulas by substituting constants for free variables. This simplifies definitions without influencing the complexity classes of queries [1].

5 Semantics of 4BDI

5.1 From Beliefs and Goals to Desires

Typically, goals in planning and agent systems are expressed by plain formulas. Given current beliefs and a goal, one can determine desires, representing desired world aspects. This, might be done in many different ways, for example exploiting the correspondence between a disjunctive normal form of the goal, the beliefs and the resulting world aspects, as shown in the following example.

³ As standard in query languages, we exclude function symbols.

Example 2 (Example 1 continued). Assume that the current Eve's goal is:

$$\exists x(\text{prepare}(x) \wedge \text{pass}(\text{test})). \quad (6)$$

Given that the domain of x is $\{\text{pres}, \text{essay}\}$, (6) is equivalent to:

$$(\text{prepare}(\text{pres}) \vee \text{prepare}(\text{essay})) \wedge \text{pass}(\text{test}),$$

whose disjunctive normal form is:

$$(\text{prepare}(\text{pres}) \wedge \text{pass}(\text{test})) \vee (\text{prepare}(\text{essay}) \wedge \text{pass}(\text{test})).$$

Since Eve believes that traveling does not leave her a sufficient amount of time for preparing presentation nor writing an essay, the resulting desire base can contain two desired world aspects:

$$\{\text{prepare}(\text{pres}), \text{pass}(\text{test}), \neg \text{travel}\}, \{\text{prepare}(\text{essay}), \text{pass}(\text{test}), \neg \text{travel}\}. \quad \square$$

Desire determination is a broad topic. In the rest of the paper we abstract from particular methods by assuming that a goal-desires bridging function is provided as an inherent part of 4BDI agent-specific frame.

5.2 From Beliefs and Desires to Intentions

Intentions are chosen desires to which the agent has committed in a course of practical reasoning. In its first phase, the agent deliberates what intention to achieve, and then, during means-end reasoning, plans how to achieve it. The resulting plan is typically chosen from a pre-assembled collection of plans constituting a plan library that creates an independent plan base external to the 4BDI framework. Otherwise, a plan leading from the current world state to a state satisfying the intention, can be constructed from scratch. The whole process of practical reasoning must be somehow addressed by an agent. The first phase of filtering intentions from desires is encapsulated in a rather complex desires-intentions bridging function given as a part of an agent-specific 4BDI frame, as illustrated in the following example.

Example 3 (Example 2 continued). In order for Eve to accomplish the desired world aspects, she has to examine her capabilities. In the scenario three actions are explicitly referred to: $\text{prepare}(\text{pres})$, $\text{prepare}(\text{essay})$, $\text{pass}(\text{test})$. Given that:

- Eve does not have enough time to write an essay (a corresponding precondition of action $\text{prepare}(\text{essay})$ is violated);
- she is capable of preparing presentation if she does not spend time on traveling;
- she can attempt the test, what is a part of a plan to accomplish a desire $\text{pass}(\text{test})$,

the resulting intention base may, among others, contain the world aspect:

$$\{\text{prepare}(\text{pres}), \text{attempt}(\text{test}), \neg \text{travel}\}. \quad (7)$$

□

Remark 2. An interesting phenomenon implicitly occurring in Example 3 is typically neglected in planning scenarios: the plan extracted from (7) does not necessarily guarantee the initial goal (6), involving $pass(test)$. Indeed, taking a test does not imply passing it. However, in realistic scenarios, one occasionally accepts plans with unknown or inconsistent results when no plan ensuring success is available. These aspects are addressed in [6], sharing with 4BDI the same logical background.⁴ \square

5.3 4BDI Frames and Structures

Let us now define 4BDI frames and structures as outlined in Sections 5.1–5.2. Let the signature $\Sigma = \langle \mathcal{P}, \mathcal{C} \rangle$ be fixed. We assume that the domain consists of constants of \mathcal{C} . Recall that a *world aspect* is a finite set of ground literals over Σ and a *cluster* is a finite set of world aspects.

By \mathbb{W} and \mathbb{C} we denote the sets of world aspects and clusters, respectively. In the next definition we use notation $\mathcal{D} \stackrel{\text{def}}{=} \mathcal{G}(\mathcal{B}, G)$ and $\mathcal{I} \stackrel{\text{def}}{=} \mathcal{A}(\mathcal{B}, \mathcal{D}) = \mathcal{A}(\mathcal{B}, \mathcal{G}(\mathcal{B}, G))$.

Definition 4 (4BDI frames). A 4BDI frame is a pair $\mathbb{F} = \langle \mathcal{G}, \mathcal{A} \rangle$, where:

- $\mathcal{G} : \mathbb{C} \times \mathcal{F}^I \rightarrow \mathbb{C}$ is a goal-desires bridging function (takes a belief base and a goal as arguments and returns a desire base);
- $\mathcal{A} : \mathbb{C} \times \mathbb{C} \rightarrow \mathbb{C}$ is a desires-intentions bridging function (takes a belief base and a desire base as arguments and returns an intention base),

such that:

1. for all $G \in \mathcal{F}^I$ and $\mathcal{B} \in \mathbb{C}$, $\underbrace{\mathcal{A}(\mathcal{B}, \mathcal{G}(\mathcal{B}, G))}_{\mathcal{I}} \subseteq \underbrace{\mathcal{G}(\mathcal{B}, G)}_{\mathcal{D}}$
 intentions are selected from desires);
2. for each positive literal ℓ ,
 - for all $w \in \mathcal{I}$, $w(\ell) \in \mathcal{T}^d$: generalization of intention consistency at the world aspects level;
 - $\text{lub}_{\leq_i} \{w(\ell) \mid w \in \mathcal{I}\} \in \mathcal{T}^d$: generalization of intention consistency at the intention bases level. \square

Given particular goals, 4BDI frames are used to determine 4BDI structures as defined below.

Definition 5 (4BDI structures). Let $G \in \mathcal{F}^I$ be a plain formula representing a goal,⁵ and $\mathcal{B} \in \mathbb{C}$ be a belief base. Let $\mathbb{F} = \langle \mathcal{G}, \mathcal{A} \rangle$ be a 4BDI frame. A 4BDI structure for G over \mathbb{F} and \mathcal{B} is a triple $\mathbb{F}(G) \stackrel{\text{def}}{=} \langle \mathcal{B}, \mathcal{D}, \mathcal{I} \rangle$, where:

1. $\mathcal{D} = \mathcal{G}(\mathcal{B}, G)$: desires are obtained from the goal assuming beliefs about the current world;
2. $\mathcal{I} = \mathcal{A}(\mathcal{B}, \mathcal{D})$: intentions are obtained from desires assuming beliefs about the current world. \square

⁴ For an open source implementation of the paraconsistent rule language 4QL with belief bases, belief shadowing and actions, see the `inter4QL` interpreter available via `4ql.org`.

⁵ If there are a (finite) number of goals, we consider G to be the conjunction of them.




Fig. 3. Sample four-valued orderings on truth values.

4BDI agents’ mental attitudes are represented by 4BDI structures. That is, for each 4BDI agent $a \in Ag$, there is a goal G_a , an associated 4BDI frame \mathbb{F}_a and a 4BDI structure $\mathbb{S}_a \stackrel{\text{def}}{=} \mathbb{F}_a(G_a) = \langle \mathcal{B}_a, \mathcal{D}_a, \mathcal{I}_a \rangle$.

5.4 Semantics of 4BDI Formulas

When multiple information sources are present, two orderings on truth values are typically considered to address distinct needs:

- *truth ordering* for evaluating the truth value of formulas to reflect their “truth contents” (as traditionally done in logics);
- *information ordering* for fusing information from multiple sources.

Sample orderings on the values \mathfrak{t} , \mathfrak{f} , \mathfrak{i} , \mathfrak{u} , are shown in Figure 3. The ordering \leq_i is commonly used in modeling knowledge gathering:

- initially there is no knowledge about a given fact what is represented by \mathfrak{u} ;
- then, evidences for the fact or against it, making it \mathfrak{t} or \mathfrak{f} ;
- finally, evidences both for and against the fact may be collected resulting in \mathfrak{i} .

In paraconsistent reasoning, different truth orderings are chosen depending on the application area. While the most popular is perhaps \leq_t [3,2], some other approaches advocate for \leq_l [13,30]. Here we choose \leq_l , \leq_i . However, \leq_l can easily be replaced by \leq_t or other ordering adequate for the application in question.

We use a four-valued logic with the set of truth values $\mathcal{T} \stackrel{\text{def}}{=} \{\mathfrak{t}, \mathfrak{f}, \mathfrak{i}, \mathfrak{u}\}$ ordered by \leq_l, \leq_i shown in Figure 3. The set of *designated truth values* that act as true (see, e.g., [40]), denoted by \mathcal{T}^d , is a set $\mathcal{T}^d \subset \mathcal{T}$ such that $\mathfrak{t} \in \mathcal{T}^d$ and $\mathfrak{f} \notin \mathcal{T}^d$.

Remark 3. Note that the set \mathcal{T}^d permits to accept rational intentions that are not fully logically justified. This has already been indicated in Remark 2. However, \mathcal{T}^d can also be used to fine-tune 4BDI in this respect. For example, there may be no exact plan to accomplish a goal “save a victim after a disaster” by making it true. However, there may be a plan whose effect on saving the victim is inconsistent or unknown. It is rational to carry out such a plan rather than doing nothing. In 4BDI this can be achieved by setting $\mathcal{T}^d = \{\mathfrak{t}, \mathfrak{i}, \mathfrak{u}\}$. Depending on the strategy, one could set \mathcal{T}^d to be $\{\mathfrak{t}, \mathfrak{i}\}$, $\{\mathfrak{t}, \mathfrak{u}\}$, or remain with $\{\mathfrak{t}\}$. \square

$w(\tau) \stackrel{\text{def}}{=} \tau, \text{ for } \tau \in \mathcal{T};$ $w(\ell) \stackrel{\text{def}}{=} \begin{cases} \mathfrak{t} & \text{when } \ell \in w \text{ and } \neg\ell \notin w; \\ \mathfrak{f} & \text{when } \ell \notin w \text{ and } \neg\ell \in w; \\ \mathfrak{i} & \text{when } \ell \in w \text{ and } \neg\ell \in w; \\ \mathfrak{u} & \text{when } \ell \notin w \text{ and } \neg\ell \notin w; \end{cases}$	$w(\neg A) \stackrel{\text{def}}{=} \neg w(A);$ $w(A \wedge B) \stackrel{\text{def}}{=} \text{glb}_{\leq_l} \{w(A), w(B)\};$ $w(A \vee B) \stackrel{\text{def}}{=} \text{lub}_{\leq_l} \{w(A), w(B)\};$ $w(\forall x(A(x))) \stackrel{\text{def}}{=} \text{glb}_{\leq_l} \{w(A(c)) \mid c \in \mathcal{C}\};$ $w(\exists x(A(x))) \stackrel{\text{def}}{=} \text{lub}_{\leq_l} \{w(A(c)) \mid c \in \mathcal{C}\}.$
--	---

Table 1. The semantics of plain formulas, where $w \in \mathbb{W}$, ℓ is a positive literal and $A, B \in \mathcal{F}^I$.

$\bar{\mathbb{S}}(\mathbb{B}_a(A)) \stackrel{\text{def}}{=} \text{lub}_{\leq_i} \{w(A) \mid w \in \mathcal{B}_a\};$ $\bar{\mathbb{S}}(\mathbb{D}_a(A)) \stackrel{\text{def}}{=} \text{lub}_{\leq_i} \{w(A) \mid w \in \mathcal{D}_a\};$ $\bar{\mathbb{S}}(\mathbb{I}_a(A)) \stackrel{\text{def}}{=} \text{lub}_{\leq_i} \{w(A) \mid w \in \mathcal{I}_a\}.$

Table 2. Semantics of 4BDI formulas, where $\bar{\mathbb{S}} = \langle \mathbb{S}_a \mid a \in Ag \rangle$ and $A \in \mathcal{F}^I$.

We will need the greatest lower bound (glb) and the least upper bound (lub) wrt considered orderings. We assume that:

$$\begin{aligned} \text{glb}_{\leq_l}(\emptyset) &\stackrel{\text{def}}{=} \mathfrak{t}; & \text{lub}_{\leq_l}(\emptyset) &\stackrel{\text{def}}{=} \mathfrak{f}; \\ \text{glb}_{\leq_i}(\emptyset) &\stackrel{\text{def}}{=} \mathfrak{i}; & \text{lub}_{\leq_i}(\emptyset) &\stackrel{\text{def}}{=} \mathfrak{u}. \end{aligned} \tag{8}$$

Let a signature $\Sigma = \langle \mathcal{R}, \mathcal{C} \rangle$ be fixed. Semantically, plain formulas are evaluated in world aspects, while 4BDI formulas are evaluated in clusters provided by 4BDI structures. Let $w \in \mathbb{W}$ be a world aspect. The *truth value of a plain formula* $F \in \mathcal{F}^I$ in w , denoted by $w(F)$, is defined in Table 1, where the negation of truth values is defined by:

$$\neg \mathfrak{t} \stackrel{\text{def}}{=} \mathfrak{f}, \quad \neg \mathfrak{f} \stackrel{\text{def}}{=} \mathfrak{t}, \quad \neg \mathfrak{i} \stackrel{\text{def}}{=} \mathfrak{i}, \quad \neg \mathfrak{u} \stackrel{\text{def}}{=} \mathfrak{u}. \tag{9}$$

Let $\bar{\mathbb{S}} = \langle \mathbb{S}_a \mid a \in Ag \rangle$ be the tuple of 4BDI structures associated with agents in Ag . For every $\mathbb{O} \in \{\mathbb{B}, \mathbb{D}, \mathbb{I}\}$, $a \in Ag$ and $\tau \in \mathcal{T}$, we set $\mathbb{O}_a(\tau) \stackrel{\text{def}}{=} \tau$. For any 4BDI formula of the form $\mathbb{O}_a(A)$, where $A \in \mathcal{F}^I$ is a plain formula, its *truth value in $\bar{\mathbb{S}}$* is defined in Table 2. In order to calculate the *truth values of an arbitrary 4BDI formula* $A \in \mathcal{F}^{bdi}$ in $\bar{\mathbb{S}}$ of the form $\mathbb{O}_a(B)$, we use Algorithm 1 which successively computes truth values of subformulas of the form $\mathbb{O}_a(F)$, where $F \in \mathcal{F}^I$, and substitutes them with the calculated truth values.

The following proposition demonstrates that the 4BDI semantics preserves commonly assumed logical properties of beliefs, desires, and intentions. Note that, in many-valued logics, implication can be defined in a variety of ways. Rather than implication, we employ truth ordering \leq_l which reflects the semantics of classical implication on the truth values $\mathfrak{t}, \mathfrak{f}$. Below \pm denotes the empty string or the negation connective.

Input: - A 4BDI formula $A \in \mathcal{F}^{bdi}$ of the form $\mathbb{O}_a(B)$
 - A tuple of 4BDI structures $\bar{\mathbb{S}} = \langle \mathbb{S}_a \mid a \in Ag \rangle$

Output: Truth value of A in $\bar{\mathbb{S}}$

```

2 set  $C = A$ 
3 repeat
4   | chose a subformula of  $C$ ,  $\mathbb{O}'_b(F)$ , where  $F \in \mathcal{F}^I$ 
5   | set  $\tau$  to be the truth value of  $\mathbb{O}'_b(F)$  in  $\bar{\mathbb{S}}$  by applying
      | a suitable clause shown in Table 2
6   | replace in  $C$  all occurrences of  $\mathbb{O}'_b(F)$  with  $\tau$ ;
7 until  $C$  is reduced to a truth value
8 return  $C$ ;
    
```

Algorithm 1: Calculating truth values of 4BDI formulas.

Proposition 1. For any formula $A \in \mathcal{F}^I$, ground literal ℓ , any tuple of 4BDI structures, $\bar{\mathbb{S}} = \langle \mathbb{S}_a \mid a \in Ag \rangle$ and any $a \in Ag$,

$$\bar{\mathbb{S}}(\pm \mathbb{B}_a(A)) \leq_l \bar{\mathbb{S}}(\mathbb{B}_a(\pm \mathbb{B}_a(A))); \quad (10)$$

$$\bar{\mathbb{S}}(\pm \mathbb{D}_a(A)) \leq_l \bar{\mathbb{S}}(\mathbb{B}_a(\pm \mathbb{D}_a(A))); \quad (11)$$

$$\bar{\mathbb{S}}(\pm \mathbb{I}_a(A)) \leq_l \bar{\mathbb{S}}(\mathbb{B}_a(\pm \mathbb{I}_a(A))); \quad (12)$$

$$\bar{\mathbb{S}}(\neg \mathbb{O}_a(f)) = \dagger, \text{ for } \mathbb{O} \in \{\mathbb{B}, \mathbb{D}, \mathbb{I}\}; \quad (13)$$

$$\bar{\mathbb{S}}(\mathbb{I}_a(\ell)) \in \mathcal{T}^d; \quad (14)$$

$$\bar{\mathbb{S}}(\mathbb{I}_a(A)) \leq_l \bar{\mathbb{S}}(\mathbb{D}_a(A)). \quad (15)$$

Note that (10)–(12) represent positive and negative introspection, (13)–(14) generalizes the consistency laws for $\mathbb{B}, \mathbb{D}, \mathbb{I}$, while (15) reflects that intentions are selected from desires. In particular, (14)–(15) reflect requirements stated in points 1–2 of Definition 4.

6 Shadowing Beliefs, Desires and Intentions

Belief revision/update/merging is an important task in agents' activities [34]. However, it typically requires deep and/or complex adjustments of belief bases even when adaptations happen to be transient. In [5], a new kind of beliefs' update, belief shadowing, is introduced. It depends on a swap of beliefs when part of one belief base is to be shadowed by a belief base of superior agents or roles. Then no changes to belief bases are needed. This substantially improves the complexity of reasoning. Because beliefs influence agents' motivational attitudes, this approach is expanded for 4BDI agents.

As an agent may play dynamically assigned roles, let Rl denote a (finite) set of *role names* ($Rl \cap Ag = \emptyset$). We assume that roles have their associated beliefs, desires and intentions, and they can also extend the abilities of agents. For example, a role “coordinator” may have a belief base that includes aspects of mission coordination, as well as desires and intentions that reflect a role-associated goal. When an agent takes on this role, it not only adds its own beliefs, but it also expands its capabilities of coordinating the activities of other agents. The beliefs and goal-desires/desires-intentions bridging functions of the role become dominant.

$$\bar{\mathbb{S}}(\mathbb{B}_{a \text{ as } r_1 \text{ as } \dots \text{ as } r_m}(F)) \stackrel{\text{def}}{=} \begin{cases} \bar{\mathbb{S}}(\mathbb{B}_{r_m}(F)) & \text{when } \bar{\mathbb{S}}(\mathbb{B}_{r_m}(F)) \neq \perp; \\ \bar{\mathbb{S}}(\mathbb{B}_{a \text{ as } r_1 \text{ as } \dots \text{ as } r_{m-1}}(F)) & \text{otherwise.} \end{cases}$$

Table 3. Semantics of belief shadowing, where $F \in \mathcal{F}^I$, $a \in Ag$, $r_1, \dots, r_m \in Rl$ and $\bar{\mathbb{S}} = \langle \mathbb{S}_i \mid i \in Ag \cup Rl \rangle$.

Definition 6 (Shadowing expressions). By a shadowing expression we mean any expression of the form:

$$a \text{ as } r_1 \text{ as } r_2 \text{ as } \dots \text{ as } r_{m-1} \text{ as } r_m, \quad (16)$$

where $a \in Ag$, $m \geq 1$ and $r_1, \dots, r_m \in Rl$. For ϵ denoting the empty sequence, we define $a \text{ as } \epsilon \stackrel{\text{def}}{=} a$. The set of shadowing expressions is denoted by \mathcal{E} . \square

The intuition behind shadowing expression $a \text{ as } r$ is that agent a takes on the role r . That is, mental attitudes of a become those of r , unless r is ignorant about the status of a given attitude, in which case the attitude of a is binding. In the general form (16), a takes on the roles r_1, \dots, r_{m-1}, r_m . Mental attitudes of a become those of r_m , unless r_m is ignorant about the status of a given attitude, in which case, inductively, the attitude of $a \text{ as } r_1 \text{ as } \dots \text{ as } r_{m-1}$ is binding.

To make use of roles in the logical language, the $\mathbb{B}, \mathbb{D}, \mathbb{I}$ operators introduced in Definition 1 have to be extended by assuming that Line (5) of the BNF grammar is replaced by:

$$\mathbb{B}_e(\langle \mathcal{F} \rangle) \mid \mathbb{D}_e(\langle \mathcal{F} \rangle) \mid \mathbb{I}_e(\langle \mathcal{F} \rangle), \text{ for } e \in \mathcal{E}. \quad (17)$$

That is, rather than using agent names, we use shadowing expressions to indicate the context of a given operator.

Definition 7 (Extended operators and formulas). For e of the form (16) with $m \geq 1$, the operators $\mathbb{B}_e()$, $\mathbb{D}_e()$, $\mathbb{I}_e()$ are respectively called extended belief, extended desire, and extended intention operators. By an extended formula we mean any formula defined by (17). An extended 4BDI formula is a formula defined as in Definition 3 with $a \in Ag$ replaced by $e \in \mathcal{E}$. The set of extended 4BDI formulas is denoted by \mathcal{F}^{ext} . \square

Notice that the language offers a rich expressiveness. For example, one can formulate properties not directly available in other logical formalisms, like:

$$\mathbb{B}_{jack \text{ as } coord}(\mathbb{I}_{eve \text{ as } vln \text{ as } med}(assist)),$$

stating that *jack*, acting as a mission coordinator (*coord*), believes that *eve*, acting as a volunteer (*vln*) being a medicine student (*med*), intends to assist medical staff in basic medical care.

The semantics of extended belief operators \mathbb{B}_e , with $e \in \mathcal{E}$, is provided in Table 3. Notice that for $i \in Ag \cup Rl$, $\mathbb{B}_i()$ is defined as in Table 2 since in these cases, \mathcal{B}_i is a belief base in the sense assumed in Section 5.

The following example illustrates belief shadowing.

Example 4 (Example 3 continued). Assume a large-scale natural disaster occurred, resulting in the displacement of a large number of evacuees. As a medicine student, Eve joined a volunteer team that was formed at her university to support evacuees in everyday matters or assist medical staff in basic medical care. The belief base associated with the role volunteer (*vln*) consists of two world aspects:

$$\{support, \neg assist, travel\}, \{\neg support, assist, travel\}. \quad (18)$$

When Eve takes the role *vln*, her beliefs are expressed by $\mathbb{B}_{eve \text{ as } vln}()$. For example, though Eve believes she should not travel, her belief as a volunteer, $\mathbb{B}_{eve \text{ as } vln}(travel)$, obtains the value \mathfrak{t} , since the volunteer belief base consisting of world aspects (18) shadows Eve's original beliefs. Using the contents of (18), the beliefs $\mathbb{B}_{eve \text{ as } vln}(support)$ and $\mathbb{B}_{eve \text{ as } vln}(assist)$ obtain the value \mathfrak{i} .

Consider now a role of a medical assistant (*med*) with its belief base containing only the world aspect:

$$\{\neg support, assist, travel\}. \quad (19)$$

According to the contents of (19):

- the value of $\mathbb{B}_{eve \text{ as } med}(assist)$ and $\mathbb{B}_{eve \text{ as } vln \text{ as } med}(assist)$ become \mathfrak{t} ;
- the value of $\mathbb{B}_{eve \text{ as } med}(support)$ and $\mathbb{B}_{eve \text{ as } vln \text{ as } med}(support)$ become \mathfrak{f} . \square

When shadowing takes place, beliefs of $\mathbb{B}_{a \text{ as } r_1 \text{ as } \dots \text{ as } r_m}(F)$ may be evaluated in any of structures $\mathbb{S}_a, \mathbb{S}_{r_1}, \dots, \mathbb{S}_{r_m}$. Indeed, when for $k < i \leq m$, the truth value of $\mathbb{B}_i(F)$ in \mathbb{S}_i is \mathfrak{u} , the value of $\mathbb{B}_k(F)$ in \mathbb{S}_k is examined. Therefore goal-desires and desires-intentions bridging functions have to be prepared for belief bases no longer understood as sets of world aspects, but as more general servers whose role is to evaluate and return answers to a queries asked to a sequence of structures.

Definition 8 (General belief bases). Let $\mathcal{F}^{bel} \subseteq \mathcal{F}^{ext}$ be a set of extended formulas where extended operators are of the form $\mathbb{B}_e()$, for $e \in \mathcal{E}$. By a general belief base we mean any mapping $\mathcal{F}^{bel} \rightarrow \mathcal{T}$ assigning truth values to extended BDI formulas in \mathcal{F}^{bel} . The set of general belief bases is denoted by \mathbb{B}^* . \square

Though Definition 8 is general, in the sequel we deal with belief bases adequate for shadowing expressions. They can be seen as sequences $\langle \mathcal{B}_1, \dots, \mathcal{B}_k \rangle$, reflecting shadowing expressions, where $k \geq 1$ and $\mathcal{B}_1, \dots, \mathcal{B}_k$ are belief bases. Their semantics follows from Table 3.

The *extended 4BDI frames and structures* are defined as follows.

Definition 9 (Extended 4BDI frames and structures). An extended 4BDI frame is a pair $\mathbb{F}^* = \langle \mathcal{G}^*, \mathcal{A}^* \rangle$, where:

- $\mathcal{G}^* : \mathbb{B}^* \times \mathcal{F}^I \rightarrow \mathbb{C}$ is a goal-desires bridging function (takes an extended belief base and a goal as arguments and returns a desire base);
- $\mathcal{A}^* : \mathbb{B}^* \times \mathbb{C} \rightarrow \mathbb{C}$ is a desires-intentions bridging function (takes an extended belief base and a desire base as arguments and returns an intention base),

such that the requirements formulated in points 1 and 2 of Definition 4 are satisfied for \mathcal{D} and \mathcal{A} substituted respectively by \mathcal{G}^* and \mathcal{A}^* .

An extended 4BDI structure for a goal $G \in \mathcal{F}^I$ over \mathbb{F}^* and \mathcal{B}^* is a triple $\mathbb{F}^*(G) \stackrel{\text{def}}{=} \langle \mathcal{B}^*, \mathcal{D}^*, \mathcal{I}^* \rangle$, defined as in Definition 5, replacing \mathcal{B} by \mathcal{B}^* . \square

While belief bases are directly shadowed, desire and intention bases are shadowed indirectly, since:

1. \mathcal{D}^* is constructed using the shadowed belief base \mathcal{B}^* ;
2. \mathcal{I}^* is constructed using the shadowed belief base \mathcal{B}^* and the desire base \mathcal{D}^* .

Extended 4BDI frames and structures behave like 4BDI frames and structures in the sense of Definitions 4, 5. In particular, Proposition 1 holds for extended operators, too.

The following example illustrates shadowing of 4BDI structures.

Example 5 (Example 4 continued). Desire bases for ‘*eve as vln*’ and ‘*eve as med*’ are computed using a goal and beliefs provided by the belief base querying interface, i.e., using beliefs of the form $\mathbb{B}_{eve\ as\ vln}()$ and $\mathbb{B}_{eve\ as\ med}()$, respectively. Intention bases are to be computed from beliefs and desires, so in addition to $\mathbb{B}_{eve\ as\ \dots}()$, one uses desire operators of the form $\mathbb{D}_{eve\ as\ vln}()$ and $\mathbb{D}_{eve\ as\ med}()$. After computing intention bases, querying about intentions will use intention operators, like $\mathbb{I}_{eve\ as\ vln}()$ and $\mathbb{I}_{eve\ as\ med}()$. Of course, the operators $\mathbb{B}_{eve\ as\ \dots}()$, $\mathbb{D}_{eve\ as\ \dots}()$ and $\mathbb{I}_{eve\ as\ \dots}()$ can occur as a part of a more complex query expressed by 4BDI formulas.

To illustrate Definition 9, consider a goal *support* \vee *help*. The value of *help* is unknown to Eve as well as to the roles *vln*, *med*. The desire and intention bases may be then computed using \mathcal{G}^* and \mathcal{A}^* , e.g., to contain the world aspect $\{\textit{support}, \textit{travel}\}$ and perhaps other literals from which one could extract a plan to achieve *support* = \mathfrak{t} and *travel* = \mathfrak{t} . Given that *help* and a plan to make it true is a part of a background knowledge available to \mathcal{G}^* and \mathcal{A}^* , another world aspect in the intention base could also contain *help*, such as being $\{\textit{help}, \textit{travel}\}$. \square

Observe that the agent’s 4BDI loop, discussed in Section 2, is realized both by 4BDI and extended 4BDI frames and structures: a frame delivers goal-desires and desires-intentions bridging functions which are then used to determine a 4BDI structure relevant for a given goal and current beliefs of the agent.

We assume that agents’ and roles’ mental attitudes are represented by (extended) 4BDI structures. That is, for each 4BDI agent and role $i \in Ag \cup Rl$, there is a goal G_i , an associated (extended) 4BDI frame \mathbb{F}_i and structure $\mathbb{S}_i \stackrel{\text{def}}{=} \mathbb{F}_i(G_i) = \langle \mathcal{B}_i^{(*)}, \mathcal{D}_i, \mathcal{I}_i \rangle$, where the superscript $(*)$ denotes a 4BDI or an extended 4BDI structure, as appropriate.

Finally, the semantics of $\mathbb{D}_e()$ and $\mathbb{I}_e()$, where $e \in \mathcal{E}$, is defined in Table 4, where $\mathcal{D}_{a\ as\ r_1\ as\ \dots\ as\ r_m}$ and $\mathcal{I}_{a\ as\ r_1\ as\ \dots\ as\ r_m}$ are a desire base and an intention base constructed as in Definition 9 with $\mathcal{B}^* = \langle \mathcal{B}_a, \mathcal{B}_{r_1}, \dots, \mathcal{B}_{r_m} \rangle$. To broaden the definitions provided in Tables 3–4 for all extended 4BDI formulas, successively eliminate 4BDI subformulas by substituting them with truth values they evaluate to (along the lines of Algorithm 1).

$$\boxed{\begin{array}{l} \mathbb{S}(\mathbb{D}_{a \text{ as } r_1 \text{ as } \dots \text{ as } r_m}(F)) \stackrel{\text{def}}{=} \text{lub}_{\leq_i} \{w(F) \mid w \in \mathcal{D}_{a \text{ as } r_1 \text{ as } \dots \text{ as } r_m}\}; \\ \mathbb{S}(\mathbb{I}_{a \text{ as } r_1 \text{ as } \dots \text{ as } r_m}(F)) \stackrel{\text{def}}{=} \text{lub}_{\leq_i} \{w(F) \mid w \in \mathcal{I}_{a \text{ as } r_1 \text{ as } \dots \text{ as } r_m}\}. \end{array}}$$

Table 4. Semantics of extended \mathbb{D} and \mathbb{I} operators, where $F \in \mathcal{F}^I$, $a \in Ag$, $r_1, \dots, r_m \in Rl$ and $\mathbb{S} = \langle \mathbb{S}_i \mid i \in Ag \cup Rl \rangle$.

7 Complexity of the Approach

In the rest of this section, by complexity, we mean the data complexity of evaluating queries on the involved structures. That is, we assume that the query is fixed, and the complexity is expressed in terms of the size of the involved databases.

World aspects are sets of literals. Data complexity of first-order queries on sets of literals (thus world aspects, too) is PTIME and LOGSPACE [1] wrt the cardinality of the sets. Let $|w|$ and $|c|$ be the cardinalities of world aspect w and cluster c , respectively. By the size of (extended) 4BDI structure $\mathbb{S} = \langle \mathcal{B}^{(*)}, \mathcal{D}, \mathcal{I} \rangle$, denoted by $|\mathbb{S}|$, we mean $|\mathcal{B}^{(*)}| + |\mathcal{D}| + |\mathcal{I}|$. For a tuple of 4BDI structures $\bar{\mathbb{S}}$, we define $|\bar{\mathbb{S}}| \stackrel{\text{def}}{=} \sum_{\mathbb{S} \in \bar{\mathbb{S}}} |\mathbb{S}|$. Note that \mathcal{G} and \mathcal{A} of (extended) 4BDI frames are used to construct 4BDI structures, not being involved in further query evaluation.

Let $\mathbb{O} \in \{\mathbb{B}, \mathbb{D}, \mathbb{I}\}$, $i \in Ag \cup Rl$ and $F \in \mathcal{F}^I$ be a plain formula. Then evaluating a query $\mathbb{O}_i(F)$ on a cluster c is $\mathcal{O}(|c| * p(\max_{w \in c} |w|))$, where $p(n)$ is the complexity of evaluating F on a set of literals of cardinality n . Using Algorithm 1, we obtain the following results.

Theorem 1 (Tractability of querying (extended) 4BDI structures). *For every (extended) 4BDI formula $A \in \mathcal{F}^{ext}$ and a tuple of (extended) 4BDI structures $\bar{\mathbb{S}}$, evaluating the query expressed by A on $\bar{\mathbb{S}}$ is in PTIME and LOGSPACE wrt $|\bar{\mathbb{S}}|$.* \square

That is, in terms of data complexity, reasoning over (extended) 4BDI structures and their shadowing is tractable wrt the size of the structures involved.

Notice that tractability is achieved wrt the size of the involved 4BDI structures. Of course, 4BDI structures can be specified using rule-based languages, like variants of DATALOG^{??} or 4QL with tractable model computation. Since SQL is a very well tested and efficient technology and our queries can easily be translated to SQL, we look at the computed models as SQL databases. That way we inherit LOGSPACE complexity of first-order/SQL queries, what does not have to hold for rule languages, where computing models typically requires PSPACE. This feature is important in applications. For example, given that a 4BDI structure is huge, even of the size like 10^{80} (being the number of atoms in the observable universe), $\log(10^{80}) = 80 * \log(10) < 270$, so the additional space to compute queries is negligible. Of course, having rule-based specification is desirable and convenient in many scenarios. However, when queries are much more frequent than updates, one can compute the models of the most used parts of rule-based programs after each substantial portion of updates, store models as SQL databases, and expose them to queries enjoying a more efficient evaluation engine.

Complexity of constructing (extended) 4BDI structures from (extended) 4BDI frames and goals depends on the complexity of frame functions \mathcal{G} and \mathcal{A} . In general,

these functions may be of exponential (or higher) complexity, reducible to acceptable levels in real-world scenarios. For example, instead of generating all desired/intended world aspects, one can focus on a sufficient number or the most important ones.

8 Summary of Design Choices

8.1 Querying vs Reasoning

One of primary goals of this research was to develop efficient reasoning machinery by identifying the BDI components that enjoy tractability of reasoning. The natural candidates for such components are belief, intention and desire bases. Of course, even in the most basic instance of classical propositional logic satisfiability is NPTIME and validity checking is CO-NPTIME . On the other hand, in many real-world applications one does not expect agents or robots to derive advanced conclusions using proof systems or SAT solvers when they reason about the environment, current beliefs, desires and intentions once they are established and stored in databases. Therefore, we decided to chose reasoning-by-querying and use database techniques with tractable model and query evaluation. For the sake of efficient query answering we emphasized the possibility of using SQL database management systems. However, our tractability result also remains true whenever one uses rule-based languages, e.g. belonging to the DATALOG^\neg or 4QL families of tractable formalisms. Of course, as discussed in Section 7, LOGSPACE is retained when rule-based specifications are compiled into SQL or other representations of first-order models.

To summarize, the possibility to compile belief, desire, and intention bases to standard database management systems with tractable query evaluation, is a minimal prerequisite for tractability.

8.2 Paracompleteness and Paraconsistency

In order to express the desired phenomena in real-world applications, we have used a four-valued formalisms with non-classical truth values representing the lack of knowledge or inconsistency of the expressed properties. As indicated in [16], to define shadowing one needs at least the three-valued logic of Kleene, K_3 [26] with the third value representing unknown. On the other hand, when a part of a belief (desire, intention) base is shadowed by superior beliefs (desires, intentions), inconsistencies are practically unavoidable. One could unify truth values \mathfrak{u} and \mathfrak{i} and use K_3 . In fact, \mathfrak{i} behaves similarly to \mathfrak{u} : inconsistency of a formula identifies that one information source claims validity while another one reveals falsity of the formula. That is, at a meta-level it is *unknown* which source is right. Therefore, three-valued logical connectives have the same semantics no matter whether the third truth value is \mathfrak{i} or \mathfrak{u} (see, e.g., [35]). In order to distinguish those cases, in addition to \mathfrak{u} , in 4BDI we have employed \mathfrak{i} which is additionally useful:

- for modeling convenience;
- to simplify the 4QL formalism;
- to simplify and strengthen the querying machinery (see, e.g., [30,17]).

Another important phenomena addressed by para-completeness and para-consistency are related to actions and planning with incomplete and/or inconsistent information, where one may need to allow inconsistent beliefs, goals, desires and intentions (see also Remark 2).

All in all, para-complete formalisms could, in principle, suffice to define shadowing. However, without para-consistency, several challenges related to expressiveness, modeling convenience and reasoning would arise.

8.3 Shadowing vs Revision

As known from the literature, belief revision/update is typically a computationally demanding task [27]. Among others, the complexity sources are related to inconsistency avoidance and deciding which parts of databases should be replaced/repaid in response to incoming contradictory data. On the other hand, agents often do not need to deeply revise their mental attitudes. Sometimes it suffices to temporarily accept superior attitudes and act accordingly. This happens when an agent joins a group or takes on a particular role. Of course, in such cases one can expect that inconsistencies will occur. However, using a para-consistent formalism agents can utilize (partially) inconsistent databases without trivializing the reasoning and obtained conclusions. Tools to react on inconsistencies using nonmonotonic/heuristic rules can also be used [17,29,30].

In summary, in order to allow agents to live with inconsistencies and react appropriately, in 4BDI we have chosen a four-valued para-consistent formalism.

8.4 Resolving Incompleteness and Inconsistencies

In the research we focused on SQL-like implementations of 4BDI databases. However, as 4BDI is close to rule-based technologies, we expect that such SQL implementations will routinely store models of beliefs, desires and intentions computed by rule-based engines. During the model computation incomplete/inconsistent information can be resolved using nonmonotonic rules. Adequate tools are available, e.g., in 4QL's module querying mechanism. They can be invoked according to a chosen strategy. As discussed in [18], as regards timing "[...] there are at least three strategies:

- killing inconsistency at the root: to solve them as soon as possible;
- living with inconsistency: to postpone disambiguation to the last possible moment (or even forever);
- intermediate: to solve inconsistency each time new relevant information appears."

In 4BDI the first strategy can be applied when building the model. The others demand model recomputation when necessary or when new information becomes available.

9 Related Work

Since the pioneering work of Bratman [8], BDI-based models have been intensively studied both from application-oriented and formal point of view. There is a broad literature discussing applied aspects of BDI (see [7,10] and numerous references there). On

the formal side, many logics supporting BDI reasoning have also been developed and investigated [11,15,19,24,25,28,31,32,33,37,42]. They primarily employ multi-modal logics with relatively high reasoning complexity. While the majority of these approaches focus on plans, actions, and temporal aspects, we abstract from them and bring pre-existing beliefs, desires and intentions into focus. Their construction is encapsulated in goal-desires and desires-intentions bridging functions, constituting essential components of 4BDI frames. The shift from general logical reasoning to the database perspective [17,38] enables us to significantly reduce the high complexity of general logical reasoning to tractable reasoning-by-querying.

Paraconsistent reasoning is a well-established area (see, e.g., [4,12] and references there). Even though inconsistency in modern AI applications is omnipresent, paraconsistent BDI models have not received sufficient attention. This aspect is pointed out in [24], where paraconsistent approach to BDI, specifically to desires, is identified as a challenging research area. In [41] the authors discuss paraconsistent logic, where BDI serves as a motivation, however focusing on the classical language. In 4BDI we use paraconsistent databases and reasoning-by-querying at all levels of mental attitudes, where queries are expressed in a BDI specific language. To the best of our knowledge, 4BDI is the first paraconsistent BDI formalism.

As agent systems act in dynamic environments, belief update and revision are in the mainstream of the area. For representative approaches see [20,34,36,39] and references there. However, belief revision/update is typically a computationally demanding task [27]. In contrast to other approaches, the paper [5] provides tractable logical formalism to deal with shallow and possibly transient belief change. Without losing tractability, 4BDI substantially extends this formalism to deal with shadowing desires and intentions, too.

10 Conclusions

The research reported in the paper introduces a novel tractable and implementation friendly model of beliefs, desires and intentions, 4BDI, developed from the database perspective. In particular, we have defined 4BDI frames and structures designed for modeling paraconsistent and paracomplete reasoning about BDI agents facing imperfect information. In comparison to other approaches, the formalism is conceptually light. It offers means for BDI revision via shadowing. This provides a level of expressiveness not available in other approaches. Importantly, a shift from general reasoning to reasoning-by-querying, renders the approach tractable. To the best of our knowledge, the reported expressive power while retaining tractability has never been achieved before. What is also important, 4BDI unifies both theoretical and practical reasoning in a single approach. This is accomplished by a uniform representation of beliefs, desires and intentions, combined with querying and shadowing as closely related logic-based mechanisms.

The 4BDI framework is implementation-oriented. Indeed, belief, desire and intention bases can almost directly be implemented with SQL or rule-based database management systems or in rule-based languages. Though a single query may refer to many

agents, 4BDI structures are meant to be allocated to individual agents and roles. Therefore, 4BDI suits better to distributed rather than centralized architectures.

Future work will include a composition of 4BDI group structures from structures of group members, aiming for a uniform treatment of individual agents and arbitrarily nested groups. This task requires a nontrivial construction of goal-desire and desire-intention bridging functions at the group level. A separate research may concern intention refinement using 4BDI structures and a suitable action and change theory. The approaches discussed in [24] combined with actions over paraconsistent belief bases presented in [6] could be good starting points in this direction.

Acknowledgments

This work has been supported by grant 2015/19/B/ST6/02589 of the National Science Centre Poland. We would also like to thank reviewers for constructive comments.

References

1. Abiteboul, S., Hull, R., Vianu, V.: Foundations of Databases. Addison-Wesley (1996)
2. Arieli, O., Avron, A.: The value of the four values. *Artif. Intell.* **102**(1), 97–141 (1998)
3. Belnap, N.: A useful four-valued logic. In: Epstein, G., Dunn, J. (eds.) *Modern Uses of Many Valued Logic*. pp. 5–37. Reidel (1977)
4. Béziau, J.Y., Carnielli, W., Gabbay, D. (eds.): *Handbook of Paraconsistency*. College Pub. (2007)
5. Białek, Ł., Dunin-Kępicz, B., Szałas, A.: Belief shadowing. In: Mascardi, V., Ricci, A., Weyns, D. (eds.) *Engineering Multi-Agent Systems. EMAS 2018. LNCS*, vol. 11375, pp. 158–180. Springer (2019)
6. Białek, Ł., Dunin-Kępicz, B., Szałas, A.: A paraconsistent approach to actions in informationally complex environments. *Ann. Math. Artif. Intell.* **86**(4), 231–255 (2019)
7. Bordini, R., El Fallah Seghrouchni, A., Hindriks, K., Logan, B., Ricci, A.: Agent programming in the cognitive era. *Auton. Agents Multi Agent Syst.* **34**(2), 37 (2020)
8. Bratman, M.: *Intention, Plans, and Practical Reason*. Harvard Univ. Press (1987)
9. Bratman, M., Israel, D., Pollack, M.: Plans and resource-bounded practical reasoning. *Comput. Intell.* **4**, 349–355 (1988)
10. Cardoso, R., Ferrando, A.: A review of agent-based programming for multi-agent systems. *Computers* **10**(2), 16 (2021)
11. Cohen, P., Levesque, H.: Intention is choice with commitment. *Artif. Intelligence* **42**(3), 213–261 (1990)
12. da Costa, N., Béziau, J., Bueno, O.O.: On the usefulness of paraconsistent logic. In: Vanderveken, D. (ed.) *Logic, Thought and Action, LEUS*, vol. 2, chap. 20, pp. 465 – 478. Springer (2005)
13. deAmo, S., Pais, M.: A paraconsistent logic programming approach for querying inconsistent databases. *Int. J. Approx. Reason.* **46**(2), 366–386 (2007)
14. d’Inverno, M., Luck, M.: *Understanding Agent Systems*. Springer (2001)
15. d’Inverno, M., Luck, M., Georgeff, M., Kinny, D., Wooldridge, M.: The dMARS architecture: A specification of the distributed multi-agent reasoning system. *Auton. Agents Multi Agent Syst.* **9**(1-2), 5–53 (2004)
16. Dunin-Kępicz, B., Szałas, A.: Shadowing in many-valued nested structures. In: *Proc. 50th IEEE Int Symp. on Multiple-Valued Logics*, pp. 230–236. IEEE (2020)

17. Dunin-Kępicz, B., Szałas, A.: Taming complex beliefs. *Trans. on Comp. Collective Intelligence XI LNCS* **8065**, 1–21 (2013)
18. Dunin-Kępicz, B., Szałas, A., Verbrugge, R.: Tractable reasoning about group beliefs. In: Dalpiaz, F., Dix, J., van Riemsdijk, M. (eds.) *EMAS'2014. LNCS*, vol. 8758, pp. 328–350. Springer (2014)
19. Dunin-Kępicz, B., Verbrugge, R.: *Teamwork in Multi-Agent Systems. A Formal Approach*. John Wiley & Sons, Ltd. (2010)
20. Fermé, E., Hansson, S.O.: Agm 25 years: Twenty-five years of research in belief change. *J. of Philosophical Logic* **40**(2), 295–331 (2011)
21. Gebser, M., Kaminski, R., Kaufmann, B., Schaub, T.: *Answer Set Solving in Practice. Synthesis Lectures on AI and ML*, Morgan and Claypool Pub. (2012)
22. Gelfond, M., Kahl, Y.: *Knowledge Representation, Reasoning, and the Design of Intelligent Agents - The Answer-Set Programming Approach*. Cambridge University Press (2014)
23. Georgeff, M., Lansky, A.: Reactive reasoning and planning. In: Forbus, K., Shrobe, H. (eds.) *Proc. 6th National Conf. on AI*. pp. 677–682. Morgan Kaufmann (1987)
24. Herzig, A., Lorini, E., Perrussel, L., Xiao, Z.: BDI logics for BDI architectures: Old problems, new perspectives. *Künstliche Intell.* **31**(1), 73–83 (2017)
25. van der Hoek, W., Wooldridge, M.: Logics for multiagent systems. *AI Mag.* **33**(3), 92–105 (2012)
26. Kleene, S.: On notation for ordinal numbers. *J. Symbolic Logic* **3**, 150–155 (1938)
27. Liberatore, P.: The complexity of belief update. *Artif. Intell.* **119**(1), 141 – 190 (2000)
28. Lorini, E., Herzig, A.: A logic of intention and attempt. *Synthese* **163**(1), 45–77 (2008)
29. Małuszynski, J., Szałas, A.: Living with inconsistency and taming nonmonotonicity. In: de Moor, O., Gottlob, G., Furche, T., Sellers, A. (eds.) *Datalog Reloaded. LNCS*, vol. 6702, pp. 334–398. Springer (2011)
30. Małuszynski, J., Szałas, A.: Partiality and inconsistency in agents' belief bases. In: Barbucha, D., Le, M., Howlett, R., Jain, L. (eds.) *KES-AMSTA. Frontiers in AI and Applications*, vol. 252, pp. 3–17. IOS Press (2013)
31. Meyer, J.J., de Boer, F., van Eijk, R., Hindriks, K., van der Hoek, W.: On programming KARO agents. *Log. J. IGPL* **9**(2), 245–256 (2001)
32. Meyer, J., Broersen, J., Herzig, A.: BDI logics. In: van Ditmarsch, H., Halpern, J., van der Hoek, W., Kooi, B. (eds.) *Handbook of Logics of Knowledge and Belief*, chap. 10, pp. 453 – 498. College Pub. (2015)
33. Novák, P., Dix, J.: Modular BDI architecture. In: Nakashima, H., Wellman, M., Weiss, G., Stone, P. (eds.) *5th Int. Joint Conf. AAMAS*. pp. 1009–1015. ACM (2006)
34. Peppas, P.: Belief revision. In: van Harmelen, F., Lifschitz, V., Porter, B. (eds.) *Handbook of KR*, pp. 317–359. Elsevier (2008)
35. Priest, G.: The logic of paradox. *Journal of Philosophical Logic* **8**(1), 219–241 (1979)
36. Priest, G.: Paraconsistent belief revision. *Theoria* **67**(3), 214–228 (2001)
37. Rao, A., Georgeff, M.: Decision procedures for BDI logics. *J. Logic and Computation* **8**(3), 293–344 (1990)
38. Shoham, Y.: Logical theories of intention and the database perspective. *J. Philos. Log.* **38**(6), 633–647 (2009)
39. Testa, R.R., Coniglio, M.E., Ribeiro, M.M.: AGM-like paraconsistent belief change. *Logic Journal of the IGPL* **25**(4), 632–672 (2017)
40. Urquhart, A.: Many-Valued Logic. In: Gabbay, D., Guentner, F. (eds.) *Handbook of Philosophical Logic*. vol. 3, pp. 71–116. Reidel (1986)
41. Villadsen, J.: Paraconsistent assertions. In: Lindemann, G., Denzinger, J., Timm, I., Unland, R. (eds.) *Proc. MATES: 2nd German Conf. Multiagent System Technologies. LNCS*, vol. 3187, pp. 99–113. Springer (2004)
42. Wooldridge, M.: *Reasoning about Rational Agents*. John Wiley & Sons, Inc. (2003)