

# An Advanced Control Architecture for Autonomous Mobile Robots: Modeling Intentions and Persuasions

Pietro Baroni<sup>\*</sup>, Daniela Fogli<sup>\*</sup>, Giovanni Guida<sup>\*</sup>, Silvano Mussi<sup>°\*</sup>

<sup>\*</sup> Università di Brescia, Dipartimento di Elettronica per l'Automazione,  
Via Branze 38, 25123 Brescia, Italy  
fax: +39 30 380014  
e-mail: {baroni, guida, fogli}@bsing.ing.unibs.it

<sup>°</sup> Consorzio Interuniversitario Lombardo per l'Elaborazione Automatica,  
Via Sanzio 4, 20090 Segrate (MI), Italy  
fax: +39 2 2135520  
e-mail: mussi@icil64.cilea.it

## Abstract

The aim of this paper is to propose a novel control architecture for autonomous mobile robots, based on the explicit representation of robot intentions and persuasions. The proposal relies on two basic assumptions: i) robot mental states (namely intentions and persuasions) are modeled as active autonomous entities; ii) the robot control architecture is conceived as a multi-agent system, where mental entities communicate, cooperate and conflict, giving rise, through their interactions, to a globally intelligent behavior. The paper discusses first the basic motivations for the proposal, introducing the concept of active mental entity and characterizing intentions and persuasions. Then a detailed description of the organization of the control architecture is provided, together with an outline of its operation. The potential of the proposed architecture is finally demonstrated through an application example concerning a department mail-delivery robot.

## 1. Introduction

Requirements for an autonomous mobile robot (AMR) include very complex and challenging issues, such as the need to cope appropriately and in a timely fashion with changes in a dynamic and unknown environment. This, in turn, entails the capability of maintaining and pursuing multiple goals, modifying them opportunistically. The design of a proper control architecture is widely recognized as a key factor to enable an AMR to effectively meet such requirements.

In recent years, there has been a progressive evolution from hierarchical paradigms, based on the sense-plan-act loop [15] [7], to behavior-based architectures featuring the explicit representation of behaviors [3] [1] [10] and, eventually, to distributed control architectures [2] [8], where robot control is distributed among autonomous interacting agents. In parallel with this evolution, other researchers [5] [14] have underlined the importance of endowing an autonomous robot with an explicit representation of internal

mental states, such as goals, beliefs, desires, etc. In fact the possibility to reason about mental states can significantly improve the capability of an intelligent autonomous agent to act opportunistically in a dynamic world. However, these research works have been developed mainly at a theoretical and formal level, whereas a structured and systematic proposal about how to integrate mental state representation in the frame of an actual robot control architecture is still lacking.

In this paper, we present a robot control architecture which is aimed to fill this gap. Starting from the proposal outlined in [2], the explicit representation of two basic types of robot mental states -namely: intentions and persuasions- is integrated in the framework of a distributed control architecture. In order to illustrate this new paradigm we first introduce the concept of active mental entity and characterize intentions and persuasions. Then we present a detailed description of the control architecture and an outline of its operation. Finally, we introduce an application example concerning a department mail-delivery robot: the example demonstrates the main features and the potential of the proposed approach in facing unusual and complex tasks.

## 2. Modeling Intentions and Persuasions as Active Mental Entities

### 2.1 From mental states to active mental entities

Intelligent behavior is often explained by using terms such as desire, inhibition, intention, hope, obligation, prohibition, etc. These terms denote *mental entities*, i.e. entities that are inside the mind of an intelligent agent and which are responsible of his/her external behavior.

Modeling agent mental activity is the subject of several works in the area of intelligent agent design. Among the most notable proposals in this field there is the BDI theory [12], to which several agent architecture models are inspired, such as IRMA [11] and PRS [6]. The BDI theory proposes an abstract architecture which models the mental activity of a rational agent in terms of three types of mental states: beliefs, desires, and intentions. These mental states are

modeled, however, as mere data structures and a central mechanism, namely the interpreter, uses these data structures to make decisions and manage agent operation.

Our main point concerns the fact that mental entities should not be regarded as passive entities. Mental activity can be better modeled by providing mental entities with a sort of "agentification". In fact, if we consider the mental processes occurring in our mind, we can reasonably understand them as the result of the cooperation and conflict between various intentions, inhibitions, hopes, desires, etc. A similar description of human reasoning has been proposed, for instance, by Minsky [9].

Therefore, mental entities should not be considered as passive entities, i.e. mere information structures, on which an agent operates according to stated and fixed procedural rules, but as active entities, endowed with the capability of autonomous operation. In fact, since mental entities have to dynamically interact to produce a globally intelligent behavior, it is essential that they are provided with individual and independent operation capabilities. Therefore, we call them *active mental entities*, stressing that they are autonomous and can operate according to their own goals and strategies.

The concept of active mental entity allows a novel agent structure to be defined: according to our approach, intelligent agents can be implemented through a distributed paradigm comprising active mental entities which cooperate in order to determine agent behavior. So, if in general "intelligence emerges from the interaction of the components of the systems" [4], in our specific proposal, agent intelligence emerges from the interaction of active mental entities.

In order to emphasize the power of a distributed approach vs. a centralized one, let us reconsider the central role played by the interpreter of the BDI-architecture. This mechanism is in charge of solving competitions between alternative plans, substituting current intentions with new ones, and choosing new courses of actions if changes in the environment occur. In a word, the interpreter is in charge of all the most challenging tasks related to agent operation. Therefore, the design of such an interpreter may be critical. Moreover, it can become a bottleneck during operation. In our proposal, there is no central structure in charge of all such crucial functions, since they are performed through the free interaction between mental entities. For example, when a conflict arises the involved mental entities engage a debate with the purpose of solving it. During this debate they can look for support to their theses or can establish alliances with other mental entities. This way, the emerging behavior is the result of the interaction of autonomous entities in a distributed environment, rather than of a complex computation carried out by a centralized reasoning activity. Of course, centralized control may be preferred, being more efficient and simple to design, in those cases where task complexity is reasonably low and the expected robot behavior can be easily specified in all situations of interest. Even in these cases however, resorting to a distributed paradigm may be definitely in favor of a more disciplined

and modular design, thus supporting incremental design and development and a gradual extension of robot capabilities.

## 2.2 Intentions and persuasions

The mental entities we take into consideration in this paper are *intention* and *persuasion*. We stress that the goal of this paper is neither to propose a general taxonomy of mental entities, nor to enter the debate about which is the minimal set of mental entities that can guarantee an intelligent behavior, but rather to show the practical advantages of providing agents with active mental entities. For the sake of brevity and simplicity, we limit therefore our analysis to the two mental entities mentioned above, as particularly important and sufficient to demonstrate the potential of our approach.

We say that an agent has an *intention* when it commits itself to pursue something, i.e. to attain a given achievement, that is called the *subject* of the intention. Intentions may rely on some *validity conditions*; for instance the intention "I want to find Mr. Smith" is valid only under the condition that it believes that it is possible for it to find Mr. Smith and, of course, Mr. Smith has not been found yet.

Intentions are characterized by persistence, that is they remain active until their validity conditions hold, otherwise they are dismissed. Some intentions are valid only until they are not achieved (for instance the intention "I want to find Mr. Smith" is dismissed once Mr. Smith has been actually found), whereas others do not depend on specific achievements (such as "I want to preserve my integrity").

We remark that the concept of intention is different from the concept of goal. An intention expresses the concept: "I want to pursue something". So, in general, an intention is characterized by a certain degree of vagueness and indefiniteness. It can not be directly and univocally translated into a precise sequence of actions. However, the existence of an intention in one's mind prompts him to recognize, in real circumstances, some favourable opportunities to exit from vagueness, identifying a precise way in which the intention may be put in concrete form. Conversely, a goal defines a precise state of the world to be reached through a suitable action plan.

Given such distinctions between the concepts of intention and goal, it seems quite natural to model intentions as active mental entities, since we assume that an intention is definitely committed to reach its achievement, possibly cooperating or conflicting with other intentions.

We say that an agent has a *persuasion* when it has some belief about the truth value of a given proposition, called the persuasion *subject*. Persuasion may be grounded on long-term knowledge (for example "Walls are unmovable obstacles"), on sensory data collected from external environment (for example "There is a wall in front of me"), or on another persuasion (for example "I can not move forward", given that "There is a wall in front of me").

Depending on their origin, persuasions can be updated or revised when long-term knowledge, sensory data or other persuasions are updated or revised. Persuasions concern

propositions, i.e. specific facts whose truth or falsity is of interest for the system. A persuasion is generated when a new interesting proposition is met and is dismissed when the interest in the proposition ceases. In general, a proposition is interesting when its truth value influences the achievement of an intention or determines the attribution of the truth value to another proposition of interest.

Persuasions are characterized by persistence, that is they remain active until their relevant propositions are considered interesting.

Persuasions are modeled as autonomous active entities since we assume that a persuasion is not just the passive result of some perceptual or reasoning activity, but it is definitely committed to reinforce itself, i.e. to find new elements to believe or disbelieve in the related proposition and to verify the already available ones. In such activity, a persuasion may cooperate or conflict with other persuasions. In particular, it can happen that several persuasions concerning contradictory propositions are generated and are active at the same time. Since they have different origins and rely on different evidences, they may be conflicting. In this case, each persuasion is stimulated to look for supports to its own thesis and for counterexamples to the theses of opponent persuasions, in order to resolve the conflict through a debate.

The following relationships between intentions and persuasions hold. When an intention is active its subject and its validity conditions become interesting propositions. Therefore, an intention is always supported by one or more persuasions which concern its validity conditions and enable the intention to exist and to persist. Moreover another persuasion is in charge of looking after the attainment of the subject. In general, different strategies may be adopted in order to pursue the same intention. Each strategy depends on a validity condition, whose truth determines the actual applicability of the strategy in a given situation. Therefore, for each strategy considered in the achievement of an intention, a persuasion is created, whose subject is the validity condition of the strategy itself.

On the other hand, in order to ascribe a truth value to its subject, a persuasion may rely on the truth value of other propositions, thus involving the creation of new persuasions, or may explicitly require some activity of information acquisition from the external world, thus involving the creation of new intentions.

In this way agents operation is characterized by a continuous intention-persuasion interaction. Intentions generate persuasions which in turn may enable or suppress other intentions, which may generate other persuasions, and so on.

### 3. A Distributed Control Architecture Supporting Active Mental Entities

#### 3.1 Agents, components and modules

According to the proposal outlined in [2], a robot control architecture can be understood as a multi-agent system where each agent is able to perform a specific task such as

managing a sonar sensor, managing a TV camera, segmenting images, controlling movement actuators, etc. In this paper we propose a substantial extension of this proposal that includes the concept of active mental entity introduced above. The proposed architecture is made up of a collection of *agents* that can communicate and cooperate in order to provide a global intelligent problem-solving behavior. Agents are assumed to be benevolent [13], i.e. they are always available to comply with cooperation requests of other agents.

An agent features a structured internal micro-organization that includes *components* and *agent knowledge*.

Let us examine components first. We distinguish two types of components:

### *operative components*, in charge of performing actions, either physical, concerning the interaction with external world through sensors and actuators, or symbolic, such as computational and reasoning activities;

### *mental components*, that is intentions and persuasions.

All agent components are understood as active entities, that can communicate and cooperate in order to produce the global agent behavior. Mental components interact among them and activate operative components which, besides performing actions, may provide a feedback to mental components about the results of the actions carried out. From outside, an agent appears, however, as a unitary active entity, featuring a definite behavior. To this end, an additional component, also an active entity, called *interface*, is included in the agent structure. It is in charge of managing interactions with other agents and the external world. It is assumed that all interactions, either among agents or among their internal components are carried out according to a message-passing paradigm.

We can further distinguish between static and dynamic components of an agent. Interface and operative components are *static entities*; they are defined at the moment of agent design and do not undergo any evolution during the agent operational life (we do not consider here learning capabilities). Mental components, instead, are *dynamic entities*; they are continuously and dynamically created and destroyed, as described in the previous section. Therefore, persuasions and intentions are designed as generic prototypes (more formally, as classes of objects) and, whenever required, specific persuasions or intentions can be generated as instances of the corresponding classes. All the components of an agent share the same basic structure. They are composed of two *modules* and possess, individually, two basic types of *component knowledge*.

The modules of a component are the kernel and the shell.

- The *kernel* is in charge of performing the specific tasks the component is capable of. In operative components the kernel is the repository of agent abilities. In mental components the kernel is devoted to carry out the typical activities defined for mental entities, such as generating strategies for intentions or looking for

support for persuasions. In the interface component, the kernel is in charge of managing message interchange.

- The *shell* is aimed at supporting communication activities. In all components, it is in charge of accepting and filtering incoming messages and of addressing messages produced in output.

Component knowledge includes:

### *Self knowledge*, that is knowledge concerning component's own specific features and capabilities, necessary to evaluate component's ability to face a given task. A particular role is played by the self knowledge of the interface, which concerns the overall agent features and capabilities, as they should be seen from outside.

### *Mutual knowledge*, that is knowledge concerning competencies and capabilities of other components, necessary to manage interaction and cooperation among components within the same agent. A particular role is played by the mutual knowledge of the interface, which concerns features and capabilities of other agents, and is used to manage interactions between different agents.

Turning now to agent knowledge, it represents the basic agent competence endowment. It is available to all agent components, which may exploit it in their operations. We distinguish two basic types of agent knowledge:

### *domain knowledge*, that is knowledge concerning the agent competence domain;

### *strategic knowledge*, that is knowledge about the criteria for evaluating pros and cons of a given action plan, necessary to mental entities in order to produce and revise their operation strategies (see below).

In order to define the global operation of our architecture we need now to proceed bottom-up.

### 3.2 Operation of mental components

Intention. Any intention is associated to a *subject*, i.e. a proposition describing a desired state of the world. The kernel of an intention consists of a problem-solving device (namely, a procedure, a knowledge base coupled with an appropriate reasoning mechanisms, or any other problem-solving system) capable of generating and dynamically refining appropriate *strategies* to achieve the subject of the intention. In order to produce and evaluate strategies, the kernel uses domain knowledge and strategic knowledge.

A strategy defines a sequence of *tasks* to be accomplished. Tasks may refer to complex or elementary activities.

Complex tasks entail the generation of other (derived) intentions. New intentions may pertain to the same agent the generating intention belongs to or fall outside its competence domain. In the former case, the new intentions generated are simply added to the agent to which the generating intention belongs, thus enriching its internal structure. In the latter case, a request of creating a new intention has to be addressed to the relevant agent. Since such request is directed outside the agent the generating intention belongs to, the intention shell transmits the

request to the agent interface. This, using its mutual knowledge, is able to correctly address the request to the relevant agent.

On the other hand, elementary tasks have to be assigned to suitable operative components, capable of achieving them. In order to address an elementary task to an operative component a request is generated by the intention kernel and assigned to the intention shell. The shell, resorting to its mutual knowledge, addresses it to the correct agent component. If the request turns out to fall outside the capabilities of agent components, it is addressed to the agent interface, and then forwarded to an appropriate agent whose capabilities match the request. To this end, the interface resorts to its mutual knowledge, and may engage a suitable dialogue with other agents in order to better focus their capabilities, availability, and general conditions for processing the request at hand.

When conflicting requests are received by an operative component (for instance "Move forward" and "Stop now") the conflict must be resolved. However it can not be handled at operative level, i.e. ignoring what intentions are behind the conflicting requests. Therefore, when a conflict is detected at operative level, it is notified and back-propagated to the intentions that (possibly indirectly) originated the conflicting requests. Then a negotiation session between the relevant intentions is started. In order to solve conflicts, contrasting intentions may resort to importance or urgency criteria, may look for alternative strategies, or can remit the conflict to the persuasions underlying them (see below).

Persuasion. Any persuasion is associated to a *subject*, (i.e. a proposition describing a state of the world) and to a *belief* about the truth value of the subject. We do not consider in this paper, for the sake of simplicity, the issue of belief quantification. Of course beliefs may be dynamically modified, if the case.

The kernel of a persuasion is a problem-solving mechanism able to:

### find, extend, and verify motivations supporting the persuasion;

### explain these motivations on request;

### debate with conflicting persuasions, either defending their own motivations or attacking those of the opponent.

In order to find support for its theses, a persuasion kernel may need to resort to external resources. Therefore it may send, through the shell, requests to operative components in charge of sensory data acquisition. Request handling is treated in the same way illustrated above for intentions.

When a conflict arises, a persuasion is able to interact with other persuasions, looking for alliances or attacking opponent's motivations. In these interactions the persuasion kernel resorts to the shell or to the agent interface in order to correctly address messages to other persuasions belonging to the same agent or to other agents. It has to be noted that a conflict does not arise just because

persuasions with the same subject and different beliefs exist. In fact, such contradictory persuasions can coexist, simply ignoring each other. A conflict arises, and needs to be solved, only in two cases:

### when contrasting operative requests are produced as a consequence of different intentions, which possibly rely on contradictory persuasions;

### when an intention, while elaborating a strategy, detects a conflict between two persuasions that can influence strategic decisions.

This choice is coherent with the main goal of AMR control architecture, that is managing and directing robot actions, rather than building a complete and correct model of the external world.

### 3.3 Operation of operative components

Operative components receive *requests* concerning specific tasks to be performed. The shell, using self knowledge, evaluates whether the request can be accepted: if acceptable, the request is forwarded to the kernel. Even if the tasks requested to an operative component are generally very specific, not necessarily they are atomic: decomposition into subtasks may be necessary before a task can be actually solved. The kernel of an operative component is therefore a problem-solving mechanism in charge of either directly carrying out an assigned task if it is atomic, or of decomposing it and then allocating the resulting subtasks to other operative components if the task is not atomic. Subtask allocation is performed through request messages, to be addressed to other operative components, by resorting to the shell and possibly to the agent interface, as described above.

### 3.4 Operation of agent interface

Agent interface play an intermediary role in communication between different agents. It is in charge of:

- examining messages coming from a component inside the agent (and that can not be addressed to another component of the same agent) and addressing them to the interface of another appropriate agent;
- examining messages coming from the interface of another agent and addressing them to the proper internal component.

To this purpose, the interface kernel exploits self and mutual knowledge, in order to establish an optimal matching between messages and addresses. The interface shell is then in charge of message handling.

### 3.5 Overall architecture operation

The overall operation of the architecture results from the autonomous operation of the agents and, in turn, from the autonomous operation of agent components. No explicit representation of behavior, either globally or at agent level, is provided. No central knowledge repository nor central control mechanism is present in the architecture.

## 4. An Example

In this section we present a simple example in order to support a better understanding of the organization and operation of our architecture and to demonstrate its advantages. For the sake of simplicity, we will omit most operational details, focusing only on the most significant aspects.

The example concerns a department mail-delivery robot, to which the user consigns an envelope to be delivered to Mr. X. Interaction with the user is managed by a specialised agent UI (User Interaction), which is characterized by a primitive intention whose subject is "obey-the-user". According to this intention, a new intention has to be generated whose subject is "deliver-mail-to Mr.-X". However, since UI has no specific competence on mail delivering, it has to address the request of creating this new intention to another competent agent. By resorting to its interface, UI identifies the MD (Mail Delivery) agent to which the request is addressed: the new intention whose subject is "deliver-mail-to Mr.-X" is therefore created within MD. MD has both general knowledge about the mail delivery task and specific knowledge about the department personnel, as far as mail delivery is concerned. So MD knows that Mr. X is actually a department employee, that his office is office number 9, and that normally, his office hours are from 9.00 a.m. to 5.00 p.m.. Since current time is 3.00 p.m. (this information is provided on request by a specialised CLOCK agent), MD generates the persuasion that Mr. X is now present in the department, (i.e. the subject "Mr. X present-now" is believed to be true).

This persuasion supports the persuasion that the intention "deliver-mail-to Mr.-X" can be actually achieved (more precisely, the subject "achievable(deliver-mail-to Mr.-X)" is believed). On these grounds, the intention "deliver-mail-to Mr.-X" can then elaborate a strategy to deliver mail to Mr. X. To this purpose, however, some more detailed persuasion, about where Mr. X is, has to be instantiated. In absence of more specific information, using default knowledge that normally an employee is in his office, the following persuasion is generated: "Mr.-X present-in-office". On the basis of this persuasion, the following simple strategy is generated:

task 1: go to office 9;

task 2: deliver the envelope to Mr. X.

Task 1 is considered first: it still concerns a quite generic and high-level task and must therefore be associated to a new intention. A request of generating such intention is therefore addressed by MD to the MM (Movement Management) agent. MM, generates the intention "go-to office-9" and, exploiting his knowledge about building topology, it generates then a strategy to reach office 9 and begins to realize it.

It has to be stressed that, meanwhile, all the mentioned mental components do not simply wait for the accomplishment of the selected strategy. On the contrary, they remain active: intentions are continuously looking for better strategies, and persuasions are continuously looking for new evidences supporting them. Moreover, it can be

assumed that some agents in charge of sensory acquisition are always active, because their output is necessary to other agents in charge of very essential primitive intentions, such as preserving robot integrity.

For instance, while the robot moves towards office 9, the intention "deliver-mail-to Mr.-X" may elaborate the following alternative strategy:

- task 1: find Mr. X around in the department
- task 2: go near Mr. X
- task 3: deliver the envelope to Mr. X.

This strategy relies on the persuasion that Mr. X is not in his office but somewhere else in the department: "Mr.-X present-but-not-in-office". However, before the strategy becomes operative, the persuasion has to find some support. To this purpose, the persuasion may generate intentions like "recognize-voice-of Mr.-X" and "recognize-face-of Mr.-X" to be addressed to agents specialized in processing audio and video input coming from sensory devices.

If sufficient resources are available, these activities can be carried out in parallel while the robot moves toward office 9, according to the first strategy generated.

Let us suppose now that, while moving towards office 9, the robot gets near a glass wall behind which there is Mr. X. Then the vision system recognizes Mr. X in front of the robot and, therefore, persuasion "Mr.-X present-but-not-in-office" gets strong support, whereas "Mr.-X present-in-office" is dismissed, since direct evidences always prevail over default knowledge.

As a consequence, also the first strategy generated on the ground of "Mr. X present-in-office" and the relevant intentions are abandoned (note however that they will revive if, for instance, it will be realized that Mr. X recognition was erroneous).

The second strategy becomes then active and, since task 1 has been achieved (find Mr. X around in the department), task 2 is pursued (go near Mr. X).

A request to create the intention "go-near Mr.-X" is therefore addressed to the agent MM (which, in the meanwhile, has dismissed all the activities implied by the first strategy). Suppose now that Mr. X is standing just in front of the robot. The task of navigating towards such a fixed target is reduced to the task go-forward. While the robot is moving forward, the VC (Video Camera) and SRS (Sonar Range Sensors) agents acquire and process data about the external world. Doing this, they continuously generate or update persuasions about the environment, whose subject is communicated to the agent CA (Collision Avoidance).

Suppose now that, while the robot is approaching the target, VC and SRS communicate to CA two contradicting persuasions: SRS has the persuasion that "there-is-an-obstacle-on-the-path", whilst VC has the persuasion "no-obstacle-on-the-path". CA recognizes that it is impossible to take a decision, given these contradicting persuasions, and therefore decides that the conflict should be solved. It puts

the two persuasions face to face by notifying each of them of the existence of the opponent persuasion.

The persuasions "there-is-an-obstacle-on-the-path" and "no-obstacle-on-the-path" enter therefore a debate in order to solve the conflict. First of all, an analysis of the motivations supporting them is carried out: "there-is-an-obstacle-on-the-path" is supported by the fact that the sonar received reflected echoes, "no-obstacle-on-the-path" is supported by the fact that in the image collected by the video camera nothing but Mr. X is seen.

The persuasions "there-is-an-obstacle-on-the-path" and "no-obstacle-on-the-path" are then in charge of searching evidence corroborating their supports or undermining the opponent's ones. For instance, "no-obstacle-on-the-path" can resort to general knowledge (this knowledge may be provided by the SRS agent itself) that sonar readings are often erroneous and notify it to "there-is-an-obstacle-on-the-path". In turn "there-is-an-obstacle-on-the-path" may reply that sonar readings are erroneous in specific conditions (near wall corners, in presence of noise sources, etc.) that are not met in the present case. At this point, "there-is-an-obstacle-on-the-path" may attack directly "no-obstacle-on-the-path" support resorting to general knowledge, provided by BT (Building Topology) agent, that, in the building, there are invisible obstacles, such as transparent glass walls.

Since "no-obstacle-on-the-path" is not able to reply to this argument, "there-is-an-obstacle-on-the-path" prevails: the presence of an obstacle is accepted and CA intervenes to modify the motion plan, going around the glass wall and eventually reaching Mr. X.

The example, although simple and presented only in a sketchy way, demonstrates two important advantages of our approach:

- ### the capability to effectively manage several different strategies in parallel pursuing multiple goals;
- ### the capability to cope with changes in a dynamic and unknown environment.

Note that both these capabilities are achieved in a way that is both technically sound and cognitively plausible. The intention-persuasion mechanism is indeed very flexible and powerful and can produce a global external behavior which is justified and explained by the mental processes occurring inside the robot.

## References

- [1] T. Balch, G. Boone, T. Collins, H. Forbes, D. Mackenzie, J. C. Santama'na, Io, Ganymede, and Callisto A Multiagent Robot Trash-Collecting Team, *AI Magazine*, Summer 1995, 39-51.
- [2] P. Baroni, G. Guida, S. Mussi, A. Vetturi, A Distributed Architecture for Control of Autonomous Mobile Robots, *Proc. of ICAR'95, 7th Int. Conf. on Advanced Robotics*, Sant Feliu de Guixols, SP, 1995, 869-877.

- [3] R. A. Brooks: A Robust Layered Control System For A Mobile Robot. *IEEE Journal of Robotics and Automation*, RA-2. April 1986, 14-23.
- [4] R. A. Brooks, Intelligence without reasoning. *Proc. 12th IJCAI International Joint Conference on Artificial Intelligence*, Sidney, Australia, 1991, 569-595.
- [5] P. R. Cohen, H. J. Levesque, Intention is choice with commitment, *Artificial Intelligence* 42(3) (1990), 213-261.
- [6] M. P. Georgeff, A. L. Lansky: Reactive Reasoning and Planning. *Proceeding of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, Seattle, WA, 1987, 268-272.
- [7] C. Isik, A. M. Meystel, Pilot Level of a Hierarchical Controller for an Unmanned Mobile Robot, *IEEE J. Robotics and Automation*, Vol. 4, no. 3, June 1988, 241-255.
- [8] R. Liscano, A. Manz, E. R. Stuck, R. E. Fayek, J. Tigli, Using a Blackboard to Integrate Multiple Activities and Achieve Strategic Reasoning for Mobile-Robot Navigation, *IEEE Expert*, April 1995, 24-36.
- [9] M. L. Minsky, *The Society of Mind*, Simon and Schuster, New York, 1986.
- [10] J. F. Montgomery, A. H. Fagg, G. A. Bekey, The USC AFV-I. A Behavior-Based Entry in the 1994 International Aerial Robotics Competition, *IEEE Expert*, April 1995, 16-22.
- [11] M. E. Pollack, The uses of plans, *Artificial Intelligence*, 57 (1), 1992, 43-68
- [12] A. S. Rao, M. P. Georgeff, Modeling Rational Agents within a BDI-Architecture, *Proc. of KR&R-91 Int. Conf. on Knowledge Representation and Reasoning*. Cambridge, MA, 1991, 473-484.
- [13] J. S. Rosenschein, M. R. Genesereth: Deals among rational agents. *Proc. 9th IJCAI International Joint Conference on Artificial Intelligence*, Los Angeles, CA, 1985, 91-99.
- [14] Y. Shoham, Agent Oriented Programming, *Artificial Intelligence*, 60, 51-92, 1993.
- [15] M.G. Slack, Planning paths through a spacial hierarchy: Eliminating stair-stepping effects, *Proc. SPIE Conference on Sensor Fusion*, 1988.