Specifying the Merging of Desires into Goals in the Context of Beliefs

Mehdi Dastani and Leendert van der Torre

The BOID Project
http://www.cs.vu.nl/~boid/
Utrecht University and Vrije Universiteit Amsterdam

Abstract. Rao and Georgeff's BDI_{CTL} logic is a popular specification and verification language for cognitive agent systems, in which desires and goals are unified into a single motivational attitude. In this paper desires and goals are distinguished to specify the merging of desires into goals, an important process in several agent systems such as Broersen *et al.*'s BOID system. We therefore introduce a BDGI_{CTL} logic. Moreover, we distinguish the merging of conflicting motivational attitudes such as desires into goals from the merging of conflicting informational attitudes such as knowledge bases or belief sources into beliefs, for which we use recent results in variants of belief revision known as semi-revision, fusion or merging. In particular, whereas belief merging is a generalization of revision, desire merging is a generalization of a kind of contraction known as (severe) withdrawal.

1 Introduction

Agent oriented software methodology [13] is the right abstraction level to analyze, design, and implement complex application domains. It is used for example for e-commerce applications and search bots that assist, actively support, and model the user. These application domains can be analyzed more intuitively and directly in terms of interacting agents or actors and high-level concepts such as beliefs, desires, goals, intentions, and plans. These agents are often called cognitive agents and the software systems are called cognitive agent systems.

Formal tools are necessary to guarantee the correctness of the analysis, design, and implementation. A popular specification and verification language for cognitive agent systems is Rao and Georgeff's BDI_{CTL} logic [10, 12], the extension of so-called computational tree logic (CTL) with normal modalities for Beliefs, Desires and Intentions (BDI). CTL is used to specify traces of a cognitive agent through time while the belief, desire, and intention modalities are used to specify the epistemic states of the agent.

A drawback of the BDI_{CTL} formalism is that desires and goals have been unified into a single motivational attitude, such that the logic cannot be used to specify or verify how goals are generated. The crucial lack of expressive power is that BDI_{CTL} does not distinguish between motivational attitudes which can conflict with beliefs (which we call desires or candidate goals) and motivational attitudes which cannot (which we call goals).

Our research question is how Rao and Georgeff's $\mathrm{BDI_{CTL}}$ has to be extended to specify or verify goal generation. In particular, we want to specify how *desires are merged into goals in the context of beliefs*. In our terminology, we are interested in the relation between beliefs, desires and goals, whereas Rao and Georgeff are interested in the relation between beliefs, goals and intentions. Our work is thus complementary to theirs.

The first step is the introduction of an extension of $\mathrm{BDI_{CTL}}$ called $\mathrm{BDGI_{CTL}}$, that makes the distinction between desires and goals explicit. We assume that the following properties hold. They are called RG weak realism, because the formula represents a kind of realism proposed by Rao and Georgeff. These properties allow conflicts between beliefs and desires but not between beliefs and goals. For example, if you believe that it is not possible to be rich and to be a scientist at the same time, then you may have the desires to be both rich and a scientist, but you may not have the goals to be both.

$$\not\models D(p) \to \neg B(\neg p) \qquad \models G(p) \to \neg B(\neg p) \qquad \text{(RG weak realism)}$$

Moreover, desires are formalized by a non-normal modal logic [3,6], because for technical reasons explained later desires and goals are also distinguished by the following two logical properties, that express that goals are a logically closed set but desires are not.

$$\not\models (D(p) \wedge D(q)) \leftrightarrow D(p \wedge q) \qquad \models (G(p) \wedge G(q)) \leftrightarrow G(p \wedge q)$$

The second step is the use of recently introduced belief revision operators known as semi-revision, fusion or merging operators [1,4,5,8] to specify goal generation in cognitive agents. The idea is to take operators that merge belief sources into beliefs in the context of integrity constraints, and use them to merge desires into goals in the context of beliefs. Due to the implicit assumption that merged desires are considered as goals and used in a planning component, desires and goals should satisfy the following properties. They are called CL realism, because the formula represents a kind of realism proposed by Cohen and Levesque (see [10]). For example, you may believe that you cannot be both rich and a scientist, without having the desire or goal that this is so.

$$\not\models B(p) \to D(p) \qquad \not\models B(p) \to G(p) \qquad \text{(CL realism)}$$

These formulas should not be theorems of the logic, because it does not make sense for the agent to plan to achieve something which it already believes to be the case. Consequently, we distinguish desire merging from belief merging. In particular, whereas belief merging is a generalization of revision, desire merging is a generalization of a kind of contraction known as (severe) withdrawal [11].

The layout of this paper is as follows. In Section 2 we introduce $\mathrm{BDGI_{CTL}}$. In Section 3 we consider belief and desire mergers.

2 Specifying the merging of desires into goals

In this section we briefly repeat Rao and Georgeff's formalism [10] and we extend it with possibly conflicting desires. In our terminology desires in Rao and Georgeff's formalism are goals instead of desires, because the motivational attitudes in this formalism are not supposed to conflict with the beliefs. We therefore call their formalism $\mathrm{BGI_{CTL}}$ instead of $\mathrm{BDI_{CTL}}$. In this paper we use an equivalent reformulation presented by Schild [12].

2.1 $\mathbf{BGI_{CTL}}$

In contrast to Schild, we only consider the single agent case.

Definition 1. *[12, Def.4,6] The admissable formulae of BGI_{CTL} are categorized into two classes, state formulae and path formulae.*

S1 Each primitive proposition is a state formula.
S2 If α and β are state formulae, then so are $\alpha \wedge \beta$ and $\neg\alpha$.
S3 If α is a path formula, then $E\alpha$ and $A\alpha$ are state formulae.
S4 If α is a state formula, then $B(\alpha), G(\alpha), I(\alpha)$ are state formulae as well.
P If α and β are state formulae, then $X\alpha$ and $\alpha \cup \beta$ are path formulae.

The semantics of $\mathrm{BGI_{CTL}}$ involves two dimensions: an epistemic and a temporal dimension. The truth of a formula depends on both the epistemic world w and the temporal state s. A pair $\langle w, s\rangle$ is called a situation in which $\mathrm{BGI_{CTL}}$ formulae are evaluated. The relation between situations is traditionally called an accessibility relation (for beliefs) or a successor relation (for time).

Definition 2. *[12, Def.2,7] A Kripke structure $M = \langle W, \mathcal{R}_1, \ldots, \mathcal{R}_n, L\rangle$ is comprised of three components. The first component is an arbitrary non-empty set W containing all worlds relevant to M. The second component is a family of relations $\mathcal{R}_i \subseteq W \times W$. The remaining third component is an assignment function L. This function assigns a particular set of worlds to each primitive proposition. $L(p)$ contains all those worlds in which p holds.*

A Kripke structure $M = \langle \mathcal{S}, \mathcal{R}, \mathcal{B}, \mathcal{G}, \mathcal{I}, L\rangle$ forms a situation structure if each of the following three conditions is met.

1. *$\mathcal{S}$ is a set of situations.*
2. *$w = w'$ whenever $\langle w, s\rangle \mathcal{R} \langle w', s'\rangle$.*
3. *$s = s'$ whenever $\langle w, s\rangle \mathcal{B} \langle w', s'\rangle$ and similarly for $\mathcal{G}$ and $\mathcal{I}$.*

Schild [12, Section 3] does not present the semantic relation of CTL, but only the one of its extension CTL* (as well as the one of the μ-calculus). This extension is not considered in this paper. A speciality of both CTL and CTL* is that some formulae are not interpreted relative to a particular situation. These are the path formulae. What is relevant here are full paths. The reference to M is omitted whenever it is understood.

Definition 3. *A full path in* $M = \langle \mathcal{S}, \mathcal{R}, \mathcal{B}, \mathcal{G}, \mathcal{I}, L \rangle$ *is an infinite sequence* $\chi = \delta_0, \delta_1, \delta_2, \ldots$ *such that for every* $i \geq 0$, δ_i *is an element of* $\mathcal{S}$ *and* $\delta_i \mathcal{R} \delta_{i+1}$. *We say that a full path starts at* δ *iff* $\delta_0 = \delta$. *We use the following convention. If* $\chi = \delta_0, \delta_1, \delta_2, \ldots$ *is a full path in* M, *then* χ^i $(i \geq 0)$ *denotes* δ_i.

Let $M = \langle \mathcal{S}, \mathcal{R}, \mathcal{B}, \mathcal{G}, \mathcal{I}, L \rangle$ *be a situation structure,* δ *a situation, and* χ *a full path. The semantic relation* $\models$ *for* BGI_{CTL} *is then defined as follows:*

S1 $\delta \models p$ *iff* $\delta \in L(p)$.
S2 $\delta \models \alpha \wedge \beta$ *iff* $\delta \models \alpha$ *and* $\delta \models \beta$.
$\delta \models \neg\alpha$ *iff* $\delta \models \alpha$ *does not hold.*
S3 $\delta \models E\alpha$ *iff there is at least one full path* χ *in* M *starting at* δ *s.t.* $\chi \models \alpha$.
$\delta \models A\alpha$ *iff for every full path* χ *in* M *starting at* $\delta, \chi \models \alpha$.
S4 $\delta \models B(\alpha)$ *iff for every* $\delta' \in \mathcal{S}$ *such that* $\delta \mathcal{B} \delta', \delta' \models \alpha$.
$\delta \models G(\alpha)$ *iff for every* $\delta' \in \mathcal{S}$ *such that* $\delta \mathcal{G} \delta', \delta' \models \alpha$.
$\delta \models I(\alpha)$ *iff for every* $\delta' \in \mathcal{S}$ *such that* $\delta \mathcal{I} \delta', \delta' \models \alpha$.
P $\chi \models X\alpha$ *iff* $\chi^1 \models \alpha$.
$\chi \models \alpha \cup \beta$ *iff there is a* $i \geq 0$ *s.t.* $\chi^i \models \beta$ *and for all* $j (0 \leq j < i), \chi^j \models \alpha$.

2.2 $\mathbf{BDGI_{CTL}}$

We extend $\mathrm{BGI_{CTL}}$ with a new modal operator.

Definition 4. *The admissable formulae of* $BDGI_{CTL}$ *are defined by the formation rules of* BGI_{CTL} *together with the following rule.*

S4 If α *is a state formula, then* $D(\alpha)$ *is a state formula as well.*

For desires, we introduce minimal or non-normal modal logic [3,6]. The accessibility relation for desires is a binary relation between situations and sets of situations $\delta \mathcal{D} \{\delta_1, \ldots, \delta_n\}$.

Definition 5. *An extended Kripke structure* $M = \langle W, \mathcal{R}_1, \ldots, \mathcal{R}_n, L \rangle$ *is defined like a Kripke model, with the exception that we have either* $\mathcal{R}_i \subseteq W \times W$ *or* $\mathcal{R}_i \subseteq W \times 2^W$.

An extended Kripke structure $M = \langle \mathcal{S}, \mathcal{R}, \mathcal{B}, \mathcal{D}, \mathcal{G}, \mathcal{I}, L \rangle$ *forms an extended situation structure if each of the following three conditions is met.*

1. $\langle \mathcal{S}, \mathcal{R}, \mathcal{B}, \mathcal{G}, \mathcal{I}, L \rangle$ *is a situation structure,*
2. $\mathcal{D} \subseteq \mathcal{S} \times 2^{\mathcal{S}}$,
3. $s = s'_i$ *for* $i = 1, \ldots, n$ *whenever* $\langle w, s \rangle \mathcal{D} \{\langle w'_1, s'_1 \rangle, \ldots, \langle w'_n, s'_n \rangle\}$

We choose the following definition of a semantic relation.[1]

Definition 6. *Let* $M = \langle \mathcal{S}, \mathcal{R}, \mathcal{B}, \mathcal{D}, \mathcal{G}, \mathcal{I}, L \rangle$ *be an extended situation structure and* δ *be a situation. The semantic relation* $\models$ *for* $BDGI_{CTL}$ *is the semantic relation* $\models$ *of* BGI_{CTL} *extended with:*

S4 $\delta \models D(\alpha)$ *iff* $\exists \{\delta_1, \ldots, \delta_n\}$ *such that* $\delta \mathcal{D} \{\delta_1, \ldots, \delta_n\}$, $\delta_i \models \alpha$ *for* $i = 1, \ldots, n$.

[1] Alternatively, we could have chosen an operator which is not closed under logical consequence, i.e. such that $\not\models D(p \wedge q) \rightarrow D(p)$. See [3,6] for details.

3 Logical relations between desires, goals and beliefs

3.1 Belief merging

In this section we briefly repeat Konieczny and Pino-Pérez operator Δ for merging belief bases with integrity constraints introduced in [4]. We first give some preliminary definitions.

Definition 7. *[4, Section 2] We consider a propositional language L over a finite alphabet P of propositional atoms. A belief base K is finite set of propositional formulae, which can be seen as the conjunction of the formulae of K. A belief set E is a multi-set of belief bases. We write $K \wedge K'$ for $K \cup K'$, $\bigwedge E$ for the conjunction of belief bases of E and $\sqcup$ for union of multi-sets. By abuse of language we write K for the multi-set $\{K\}$.*

A belief set E is consistent iff $\bigwedge E$ is consistent. Two belief sets E_1 and E_2 are equivalent, noted $E_1 \leftrightarrow E_2$, iff there exists a bijection f from $E_1 = \{K_1^1, \dots, K_1^n\}$ to $E_2 = \{K_2^1, \dots, K_2^n\}$ such that $f(K) \leftrightarrow K$.

Definition 8. *[4, Section 3] Let E be a belief set, IC a belief base (integrity constraints), and Δ an operator that assigns to each belief set E and belief base IC a belief base $\Delta_{IC}(E)$. Δ is a belief merging operator if and only if it satisfies the following properties:*

B0 $\Delta_{IC}(E) \vdash IC$
B1 *If IC is consistent, then $\Delta_{IC}(E)$ is consistent*
B2 *If $\bigwedge E \wedge IC$ is consistent, then $\Delta_{IC}(E) = \bigwedge E \wedge IC$*
B3 *If $E_1 \leftrightarrow E_2$ and $IC_1 \leftrightarrow IC_2$, then $\Delta_{IC_1}(E_1) \leftrightarrow \Delta_{IC_2}(E_2)$*
B4 *If $K \vdash IC$, $K' \vdash IC$, and $\Delta_{IC}(K \sqcup K') \wedge K$ is consistent, then $\Delta_{IC}(K \sqcup K') \wedge K'$ is consistent*
B5 $\Delta_{IC}(E_1) \wedge \Delta_{IC}(E_2) \vdash \Delta_{IC}(E_1 \sqcup E_2)$
B6 *If $\Delta_{IC}(E_1) \wedge \Delta_{IC}(E_2)$ is consistent, then $\Delta_{IC}(E_1 \sqcup E_2) = \Delta_{IC}(E_1) \wedge \Delta_{IC}(E_2)$*
B7 $\Delta_{IC_1}(E) \wedge IC_2 \vdash \Delta_{IC_1 \wedge IC_2}(E)$
B8 *If $\Delta_{IC_1}(E) \wedge IC_2$ is consistent, then $\Delta_{IC_1 \wedge IC_2}(E) \vdash \Delta_{IC_1}(E)$*

Most of these postulates are generalizations of belief revision postulates. (B0) states that the result of merging complies with the integrity constraints. (B1) ensures that, when the integrity constraints are consistent it is always possible to abstract a coherent piece of information from the belief set. (B2) says that, if possible, the result of the merging is simply the conjunction of the belief bases of the belief set with the integrity constraints, (B3) is the principle of irrelevance of syntax. The purely 'merging' postulates are (B4), (B5) and (B6). (B4) is what is called the fairness postulate. It ensures that when merging two belief bases, it cannot give full preference to one of them. (B5) and (B6) correspond to Pareto's conditions in social choice theory. Finally (B7) and (B8) state conditions on the conjunction of integrity constraints and make sure that 'closeness' is well-behaved. See [4, 5] for further details and motivations.

3.2 Merging operators defined on situation structures

Intuitively, the operator $\mathcal{S}$ is the merging operator of an extended situation structure $M = \langle \mathcal{S}, \mathcal{R}, \mathcal{B}, \mathcal{D}, \mathcal{G}, \mathcal{I}, L \rangle$ iff for each situation δ, integrity constraints $IC = \{\alpha \mid \delta \models B\alpha\}$ and set $E = \{\{\alpha \mid \delta_i \models \alpha \text{ for } i = 1 \dots n\} \mid \delta \mathcal{D} \{\delta_1, \dots, \delta_n\}\}$ we have that $\Delta_{IC}(E) = \{\alpha \mid \delta \models G\alpha\}$.

This intuition explains our motivation to use non-normal modal logic. First, note that the multi-set E as defined above can contain several identical elements, because distinct situations can satisfy the same propositions. Second, suppose that $\mathcal{D}$ is a normal accessibility relation between worlds. In that case, we could only define a logically closed set of propositions like we defined IC and $\Delta_{IC}(E)$, not a set of conflicting ones.

However, this definition is still preliminary. One of the reasons is that postulate (B0) implies $B(p) \rightarrow G(p)$. As discussed in the introduction, this property called CL realism is considered counterintuitive [10].

3.3 Desire merging operators

The following definition defines merging for desires. Postulate B0 has been weakened, and the other postulates B1-B8 have been updated to incorporate this difference. We define desire bases and desire sets analogous to belief bases and belief sets.

Definition 9. *A desire base K is finite set of propositional formulae, which can be seen as the conjunction of the formulae of K. A desire set E is a multi-set of desire bases. We write $\bigwedge E$ for the conjunction of desire bases of E, etc.*

Definition 10. *Let E be a desire set, IC a belief base (integrity constraints), and ∇ an operator that assigns to each desire set E and belief base IC a desire base $\nabla_{IC}(E)$. ∇ is a desire merging operator if and only if it satisfies the following properties:*

D0 *If IC is inconsistent, then $\nabla_{IC}(E) = IC$*
D1 *If IC is consistent, then $\nabla_{IC}(E) \wedge IC$ is consistent*
D2a $\bigwedge E \vdash \nabla_{IC}(E)$
D2b *If $\bigwedge E \wedge IC$ is consistent, then $\nabla_{IC}(E) \vdash \bigwedge E$*
D3 *If $E_1 \leftrightarrow E_2$ and $IC_1 \leftrightarrow IC_2$, then $\nabla_{IC_1}(E_1) \leftrightarrow \nabla_{IC_2}(E_2)$*
D4 *If $K \vdash IC$, $K' \vdash IC$, and $\nabla_{IC}(K \sqcup K') \wedge K$ is consistent, then $\nabla_{IC}(K \sqcup K') \wedge K'$ is consistent*
D5 $\nabla_{IC}(E_1) \wedge \nabla_{IC}(E_2) \vdash \nabla_{IC}(E_1 \sqcup E_2)$
D6 *If $\nabla_{IC}(E_1) \wedge \nabla_{IC}(E_2) \wedge IC$ is consistent, then $\nabla_{IC}(E_1 \sqcup E_2) = \nabla_{IC}(E_1) \wedge \nabla_{IC}(E_2)$*
D7 *If $IC_1 \wedge IC_2$ is consistent, then $\nabla_{IC_1}(E) \vdash \nabla_{IC_1 \wedge IC_2}(E)$*
D8 *If $\nabla_{IC_1}(E) \wedge IC_1 \wedge IC_2$ is consistent, then $\nabla_{IC_1 \wedge IC_2}(E) \vdash \nabla_{IC_1}(E)$*

We choose for the borderline case in which the integrity constraints are inconsistent, that the output is a contradiction (D0). Moreover, we split the postulate

D2 in two postulates D2a and D2b. The meaning of D1-D8 is analogous to the meaning of B1-B8. The differences are that it is no longer implied that IC follows from ∇ (D2, D7, D8), but IC is added to every the consistency check (D1, D6, D8). If a belief set consists of only one belief base K, then merging the belief set under integrity constraints IC is equivalent to revising K by IC, expressed by $K * IC$. Under similar conditions, merging a single desire base under integrity constraints IC is an operation known in the literature as 'severe withdrawal' [11] and expressed by $K - \neg IC$, where we write $\neg IC$ for $\{\neg \bigwedge IC\}$. This is basically a contraction operator that does not satisfy the recovery postulate, in our terminology $\nabla_{IC}(K) \wedge \neg IC \vdash K$.

The following proposition generalizes the well known Levi identity for merging operators and shows the close relation between belief merging and desire merging.

Theorem 1. *If ∇ is a desire merging operator, then we have that Δ defined by $\Delta_{IC}(E) = \nabla_{IC}(E) \wedge IC$ is a belief merging operator.*

4 Summary

Rao and Georgeff's $\mathrm{BDI}_{\mathrm{CTL}}$ is one of the most popular specification or verification languages for cognitive agent systems. It describes the balance between beliefs, goals and intentions, but it does not describe where the beliefs and goals come from. However, in most agent systems these beliefs and goals are the result of merging respectively belief bases and candidate goals such as desires, obligations and norms. Logical analyses of the derivation of goals from desires using Reiter's default logic has been given in [2], a logical analysis of conflicts between desires has been given in [7], and the merging of candidate goals has been addressed in [5]. In this paper we have introduced merging in the specification language by introducing a new non-normal modal operator for desires. In a similar way we can build the extension of $\mathrm{BDGI}_{\mathrm{CTL}}$ to $\mathrm{KBDGI}_{\mathrm{CTL}}$ with an additional operator K for belief sources, which have to be merged into beliefs. Another extension of $\mathrm{BDI}_{\mathrm{CTL}}$ with belief revision operators can be found in [9].

As an example we have considered a recent merging operator proposed by Konieczny and Pino-Perez. We discussed one postulate of their logic, namely (B0), and we argued that this postulate is too strong, because it implies the Cohen-Levesque notion of realism expressed by $B(p) \rightarrow G(p)$. Moreover, we introduced a weakened version of their postulates. Whereas their operator is a generalization of revision, our operator is a generalization of (severe) withdrawal.

We believe that this paper is the beginning of a study of merging desires into goals in the context of beliefs. First, interesting properties have to be identified that can be expressed in the $\mathrm{BDGI}_{\mathrm{CTL}}$ language. Second, the other belief merging postulates have to be considered, as well as other interesting merging operators proposals in the literature such as [1, 8]. Third, the merging operator should not be fixed for the whole model but allowed to vary. Fourth, merging operators should not only be defined for finite sets of propositions, because even for a finite set of propositional atoms there are infinite sets of CTL formulas.

Fifth, the language $\mathrm{BDGI}_{\mathrm{CTL}}$ must be extended such that it can distinguish – like the extended situation structure – the number of occurrences of a desire base. Sixth, the merging constraints should be axiomatized in an extension of $\mathrm{BDGI}_{\mathrm{CTL}}$. Seventh, to cover the analysis in [2], the merging operators have to be extended to conditionals. Eighth, as suggested by Richard Booth, sources should be represented explicitly in the formal framework such that we can distinguish between a belief and desire base from the same source, and from distinct sources.

Acknowledgements

Thanks to Richard Booth for pointing us to the relation between desire merging and (severe) withdrawal.

References

1. R. Booth. Social contraction and belief negotiation. In *Proceedings of the Eighth International Conference on Principles of Knowledge Representation and Reasoning (KR2002)*, pages 375–384. Morgan Kaufmann, 2002.
2. J. Broersen, M. Dastani, and L. van der Torre. Realistic desires. *Journal of Applied Non-Classical Logics*, To appear.
3. B. Chellas. *Modal Logic: An Introduction*. Cambridge University Press, 1980.
4. S. Konieczny and R. Pino Pérez. Merging with integrity constraints. In *Proceedings of the Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'99)*, volume 1638 of *LNCS*, pages 135–144, 1999. Extended version: S. Konieczny and R. Pino Pérez. Merging information under constraints: a qualitative framework. *Journal of Logic and Computation*, to appear.
5. S. Konieczny and R. Pino Pérez. On the frontier between arbitration and majority. In *Proceedings of the Eigth International Conference on Principles of Knowledge Representation and Reasoning (KR'02)*, 2002.
6. S. Kraus, K. Sycara, and A. Evenchik. Reaching agreements through argumentation; a logical model and implementation. *Articial Intelligence*, 104:1–69, 1998.
7. J. Lang, L. van der Torre, and E. Weydert. Two types of conflicts between desires (and how to resolve them). In *Proceedings of GTDT2001, AAAI Spring Symposium*, 2001.
8. T. Meyer. On the semantics of combination operations. *Journal of Applied Non-Classical Logics*, 11(1-2):59–84, 2001.
9. A. Rao and M. Georgeff. The semantics of intention maintenance for rational agents. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 704–710, 1995.
10. A.S. Rao and M.P. Georgeff. Decision procedures for BDI logics. *Journal of Logic and Computation*, 8(3):293–342, 1998.
11. H. Rott and M. Pagnucco. Severe withdrawal (and recovery). *Journal of Philosophical Logic*, 28:501–547, 1999.
12. K. Schild. On the relationship between BDI logics and standard logics of concurrency. *Autonomous Agents and Multi agent Systems*, 3(3):259–283, 2000.
13. Wooldridge M. J. and Jennings N. R. Software Engineering with Agents: Pitfalls and Pratfalls. *IEEE Internet Computing*, 3(3):20–27, May/June 1999.