A Soft COP Model for Goal Deliberation in a BDI Agent

John Thangarajah[1], James Harland[1], and Neil Yorke-Smith[2]*

[1] RMIT University, Melbourne, Australia. {johthan, jah}@cs.rmit.edu.au
[2] Artificial Intelligence Center, SRI International, USA. nysmith@ai.sri.com

Abstract. Agent systems, such as those used to control robots, make decisions about their actions and take into account changes in the surrounding environment. The agent's reasoning includes deliberating about its goals, such as whether to adopt an additional goal, to prioritize or reprioritize its goals, and to suspend some goals. In popular agent systems, such as those based around the Belief-Desire-Intention (BDI) architecture, deliberation is usually qualitative only, in that goals are dropped when they are found to be in conflict with other goals, or no longer believed to be possible, rather than as a means of increasing a measure of utility. In this paper we add a quantitative dimension to this reasoning process by formulating it as a Constraint Optimization Problem (COP). This allows us to incorporate preferences and other utility measures. We describe some criteria relevant to the reasoning process. The resulting model is able to encompass multiple aspects of agent deliberation, enabling the agent to make decisions that take into account more options and sources of information than it could by breaking the deliberation into components across its decision cycle.

1 Introduction

Agent-oriented programming is often used for devices such as robots that must operate in complex and dynamic environments. A key element of agent systems is the balance between *proactive* behaviour, i.e., pursuing goals, and *reactive* behaviour, i.e., responding to environmental changes. Accordingly, the execution cycle of such agents involves the interleaving of performing actions to achieve goals, sensing environmental changes, and deliberating over the right actions to perform.

Constraint programming can, on one hand, benefit from agents developed for distributed constraint solving, and, on the other, serve the development of agent systems, particularly in modelling multiagent coordination [9, 3].

Our work centres on the Belief-Desire-Intention (BDI) model of agency [19], which has become the predominant architecture for the design of cognitive agents. The BDI model provides an explicit, declarative representation of three key mental structures: informational attitudes about the world (beliefs), motivational attitudes on what to do (desires), and deliberative commitments to act (intentions). This explicit representation enables ready inspection of an agent's operation, both by an observer and by the agent itself, thus supporting reflective capabilities such as explanation and redirectability.

One important decision that a rational agent will deliberate over is that of *goal adoption*: under what circumstances to add a potential goal g to the set of its existing adopted goals G^A. Whereas the goals a rational agent would adopt and the intentions that would follow from them, must be consistent — for an agent to act consistently and

* Corresponding author

effectively in the world — its potential *candidate goals* may be inconsistent with one another, the state of the world, and the means of the agent to act upon them.

The desires and also the candidate goals of an agent can be inconsistent for internal or external reasons. Internally, the agent may have conflicting desires (as humans often have), and potentially or necessarily conflicting plans to achieve certain candidate goals. Externally, the agent may have conflicts between its desires and those of other agents, and between its desires and societal obligations and norms. Consistent goals may become inconsistent because of changes in the environment. Not least, goals may be potentially or necessarily inconsistent because of limited resources. Goal adoption is one component of *goal deliberation*, the process by which an agent balances its desires and candidate goals with its current set of goals and intentions.

Previous work has most commonly modelled aspects of the goal deliberation process using logical or decision-theoretic formulations. In this paper we present a broadly based model able to encompass multiple aspects of the deliberation, by characterizing the agent's reasoning as generation of and choice among possible future mental states. We develop the reasoning process as a soft Constraint Optimization Problem (COP) [22]. By formulating such a soft COP and solving it with objective criteria according to its nature, an agent can perform its goal deliberation, coming to decisions that take into account more options and sources of information than it could by breaking the deliberation into components across its decision cycle.

2 Background and Related Work

Of the numerous agent frameworks published in the literature in recent years, many, although not all, are based on the BDI model. Three key aspects of agency are captured by the BDI model: informational attitudes about the world (beliefs), motivational attitudes on what to do (desires), and deliberative commitments to act (intentions) [19].

Recent developments have begun to address the gap between theoretical BDI models and practical BDI systems. An important aspect of the theory–practice gap is between motivational attitudes that may be not consistent and feasible (*candidate goals*, we will call them) and those that must be (*goals*, we will call them). As discussed by [7] and others, many formalizations of BDI equate these two types of goals.

A critical aspect of the representation of goals is whether goals are viewed as declarative or procedural (or both), and there are various examples of both aspects in the literature on goals in agent systems [20, 34, 6]. The distinction is sometimes described as *goals to perform* versus *goals to be achieved*. The latter are naturally represented as desired states of the system, and accordingly theoretical BDI frameworks represent desires as *states* to be accomplished. On the other hand, most implemented systems represent goals as *tasks* to be performed. For consistency with these conventions, we will represent desires as states and goals as tasks. However, critically for goal deliberation, we include a declarative representation of goals. This allows us to form a direct link between the agent reasoning process and its implementation. State-to-task translation, while reasonably straightforward, is not the focus of this paper.

The decision cycle in BDI architectures follows a three-step pattern: Observe (Perceive), Decide (Deliberate), and Act, as shown in Figure 1. In general, an agent will be pursuing a particular course of action (Act) before pausing to gather input from the environment (Perceive). After processing this input, the agent may deliberate over whether

```
while( true )
   events := observe-external-events()
   drop-successful-attitudes( events, B, D, I )
   adopt-new-attitudes( events, B, D, I )
   options := generate-options( events, B, D, I )
   selected-options := option-deliberate( options, B, D, I )
   update-intentions( selected-options, I )
   execute( I )
```

Fig. 1. BDI decision cycle

to continue with its current course of action (i.e., intentions), or to modify it in response to environmental changes. Hence there is a distinction between decisions made as part of its course of action, and meta-decisions about what course of action is appropriate.

This tension between responding to environmental changes (being *reactive*) and pursuing a course of action (being *proactive*) is fundamental to agent systems. The balance between reactive behaviour and proactive behaviour will depend on the application domain, and in particular how often changes in the environment occur compared to actions of the agent.

The objective of our work is to consider the broad aspects of reasoning encompassed by a rational agent's self-reflection and deliberation over its commitments. This reasoning includes goal adoption and intention reconsideration, but encompasses more than either of these important single aspects on their own.

We are by no means the first to recognize the importance of commitment deliberation in the BDI model. Previous work has most commonly modelled aspects of the commitment deliberation process using logical or decision-theoretic formulations.

Dastani and van der Torre [5] distinguish desires from goals, formalizing a $\mathrm{BDGI}_{\mathrm{CTL}}$ logic and considering a belief-theoretic approach to merging desires into goals. The role of desires in the agent's adoption of goals is a precursor to the reasoning we include within goal deliberation. Dignum et al. [7] consider a multiagent context, in which an agent has desires, obligations and norms, and from these derives and maintains a set of goals: the B-DOING agent framework. Since the three motivational attitudes are distinguished in this framework, the problem of goal adoption now involves reconciling three potentially conflicting sets of attitudes. Our work leaves aside the origin of the agent's motivational attitude, which we collate as simply *desires*; our focus is on the deliberation of a single agent over its commitments. Note that one source of commitments can be other agents, in that an agent may take on a commitment at the request of another. Such interactions are outside the scope of this paper.

The BOID agent framework [2, 6] sees goals arising out of the interaction between different cognitive aspects — obligations as well as beliefs, desires and intentions — and the resolution of conflicts among them according to the nature (character) of the agent. Sets of goals are selected according to rule-based reasoning with priorities. Dastani and van der Torre [6] present the syntax and semantics of a language for programming BOID agents. It features a notion of the agent's cognitive state of current beliefs, goals, intentions, and plans, but not desires. As fits the BOID framework, goal generation, goal adoption, plan generation, and plan selection are separate steps that the programmer can specify in a decision cycle.

Schut et al. [24] develop a decision-theoretic approach to intention reconsideration, in which the agent meta-deliberates over the expected cost of the computation versus the expected utility to be gained. Raja and Lesser [18] present a similar utility-based formulation in terms of a Markov Decision Process.

Stroe et al. [27] consider a problem subsequent to ours, namely, action (plan) selection, but adopt a similar approach. They associate an objective function over agent status sets, and compute the ideal agent action as that induced by the status set optimizing the objective. Natarajan et al. [16] also consider optimal agent activity, in the context of an assistive agent that seeks to minimize the user's effort in a joint activity. Selection of the optimal action is according to a combined logical and probabilistic model.

Meneguzzi et al. [12] explore the BDI reasoning process by mapping BDI mental states to propositional planning problems. Like our work, their intent is to leverage established AI representations and models to give a global view of BDI agent deliberation.

Informing Goal Deliberation Goal deliberation cannot operate in a vacuum of information. Thangarajah et al. [30, 32, 31] provide information for an agent to deliberate over whether to adopt a new goal g, by computing *summary information* that captures goal effects. These summaries are computed at each node of a *goal-plan tree*. An instance of a goal-plan tree consists of the possible subgoals of g that can arise according to the plans that the agent can apply, in order to achieve an intention that fulfills g.

A similar approach is applied to analyze potential resource consumption and production by goals [30, 33]. Morley et al. [13] extend the representation to include task parameters and a richer set of plan language constructs, and dynamically update the resulting resource estimates as execution proceeds.

Shaw and Bordini [25] cast the goal adoption problem as a reachability problem in Petri nets. They are developing an automated translation from BDI agent mental state to Petri net formulation. This formulation enables the agent to consider potential negative and positive interactions between a postulated new adopted goal g and the existing goal set G^A without maintaining the explicit goal summaries of Thangarajah et al.

Shaw and Bordini [25] also describe in-progress work to develop a method based on compiling the goals and possible plans of an agent into a Constraint Satisfaction Problem (CSP), exploiting ideas from the AI planning community [28]. This formulation enables a method, given an agent's current goals G^A, candidate goal g, and current plans, for generating a CSP whose solutions correspond to valid ordering of plans that minimize goal effect conflicts (and exploit serendipitous opportunities) when attempting to achieve $g \cup G^A$. We share the use of a constraint-based model in agent cognition, not to analyze potential goal effects (one aspect of information that informs goal deliberation), but to model the whole reasoning process itself.

Whereas agents have been developed for distributed constraint solving, constraints in service of agent research have focused on modelling multiagent consultation, cooperation, and competition [9, 3]. Constraints are employed by Chalmers et al. [4] not in service of the agent's reasoning, but to specify a desired organization state. A BDI agent is constructed such that the plans it decides upon correspond to generation of possible organizations matching the specification, via solving of a Constraint Logic Program that represents domain knowledge. Perhaps the closest to our work in the CP community are developments along the lines of the agent framework of Mackworth [11], in which constraint-based controllers are developed for the heart of agent reasoning engines.

3 Goal Deliberation

Goal deliberation lies at the heart of BDI architectures: how should an agent balance its desires and candidate goals with its current set of goals and intentions? This balancing problem embraces two fundamental decisions: determining when to adopt a new goal (*goal adoption*) and when to terminate an existing intention (*intention reconsideration*). Potential conflicts between commitments are the reason a rational agent will deliberate.

As we have presented, the agents literature is for the large part silent on these questions [30, 5]. In practice, most agent systems skirt the goal adoption problem by making the simplifying assumption that desires and goals can be equated, and by not distinguishing candidate and adopted goals over an extended life cycle. Every candidate goal is adopted. Potential conflicts and synergies between goals are ignored. Intention reconsideration then is driven solely by problems encountered as execution proceeds. At best, goals are dropped when they are found to be in conflict with other goals, or no longer believed to be possible.

We view both activities as components of the larger cognitive function of deliberating over mental states. For this reason, we propose an approach in which goal adoption and intention reconsideration are inextricably linked in a *goal deliberation* process.

To express the mental attitudes that bear upon this process, we adopt an extended BDI model that distinguishes desires, candidate and adopted goals, and intentions [15]. We define the mental state of the agent as a tuple of sets of these elements. Next, we characterize the agent's goal deliberation as generation of and choice among possible future mental states. We then present, in the next section, a general proposal for implementing goal deliberation as a soft Constraint Optimization Problem.

3.1 BDGI Cognitive Model

We relate an informal description of the BDGI framework [15] which will be the basis for our model of agent cognition.

- **Beliefs** are the agent's accepted knowledge about the world and its current state. Example: "Alice and Bob are the best fits on paper for the open position".
- **Desires** are states that describe the agent's motivations. As discussed earlier, we consider only a single type of motivational attitude; we do not distinguish between (internally motivated) desires, and (externally motivated) obligations and norms.
- **Candidate Goals** are tasks that provide the agent's motivation. That is, candidate goals characterize tasks that the agent would like to accomplish, but may not be able to because of extenuating factors. Candidate goals need not be consistent with each other, with the state of the world, or with adopted intentions. For example, "I wish to interview Alice at 9 a.m. on Monday", and "I wish to interview Bob at 9 a.m. Monday" constitute conflicting candidate goals.
- **Goals** are a consistent, feasible set of tasks, which are derived from the agent's candidate goals. As noted, implemented BDI systems suppose candidate goals to be consistent. We below explain what we mean by consistent and feasible. Example: "I want to invite Alice" (and it is possible).
- **Intentions** are committed goals together with the means to achieve them (plans). Intentions represent the tasks that have successfully passed the conditional aspects; as far as it knows, the agent can achieve these tasks, and it has committed to doing so. Example: "I am inviting Alice" (by emailing her three possible dates).

- **Goal-Advice** are constraints on the adoption of candidate goals as goals. Examples: “Don’t invite multiple candidates on the same day”, “Attend only one conference this summer”.
- **Execution-Advice** are constraints for user directability of problem solving. Whereas Goal-Advice advises the agent what to do, Execution-Advice advises the agent how to do it. Examples: “Use email to invite candidates”, “Fly on a US carrier”.
- **Plans** are the means to achieve intentions (the standard BDI interpretation). Plans are selected from a library of alternative recipes for achieving an intention.

We define four conditions that a set of adopted goals G^A must satisfy. Here, consistency implies consistency w.r.t. a fixed background theory of domain constraints.

- **Self-consistency:** G^A must be mutually consistent
- **Coherence:** G^A must be mutually consistent relative to the current beliefs *B*
- **Feasibility:** G^A must be mutually satisfiable relative to current intentions *I* and available plans
- **Reasonableness:** G^A should be mutually reasonable w.r.t. current *B* and *I*

The requirements of self-consistency and coherence are common to most BDI frameworks [21]. Feasibility requires that goals be adopted only if they can be achieved.

The requirement of reasonableness requires a background domain theory. It plays an important role in personal assistant agents [14], where for example the user may pose desires that are consistent, coherent, and feasible, but not in her best interest. This type of situation can arise because a user lacks awareness of problem-solving history, commitments, or constraints. For example, suppose that the user asks her assistive agent to purchase a laptop computer; the next day, having forgotten about the request issued the prior day, she requests the agent again to purchase a computer. If the user worked in procurement, such a request may seem reasonable. However, for the typical office worker, the purchase of a computer is a relatively rare event. In this case, it would be helpful for the agent to recognize that the request is unreasonable, and check with the user as to whether she really wants the agent to proceed with the purchase.

3.2 The Goal Deliberation Process

Many formal models cast goal adoption as a filtering problem, where some maximal subset of stated desires is identified that satisfies designated requirements (such as those above). We argue that this approach is inappropriate in many situations, as it places undue emphasis on current goals. More generally, the goal adoption process should admit the possibility of modifying beliefs (through acting to change the state of the world), adopted goals, and intentions in order to enable adoption of new candidate goals.

The desires and other motivational attitudes of an agent lead it to accept or generate new candidate goals. For the remainder of the paper, we focus on the decision making process from candidate goals through to intentions.

There is little reason for an agent not to adopt a candidate goal, provided its inclusion as an adopted goal satisfies the conditions of consistency (i.e., non-conflicting), coherence, feasibility and reasonableness as stated earlier. Goal deliberation becomes significant when the candidate goal g would lead to an adopted goal set G^A that no longer satisfies the above conditions.

Among its possible decisions, the agent may choose to not adopt g at present, or to *reject* it entirely. If the new candidate goal causes conflict with an existing goal (or goals), for example because of conflict over a reusable resource, then the agent may delay adopting g or *suspend* the currently adopted goal(s), depending on which it believes it should be more committed. On the other hand, the nature of the conflict could lead to the agent *dropping* the existing goal(s), or to it creating a new candidate goal and *directly adopting* it in order to enable g to be addressed at a future stage.

As a concrete example, suppose that the user of an assistive agent desires to attend the AAMAS conference in Europe but lacks sufficient travel funds. To enable attendance, the user could shorten a previously scheduled trip for a different meeting, or she could cancel the planned purchase of a new laptop. Alternatively, the user could apply for a travel grant from the department to cover the costs of the European trip.

A typical BDI agent executes a tight control loop for determining the actions that it will perform, as shown earlier in Figure 1. Let $S = (B, D, G^C, G^A, I)$ denote the agent's current mental state at the start of a cycle through this loop. The five cognitive elements are, respectively, Beliefs, Desires, Candidate Goals, Adopted Goals, and Intentions. Recall that the control flow involves identifying modifications to S from the prior cycle (Observe), deciding what to do in response to those changes (Decide), and then performing an appropriate set of actions (Act).

Goal deliberation fits in neatly as an initial step within this control loop. Thus, to start a cycle of the control loop, the agent first performs some deliberation to determine whether it needs to revise its goals in light of the changes during the past cycle. In general, this deliberation can transition the agent's mental state to a new state $S' = (B', D', G^{C'}, G^{A'}, I')$. We next consider the possible transitions in a constraint formulation of the reasoning process.

4 Constraint Model of Goal Deliberation

Informally, we take a *Constraint Satisfaction Problem* (CSP) to consist of a set of variables V, a set of corresponding domains D that specify the possible values for the variables, and a set of constraints C [22]. A solution is a complete assignment s to all variables in V that satisfies all constraints in C. A *Constraint Optimization Problem* (COP) is a CSP together with an objective function f that evaluates tuples of variables. An optimal solution is one that maximizes $f(s)$.

In a *soft* CSP, the requirement that a solution satisfy all constraints is relaxed. We follow the semiring-based soft CSP [22], in which a weight function w_i is attached to each constraint $C_i \in C$. The weight functions ascribe values from a set E to assignments of variables in the scope of the constraint. These values are compared according to the structure of an associated semiring. For our purposes the *fuzzy semiring* will suffice. In the corresponding fuzzy CSP, weight functions take real values in $E = [0, 1]$, are aggregated by $\min$ and $\max$, and are compared with $\geq$.

By its nature, a soft CSP is an optimization problem. In a fuzzy CSP, maximally preferred solutions are those that maximize the aggregated weight functions, i.e., maximize the minimum preference level achieved over all constraints. We write v to denote the aggregated valuation function from the set of soft constraints Thus, any additional objective function f added to a soft CSP results in a *multi-criteria* optimization problem. We discuss the optimization below.

We develop a soft constraint optimization model of a BDGI agent's reasoning over mental states as follows. Consider one cycle through the control loop. For a state $S = (B, D, G^C, G^A, I)$, define *set-valued* variables B_S, D_S, C_S, G_S, and I_S for the five components of S, respectively. Each variable takes values as subsets of a domain of possible beliefs, desires, and so on, respectively. For instance, the domain of values for B_S consist of subsets of a set $\mathcal{B}$ of possible beliefs. For example, at initialization, an accounting agent might have $B_S = \texttt{balance(\$100)}$ and $\mathcal{B} = \texttt{balance(\$X)}\ \forall X \in \mathbb{Z}$. Note that the domains may not be finite.

These underlying domains — $\mathcal{B}$, $\mathcal{D}$, $\mathcal{C}$, $\mathcal{G}$, $\mathcal{I}$ — are obtained from the background theory and the recipe library of the agent, and evolve according to perception and actions. At least for beliefs, the domain is *open*: the possible values in the domains are not fixed and known.[3] However, upon entry of the control loop, the current mental state and the underlying domains are fixed and known for the remainder of the loop. Thus the constraint problem that the agent will formulate for any one iteration through its decision cycle is closed.

We encapsulate the evolution of the underlying domains by *domain update functions*, one for each element of the mental state. These functions are invoked when the agent executes an action and when it receives a perception, and when internal events occur, such as an intention succeeding or failing. The functions serve to update mental state (e.g., assert new beliefs) and underlying domain theories (e.g., enlarge the domain of possible beliefs). They take as input the current domain, together with the current mental state, the executed action or world state change, and they produce as output an updated mental state and domain. For instance, the domain update function for beliefs thus encapsulates the beliefs that the agent could *potentially* hold, as a result of the completion of the execution cycle. In the example, after executing a web query for the cost of a purchase of a laptop, the agent's domain of possible beliefs could be $\mathcal{B} = \{\texttt{balance(\$X)}\ \forall X \in \mathbb{Z}, \texttt{laptop\_cost(\$Y)}\ \forall Y \in \mathbb{Z}^+\}$. Note that there is no reason for the underlying domains to be enumerated extensionally.

The role of the domain update functions is non-trivial. The outcome of goal deliberation is a decision over what commitments to pursue. The underlying domains for candidate goals and goals, therefore, describe the choices the agent can consider for its goals. For some agents, these choices can include goal modifications, such as modifying the specifications requested of the laptop to be purchased. Thus $\mathcal{G}$ includes all possible such modifications that the agent can consider; computation of these possibilities is highly dependent on the application domain.

Constraints stipulate the permissible transition from the current mental state S to the next state S'. They describe the possible moves $S \rightarrow S'$ in terms of the evolution of each mental attitude. Formally, the semantics of the transitions can be captured with conditioned proof rules. We consider the following representative transitions:

- *Rejection* of a new candidate goal. $S' = (B, D, G^C, G^A, I)$
- *Consideration* of a new candidate goal. $S' = (B, D, G^C \cup g, G^A, I)$; g may have awakening conditions attached.
- *Rejection* of an existing candidate goal. $S' = (B, D, G^C - g, G^A, I)$
- *Expansion* of the adopted goals by adopting a candidate goal. $S' = (B, D, G^C - g, G^A \cup g, I)$

[3] Only in simplistic application domains can we expect the closed world assumption to hold.

– *Direct adoption* of a new goal. $S' = (B, D, G^C, G^A \cup g, I)$
– *Suspension* of an adopted goal. $S' = (B, D, G^C \cup g, G^A - g, I - I_g)$, where I_g is the set of intentions associated with or descended from goal g.[4]
– *Resumption* of a suspended goal. $S' = (B, D, G^C - g, G^A \cup g, I \cup I_g)$
– *Modification* of an existing adopted. $S' = (B, D, G^C, G^A - g \cup g', I)$, where g' is a modification of the goal g.
– *Revocation* of an adopted goal. $S' = (B, D, G^C, G^A - g, I - I_g)$

We state set constraints [22, Ch. 17] corresponding to each type of transition. For instance, the constraints that describe an expansion transition are $B_{S'} = B_S$, $D_{S'} = D_S$, $C_{S'} = C_S - g$, $A_{S'} = A_S \cup g$, and $I_{S'} = I_S$. Such constraints can be stated for every type of transition that the agent's deliberation will consider. Transitions not considered by the agent need not be stated. For an agent that cannot modify an adopted goal, for example, we need not state constraints corresponding to the modification transition. The constraints form a disjunction: *rejection conjuncts* $\vee$ *consideration conjuncts* $\vee \ldots$.

The transition constraints must be satisfied, but they may have differing levels of importance according to the nature of the agent, such as its level of commitment to existing intentions [2]. For instance, a more conservative agent may adopt new goals only with reluctance; it would have a low weight on consideration, expansion, and direct adoption transitions. These constraint importances can be modelled dually in a fuzzy CSP as weights on variable assignments.

The possible transitions $S \to S'$ are constrained by the user's goal advice A^G. Such advice defines user preferences over permissible transitions. For example, the advice "Don't invite multiple candidates on the same day" gives a constraint: $\mathtt{sameday}(\mathtt{g}_1, \mathtt{g}_2) \implies g_1, g_2 \not\subseteq G$, for any pair of goals g_1 and g_2 that correspond to inviting candidates for interview. These advice-derived constraints have a lesser level of importance than the transition constraints. Since advice is soft, the agent may consider advice-infeasible transitions, with a suitable penalty in the overall assessment. Indeed, to complement goal- and execution-advice, we define *meta-advice* on how to approach the optimization problem. Whereas the other forms of advice restrict the possible transitions, the meta-advice indicates what transitions the user prefers: e.g., "Refrain from dropping any intention that is more than 70% complete".

The constraint model corresponds to one cycle through the control loop, akin to a constraint model of one transition of a deterministic finite state machine. While we could develop a non-myopic model that considers the (non-deterministic and partially-observed) future consequences — just as lookahead planning is married into BDI architectures by [23] — the single cycle model fits the reactive spirit of BDI agents.

4.1 Optimality and Solving the COP

The model described so far is a soft CSP. We now consider the optimization criteria that direct the agent's deliberation. The goals and intentions for a mental state may have associated criteria that inform the deliberation process, including

– **Goals:** Value or utility (time-varying), priority, and deadline; estimated cost of achievement. For adopted goals, there is also the level of commitment.

[4] Computation of the intentions I_g that must be suspended or aborted, as a consequence of suspending or aborting a goal g, is addressed in [29].

– **Intentions:** Cost of change (deliberative effort, loss of utility, delay); level of commitment; level of effort so far (e.g., percentage complete); estimated cost to complete; estimated probability of success; and the value or utility implied by the goal from which the intention arises.

In selecting goals and intentions, as well as plans to execute intentions, the agent seeks to maximize some combination — sophisticated or otherwise — of these and other criteria. In general, the multiple criteria, together with the soft constraint satisfaction valuation v, may not necessarily be aggregated into a single objective function. Indeed, it is well known that there is no way to satisfy all the properties one would like in multi-criteria optimization; utility functions on the components cannot be combined to a single, ideal ranking function on the possible S' [8].

Hence, the transition from S to S' is a multi-criteria optimization problem not only because of the presence of both a soft constraint preference/relaxation objective (i.e., valuation function v) and an optimality objective (i.e., f), but because the optimality objective itself most likely stems from multiple criteria. While some criteria are user specified, others derive from the agent's cognitive processes, as we discuss below.

An optimal solution to the soft COP must be developed according to the literature on constrained multi-criteria optimization [26], and the algorithms for COPs developed from it, such as [10]. A simple method to balance the preference/relaxation (coming from the soft constraints) and the optimality (coming from the objective function) is a weighted sum of the v and f. A lexicographical ordering of v before f emphasizes constraint satisfaction before optimality.

Solution methods for the constraint model of agent deliberation, while of prime practical importance if the model is to be applicable in an agent framework, are outside the focus of this paper. We note, however, that an agent may be endowed with a simple strategy for selecting an optimal S' (i.e., there may be multiple such options), and with the capability for it to be overridden by more elaborate strategies specified by the agent programmer or the user for a given situation. Second, we note that the solving may exploit similarity between the COPs of successive iterations through the decision cycle.

The agent's deliberation is expressed simply with the model, as follows. Let the variables in S be assigned according to the agent's current mental state. Then a satisfying assignment for the variables in S' corresponds to a possible future mental state. An assignment that optimizes the value of S' according to the form of optimality chosen corresponds to a decision that is maximally preferred by the agent.

Within the decision cycle of a BDGI agent, goal deliberation is as follows:

1. Observe, update current mental state to S
2. Meta-deliberate: is there need for goal deliberation, and does its expected utility outweigh the expected cost?
3. If warranted, formulate COP based on S
4. Solve COP for optimal solution(s)
5. Pick a solution and make corresponding transitions to S'
6. Act according to new mental state and current intentions

An interactive agent can benefit from collaboration with the user, according to the level of adjustable autonomy. For example, an assistive agent might consult the user if it cannot establish a clear best S', given the meta-advice. For example, it might ask "Should I give up on purchasing a laptop, in order to satisfy your decision to travel to both CP and AAMAS conferences?", if it finds the two options are in Pareto trade-off.

4.2 Example

Let us continue the example of a user being aided in arranging conference travel and laptop purchase by a personal assistant agent. Let c_1 be the candidate goal "Purchase a laptop", c_2 the candidate goal "Attend CP", and c_3 the candidate goal "Attend AAMAS". Available funds are the resource *money*, expressed by a belief `balance($X)` where X is the current amount. Suppose c_1 and c_2 have led to the following adopted goals and intentions: g_1 with intention i_1: "Purchase a high-specification laptop using general funds"; and g_2 with intention i_2: "Attend CP and its workshops, staying in the conference hotel". The agent estimates it is 90% through completion of intention i_1 and 25% through intention i_2; to change i_2 would incur a financial cancellation penalty.

Suppose the user tasks the agent with the new candidate goal c_3 to attend AAMAS in Europe. Thus, the current cognitive state is $S = \{B, \{c_1, c_2, c_3\}, \{g_1, g_2\}, \{i_1, i_2\}\}$. (For simplicity, since they do not impact this example, we will not describe agent desires.) As before, suppose adopting c_3 is infeasible because of resource contention, i.e., insufficient general funds. The agent has several alternatives, including one or more of

1. Do not adopt c_3 (i.e., don't attend AAMAS)
2. Drop c_1 or c_2 (i.e., laptop purchase or CP attendance)
3. Modify g_2 to attend only the main CP conference
4. Adopt a new candidate goal c_4 to apply for a departmental travel grant

Suppose now that the user has stated meta-advice forbidding the agent from dropping any intention, thus disallowing the second alternative above. Finally, suppose the user also gave a high priority when tasking the agent with c_3.

We have variables $B_S = \texttt{balance(\$850)}$, $C_S = \{c_1, c_2, c_3\}$, $G_S = \{g_1, g_2\}$, and $I_S = \{i_1, i_2\}$. Let $\mathcal{P}$ denote the power set of a set. The underlying domains are $\mathcal{B} = \{\texttt{balance(\$X)}\ \forall X \in \mathbb{Z}\}$, $\mathcal{C} = \mathcal{P}(\{c_1, c_2, c_3, c_4\})$, $\mathcal{G} = \mathcal{P}(\{g_1, g_2, g_3, g_4, g_2^{'}\})$, and $\mathcal{I} = \mathcal{P}(\{i_1, i_2\})$, where $g_2^{'}$ denotes the modification of goal g_2. The transition constraints include those for rejection, consideration, expansion, and so on.

The constraint derived from user advice is $I_S \subseteq I_{S'}$ ("don't drop an intention"). The user's emphasis on attending AAMAS can be modelled via the preference function of the adoption transition constraint, i.e., greater preference on assignments that include c_3 in $G_{S'}$. Alternatively, it can be modelled via an optimality criterion in the optimization formulation, i.e., a weight on c_3 and an objective function that goal sets with higher weight are preferred.

According to its nature, the agent builds the optimization problem from the criteria it considers in its goal deliberation. In this case, suppose the agent's simple behaviour is to neglect measures of commitment to existing intentions (i.e., here, the estimated percentage complete) and the cost of changing an intention, and to consider only the weight the user has given to candidate goals together with her advice. This gives an objective function $f(s) = \sum_{g \in A_S} w(g)$, where $w(g)$ is the weight on a goal. Suppose the agent resolves the multiple criteria by the aggregation into $F(s) = v(s) + f(s)$.

Solving the constraint problem gives optimal values for variables as follows: $B_{S'} = B_S$, $C_{S'} = C_S$, $A_{S'} = A_S - g_2 \cup \{g_2^{'}, g_4\}$, and $I_{S'} = I_S$. Despite the penalty incurred, the agent therefore seeks to move to an S' that corresponds to taking both the third and fourth options above: to modify g_2 and adopt g_4, i.e., $S' = \{B,\ \{c_1, c_2, c_3, c_4\},\ \{g_1, g_2, g_4\},\ \{i_1, i_2\}\}$. Thus, on successful completion of g_4, the agent is able to adopt g_3 also and maintain goal feasibility,

4.3 Criteria for the Agent's Decision Making

We have shown how modelling and solving a COP can be incorporated into the agent execution cycle. Recall from Section 3 that an agent may need to deliberate over the choice of goals to pursue if a candidate goal is to be adopted and its inclusion as an adopted goal violates any of the consistency, coherence, feasibility, and reasonableness properties. An important factor that guides this deliberation process is the commitment the agent has towards a particular goal. In the following we consider some of the criteria that may be used to determine this commitment.

Utility (time-varying), priority, and deadline Utilities can derive from a combination of two sources: explicit statement for externally tasked goals, or inference by the agent. Utility may be a time-varying quantity. For example, the utility of a conference registration task might increase the closer to the early registration deadline, then fall off until nearer the final deadline.

Estimated cost of achievement, including estimated amount of resources required In resource-critical domains, an important component of the cost of achievement is the amounts of resources required. An estimate of this can be derived and updated as execution proceeds using the methods of Morley et al. [13]. Distinguishing between consumable and reusable resource consumption is also useful as the former has a greater impact than the latter. In domains where resources are less significant, the cost of achievement can be determined by a simpler measure such as the number of steps (plans, subgoals, actions) that are required to satisfy a goal.

The dependency by other goals Internal dependencies concern how many other goals depend on the successful completion of the goal, whereas external dependencies concern how many other agents depend on it. For example, suppose the goal is to book a flight as part of a top-level goal of attending AAMAS. If cancelled, the penalty, apart from the local financial penalty, is that the conference attendance goal fails, and so forth. Similarly, suppose the goal is a delegated goal by another agent. If dropped, then any obligation to fulfill the goal may be violated.

The interactions with other goals A goal may potentially interact with the pre/post/in-conditions and the effects of other goals. An interaction can be positive (cooperates with) or negative (hinders). The computation will also consider the importance of the goals that are involved, be it positive or negative. An assessment of these interactions can be computed by the methods of Thangarajah et al. [30, 33, 32, 31]. For example, the existing conflicting goal may have steps that contribute to the achievement of many other goals, or the new goal may cause potential conflict with several existing goals.

Intentions, similarly, depend upon and affect other intentions. Note that while a goal may not cause conflict, the way the agent achieves the goal (i.e., the intention) may do so; the same applies when the interaction is positive.

The penalty of retracting the goal, or the intention Penalties are linked to commitment, since performing steps that incur a high penalty to undo should increase commitment towards pursuing this same path of achieving the goal. The penalty for retracting a goal — other than causing the goals that it depends on to fail — is the cost of aborting the goal and its current plans (intentions) [29]. For example, to retract a goal to travel to AAMAS may require forfeiting flights booked, registrations paid, and so on.

The level of effort to date The effort from the conception of an intention to the current moment is captured by a sum of the computation time and resources used in service of all plans so far attempted: the more that has been invested into the goal, the more committed the agent is to it. An alternate and simpler measure of effort is to use the number of plans and subgoals completed for a given goal.

However the agent assesses the level of effort, the import of this criterion depends on the nature of the agent. A blindly committed agent, for instance, will not give up even if it is almost certain that it cannot achieve the goal.

The estimate cost to completion An estimate can be computed from the total estimated cost to achieve the intention by means of the selected plan (see earlier remarks), together with an estimate of the progress so far derived from the level of effort to date.

The likelihood of success When conflicts are detected using the mechanisms developed by Thangarajah et al. [32, 31], they can be classified as definite or potential conflicts, and as uncertain or definite success. Of course, goals in the definite success category are of higher preference, and goals with definite conflict are what causes the agent to perform a deliberation of choice between goals.

It is reasonable to infer that the more potential conflicts that a goal is involved in, the greater the probability that it may fail. Whilst this is also true for goals in the uncertain category, the probability of failure is lesser than if the conflict was potential. More principled, Pfeffer [17] presents an approach to compute, and update as execution proceeds, an estimated probability of success for any intention. The initial value of such an estimation can take into account the definite or possible positive and negative goal interactions of Thangarajah et al.

The criteria described above are domain independent. However, their impact on the commitment towards a goal may depend on the type of domain. For example, deadlines have more of an effect on time-critical domains, while resource consumption has greater impact on resource-critical domains. In addition to the above, domain-specific criteria may be provided, for example, as user advice for a personal assistant agent.

5 Conclusion

We have investigated how to extend the usual methods of agent reasoning by incorporating the formulation and solving of a soft COP into the deliberation process. This provides a richer means of reasoning about appropriate courses of action for an agent, in that qualitative measures, such as utility measures, can be embedded into the decision-making process. In addition, we have introduced a finer granularity of goals and related concepts (desires, candidate goals, goals and intentions), which can be neatly reflected in a hierarchy of constraints, in that a constraint associated with a desire, for example, is considered softer than one associated with an intention.

From a practical standpoint, empirical investigations are necessary to understand the affordability of embedding a soft constraint optimization solver within an autonomous, reactive agent. We must understand how to effectively solve our soft COP model.

One key area of future research is to investigate particular mechanisms for determining the level of commitment an agent should have to a particular goal. In practice, most conflicts between goals in agent systems occur because of resource issues. Accordingly, the vast majority of the constraint solving will most likely be focused on resource optimization (compare [25]). To able to perform the measurements necessary for this, the relative importance of the goals will need to be established.

A second area of further work is to investigate the generation of candidate goals from desires. In principle, this is a matter of finding an optimal set of consistent goals from a given set; in practice, this may also involve prioritizing goals that have already consumed substantial resources and are considered 'almost finished', as well as estimating resource usage before accepting a goal as a candidate goal.

A third area is that of scheduling. It is often the case that resource conflicts can be resolved by rescheduling goals, rather than having to choose between them (compare [1]). This can be considered another dimension of the COP, and an area where constraint-based techniques have proved effective.

Acknowledgments The authors thank David Morley and Karen Myers for discussions, and the reviewers for their constructive comments. This material is based in part upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. NBCHD030010. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA or the DOI-NBC. The work of one author from RMIT was supported by the Australian Research Council and Agent Oriented Software under grant LP0453486.

References

1. R. H. Bordini, A. L. Bazzan, R. de O. Jannone, D. M. Basso, and R. M. Vicari. AgentSpeak(XL): Efficient intention selection in BDI agents via decision-theoretic task scheduling. In *Proc. of AAMAS'02*, pages 1294–1302, Bologna, Italy, 2002.
2. 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.
3. M. Calisti and N. Neagu. Constraint satisfaction techniques and software agents. In *Proc. of AIIA 2004 workshop on Agents and Constraints*, Perugia, Italy, Sept. 2004.
4. S. Chalmers, P. Gray, and A. Preece. Supporting virtual organisations using BDI agents and constraints. In *Proc. of 6th Intl. Workshop on Cooperative Information Agents (CIA'02)*, pages 226–240, Madrid, Spain, 2002.
5. M. Dastani and L. van der Torre. Specifying the merging of desires into goals in the context of beliefs. In *Proc. of First Eurasian Conf. on Advances in Information and Communication Technology (EurAsia ICT 2002)*, pages 824–831, Shiraz, Iran, 2002.
6. M. Dastani and L. van der Torre. Programming BOID-plan agents: Deliberating about conflicts among defeasible mental attitudes and plans. In *Proc. of AAMAS'04*, 2004.
7. F. Dignum, D. Kinny, and E. Sonenberg. From desires, obligations and norms to goals. *Cognitive Science Quarterly*, 2(3–4):407–430, 2002.
8. J. Doyle and M. P. Wellman. Impediments to universal preference-based default theories. *Artificial Intelligence*, 49(1–3):97–128, 1991.
9. P. S. Eaton, E. C. Freuder, and R. J. Wallace. Constraints and agents: Confronting ignorance. *AI Magazine*, 19(2):51–65, Summer 1998.
10. U. Junker. Preference-based search and multi-criteria optimization. In *Proc. of AAAI-02*, pages 34–40, Edmonton, Canada, 2002.
11. A. K. Mackworth. Constraint-based agents: The abc's of cba's. In *Proc. of CP-2000*, pages 1–10, Singapore, 2000.

12. F. R. Meneguzzi, A. F. Zorzo, and M. da Costa Móra. Propositional planning in BDI agents. In *Proc. of SAC'04*, pages 58–63, Nicosia, Cyprus, 2004.
13. D. Morley, K. L. Myers, and N. Yorke-Smith. Continuous refinement of agent resource estimates. In *Proc. of AAMAS'06*, Hakodate, Japan, 2006.
14. K. Myers, P. Berry, J. Blythe, K. Conley, M. Gervasio, D. McGuinness, D. Morley, A. Pfeffer, M. Pollack, and M. Tambe. An intelligent personal assistant for task and time management. *AI Magazine*, 28(2):47–61, Summer 2007.
15. K. L. Myers and N. Yorke-Smith. A cognitive framework for delegation to an assistive user agent. In *Proc. of AAAI 2005 Fall Symposium on Mixed-Initiative Problem-Solving Assistants*, pages 94–99, Arlington, VA, 2005.
16. S. Natarajan, P. Tadepalli, and A. Fern. A relational hierarchical model for decision-theoretic assistance. In *Proc. of ILP'07*, Corvallis, OR, June 2007.
17. A. Pfeffer. Functional specification of probabilistic process models. In *Proc. of AAAI-05*, pages 663–669, Pittsburgh, PA, 2005.
18. A. Raja and V. Lesser. A framework for meta-level control in multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, Jan. 2007.
19. A. S. Rao and M. P. Georgeff. Modeling agents within a BDI-architecture. In *Proc. of KR'91*, pages 473–484, Cambridge, MA, 1991.
20. A. S. Rao and M. P. Georgeff. BDI agents: from theory to practice. In *Proc. of ICMAS-95*, pages 312–319, San Francisco, CA, 1995.
21. A. S. Rao and M. P. Georgeff. Decision procedures for BDI logics. *J. Logic and Computation*, 8(3):293–343, June 1998.
22. F. Rossi, P. van Beek, and T. Walsh. *Handbook of Constraint Programming*. Foundations of Artificial Intelligence. Elsevier Science Publishers, New York, NY, 2006.
23. S. Sardiña, L. de Silva, and L. Padgham. Hierarchical planning in BDI agent programming languages: A formal approach. In *Proc. of AAMAS'06*, pages 1001–1008, 2006.
24. M. Schut, M. Wooldridge, and S. Parsons. The theory and practice of intention reconsideration. *JETAI*, 16(4):261–293, Oct.–Dec. 2004.
25. P. H. Shaw and R. H. Bordini. Towards alternative approaches to reasoning about goals. In *Proc. of AAMAS'07 Workshop on Declarative Agent Languages and Technologies (DALT'07)*, Honolulu, HI, 2007.
26. R. Steuer. *Mulitple Criteria Optimization: Theory, Computation, and Application*. Wiley, New York, 1986.
27. B. Stroe, V. S. Subrahmanian, and S. Dasgupta. Optimal status sets of heterogeneous agent programs. In *Proc. of AAMAS'05*, pages 709–715, Utrecht, The Netherlands, 2005.
28. P. Surynek and R. Barták. Encoding HTN planning as a dynamic CSP. In *Proc. of CP'05*, page 868, Sitges, Spain, 2005.
29. J. Thangarajah, J. Harland, D. Morley, and N. Yorke-Smith. Aborting tasks in BDI agents. In *Proc. of AAMAS'07*, pages 8–15, Honolulu, HI, 2007.
30. J. Thangarajah, L. Padgham, and J. Harland. Representation and reasoning for goals in BDI agents. In *Proc. of 25th Australasian Computer Science Conference (ACSC'02)*, pages 259–265, Melbourne, Australia, 2002.
31. J. Thangarajah, L. Padgham, and M. Winikoff. Detecting and avoiding interference between goals in intelligent agents. In *Proc. of IJCAI'03*, pages 721–726, Acapulco, Mexico, 2003.
32. J. Thangarajah, L. Padgham, and M. Winikoff. Detecting and exploiting positive goal interaction in intelligent agents. In *Proc. of AAMAS'03*, pages 401–408, 2003.
33. J. Thangarajah, M. Winikoff, L. Padgham, and K. Fischer. Avoiding resource conflicts in intelligent agents. In *Proc. of ECAI-02*, pages 18–22, Lyon, France, 2002.
34. M. Winikoff, L. Padgham, J. Harland, and J. Thangarajah. Declarative and procedural goals in intelligent agent systems. In *Proc. of KR'02*, pages 470–481, Toulouse, France, 2002.