Robots with the Best of Intentions

S. Parsons[1] O. Pettersson[2] A. Saffiotti[2] M. Wooldridge[1]

[1]Department of Electronic Engineering
Queen Mary and Westfield College
University of London, London E1 4NS, U.K.
{s.d.parsons, m.j.wooldridge}@elec.qmw.ac.uk
http://www.elec.qmw.ac.uk/isag

[2]Applied Autonomous Sensor Systems
Department of Technology and Science
Örebro University, S-70182 Örebro, Sweden
{ola.pettersson, alessandro.saffiotti}@ton.oru.se
http://www.oru.se/forsk/aass

26.3.99

Abstract

Intelligent mobile robots need the ability to integrate robust navigation facilities with higher level reasoning. This paper is an attempt at combining results and techniques from the areas of robot navigation and of intelligent agency. We propose to integrate an existing navigation system based on fuzzy logic with a deliberator based on the so-called BDI model. We discuss some of the subtleties involved in this integration, and illustrate it on a simulated example. Experiments on a real mobile robot are under way.

1 Introduction

> Milou works in a food factory. He has to regularly go and fetch two food samples (potato crisps) from two production lines in two different rooms, A and B, and bring them to an electronic mouth in the quality control lab. Milou must now plan his next delivery. He decides to get the sample from A first, since room A is a little bit nearer than B. While going there, however, he finds the main door to that room closed. Milou knows that there is another door that he could use, but he considers the desirability of doing so. The alternative way to A is hard for Milou, since it goes through a narrow corridor which is usually cluttered with boxes. Besides, doors usually do not stay closed long. Hence, Milou decides to first go to B, and come back to A later on. He goes to room B, picks up the potato crisp and returns. The door to A is still closed, and this time Milou has no other choice than taking the difficult route. He does so, obtains the desired crisp, and finally rolls over to the lab and completes his task.

Performing the above task requires the ability to navigate robustly in real-world, unmodified environments. The robot Milou must be able to reliably find his way, keep track of his own position, avoid the obstacles in the cluttered corridor, and so on. However, this task also requires some higher level capabilities, like reasoning about alternative ways to perform a given task, and reconsidering the available options in face of new events. Our ability to develop intelligent mobile robots and to deploy them in real-world environments will critically depend on our ability to integrate these two aspects of the autonomous navigation problem.

Today's research on mobile robotics has produced a large number of techniques for robust robot navigation in real environments in the presence of uncertainty. These techniques typically focus on the navigation problem, and do not engage in abstract reasoning processes of the type encountered in the above scenario. On the other hand, research in intelligent agency has resulted in a number of interesting theories for reasoning about actions and plans. Unfortunately, these theories are typically stated at a very abstract level, and ignore the oddities and uncertainties that arise from operating in a real, physical environment.

This paper is a preliminary attempt at integrating results and techniques from the areas of robot navigation and of intelligent agency. Our two main ingredients are: (i) a theory of intelligent agency based on the interplay between *beliefs, desires* and *intentions,* commonly referred to as 'BDI'; and (ii) a behaviour-based robot navigation system grounded in fuzzy logic, called 'Thinking Cap'. In what follows, we outline the characteristics of these ingredients that are relevant to this paper, and discuss how we can integrate them. We also show an illustrative example based on the above scenario.

2 The BDI model

In the past few years there has been a lot of attention given to building formal models of autonomous software agents; pieces of software which operate to some extent independently of human intervention and which therefore may be considered to have their own goals, and the ability to determine how to achieve their goals. Many of these formal models are based on the use of mentalistic attitudes such as beliefs, desires and intentions. The beliefs of an agent model what it knows about the world, the desires of an agent model which states of the world the agent finds preferable, and the intentions of an agent model those states of the world that the agent actively tries to bring about.

The development of the BDI paradigm was to a great extent driven by Bratman's theory of (human) practical reasoning [1], in which *intentions* play a central role. Put crudely, since an agent cannot deliberate indefinitely about what courses of action to pursue, the idea is it should eventually *commit* to achieving certain states of affairs, and then devote resources to achieving them. These chosen states of affairs are intentions, and once adopted, they play a central role in future practical reasoning [2, 3].

It should be noted that the current popularity of the BDI paradigm in the area of software agents is due to more than just an anthropomorphic desire to attribute mental states to inanimate objects. On the contrary, the use of such ideas has strong justification from a software engineering perspective, allowing the modular development of systems by partitioning the information about a domain into different categories which are handled in different ways. In addition, the BDI approach has proved effective as the basis of a number of exacting applications, including the monitoring and control of spacecraft systems [5], and managing the flow of aircraft arriving at an airport [6].

A major issue in the design of agents that are based upon models of intention is that of when to *reconsider* intentions. An agent cannot simply maintain an intention, once adopted, without ever stopping to reconsider it. From time-to-time, it will be necessary to check, for example, whether the intention has been achieved, or whether it is believed to be no longer achievable [3]. In such situations, it is necessary for an agent to deliberate over its intentions, and, if necessary, to *change focus* by dropping existing intentions and adopting new ones.

In [15] we started the formal analysis of this problem. In particular we proposed a notion of optimality of deliberation, which can be glossed as "an agent is optimal if it always deliberates when deliberation will change its intentions and never deliberates when deliberation would not change its intentions", and showed that this can be used to develop a formal description of agents which are *bold* and *cautious* in the sense of Kinny and Georgeff [8]. The idea is that different types of environment require different types of strategies. In rapidly changing environments it makes sense for an agent to spend a lot of time deliberating in order to avoid spending time trying to achieve things which have become impossible. In more static environments there is much less call for agents to deliberate because once they have adopted an intention there is only a small chance that the world will change so as to make that intention impossible to achieve.

3 The 'Thinking Cap'

The 'Thinking Cap' (TC) is a system for autonomous robot navigation based on fuzzy logic which has been implemented and validated on several mobile platforms. A full description of the TC can be found in [11]. Parts of the TC were previously reported in [13, 12, 14]. The main ingredients of TC are:

- a library of *fuzzy behaviours* for indoor navigation, like obstacle avoidance, wall following, and door crossing;
- a *context-depending blending* mechanism that combines the recommendations from different behaviours into a tradeoff control;
- a set of *perceptual routines*, including sonar-based feature extraction, and detection of closed doors and blocked corridors;
- an approximate *map* of the environment, together with a positioning mechanism based on natural landmarks;
- a *navigation planner* that generates a behaviour combination strategy, called a B-plan, that achieves the given navigation goal; and
- a *monitor* that reinvokes the planner whenever the current B-plan is no more adequate to the current goal.

For the goals of this paper, we regard TC as a black box that provides a robust navigation service, and that accepts goals of the form '(goto X)'. There are however two peculiar characteristics of TC that are important here.

Firstly, navigation goals in TC are fuzzy: in '(goto X)', 'X' is a fuzzy location in the robot's map. (More precisely, a goal is formally defined in the TC framework as a fuzzy set of trajectories.) This means that a goal in TC can be more or less satisfied, as measured by a *degree of satisfaction*, a real number in the $[0, 1]$ interval. Typically, this degree depends on the distance between the robot and the desired location, but more complex goals may have more complex degrees of satisfaction.

Secondly, the 'adequacy' of the current B-plan which is monitored by the TC is in fact a *degree of adequacy*, again measured by a number in $[0, 1]$. This degree of adequacy is the result of the composition of three terms. (i) A degree of 'goodness', that takes into account the prior information available about the environment; for example, a B-plan that includes passing through a long and narrow corridor has a small degree of goodness. (ii) A degree of 'competence', that dynamically considers the truth of the preconditions of the B-plan in the current situation; for example, if a door that has to be crossed is found closed this degree drops to 0. And (iii) a degree of 'conflict', that measures the conflict between the behaviours which are currently executing in parallel. Both the degrees of satisfaction of the current goal and the degree of adequacy are recomputed by the TC at each control cycle (100 ms).

4 Integrating the BDI model and the Thinking Cap

Our work is based on the premise that the BDI model and the Thinking Cap represent two ends of the spectrum as far as the mental abilities of an autonomous robot are concerned. The TC can construct plans to achieve a single high level intention (like "go to the lab"), but has little to say about when such a plan has either failed or should be reconsidered because it might now be impossible to carry out. In contrast, the BDI model (at least in so far as we have analysed it with respect to intention reconsideration) is only concerned with high level intentions and whether or not they should be reconsidered as its beliefs about the world themselves change.

A consequence of this premise is that it might be profitable to combine the TC with a BDI architecture of the kind proposed in [15]. Our first attempt to integrate these two different systems is to consider them as separated blocks with a minimal interface between them, as shown in Fig. 1.

The BDI deliberator generates high-level intentions of the type (goto X) and sends them to the TC. (In future versions, intentions may include manipulation or observation activities.) The TC receives these intentions and considers them as goals. For each goal, it generates a B-plan, and starts execution. It also monitors this execution, and switches to a new B-plan if the current one turns out to be inadequate.

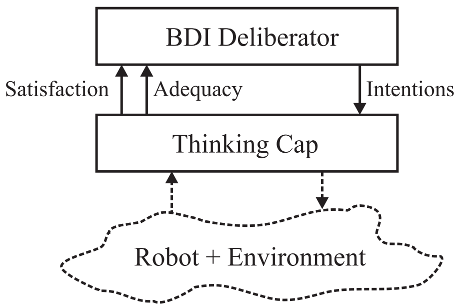

Figure 1: Integration between a BDI deliberator and the Thinking Cap.

During execution, it constantly computes the current degrees of satisfaction and of adequacy, as mentioned above.

These degrees are sent back to the BDI deliberator. From the point of view of the deliberator, the degree of satisfaction measures how much the current intention has been achieved, and the degree of adequacy measures how much this intention is considered achievable. Differently from the standard BDI model, however, this information is not given by binary values, but by continuous measures. It is these indicators of the state of the world *vis à vis* the current intention which help the deliberator determine when it is appropriate to reconsider its intentions.

More precisely, the deliberator uses these values in two ways. Firstly, to decide *when* it is time to deliberate. Two of the possible causes that lead the deliberator to reconsider its intentions are: (i) an increase in the value of satisfaction; and (ii) a drop in the value of adequacy. Secondly, to actually *deliberate*, that is to reconsider its intentions in light of the new information. Deliberation may involve comparing the available options, and possibly adopting a new intention which is then sent to the TC. As we shall see below, considering degrees instead of binary values allows the deliberator to take more informed decisions.

5 Example

We report a simple experiment where we execute the potato crisp scenario in a simulated environment. We have used the Nomadic simulator, which includes simulation of the sensors and some moderate sensor and positioning noise. This experiment is meant to illustrate the concepts and mechanisms involved in our integrated approach to robot deliberation and navigation. The successive phases of the simulated run are shown in Figures 3 and 4. Fig. 5 shows the values of adequacy and of satisfaction of the currently executing intention at each moment of the run.

Initially, the BDI deliberator considers the new task and decides a strategy, represented by the intention tree shown in Fig. 2 (left). (The details of how this is done are not relevant here; the dots indicates other intentions, like picking up the crisp, which we ignore.) The deliberator then passes the first intention (goto A) to TC, which generates a suitable B-plan for it. In this case there are two possible B-plans, one for each possible door leading to A, and the TC selects the one with the highest degree of (expected)

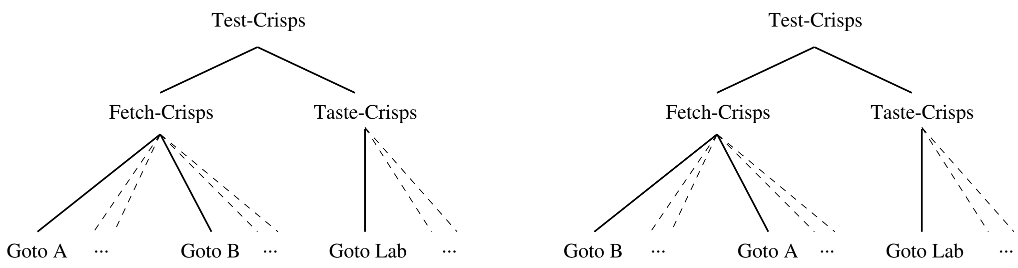

Figure 2: Two intention trees for our example task.

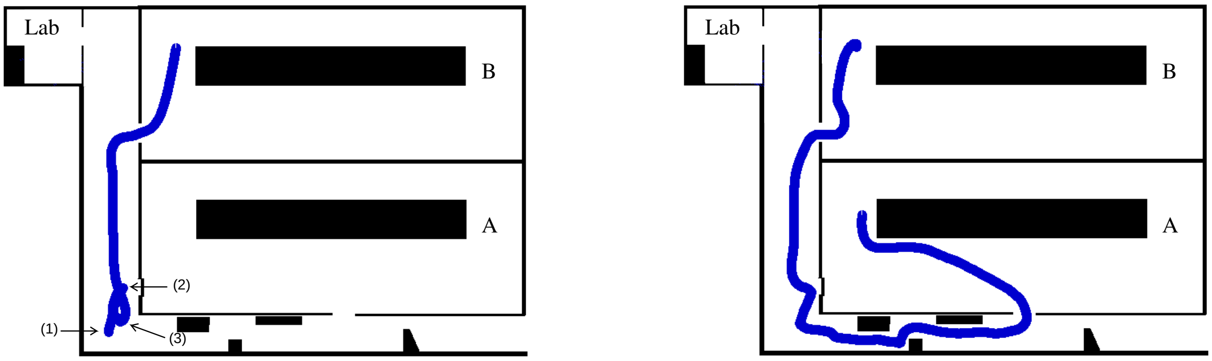

Figure 3: Left: Milou has the intention (goto A), but this turns out to be difficult to achieve, and adopts the new intention (goto B). Right: Milou again has the intention (goto A).

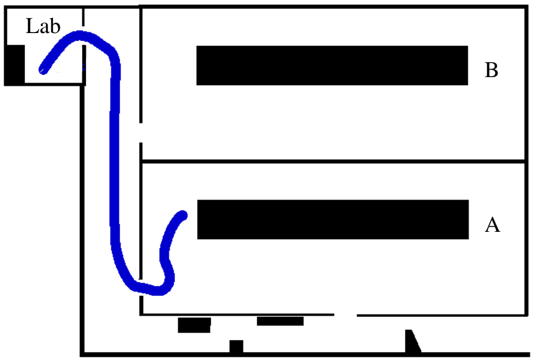

Figure 4: Both previous intentions are fulfilled, and Milou adopts the intention (goto Lab).

goodness. Since the TC knows about the low degree of traversability of the lower corridor,[1] the selected B-plan is the one that goes through the main door of A, the one on its left wall. Milou starts executing this B-plan from the lower left corner, as indicated by (1) in Fig. 3 (left).

When Milou arrives to this door (2), the sonars detect that the door is closed. Since one of the assumptions in the B-plan is that the door must be traversable, the degree of adequacy of this plan drops to 0 (Fig. 5 at about 20 s). The TC notices the problem, generates a new B-plan that goes through the second door, and starts executing it. However, this B-plan has a low degree of goodness because it includes passing through the cluttered corridor. This causes a drop of the adequacy level to a low 0.2. This is soon[2] noticed by the BDI deliberator, which reconsiders its options. Since the current intention turns out to be difficult (although not impossible) to achieve, and since there is an alternative way to perform the task (Fig. 2 right), the BDI deliberator decides to swich to this alternative way and to reverse the order of visiting the two production lines. Hence, it sends the new intention (goto B) to the TC (Fig. 5 at 30 s). The TC generates a new B-plan for this intention and swaps it in. Poor Milou then stops its travel to the lower corridor (point (3) in Fig. 3 left), turns around, heads to room B, and eventually reaches the collection point in front of conveyer belt B.

The achievement of the intention (goto B) is reflected in the rise of the satisfaction level (Fig. 5 at 75 s). This is noticed by the BDI deliberator, which then sends the next intention to the TC: in our case, this is again the intention (goto A). Since the information about closed doors inside the TC is transient, the TC again generates a B-plan for this intention that goes through the main door. Milou finds its way

[1] Currently, this information is stored in the map; in the future, the robot may acquire this knowledge during exploration.

[2] We add some delay on the adequacy level in order to leave the TC the time to try and fix problems (like the closed door) before the deliberator does something.

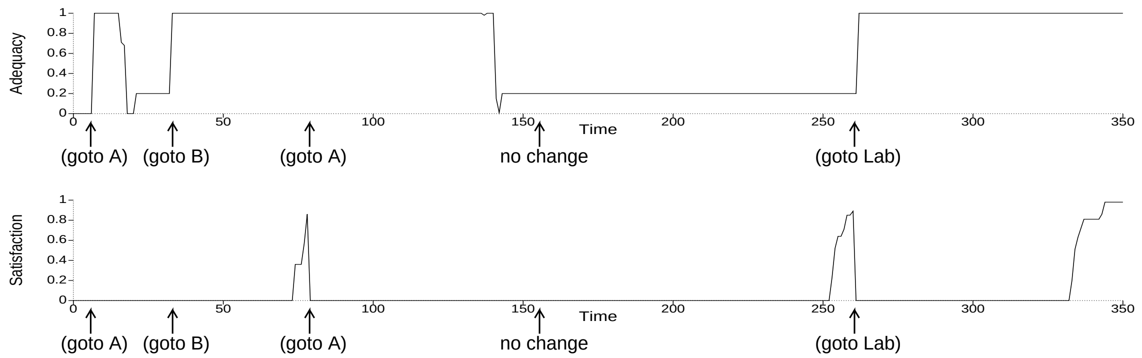

Figure 5: Measures of adequacy (top) and satisfaction (bottom) sent by the TC to the deliberator during the run. The arrows indicate the deliberation points, and the new intentions generated.

from room B, but unfortunately it finds that the door is still closed (Fig. 3 right). As before, the TC generates an alternative B-plan going through the lower corridor and starts to execute it. This produces a drastic drop in the adequacy level, which is noticed by the BDI deliberator (Fig. 5 at 160 s). However, this time there is no alternative option, so the deliberator decides to keep with the current intention, even if it is difficult to achieve. The navigation functionalities of the TC allow Milou to safely, if slowly, get around the obstacles, and reach the collection point in front of conveyer belt A.

The first two intentions are now fulfilled, and the BDI deliberator sends the last one (goto Lab) to the TC. Again, the TC tries the main door first. This time we are lucky, since someone has actually opened this door, and Milou eventually finds its way to the lab, thus completing the mission (Fig. 4).

6 Discussion

The problem of how to integrate the execution of low-level navigation primitives to high-level reasoning processes is at the heart of autonomous robot navigation. Several proposals have already appeared in the literatures that use a BDI approach for this goal. For example, the Saphira architecture [9] uses a simplified version of PRS [4], a computational incarnation of the BDI model, at the higher level, and fuzzy navigation behaviours at the lower level. In that architecture, the PRS system arbitrates the on-off activation of individual fuzzy behaviours, which are seen as ground level intentions. A similar approach is taken in [7] and in [10], where PRS-like systems are used to arbitrate low-level processes.

Our proposal departs from these approaches in the way we partition the responsibilities between the Thinking Cap and the BDI deliberation system. We rely on the underlying navigation abilities of the TC to take care of fuzzy behaviour arbitration and blending in a sophisticated way. And we limit the role of the deliberation system to take care of higher level decisions about which overall navigation goal should be pursued next. This partition allows us to make a better use of the respective powers of the TC and of the BDI level. By passing the adequate performance measures from the lower to the upper level we allow the latter to take more abstract, yet still fully informed decisions. We have shown that the use of measures instead of crisp values helps the higher level processes to generate the best possible intentions given the oddities and uncertainties that are inherent in real-world operation. We believe that a careful integration between these two levels in face of uncertainty is pivotal to our ability to deploy fully autonomous mobile robots.

The work presented above is still preliminary, and should be taken as a feasibility study more than a report of assessed results. Many variations of, and extensions to, the simple ideas presented here are possible, and their investigation is part of our current work. Firstly, the information passed by the TC to the BDI level can be much richer. For example, it may include the reasons why a B-plan has (partially) failed, the conditions that would increase its level of adequacy, or indications about the existence of alternative B-plans and their degrees of adequacy. Secondly, in our framework the BDI level does not have any way to recognise new opportunities that arise at the navigation level, like an open door that offers an unanticipated shortcut. Thirdly, and related to the previous point, we have not addressed the important issue of which information about the environment is available to the BDI level. Currently, no perceptual information is passed to this level by the TC, but this will clearly have to be changed in the future. Fourthly, more measures about the quality of execution could be communicated between the TC and the BDI deliberator, e.g., a measure of the current positional uncertainty. Finally, the choice of the strategy used to decide when the BDI should deliberate and when it should let the TC do its job depends on the characteristics of the environment, and it may itself be the result of another, higher level deliberation. Including this idea in our framework would lead to a "tower of meta-controllers" similar to the one suggested in [15]. Such an approach would allow the robot to dynamically adjust its policy for redeliberation if it finds that its current policy is incorrect with respect to its current environment.

In closing, we note that the example shown above has only been run in simulation — although the navigation system alone has been extensively validated on several real robots [13, 12, 11]. We are aware that the actual verification of the ideas sketched in this paper will only come from intensive testing in real and challenging environments. We are currently in the process of implementing our integrated system on a Nomad 200, and we hope to be able to show the first experimental results at the workshop.

Acknowledgements

The work reported here has benefit from discussions with Dimiter Driankov. Partial financial support for this project was provided by a grant from the Faculty of Engineering at QMW, and by the Swedish KK Foudation.

References

[1] M. E. Bratman. *Intention, Plans, and Practical Reason.* Harvard University Press, Cambridge, MA, 1987.

[2] M. E. Bratman, D. J. Israel, and M. E. Pollack. Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4:349–355, 1988.

[3] P. R. Cohen and H. J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.

[4] M. P. Georgeff and F. F. Ingrand. Decision-making in an embedded reasoning system. In *Procs. of the AAAI Conf.*, pages 972–978, Detroit, MI, 1989.

[5] M. P. Georgeff and F. F. Ingrand. Monitoring and control of spacecraft systems using procedural reasoning. In *Proceedings of the Space Operations Automation and Robotics Workshop*, 1989.

[6] M. P. Georgeff and A. S. Rao. Profile of the Australian Artificial Intelligence Institute. *IEEE Expert*, 6, December:89–92, 1996.

[7] F. F. Ingrand, R. Chatila, R. Alami, and F. Robert. PRS: a high level supervision and control language for autonomous mobile robots. In *Procs. of the Int. Conf. on Robotics and Automation*, Minneapolis, MN, 1996.

[8] D. Kinny and M. Georgeff. Commitment and effectiveness of situated agents. In *Procs. of the Int. Joint Conf. on Artificial Intelligence*, pages 82–88, Sydney, Australia, 1991.

[9] K. Konolige, K.L. Myers, E.H. Ruspini, and A. Saffiotti. The Saphira architecture: A design for autonomy. *Journal of Experimental and Theoretical Artificial Intelligence*, 9(1):215–235, 1997.

[10] J. Lee, M. J. Huber, E. H. Durfee, and P. G. Kenny. UM-PRS: an implementation of the procedural reasoning system for multirobot applications. In *AIAA/NASA Conf. on Int. Robots in Field, Factory, Service and Space.* American Institute of Aeronautics and Astronautics, 1994.

[11] A. Saffiotti. *Autonomous robot navigation: a fuzzy logic approach.* PhD thesis, Université Libre de Bruxelles, Brussels, Belgium, 1998.

[12] A. Saffiotti, K. Konolige, and E. H. Ruspini. A multivalued-logic approach to integrating planning and control. *Artificial Intelligence*, 76(1-2):481–526, 1995.

[13] A. Saffiotti, E. H. Ruspini, and K. Konolige. Blending reactivity and goal-directedness in a fuzzy controller. In *Proc. of the 2nd IEEE Int. Conf. on Fuzzy Systems*, pages 134–139, San Francisco, California, 1993. IEEE Press.

[14] A. Saffiotti and L.P. Wesley. Perception-based self-localization using fuzzy locations. In M. van Lambalgen L. Dorst and F. Voorbraak, editors, *Reasoning with Uncertainty in Robotics — Proc. of the 1st Int. Workshop*, number 1093 in LNAI, pages 368–385. Springer, Berlin, DE, 1996.

[15] M. Wooldridge and S. Parsons. Intention reconsideration reconsidered. In *Proc. of the 5th Int. Workshop on Agent Theories Architectures and Languages (ATAL)*, Paris, F, 1998.