

Beliefs, Obligations, Intentions, and Desires as Components in an Agent Architecture

Jan Broersen Mehdi Dastani Leendert van der Torre

December 24, 2004

Abstract

In this paper we discuss how cognitive attitudes like beliefs, obligations, intentions and desires can be represented as components with input/output functionality. We study how to break down an agent specification into a specification of individual components and a specification of their coordination. A typical property discussed at the individual component specification level is whether the input is included in the output, and a typical property discussed at the coordination level is whether beliefs override desires to ensure realism. At the individual level we show how proof rules of so-called input/output logics correspond to properties of functionality descriptions, and at the coordination level we show how global constraints coordinating the components formalize coherence properties.

1 Introduction

The structured approach to the specification of agent systems based on a notion of compositional architecture has been called *compositional design*, and an example of a compositional design method is the DEsign and Specification of Interacting REasoning components (DESIRE) developed by Treur and colleagues [7]. The reasoning process is structured according to a number of reasoning components that interact with each other. Components may or may not be composed of other components, where components that cannot be decomposed are called primitive components. The functioning of the overall agent system is based on the functionality of these primitive components plus the composition relation that coordinates their interaction. Specification of a composition relation may involve, for example, the possibilities for information exchange among components and the control that activates the components. The functionality of a (primitive) reasoning component within a compositional architecture can be described by Treur's functionality descriptions [33]. Typical properties are whether the input of a component is included in its output, or whether it supports reasoning by cases.

Our research question is how to represent cognitive attitudes with components such that an agent specification is broken down into a specification of individual components and the specification of their coordination (following a divide and conquer strategy). In this paper we relate properties of Treur's functionality descriptions to the validity of

proof rules in Makinson and van der Torre’s input/output logics [25, 26]. These logics have been developed as a general theory of propositional input/output operations as processes resembling inference, but where input propositions are not in general included among outputs, and the operation is not in any way reversible. We show that an input/output logic, characterized by a set of proof rules, specifies the input/output behavior of reasoning components as described by the functionality descriptions of those components. Functionality descriptions can be specified in input/output logic, thus enabling a *logical* specification of reasoning components. The correspondence between Treur’s functionality descriptions and the input/output logics relates functionality descriptions to cognitive agent architectures, because, according to Makinson and van der Torre, examples of the input/output operations they study “arise in contexts of conditional obligations, goals, ideals, preferences, actions, and beliefs” [25].

We also study the coordination of components that represent cognitive attitudes. The coherence among cognitive components can be described by various kinds of coordination principles. For example, it may make sense for belief sets to include the observations, but it makes less sense for obligation or desire sets to include the beliefs. You may desire to go to the dentist and believe that going to the dentist implies pain, but this does not imply the desire for pain [6]. Other coordination principles are encoded in the overall conflict resolution strategy in the architecture’s control specification. A typical property is realism, which has been formalized as the overriding of desires by beliefs [32, 9].

Our motivation to study how cognitive attitudes can be allocated in an agent architecture is to give formal foundations for several agent architectures such as Castelfranchi *et al.*’s deliberative normative agent architecture [11] and Broersen *et al.*’s BOID architecture [8]. The BOID architecture extends Thomason’s BDP logic [32], based on a planning component and a conflict resolution component for conditional beliefs and desires, with conditional obligations and prior intentions, borrowed from respectively deontic action programs [17] and BDI systems [12, 27].

The layout of this paper is as follows. In Section 2 we discuss the design of component-based agents, with the BOID agent as an example. In Section 3 we consider the specification of individual components, and in Section 4 we consider the coordination of the components.

2 Designing component-based agents

In the study of architectures for autonomous agents, cognitive attitudes have been represented as components. A typical example of this allocation is the BOID architecture [8], which contains components outputting beliefs (B), obligations (O), intentions (I) and desires (D). In this section we discuss how to design a component-based BOID agent. Design problems are iteratively replaced by a number of smaller design problems together with the problem of how the solutions are composed based on coordination principles, which leads to a transparent structure of the design and to support for reuse and maintainability of components and generic models.

Figure 1 visualizes the general goal-based agent architecture studied in [14], that consists of a goal generation and a plan generation component. The agent can be

seen as a black box with observations (obs) as input and intended plans or actions as output. As usual, the agent is assumed to sense the environment by detectors and execute actions by effectors. The design problem of the agent is now reformulated as designing the goal generation component, the plan generation component, and their coordination. This design suggests sequential composition, which may or may not be possible for a specific problem or agent. We therefore do not restrict ourselves to sequential composition, and also allow for other compositional relations.

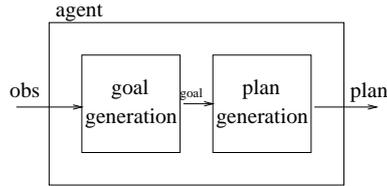


Figure 1: Goal-based agent

Moreover, Figure 2 visualizes the goal generation component studied in [8, 9], which contains the four components *beliefs* (B), *obligations* (O), *intentions* (I), and *desires* (D). The design problem of the goal generation module decomposes into the problem of how to model the input/output behavior of the agent’s cognitive attitudes, and the problem of their interaction. Again we do not restrict ourselves to sequential composition. This is visualized by feedback loops, which have subtle implications.

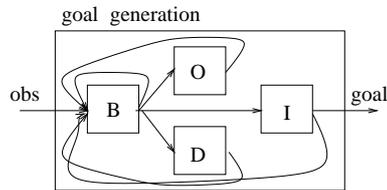


Figure 2: Goal generation component of goal-based agent

Finally, these four cognitive components are primitive components, because they are not further broken down. We could have further distinguished among distinct types of beliefs, distinct types of obligations, etc. The components can be implemented in a variety of ways. For example, the belief component may maximize cross entropy or apply AGM belief revision [3], and its output may be a probability distribution, a set of them, plausibility measures, a belief set, etc. In a logic-based goal generation component like the BOID architecture, the input (the observations) and the output (the goals) are sets of logical formulas. The behavior of each component as well as the interaction among the components are described as follows.

Component specification. Each component can be described by (or may even consist of) a set of rules – pairs of logical formulas – and the input and the output of each component are sets of formulas. For example, the belief component may be

described by the set of rules $\{\top \leftrightarrow \text{rain}, \text{rain} \leftrightarrow \text{wet}\}$ representing the beliefs that ‘it rains’ and ‘if it rains, then streets are wet’. The belief component outputs $\{\text{rain}\}$ without any input (empty observation, \emptyset) and it outputs $\{\text{rain}, \text{wet}\}$ if it receives $\{\text{rain}\}$ as input. In the BOID architecture, the output of a BOID component is its input together with the body of applicable rules.

Component coordination specification. At any point in time, the feedback connections guarantee that the components can receive outputs, which are generated so far, in order to generate new outputs. For example, in Figure 2, there is a feedback connection from the output interface of the belief component to its input interface guaranteeing that the output of the belief component is used to generate new belief outputs. Similarly, the output interface of O, I, and D components are connected to the input interface of the B component by feedback connections. The outputs of O, I, and D components are sent back to the belief component through these feedback connections.

A tool for the design of individual components is Treur’s so-called functionality description [33], described in more detail in Section 3, which according to Treur is independent of the component’s specific internal knowledge representation, inference relations or implementation. The component’s internal structure can be changed as long as its functionality remains the same. In this way a functionality description not only obtains a clear and well-defined analysis of the component’s behavior, but it also supports reusability and maintainability, as well as information hiding.

A typical example of a phenomenon which can be defined on the individual component level as well as on the composition level is the resolution of conflicts. It is well known from the work on BDI agents that the type of an agent determines its behavior, we can thus refer to it as an abstract behavior type [4]. The agent type is governed by the specific way it handles the rational balance among its cognitive attitudes such as beliefs, obligations, intentions and desires. For example, the type of an agent determines how it resolves conflicts within and among its cognitive attitudes. In the BOID architecture the rational balance such as resolving conflict among cognitive attitudes is determined based on the properties of their architectures which are in turn defined in terms of data representation and a reasoning mechanism. Two types of conflicts can be distinguished in the goal generation component in Figure 2.

Component specification: conflicts within a component. Conflicts and their resolutions are analyzed within the component itself and therefore not visible from the outside of the components. The conflict resolution behavior within a component can be specified in terms of properties of functionality descriptions, i.e., properties of input/output behavior of the component.

Component coordination specification: conflicts among components. Conflicts can be formalized as conflicts among outcomes of functionality descriptions. These conflicts are addressed in the possibly dynamic composition relation. There are six total orders, the number in which the three letters O/I/D can be ordered, which expresses that:

- The agents are realistic and the belief therefore overrides all other mental attitudes.
- Intentions may override desires and obligations (stable agents), or the other way around (unstable agents).
- Desires may override obligations (selfish), or the other way around (social).

Besides the six total orders there are also many partial orders, in which for example there is no general preference of desires over obligations or vice versa.

For further information on the BOID project see the BOID's home page [5].

3 Individual component specification

In Section 3.1 we discuss the use of functionality descriptions in compositional design to formalize the functionality of a primitive component within a compositional architecture, and in Section 3.2 we relate this notion to proof rules of input/output logics.

3.1 Functionality descriptions ...

As part of the foundations of compositional architectures the notion of functionality description has been introduced in [33]. A functionality description is the set of input/output combinations provided by the component, that

- gives a *semantical description* of the reasoning that provides conclusions ‘for the outside’ – for other components outside the reasoning component – as output, but at the same time uses information ‘from the outside’ as input. In this sense we can say that such functionality describes the *interactive* role of a given component.
- may be seen as a *declarative description*, in the sense that a functionality description abstracts from the dynamic aspects of the component's reasoning process. It only describes what facts can be derived in a given situation, and not at what time and in which order specific facts actually are derived.

The remainder of this section repeats some definitions of [33]. The signature is based on a distinction between input, intermediate and output signature, which are all assumed to be finite.

Definition 1 ([33], Def. 3.1.) *A propositional signature Σ is a three-tuple, $\langle InSig(\Sigma), IntSig(\Sigma), OutSig(\Sigma) \rangle$, where $InSig(\Sigma)$, $IntSig(\Sigma)$, and $OutSig(\Sigma)$ are ordered sets of atom names, respectively called the input signature, the internal signature, and the output signature. The input and output signature can contain common names, but the internal signature is disjoint from them. All signatures are assumed to be finite.*

The semantics is based on partial models. Reasoning components are able to draw partial conclusions if a partial input information state is given. Therefore the formal description of the declarative functionality of a component also treats partiality of information both at the input side and the output side of the component.

Definition 2 ([33]) A partial model M of signature Σ is an assignment of truth values from $\{1, 0, u\}$ to the atoms of Σ . By $M(a)$ we denote the truth value assigned to atom a . We call M a complete model if for all atoms a the truth value $M(a)$ is not u . Moreover, we call M a complete input model if for all atoms a of $InSig(\Sigma)$ the truth value of a is not u , but for all atoms a of $\Sigma \setminus InSig(\Sigma)$, the truth value of a is u . The model $Out(M)$ is obtained from M such that for all atoms a of $OutSig(\Sigma)$ the truth value of a is its truth value in M , but for all atoms a of $\Sigma \setminus OutSig(\Sigma)$, the truth value of a is u .¹

An important notion is that a model can refine another model of the same signature, which means that if the less refined model assigns true or false to propositions the more refined model assigns the same truth value to those propositions, but when the less refined model assigns unknown to propositions, the more refined model may assign true or false. This notion is used to define the greatest common information state, which may be seen as the maximal information on which all members of the set of information states agree. The greatest common information state of the set of information states $gci(V)$ can be constructed as follows: for any atom a on which all members of V agree, take this truth value, and if the members of V disagree, take the truth value u .

Definition 3 ([33]) The refinement relation \leq between partial models of the same signature is defined by $M \leq N$, if for every atom a it holds that $M(a) \leq N(a)$ (i.e., point by point), where the partial order of truth values is defined by $u < 1$ and $u < 0$. Let V be a non-empty set of partial models of signature Σ . By $P(V)$ we denote the set of all partial models that can be refined to a model in V . The greatest common information state of V is the partial model $N \in P(V)$, denoted by $N = gci(V)$, such that for all $M \in V$, it holds that $N \leq M$, and for any N' satisfying this condition, it holds that $N' \leq N$.

We have now given the necessary machinery to define functionality descriptions. Treur also defines additional properties of functionality descriptions, such as regularity, which we do not discuss in this paper.

Definition 4 ([33], Def. 4.1.) Suppose a signature Σ is given, a non-empty set of complete input models W_{in} for Σ , and a mapping $\alpha : P_{in} \rightarrow P$, where $P_{in} = P(W_{in})$ and P is the set of partial models for Σ .²

1. The mapping α is called conservative if for all $M \in P_{in}$ it holds that

$$M \leq \alpha(M)$$

2. The mapping α is called monotonic if for all $M, N \in P_{in}$ it holds that

$$M \leq N \Rightarrow \alpha(M) \leq \alpha(N)$$

¹Treur's definitions of complete input model and the model $Out(M)$ are obtained via reductions and expansions, but since we are only interested in the notions we need to define functionality descriptions, we do not give the details.

²Treur says " P is a set of partial models for Σ ".

3. The mapping α is called self-bounded if for all $M, N \in P_{in}$ it holds that

$$M \leq \alpha(N) \Rightarrow \alpha(M) \leq \alpha(N)$$

4. The mapping α is called well-informed if for all $M \in P_{in}$ it holds that

$$out(\alpha(M)) = gci(\{out(\alpha(N)) \mid N \in W_{in} \& M \leq N\})$$

A partial model can be represented by a finite set of propositional literals, and therefore by a propositional formula. In this representation, the refinement relation represents propositional derivability. For example, a complete model that assigns true (1) to a and false (0) to b refines a partial model that assigns unknown (u) to a and false (0) to b , because $a \wedge \neg b$ logically implies $\neg b$.

Moreover, in this interpretation we can interpret α as a kind of logical consequence operator. However, it is a consequence operator of an unusual kind, as α does not satisfy the Tarskian properties. For example, α is not necessarily conservative and the associated consequence relation therefore does not necessarily satisfy identity. In the following section we therefore relate functionality descriptions to input/output logic, as this logic does not necessarily satisfy the identity rule either.

3.2 ... and input/output logics

Input/output logic (IOL) [25, 26] is a theory of input/output operations resembling inference, but where input propositions are not in general included among outputs, and the operation is not in any way reversible. Makinson and van der Torre write *out* instead of *output*, but we write *output* to avoid confusion with Treur's notion of *out*. Moreover, Makinson and van der Torre define input and output of their operations as arbitrary formulas, but to represent the functionality descriptions, here we restrict ourselves to conjunctions of sets of literals. To facilitate the semantics of the output operations Makinson and van der Torre extend the generating set G with (\top, \top) such that it is never empty. Since we do not consider semantics of input/output logics in this paper we also do not consider this borderline case.

Definition 5 (Input/output logic) [25] *Let L be the fragment of propositional logic that contains all conjunctions of literals, \vdash derivability in L , G be a set of pairs of conjunctions of propositional literals $\{(a_1, x_1), \dots, (a_n, x_n)\}$, read as 'if input a_1 then output x_1 ', etc., and consider the following proof rules strengthening of the input (SI), weakening of the output (WO), conjunction for the output (AND), disjunction of the input (OR), cumulative transitivity (CT), transitivity (TRANS) and identity (Id) defined as follows:*

$$\begin{array}{cc} \frac{(a, x)}{(b, x)} b \vdash a \text{ (SI)} & \frac{(a, x)}{(a, y)} x \vdash y \text{ (WO)} \\ \frac{(a, x)(a, y)}{(a, x \wedge y)} \text{ (AND)} & \frac{(a \wedge b, x)(a \wedge \neg b, x)}{(a, x)} \text{ (OR)} \end{array}$$

$$\frac{(a, x)(a \wedge x, y)}{(a, y)} (CT) \quad \frac{(a, x)(x, y)}{(a, y)} (TRANS)$$

$$\frac{}{(a, a)} (Id)$$

Each subset of the above rules together with a ‘silent’ rules that replaces propositional formulas by logical equivalent ones, defines an output operation $output(G)$ as a closure operation on G . Makinson and van der Torre focus on the following eight output operators.

- $output_1$: $SI+AND+WO$ (simple-minded output)
- $output_2$: $SI+AND+WO+OR$ (basic output)
- $output_3$: $SI+AND+WO+CT$ (simple-minded reusable output)
- $output_4$: $SI+AND+WO+OR+CT$ (basic reusable output)
- out_i^+ : out_i+Id (throughput)

Moreover, a mapping from propositions to propositions is defined indirectly by $x \in output(G, a)$ iff $(a, x) \in output(G)$.

To relate the functionality description with input/output logic, we assume that if $N = \alpha(M)$, then M is a description of the input, and N is a *complete* description of the output. We ignore the fact that functionality descriptions have been restricted to a set of complete input models called a domain descriptions W_{in} .³

Definition 6 Let Σ be a propositional signature as introduced in Def. 1, L the set of conjunctions of literals from Σ , and L_{in} the set of conjunctions from literals built from $InSig(\Delta)$. For a partial model M , let $\tau(M)$ be the following conjunction of literals:

$$\tau(M) = \bigwedge \{a \mid M(a) = 1, a \in \Sigma\} \wedge \bigwedge \{\neg a \mid M(a) = 0, a \in \Sigma\}$$

The input/output relation associated with α is defined as follows. $(a, x) \in output_\alpha$ iff $\exists M \in P(W_{in}), N \in P$ such that $N = \alpha(M)$, $a \equiv \tau(M)$ and $\tau(N) \vdash x$. Moreover, a mapping from propositions to propositions is defined indirectly by $x \in output_\alpha(a) \Leftrightarrow (a, x) \in output_\alpha$.

The $output_\alpha$ operation is an input/output relation closed under the two rules AND and WO. It can be interpreted as a non-monotonic weakening of Makinson and van der Torre’s input/output logics, as compared to their systems the closure operator does not satisfy strengthening of the input.

Theorem 1 We have

1. $output_\alpha$ satisfies AND and WO.

³The extension of the results below for a given domain description is straightforward.

2. For any input/output operation satisfying AND and WO, there is a functionality description α such that this input/output operation is identical to $output_\alpha$.

Proof. (sketch). 1. Follows from Definition 6 where " $\tau(N) \vdash x$ " defines output as logically closed set. 2. follows from lack of properties of α , and can be proven by construction: assume an arbitrary output relation, and construct α .

Before we present our main theorem, we first relate the refinement relation on partial models to propositional derivability. For conservative, monotone and well-founded, we only need the simple properties of Lemma 1.

Lemma 1 We have:

1. $M \leq N$ iff $\tau(N) \vdash \tau(M)$
2. $M \leq \alpha(N)$ iff $\tau(M) \in output_\alpha(\tau(N))$
3. $\alpha(M) \leq \alpha(N)$ iff $output_\alpha(\tau(M)) \subseteq output_\alpha(\tau(N))$

Proof. Follows directly from the definition of τ in Definition 6.

For the property of well-informed we need an inductive argument, because it is based on the greatest common information state of a set of models.

Lemma 2 We have

1. $\forall M \in P_{in} : Out(\alpha(M)) \leq gci(\{Out(\alpha(N)) \mid N \in W_{in} \wedge M \leq N\})$
iff $\forall M \in P_{in} : Out(\alpha(M)) \leq gci(\{Out(\alpha(N)) \mid N \in P_{in} \wedge M \leq N\})$
2. $\forall M \in P_{in} : Out(\alpha(M)) \leq gci(\{Out(\alpha(N)) \mid N \in P_{in} \wedge M \leq N\})$
iff $\forall M, N \in P_{in} : Out(\alpha(M)) \leq Out(\alpha(N))$
3. $\forall M \in P_{in} : gci(\{Out(\alpha(N)) \mid N \in P_{in} \wedge M \leq N\}) \leq Out(\alpha(M))$ iff $\forall M \in P_{in} \forall p \in InSig(\Delta) : \tau(Out(\alpha(M))) \vdash output_\alpha(\tau(M) \wedge p) \vee output_\alpha(\tau(M) \wedge \neg p)$

Proof. (sketch). 1. and 2. follow directly from the definitions. Item 3. follows from fact that \leq is a semi-lattice, and gci is the infimum⁴ of this semi-lattice.

We now relate the four properties of functionality descriptions to proof rules of input/output logic. There is one final complication. Functionality descriptions are only defined on partial models. Consequently, $output_\alpha$ is only defined for *consistent* formulas, that is, sets of literals that do not contain a propositional atom and its negation.

⁴The infimum is $inf(V) = \{M \mid \forall N \in V, M \leq N, \nexists M' > M, \forall N \in V, M' \leq N\}$.

Theorem 2 *Let the set of consistent propositional formulas contain the conjunctions of sets of literals where these sets do not contain an atom and its negation. For any functionality description α , the following hold:*

1. α is conservative iff output_α satisfies identity (Id) for consistent input formulas.
2. α is monotonic iff output_α satisfies strengthening of the input (SI) for consistent input formulas.
3. α is self-bounded iff output_α satisfies transitivity (TRANS).
4. α is well-informed iff output_α satisfies the strengthening of the input (SI) and the disjunction rule (OR) for input formulas.

Proof.

Let L and L_{in} be set of propositional formulae as defined in Def. 6.

1. α is conservative iff output_α satisfies identity (Id) for consistent input formulas.

$$\frac{\frac{\forall M \in P_{in} : M \leq \alpha(M)}{\forall M \in P_{in} : \tau(M) \in \text{output}_\alpha(\tau(M))} \text{Lem.1}}{\frac{\forall a \in L_{in} : a \in \text{output}_\alpha(a)}{\forall a \in L_{in} : (a, a) \in \text{output}_\alpha} \text{Def.6}} \text{Def.6}$$

2. α is monotonic iff output_α satisfies strengthening of the input (SI) for consistent input formulas.

$$\frac{\frac{\frac{\forall M, N \in P_{in} : M \leq N \Rightarrow \alpha(M) \leq \alpha(N)}{\forall M, N \in P_{in} : \tau(N) \vdash \tau(M) \Rightarrow (\text{output}_\alpha(\tau(M)) \subseteq \text{output}_\alpha(\tau(N)))} \text{Lem.1}}{\frac{\forall a, b \in L_{in} : b \vdash a \Rightarrow (\text{output}_\alpha(a) \subseteq \text{output}_\alpha(b))} \text{Def.6}}{\forall a, b \in L_{in} : b \vdash a \Rightarrow (\forall x \in L : (a, x) \in \text{output}_\alpha \rightarrow (b, x) \in \text{output}_\alpha)} \text{Def.6}}$$

3. α is self-bounded iff output_α satisfies transitivity (TRANS)

$$\frac{\frac{\frac{\forall M, N \in P_{in} : M \leq \alpha(N) \Rightarrow \alpha(M) \leq \alpha(N)}{\forall M, N \in P_{in} : \tau(M) \in \text{output}_\alpha(\tau(N)) \Rightarrow (\text{output}_\alpha(\tau(M)) \subseteq \text{output}_\alpha(\tau(N)))} \text{Lem.1}}{\frac{\forall a, b \in L_{in} : b \in \text{output}_\alpha(a) \Rightarrow (\text{output}_\alpha(b) \subseteq \text{output}_\alpha(a))} \text{Def.6}}{\forall a, b \in L_{in} : (a, b) \in \text{output}_\alpha \Rightarrow (\forall x \in L : (b, x) \in \text{output}_\alpha \rightarrow (a, x) \in \text{output}_\alpha)} \text{Def.6}}$$

4. α is well-informed iff output_α satisfies the following two rules. First it satisfies the strengthening of the input (SI) for out formulas.

$$\begin{array}{l}
\frac{\forall M \in P_{in} : Out(\alpha(M)) \leq gci(\{Out(\alpha(N)) \mid N \in W_{in} \wedge M \leq N\})}{\forall M \in P_{in} : Out(\alpha(M)) \leq gci(\{Out(\alpha(N)) \mid N \in P_{in} \wedge M \leq N\})} \text{Lem.1} \\
\frac{\forall M, N \in P_{in} : M \leq N \Rightarrow Out(\alpha(M)) \leq Out(\alpha(N))}{\forall M, N \in P_{in} : \tau(N) \vdash \tau(M) \Rightarrow output_{\alpha}(\tau(M)) \subseteq output_{\alpha}(\tau(N))} \text{Lem.1} \\
\frac{\forall a, b \in L_{in} : b \vdash a \Rightarrow output_{\alpha}(a) \subseteq output_{\alpha}(b)}{\forall a, b \in L_{in} : b \vdash a \Rightarrow \forall x \in L : output_{\alpha}(a, x) \Rightarrow output_{\alpha}(b, x)} \text{Def.6}
\end{array}$$

Secondly it satisfies the disjunction rule (OR) for input formulas.

$$\begin{array}{l}
\frac{\forall M \in P_{in} : gci(\{Out(\alpha(N)) \mid N \in W_{in} \wedge M \leq N\}) \leq Out(\alpha(M))}{\forall M \in P_{in} \forall p \in InSig(\Delta) : \tau(Out(\alpha(M))) \vdash output_{\alpha}(\tau(M) \wedge p) \vee output_{\alpha}(\tau(M) \wedge \neg p)} \text{Lem.2} \\
\frac{\forall a \in P_{in} \forall b \in InSig(\Delta) : output_{\alpha}(a) \subseteq output_{\alpha}(a \wedge b) \cap output_{\alpha}(a \wedge \neg b)}{\forall a \in P_{in} \forall b \in InSig(\Delta) : \forall x \in L : (a \wedge b, x) \in output_{\alpha}, (a \wedge \neg b, x) \in output_{\alpha} \Rightarrow (a, x) \in output_{\alpha}} \text{Def.6}
\end{array}$$

Finally we relate functionality description α to the set of generators G , that is, we define a mapping π from functionality descriptions to sets of generators such that $output_{\alpha} = output(\pi(\alpha))$. In the general case, we can simply define $G = \{(\tau(M), \tau(N)) \mid M \in P_{in}, N = Out(\alpha(M))\}$, but for functionality descriptions satisfying additional properties, more concise representations for the set of generators can be given.

Theorem 3 *If $G = \{(\tau(M), \tau(N)) \mid M \in P_{in}, N = Out(\alpha(M))\}$, then $output_{\alpha} = output(\pi(\alpha))$.*

The logical characterization of the properties of functionality descriptions suggests various other properties of functionality descriptions which can be studied. For example, a notion of well-informed can be defined as simply the validity of the OR rule, and a notion of self-boundedness can be defined using CT instead of Trans. Moreover, the logical characterization may be used to extend the notion of functionality description to more expressive languages, for example languages that include disjunction. We do not further consider these issues in this paper.

Summarizing, Theorem 2 relates compositional design to cognitive attitudes, in the sense that typical examples of input/output processes studied in input/output logic arise in contexts of conditional beliefs, obligations, goals, ideals, preferences, and actions. This is relevant for the application of compositional design of agent architectures such as the BOID architecture considered in this paper. The major issue for further research for the foundations of compositional architectures is the logical analysis of the *interaction* of input/output components [22], such as the distinction between object and meta level interactions (for which hierarchical logics may be used) and feedback loops among components (for which generalized self-boundedness or generalized CT rules may be used). A first step towards such issues is discussed in the following section.

4 Coordination of components

In this section we discuss the coordination of components. In Section 4.1 we discuss how functionality descriptions of individual components can be combined, and in Section 4.2 we discuss their formalization in logic.

4.1 Composition relations ...

Functionality descriptions can be combined by composing functions in the usual way. For example, consider the goal-based agent in Figure 1 with input F (observations), and two functionality descriptions α_p (plan generation) and α_g (goal generation). They can be combined to β by first applying α_g on F and thereafter α_p , i.e., $\beta(F) = \alpha_p(\alpha_g(F))$ or

$$\beta(F) = \alpha_p \circ \alpha_g(F)$$

Moreover, consider the more complex example of the goal generation component in Figure 2. As a first naive approximation, any subset of the four components can be applied in any order, i.e.

$$\alpha_g = \alpha_b \cup \alpha_o \cup \dots \cup \alpha_b \circ \alpha_o \cup \dots \cup \alpha_b \circ \alpha_o \circ \alpha_i \circ \alpha_d$$

This is obviously too strong, but in another sense, it is also too weak.

A first reason for calling this composition too strong, is that it connects every component to every other component. This is not in accordance with, for instance, the architecture as shown in Figure 2; the input of the belief component contains the observations as well as the output of the obligation, intention and desire components. Moreover, the output of the belief component is itself connected to the input of the obligation, intention and desire component. As explained in Section 2, this means that the motivational attitudes induce new beliefs, and that the beliefs may induce new motivational attitudes, respectively. Which beliefs and motivational attitudes are derived depends on the rules and the properties of the components, i.e., on their input/output logic. The behavior of the BOID component depends on rules of the individual components, properties of the individual components (i.e., their input/output logic), and the links among components (which are assumed to be fixed as in Figure 2). The most interesting property is the identity or ID rule, which states that the input is included in the output. For some components this may lead to desired results, but for other components this leads to counterintuitive results. For example, in the BOID architecture it may make sense for beliefs to include the observations, but it makes less sense for obligations or desires to include the beliefs. You may desire to go to the dentist and believe that going to the dentist implies pain, but this does not imply the desire for pain [6]. We can implement these differences in input-output behavior simply by accepting the identity rule for the belief component and not for the obligation, intention and desire components.

A second reason for calling the composition too strong, is that it neglects the control loop, and in particular the conflict resolution mechanisms. For example, Thomason [32] argues that beliefs should override desires with the following example. If you think it is going to rain and you believe that if it rains, you will get wet, and you would

not like to get wet, then you have to conclude that you get wet. Beliefs should therefore prevail in conflicts with desires. This issue is further discussed in the following section.

But, the composition is also too weak, since it does not take complex mechanisms like feedback loops and reasoning by cases into account. These are exactly the properties already discussed in Section 3 of self-boundedness and well informedness, i.e., of the cumulative transitivity (CT) and disjunction (OR) rule! The CT rule, or its variant the transitivity (TRANS) rule, is valid within the belief component (rules can be applied one after the other) and it is valid among the components via the feedback loops. The OR rule states that the component satisfies reasoning by cases. For example, suppose there are two belief rules:

- If a , then the agent believes x ;
- If $\neg a$, then the agent believes x .

Can we conclude that the agent believes x if there are no observations? We can if the agent's belief component supports reasoning by cases. It is usually considered to be a reasonable conclusion, but it is seldom implemented due to the complexity of generating cases. More complex variants of this reasoning scheme among components are, e.g., the following two rules:

- If a , then the agent desires x ;
- If $\neg a$, then the agent is obliged x .

Can we conclude that the agent is motivated for x if there are no observations or beliefs? Only if the agent's control loop supports reasoning by cases.

4.2 ... and global constraints

The general approach we advocate to combine functionality specifications is:

1. to define a very general composition that allows all possible behaviors of the system such that it is not too weak, and thereafter
2. add constraints on this composition such that it is not too strong.

The first step of this approach is usually straightforward (e.g., introduce some operators for iteration) but the second may involve many different techniques. For example, a constraint that can be added on the composition is the realism constraint. One way to express this constraint is to construct a conflict resolution mechanism in which beliefs override desires, such as the mechanisms proposed in Thomason's BDP logic [32] and in Broersen *et.al.*'s BOID architecture [8]. An alternative, more general and more insightful approach is to formulate the intuition as a property of the functional agent description, i.e., as the relation between observations and extensions [9]. Consider the following conflict:

1. The agent believes the car will be sold
2. The agent believes the car will not be sold

The agent does not know what to believe: it is confused. Alternatively, the agent has two incompatible belief sets, one which argues that the car will be sold, and another which argues that the car will not be sold. Confusion can be formalized as an inconsistent belief set, whereas multiple belief states can be formalized by multiple extensions in for example Reiter's default logic. In this paper we follow the latter approach. Moreover, consider the following conflict:

1. The agent desires the car to be sold
2. The agent desires the car not to be sold

The agent has two conflicting desires, which may both become candidate goals. We call this an internal desire conflict. In this paper we assume that also a conflict among desires leads to multiple extensions. Finally, consider the following conflict:

1. The agent believes the car will be sold
2. The agent desires the car not to be sold

The agent's desire conflicts with its belief. Such mixed conflicts can be interpreted in various ways. One way, which we adopt in this paper, is due to Thomason [32]. He argues that it is unrealistic to allow the agent's desire to become a goal when it conflicts with its beliefs, and that therefore beliefs should override desires, with the following example. If the agent believes it is raining and it believes that if it rains, it will get wet, and it desires not to get wet, then the agent cannot pursue the goal of not getting wet. This example shows that it is unrealistic (i.e., it allows for wishful thinking) to allow the desire of not getting wet to become a goal. Thus, beliefs prevail in conflicts with desires. Thomason's interpretation can be contrasted with the following example:

1. The agent believes the fence is white
2. The agent desires the fence to be green

In this example the agent can see to it that the fence becomes green by painting it, so pursuing the goal that the fence is green is not wishful thinking. The difference between this example and the previous one is that this is not a conflict due to implicit temporal references. The belief implicitly refers to the present whereas the desire refers to the future:

1. The agent believes the fence is white *now*
2. The agent desires the fence to be green *in the future*

In this paper we only use abstract examples in which we do not give an interpretation for the propositional atoms. If there is a conflict between a belief and a desire, then there always is a real conflict (as in the car selling example), never an apparent conflict (as in the fence example) that can be solved by considering time. We also do not discuss the kind of revision or updating involved in the fence example.

An intriguing question is how to resolve conflicts among beliefs and desires, in case more than two rules are involved. Consider the example in Figure 3.

Figure 3 represents the following four rules:

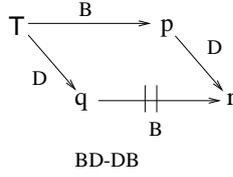


Figure 3: Is it realistic to desire q ? Is it realistic to desire r ?

1. The agent believes p
2. The agent desires q
3. If p , then the agent desires r
4. If q , then the agent believes $\neg r$

In the sequel we give some definitions to determine whether it is realistic to desire q or r .

We introduce two properties that characterize realistic desires. They are not restricted to one particular logic or architecture, but they can be applied to any extension-based approach.

The agent is characterized by a function, which we denote by Δ , from so-called BD theories (sets of belief and desire rules) to extensions (logically closed sets of propositional sentences). This terminology is inspired by Reiter's default logic [30]. However, for now we do not assume any further properties on the relation between BD theories and their extensions. For example, we do not assume that observations are included in the extensions. In this sense the phrase 'extension' may be slightly misleading.

Definition 7 (Agent, BD theory, extension) A BD theory is a tuple $T = \langle W, B, D \rangle$, where W is a set of propositional sentences of a propositional language L and B and D sets of ordered pairs of such sentences. An extension is a logically closed set of L sentences. Δ is a function which returns for each BD theory a set of extensions. We write $\Delta(T)$ for the set of all extensions of a BD theory T (there may be none, one or multiple extensions), and we write $Th_L(S)$ for all propositional consequences of the set of propositional formulas S . For representational convenience we write $\Delta(W, B, D)$ for $\Delta(\langle W, B, D \rangle)$, and we write $Th_L(\alpha)$ for $Th_L(\{\alpha\})$.

Using definition 1, the example of Figure 3 can be represented by a BD theory $T = \langle \emptyset, \{\top \xrightarrow{B} p, q \xrightarrow{B} \neg r\}, \{\top \xrightarrow{D} q, p \xrightarrow{D} r\} \rangle$. Note that this definition allows us to use any pairs of propositional formulas, which means that we consider a more general setting than in the examples thus far.

Realisms concerns the rational balance in case of conflict. Therefore we first define what a conflict is. Conflicting theories lead to an inconsistent extension if rules are applied.⁵

⁵An alternative stronger definition is that T is a conflict if the following is inconsistent:

$$W \cup \{a \rightarrow x \mid a \xrightarrow{B} x \in B \text{ or } a \xrightarrow{D} x \in D\}$$

Definition 8 (Conflict) Let $T = \langle W, B, D \rangle$ be a BD theory. T contains a conflict iff there is no consistent set E of L sentences such that:

- $W \subseteq E$
- If $a \xrightarrow{B} x \in B$ or $a \xrightarrow{D} x \in D$ and $a \in Th_L(E)$ then $x \in Th_L(E)$

One way to proceed is to define for each BD theory when a desire is realistic and when it is unrealistic. A drawback of this approach is that it has to commit to a logic of rules for the belief and desire rules. We therefore follow another approach, which may be called comparative. The basic pattern is as follows. If a realistic function Δ returns for a BD theory T a set of extensions S , then we can deduce that it does not return for other BD theories T' extensions S' . The latter extensions S' would be unrealistic, i.e., based on unrealistic desires.

The realism properties are defined in terms of sets of belief and desire rules. However, BD theories are not mapped on a set of belief and desire rules, but on extensions generated by such rules. We therefore have to associate with each extension a set of belief and desire rules. The following definition associates an extension with the set of rules which are applied in it (sometimes called its generating set [30]).

Definition 9 (Applied rules) Let $T = \langle W, B, D \rangle$ be a BD theory and let the set E be one of its extensions. The set of applied belief rules in extension E is $R_B(T, E) = \{\alpha \xrightarrow{B} w \in B \mid \alpha \wedge w \in E\}$, and the set of applied desire rules is $R_D(T, E) = \{\alpha \xrightarrow{D} w \in D \mid \alpha \wedge w \in E\}$.

The intuition behind a priori realism in Definition 10 is as follows. Consider a BD theory $(\langle W, B, D \rangle)$ and an extension of this BD theory (E) in which at least one desire has been applied. We call this the a posteriori state. We want to ensure that these applied desires result in realistic extensions. We therefore consider the state in which this rule has not been applied (BD theory $\langle W, B, D' \rangle$ with extension E'). We call this state the a priori state. We now say that the desire rule is realistic if the set of applied belief rules in the a priori state is a subset of the set of applied belief rules in the a posteriori state. This implies that the removal of realistic desires from the BD theory can only decrease the extension, not increase it or remove it.⁶

Definition 10 (A priori realism) Δ is a priori realistic iff for each $E \in \Delta(W, B, D)$ and $D' \subseteq R_D(\langle W, B, D \rangle, E)$ there is an $E' \in \Delta(W, B, D')$ such that we have $R_B(\langle W, B, D' \rangle, E') \subseteq R_B(\langle W, B, D \rangle, E)$. We also say that each $E \in \Delta(W, B, D)$ that satisfies the above condition is realistic, and we say that all applied desire rules of a realistic extension are realistic desire rules.

In the remainder of this section we illustrate a priori realism by some examples.

Which definition of conflict is used depends on the underlying logic of rules, see, e.g., [25] for some possibilities. For the definitions of realism in this paper the exact definition of conflict is not important.

⁶An alternative closely related definition of a priori realism is as follows. For each $E \in \Delta(W, B, D)$ and $D' \subseteq R_D(\langle W, B, D \rangle, E)$ there is an $E' \in \Delta(W, B, D')$ such that $E' \subseteq E$. A simple instance of this property, which we may call ‘restricted a priori realism,’ is the case where D is the empty set. This property says that every BD extension extends a B extension. For each $E \in \Delta(W, B, D)$ there is an $E' \in \Delta(W, B, \emptyset)$ such that $E' \subseteq E$.

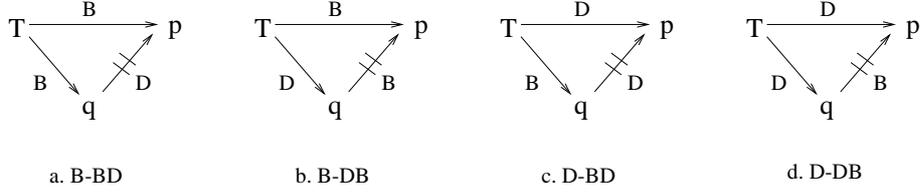


Figure 4: Desire-belief triangles

Example 1 Consider the four triangles in Figure 4. Intuitively, we have a: $\{p, q\}$, b: $\{p\}$ or $\{p, q\}$, c: $\{\neg p, q\}$ or $\{p, q\}$, d: $\{p\}$ or $\{\neg p, q\}$.

Case a. Let $T = \langle \emptyset, \{\top \xrightarrow{B} p, \top \xrightarrow{B} q\}, \{q \xrightarrow{D} \neg p\} \rangle$ and
 $T' = \langle \emptyset, \{\top \xrightarrow{B} p, \top \xrightarrow{B} q\}, \emptyset \rangle$.

T contains a conflict, because any set E as defined in Definition 8 contains p as well as $\neg p$. Moreover, assume $\Delta(T') = \{Th_L(p \wedge q)\}$ with $R_B(T', Th_L(p \wedge q)) = \{\top \xrightarrow{B} p, \top \xrightarrow{B} q\}$. Due to a priori realism we have for each element E of $\Delta(T)$ that $R_B(T, E)$ contains $\{\top \xrightarrow{B} p, \top \xrightarrow{B} q\}$, and consequently E has to contain $Th_L(p \wedge q)$. $\Delta(T)$ thus cannot contain for example $Th_L(q \wedge \neg p)$. In other words, according to Property 10 we have that the desire for $\neg p$ is unrealistic.

Case b. Let $T = \langle \emptyset, \{\top \xrightarrow{B} p, q \xrightarrow{B} \neg p\}, \{\top \xrightarrow{D} q\} \rangle$ and
 $T' = \langle \emptyset, \{\top \xrightarrow{B} p, q \xrightarrow{B} \neg p\}, \emptyset \rangle$.

T contains a conflict, because any set E as defined in Definition 8 contains p as well as $\neg p$. Assume $\Delta(T') = \{Th_L(p)\}$ with $R_B(T', Th_L(p)) = \{\top \xrightarrow{B} p\}$. Due to a priori realism, each element of $\Delta(T)$ has to contain $Th_L(p)$. Consequently, $\Delta(T)$ thus can contain $Th_L(p \wedge q)$, but it cannot contain for example $Th_L(q \wedge \neg p)$. In other words, according to Property 10 we have that $\neg p$ is unrealistic. Note that there is not a desire for $\neg p$, but that $\neg p$ would be a believed consequence of a desire (for q). However, it does not imply that the desire for q is unrealistic.

Case c. Let $T = \langle \emptyset, \{\top \xrightarrow{B} q\}, \{\top \xrightarrow{D} p, q \xrightarrow{D} \neg p\} \rangle$ and
 $T' = \langle \emptyset, \{\top \xrightarrow{B} q\}, \emptyset \rangle$.

If $\Delta(T') = \{Th_L(q)\}$, then each element of $\Delta(T)$ has to contain $Th_L(q)$, but $\Delta(T)$ still can contain for example $Th_L(p \wedge q)$ and $Th_L(\neg p \wedge q)$. In other words, according to Property 10 in T' neither p nor $\neg p$ would be unrealistic. It illustrates that Property 10 does not classify conflicts among desires as unrealistic.

Case d. Let $T = \langle \emptyset, \{q \xrightarrow{B} \neg p\}, \{\top \xrightarrow{D} p, \top \xrightarrow{D} q\} \rangle$,
 $T' = \langle \emptyset, \{q \xrightarrow{B} \neg p\}, \{\top \xrightarrow{D} p\} \rangle$, and
 $T'' = \langle \emptyset, \{q \xrightarrow{B} \neg p\}, \{\top \xrightarrow{D} q\} \rangle$.

If $\Delta(T') = \{Th_L(p)\}$ and $\Delta(T'') = \{Th_L(q, \neg p)\}$, then a priori realism implies $Th_L(p \wedge q) \notin \Delta(T)$. However, note that $Th_L(p)$ and $Th_L(q \wedge \neg p)$ may be in $\Delta(T)$ (verification left to the reader).

5 Further research

There are several consequences of Theorem 2 for further logical investigations.

First, the declarative functionality descriptions can be compared to other logical formalisms. For example, out_3^+ is the logic of Reiter's normal default logic without constraints, and out_4^+ is material implication (and thus satisfies cut elimination, in the sense that it is equivalent to out_2^+) [25]. Moreover, some lemmas shown for the mapping α now correspond to well-known logical properties. For example, the fact that α is monotonic if it is conservative and self-bounded, now corresponds to the fact that Id and TRANS imply SI (in the context of WO), proven as follows.

$$\frac{\frac{\frac{-}{(a \wedge b, a \wedge b)} \text{ID}}{(a \wedge b, a)} \text{WO} \quad (a, x)}{(a \wedge b, x)} \text{TRANS}$$

Second, the declarative functionality descriptions can be compared to other possible mappings. A possible extension is to allow for defeasible inference. Systems without SI have been studied in conditional logic [23] and more recently by [21]: Id+CT is cumulative inference, Id+CT+OR is preferential inference.

Third, with the partition of the signature, we can use weaker logical systems to obtain the same output. We do not need the Id rule if the input and output variables do not overlap, i.e., we can use out_3 and out_4 instead of out_3^+ and out_4^+ . Moreover, if we also do not have intermediate variables then we can drop CT, i.e., we can use out_1 and out_2 . This shows the use of all these different systems: The more complex combination of variables is allowed, the more extended proof system (input/output logic) is required.

Fourth, another issue of relevance for complex compositional systems is their computational complexity. Note, for instance, that an extended set of proof rules does not necessarily mean increased computational complexity (e.g., out_2 and out_3 are probably computationally more complex than out_1 or out_4^+). Again the logical characterization may be of use, because complexity results may already be known or such results are easier to obtain in a standard logic setting (thus far no complexity results on input/output logics are known, but results are known on instances like Reiter's normal default logic).

Moreover, in this paper we have just made a start with the formal analysis of the coordination of cognitive components. There are many issues to be studied.

6 Related research

In agent technology, the relation between architectures and formal logic is a topic of frequent discussion. The most discussed example is the relation between Georgeff and Rao's BDI architectures [19, 28, 29] and their BDI logic [27]. Typical questions are 'how to use formal languages in the design of such architectures?', 'what is the relation between the architecture and its logic?', and more generally, 'how to exploit computational logic for multi agent systems?' There are at least two fundamentally

different uses of logics for agents. Languages for analysis, requirement specification and verification enable reasoning *about* agents, and knowledge specification languages are used *within* agent architecture to represent agent’s mental attitudes. Despite this clear distinction in the *uses* of agent logics, there is no such clear distinction in the agent logics themselves, because some logics can be used as specification as well as knowledge representation language. An example is the modal BDICTL logic [27, 31] – though it seems better fit for the former task.

In practical reasoning monadic modal logics were introduced to formalize various cognitive notions, most notably epistemic operators (knowledge and belief) and deontic operators (obligation and permission). However, the use of monadic modalities was soon seen as too restrictive. The area of epistemic logic transformed into the area of belief revision, in particular under the Alchourrón-Gärdenfors-Makinson paradigm [3]. The area of deontic logic studied an issue related to belief revision, which we may call obligation revision (though it is usually called contrary-to-duty reasoning [35]). Von Wright [38], Alchourrón [2, 1] and Makinson [24] started to look for general characterizations of the problems, and they concluded that modal logics with their focus on nested modalities did not touch the core of the problem, which was to be found in the notion of conditionality. Input/output logic, as used in this paper, is one of the logics developed in this new tradition. Moreover, several other dynamic approaches have been introduced lately in the area of logic, in particular in relation to language and philosophy, such as Veltman’s update semantics [36] and van Benthem’s work [34] on logical dynamics.

In general, previous theoretical research has often neglected to show possible implementations. Like [16, 11], we not only provide a theoretical framework, but we also provide an architecture which includes a control procedure in the style of Rao and Georgeff’s Procedural Reasoning System (PRS) [28, 19]. In particular, we provide a theoretical framework to specify agent’s properties by extending the BDP formalism with prior intentions and obligations, and use BOID as a generic agent architecture to design and implement agents. In this section, we discuss three existing formalisms and explain how our formalism is related to them.

6.1 Deontic Action programs

Deontic action programs, introduced by Eiter *et al.* [17], are sets of horn-type rules of the form $L_1, L_2, \dots, L_n \rightarrow A$, where the L_i are literals (possibly negated atoms), either of a Boolean form, or from the set $P(\alpha), F(\alpha), O(\alpha), W(\alpha), Do(\alpha)$, which stand for permission, prohibition, obligation, waiving, and performance of α , respectively. In the head A , negations and boolean combinations are not allowed. The ‘semantics’ of rules is defined in terms of status sets, which are application sets of the above operators to ground instances of actions α . The status sets correspond to what we call ‘extensions’. Eiter *et al.* discuss many different choices for calculating with action programs, resulting in many different types of status sets. They propose closure of status sets S under the deontic rules:

- if $O(\alpha) \in S$ then $Do(\alpha) \in S$
- if $Do(\alpha) \in S$ then $P(\alpha) \in S$

Furthermore, status sets have to obey consistency constraints:

if $O(\alpha) \in S$ then $W(\alpha) \notin S$
 if $P(\alpha) \in S$ then $F(\alpha) \notin S$
 if $P(\alpha) \in S$ then $Pos(\alpha)$ (a rule ensuring that α is only permitted if its precondition is satisfied)

Status sets closed under the program rules, the above deontic rules, obeying the consistency constraints and other, independently specified constraints, are called ‘feasible status sets’. But feasible status sets lack a criterion of minimality: there may be unsupported elements. This leads to the concepts of ‘groundedness’ and ‘rational status set’. The class of rational status sets is a subset of the class of feasible status sets. The criterion that it should not be possible to use rules contrapositively leads to the concept of ‘reasonable status set’, the class of reasonable status sets being a subset of the class of rational status sets.

Deontic action programs and BOID theories have several aspects in common. Both formalizations are rule-based, encompass normative concepts (obligation), and do not allow for contrapositive rule application. But there are also many notable differences. A technical difference follows from the comparison with Reiter’s default theories. BOID rules are like normal defaults: information in the heads of rules is allowed to occur negated and may give rise to conflicts, while negation symbols in the body are interpreted in the standard way. In the rules of action programs, negations in the heads are not allowed, while negation in the bodies is interpreted according to ‘negation as failure’. Also, both approaches have different subject domains: BOID rules do not deal with permission, prohibition, waiving and executability, while action programs do not deal with beliefs, desires and intentions. From this difference in subject domain follows a difference in research questions. The central problem for BOID theories is conflict resolution. BOID aims at selecting actions for execution, given possibly conflicting information about informational, internal and external motivational attitudes. Action programs, on the other hand, aim at selecting actions for execution, given information about action and external motivational attitudes. Possible conflicts with internal or informational attitudes are not considered. In the action selection process, BOID always comes to an answer, while action programs may return the empty action.

From a deontic perspective, it is possible to raise some objections against the approach taken by Eiter *et al.*. A first observation is that status sets fail to establish the important deontic relation that in standard deontic logic (SDL [37]) is expressed as $O(\phi) \leftrightarrow F(\neg\phi)$, and that in deontic action logics is expressed as $O(\alpha) \leftrightarrow F(\sim\alpha)$ (see [10]). In other words, this property says that ‘obligation to do α ’ equals ‘prohibition to do anything other than α ’. If Eiter *et al.* want to be faithful to this deontic principle, they should consider incorporation of the rules: (1) if $O(\alpha) \in S$ then $F(\beta) \in S$ for all actions β other than α , and (2) if $F(\beta) \in S$ for all actions β other than α then $O(\alpha) \in S$. Especially this second clause might be important, since it may lead to selection of actions for execution only from prohibitions for ‘alternative’ actions.

A more fundamental objection concerns the motivation for including deontic operators in action programs in the first place. The main reason for viewing deontic notions as useful for modeling behavior, is their capacity to specify violations [39]. In the use Eiter *et al.* make of deontic operators, such violations do not occur. The central elements of status sets are the elements $Do(\alpha)$, that determine the action that is undertaken. The other deontic elements of status sets have no independent operational mean-

ing in themselves and are only important in the sense that they impose restrictions on occurrences of elements $Do(\alpha)$. This means that no violations of deontic constraints are considered: status sets always obey all specified norms. The authors notice this weakness themselves and suggest that it can be remedied by allowing and deciding internal conflicts among obligations and by the introduction of violation costs. But this does not solve the problem how to specify the way an agent is supposed to behave when a certain norm violation has occurred (see [10]).

6.2 BDI Approach

The BDI approach [19, 28, 29, 27] is one of the most popular approaches to capturing the rational balance between mental attitudes of cognitive agents. In the BDI framework, beliefs, desires and intentions are considered as independent modalities formalized by separate modal operators. The specific relationships among these modalities are given by two types of constraints called static and dynamic constraints. The static constraints are formalized by special axioms which, for example, ensure stable decision making behavior. The dynamic constraints, known as *commitment strategies* capture the dynamic of mental attitudes, i.e. the ways in which the (relations between) attitudes change over time. An important commitment strategy concerns the perseverance with which an agent maintains its goals. In Rao and Georgeff’s formalism $Int_a(A\Diamond\alpha)$ describes the pursuit of a goal: agent a intends to eventually achieve α , no matter how the future evolves. The condition that a proposition α will eventually come true is expressed by $\Diamond\alpha$, while A is to be read as “for all possible courses of future events”. The question here is how long an agent should maintain $Int_a(A\Diamond\alpha)$. It seems to be reasonable for an agent to maintain a goal α as long as the agent believes that α is attainable. This is exactly what $Bel_a(\alpha)$ denotes. The agent should then give up to maintain a goal as soon as the agent believes that α is not attainable anymore. This condition can be expressed with the help of the E operator, which reads “for at least one possible course of events.” This condition is translated into $\neg Bel_a(E\Diamond\alpha)$. The until operator U then allows us to formulate an axiom of change:

$$A(Int_a(A\Diamond\alpha)U(Bel_a(\alpha) \vee \neg Bel_a(E\Diamond\alpha)))$$

This is what Rao and Georgeff call ‘single-minded commitment’. Other commitment strategies are defined in a similar way. Commitment strategies of this type are used to characterize different types of agents.

Like BDI, BOID aims at a formalization of the rational balance between mental attitudes of agents. However, the BOID formalism extends the mental attitudes of BDI with obligations. The gap between BOID architectures and their underlying logics is much smaller than the large gap between modal BDI logics and their architectures. This gap is due to the fact that the BDI formalism is developed to *specify* agent behavior while the BOID formalism is developed to *design* agents. Input/output logic can employ the BOID formalism to specify agent behavior in a straightforward way. In fact, in this paper we have shown how input/output logic can be used to specify the behavior of individual components.

Finally, in modal logics *nested* modalities can be represented, such as: ‘You are obliged not to desire to smoke.’ This cannot be represented in input/output logics or

in Reiter’s default logic. However, we can replace the propositional base language of input/output logics by a modal language containing B/O/I/D modalities. Now rules may be added such that the desired extensions can be derived which can be used as goals in the BOID architecture [13].

6.3 BDP Approach

A rather different formalism to specify the behavior of cognitive agents is known as the Beliefs, Desire, and Planning (BDP) formalism proposed by Thomason [32]. BDP is designed to integrate reasoning about beliefs and desires with planning. The basic idea is to extend the planning formalisms with goals that are not given but derived from agent’s beliefs and desires. Thomason first discusses the BD formalism without planning and only thereafter BDP, such that in the first part he can focus on the interaction between beliefs and desires.

The BDP formalism is based on Reiter’s default logic [30]. The basic idea is to model beliefs and desires both as Reiter defaults, *without modalities for belief or desire*, such that the extensions contain all the indirect effects of actions. That is, a BD-basis is a tuple $\langle M, NB, ND \rangle$ with M a set of formulas, NB a set of B-rules ‘if a , then I believe x ’ written as $a \xrightarrow{B} x$, and ND a set of D-rules ‘if a , then I desire x ’ written as $a \xrightarrow{D} x$. Extensions are built in the usual way without distinguishing between beliefs and desires, so for example, the BD-basis $a, a \xrightarrow{B} b, b \xrightarrow{D} c$ has as an extension $Th(\{a, b, c\})$. But then, there are two types of conflicts:

- Conflicts between a belief and a desire leads to overriding of desire by belief to block wishful thinking.
- Other conflicts like between two desires lead to multiple extensions.

BDP is partly inspired by computational BDI systems, but there are important differences. First, in BDI logics decisions are restrained by previous intentions. Second, whereas BDI logics only formalize unconditional beliefs and desires, BDP also formalizes conditional ones. Third, whereas BDI logics are based on modal logic, BD logic (sub-formalism of the BDP logic) is based on Reiter’s normal default logic.

The BOID formalism extends the BDP with Obligations and Intentions. In particular, the BDP formalism is based on conflict resolution for conditional beliefs and desires which is extended in the BOID formalism with conditional obligations and desires borrowed from respectively deontic action programs [17] and BDI logic [12, 27]. However, in contrast to the BDP formalism, the BOID formalism does not account for planning though this is not a principle shortcoming. We believe that the computed and selected extension in the BOID formalism, which can be considered as a set of goals, should be mapped into a sequence of actions (i.e., a plan) through which the goals can be reached. The mapping from a set of goals to a plan can be accomplished by an additional planning component in the BOID architecture [13].

Extending the BDP logic with obligations and intentions increases the number of possible conflicts dramatically. In both approaches conflict resolution is defined as selecting a subset of the conditionals that does not derive a contradiction. As different

types of rules are formalized as normal Reiter defaults, this means that a Reiter extension should be selected. The fact that extensions are built represents that all the effects of the decisions are taken into account.

In order to design cognitive agents based on the BDP formalism, Thomason sketches an agent architecture [32] which is far from being used to design cognitive agents. For this reason, we have introduced the BOID architecture [8] which can be used to design cognitive agents based on an extension of the BDP formalism.

7 Concluding remarks

To scale up agent applications from toy examples to industrial applications, agent logics need to support compositional or component-based design, and include related tools like abstraction and refinement operators. In compositional architectures for reasoning systems or agents the functioning of the overall system is based on the functionality of the components, in addition to the composition relations by which they are glued together within the compositional architecture. Specification of a composition relation may involve, for example, the possibilities for information exchange among components and control.

In this paper we considered the compositional design of the cognitive BOID architecture. In a component-based BOID agent questions concerning the behavior of the whole agent or satisfaction of the overall agent's functionality, can be decomposed into questions for the case of a single primitive interactive reasoning component within the agent, and questions related to the agent's compositional structure. For the former type of questions the notion of declarative functionality description comes in the play. At the level of the composition as a whole, the (implementation) internal details of a component are irrelevant, as long as an adequate overall description of the component's behavior is available. Since functionality descriptions abstract from these internal details, this notion is well-suited for these purposes. As cognitive attitudes are represented by components, the cognitive attitudes have to be represented by conditional structures – for which we used input/output logics. Thus, compositionality forced us to use a more complex logic to describe the behavior of components as well as the composition of components.

For the specification of individual components we have related functionalities of reasoning components to input/output logic [25, 26], in the sense that it has been shown how the formalization of the functionality description of a reasoning component, introduced in [33], can be mapped onto an input-output logic specification, and, moreover, how specific properties of such functionality descriptions correspond to specific proof rules in input-output logic. This introduces a structural relation among components within architectures of reasoning systems, as specified from a software engineering perspective (e.g., in a design language as DESIRE), to logic. It shows that input-output logics can be used as the basis of an adequate formalization of the semantics of a component in a compositional agent-based reasoning system.

Compositionality has not only some important advantages for design, but it also may be a crucial factor for formal analysis like verification. This is a topic for further research. Compositional verification extends the main advantages of compositionality

for designing to:

A well-structured verification process. Both conceptually and computationally the complexity of the verification process can be handled by compositionality at different levels of abstraction. The verification problem is reduced to a number of much simpler verification problems, and the problem of defining the logical relation that compose properties of components together to properties of the whole system.

The reusability of proofs for properties of reused components. If a modified component satisfies the same properties as the previous one, then the proof of the properties at the higher levels of abstraction can be reused to show that the new system has the same properties as the original. This has high value for a library of reusable generic models and components. A library of reusable components and task models may consist of both specifications of the components and models, and their design rationale. As part of the design rationale, at least the properties of the components and their logical relations can be documented. For more details about compositional verification, see [18, 20, 15].

Acknowledgment

Thanks to Zhisheng Huang, Joris Hulstijn and Jan Treur for discussions on the issues discussed in this paper.

References

- [1] C.E. Alchourrón. Philosophical foundations of deontic logic and the logic of defeasible conditionals. In J.-J. Meyer and R. Wieringa, editors, *Deontic Logic in Computer Science: Normative System Specification*, pages 43–84. John Wiley & Sons, 1993.
- [2] C.E. Alchourrón and Bulygin. The expressive conception of norms. In R. Hilpinen, editor, *New Studies in Deontic Logic: Norms, Actions and the Foundations of Ethics*, pages 95–124. D. Reidel, 1981.
- [3] C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [4] F. Arbab. Abstract behavior types: A foundation model for components and their composition. In *Formal Methods for Components and Objects: First International Symposium, FMCO 2002, Leiden, The Netherlands, November 5-8, 2002, Revised Lectures*, volume 2852 of *LNCS*, pages 33 – 70, 2003.
- [5] The BOID’s home page. <http://www.cs.vu.nl/~boid/>.

- [6] M.E. Bratman. *Intention, plans, and practical reason*. Harvard University Press, Cambridge Mass, 1987.
- [7] F.M.T. Brazier, C.M. Jonker, and J. Treur. Principles of compositional multi-agent system development. In J. Cuena, editor, *Proceedings of the 15th IFIP World Computer Congress, WCC'98, Conference on Information Technology and Knowledge Systems, IT&KNOWS'98*, pages 347–360, 1998. To be published by IOS Press.
- [8] J. Broersen, M. Dastani, J. Hulstijn, and L. van der Torre. Goal generation in the BOID architecture. *Cognitive Science Quarterly*, 2(3-4):428–447, 2002.
- [9] J. Broersen, M. Dastani, and L. van der Torre. Realistic desires. *Journal of Applied Non-Classical Logics*, 12(2):287–308, 2002.
- [10] J.M. Broersen, R.J. Wieringa, and J.-J.Ch. Meyer. A fixed-point characterization of a deontic logic of regular action. *Fundamenta Informaticae*, 49(4):107–128, 2001. Special issue, proceedings DEON2000.
- [11] C. Castelfranchi, F. Dignum, C. Jonker, and J. Treur. Deliberative normative agents: principles and architecture. In *Proceedings of the ATAL'99*, LNCS, pages 364–378, 1999.
- [12] P.R. Cohen and H.J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.
- [13] M. Dastani and L. van der Torre. Programming boid agents: a deliberation language for conflicts between mental attitudes and plans. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS'04)*, pages 706–713, 2004.
- [14] M. Dastani and L. van der Torre. What is a normative goal? Towards goal-based normative agent architectures. In *Regulated Agent-Based Systems. Postproceedings RASTA02*, LNAI 2934, pages 210–227. Springer, 2004.
- [15] W.P. de Roeper, H. Langmaack, and A. Pnueli (eds.). *Proceedings of the International Workshop on Compositionality, COMPOS'97*, volume 1536 of *Lecture Notes in Computer Science*. Springer Verlag, 1998.
- [16] F. Dignum, D. Morley, E. Sonenberg, and L. Cavedon. Towards socially sophisticated BDI agents. In *Proceedings of the Fourth International Conference on MultiAgent Systems (ICMAS-2000)*, pages 111–118, 2000.
- [17] Thomas Eiter, V.S. Subrahmanian, and G. Pick. Heterogeneous active agents I: Semantics. *Artificial Intelligence*, 108 (1-2):179–255, 1999.
- [18] J. Engelfriet, C.M. Jonker, and J. Treur. Compositional verification of multi-agent systems in temporal multi-epistemic logic. *Journal of Logic, Language and Information*, 11(2):195–225, 2002.

- [19] M. Georgeff and A. Lansky. Reactive reasoning and planning. In *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, pages 677–682, 1987.
- [20] C.M. Jonker and J. Treur. Compositional verification of multi-agent systems: a formal analysis of pro-activeness and reactiveness. In W.P. de Roever, H. Langmaack, and A. Pnueli, editors, *Proceedings of the International Workshop on Compositionality, COMPOS'97*, volume 1536 of *Lecture Notes in Computer Science*, pages 350–380. Springer Verlag, 1998.
- [21] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.
- [22] P. Leemans, J. Treur, and M. Willems. On the verification of knowledge-based reasoning modules. Technical Report IR-346, Vrije universiteit, faculteit der wiskunde en informatica, 1993.
- [23] D. Lewis. *Counterfactuals*. Blackwell, Oxford, 1973.
- [24] D. Makinson. On a fundamental problem of deontic logic. In P. McNamara and H. Prakken, editors, *Norms, Logics and Information Systems. New Studies on Deontic Logic and Computer Science*, pages 29–54. IOS Press, 1999.
- [25] D. Makinson and L. van der Torre. Input-output logics. *Journal of Philosophical Logic*, 29:383–408, 2000.
- [26] D. Makinson and L. van der Torre. Constraints for input-output logics. *Journal of Philosophical Logic*, 30:155–185, 2001.
- [27] A. Rao and M. Georgeff. Modeling rational agents within a BDI architecture. In *Proceedings of the KR'91*, pages 473–484, 1991.
- [28] A. Rao and M. Georgeff. An abstract architecture for rational agents. In *Proceedings of the KR'92*, pages 439–449, 1992.
- [29] A. Rao and M. Georgeff. BDI agents: From theory to practice. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95)*, pages 312–319, 1995.
- [30] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1-2):81–132, 1980.
- [31] K. Schild. On the relationship between BDI logics and standard logics of concurrency. *Autonomous Agents and Multi agent Systems*, 3(3):259–283, 2000.
- [32] R. Thomason. Desires and defaults:. In *Proceedings of the KR'2000*, pages 702–713. Morgan Kaufmann, 2000.
- [33] J. Treur. Semantic formalisation of interactive reasoning functionality. *International Journal of Intelligent Systems*, 17:645–686, 2002.

- [34] J. van Benthem. *Exploring logical dynamics*. Cambridge University Press, 1996.
- [35] van der Torre and Tan. Contrary-to-duty reasoning with preference-based dyadic obligations. *Annals of Mathematics and Artificial Intelligence*, 1999.
- [36] F. Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25:221–261, 1996.
- [37] G.H. von Wright. Deontic logic. *Mind*, 60:1–15, 1951.
- [38] G.H. von Wright. Deontic logic: as I see it. In P. McNamara and H. Prakken, editors, *Norms, Logics and Information Systems. New Studies on Deontic Logic and Computer Science*, pages 15–25. IOS Press, 1999.
- [39] R.J. Wieringa and J.-J.Ch. Meyer. Applications of deontic logic in computer science: A concise overview. In J.-J.Ch. Meyer and R.J. Wieringa, editors, *Deontic Logic in Computer Science: Normative System Specification*, pages 17–40. Wiley, 1993.