In 32nd International Symposium on Robotics (ISR 2001), Seoul, Korea April 19-21, 2001.

Intelligent Inhabited Environments: Cooperative Robotics & Buildings

Martin Colley, Graham Clarke, Hani Hagras, Vic Callaghan, Anthony Pounds-Cornish

Department of Computer Science
University of Essex, Wivenhoe Park
Colchester, Essex, CO4 3SQ, UK
Email: robots@essex.ac.uk

Abstract

In this paper we describe an innovative multi-agent environment consisting of an Intelligent-Building (IB) inhabited by a variety of agents ranging from mobile robots, embedded-agents to people some of which carry smart wearable gadgets. These cooperate together to form what we term an *intelligent inhabited environment*. We discuss the high-level multi embedded-agent model, explaining how it facilitates inter-agent communication and cooperation between heterogeneous sets of agents. We present an application aimed at establishing an agent based care/rehabilitation system in which a collection of building and robot agents cooperates to care for human occupants. The technological basis of our solutions stem from a genetic-fuzzy technique that has already been developed and successfully applied to the control of autonomous outdoor vehicles. We discuss in detail how we might use this system to identify significant variations from normal behaviour or identify emergency situations in which specialised robots might be summoned to assist. Finally we report results from earlier experiments on autonomous mobile robot navigation and embedded-agent based environment control in intelligent buildings.

1. Vision

In our previous work we have looked at the problems involved in developing intelligent buildings [5], mobile robotics [10] and the development of agent based technologies embedded in ubiquitous and mobile gadgets [4]. In this paper, we describe an innovative multi-agent environment consisting of an intelligent-building inhabited by a variety of agents ranging from mobile robots and embedded-agents as shown in Figure (1) to smart wearable gadgets. An essential feature that characterises all our work is that intelligent habitat technology needs to be centred on the individual. The agents should tailor their behaviour base to an individual wherever possible rather than generalise across a group of individuals. We have proposed that an elegant solution to producing such a system would be to embed agents and sensors into wearable devices (e.g. mobile phones, watches, smart-clothing etc). Here agents and sensors reside in both body-wearable artefacts and buildings. The agents share common functionality and are enabled to interact and work together as shown in Figure (2).

The underlying engineering infrastructure is based on a combination of: DAI (Distributed Artificial Intelligence), Hierarchical Fuzzy-Logic/GA based embedded-agents and network technology. We discuss the high-level multi embedded-agent model, explaining how it facilitates inter-agent communication and cooperative operation between heterogeneous sets of agents.

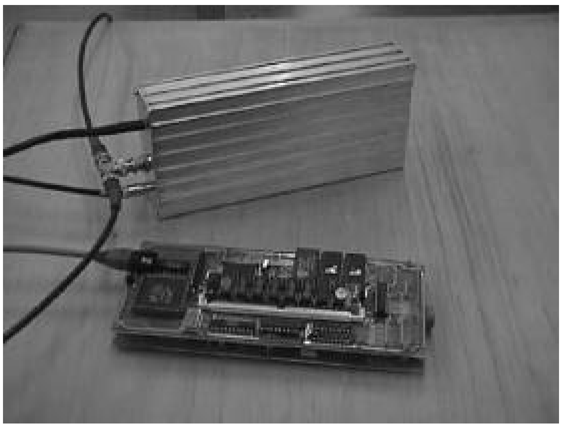

Figure (1): Prototype Essex IB Agent

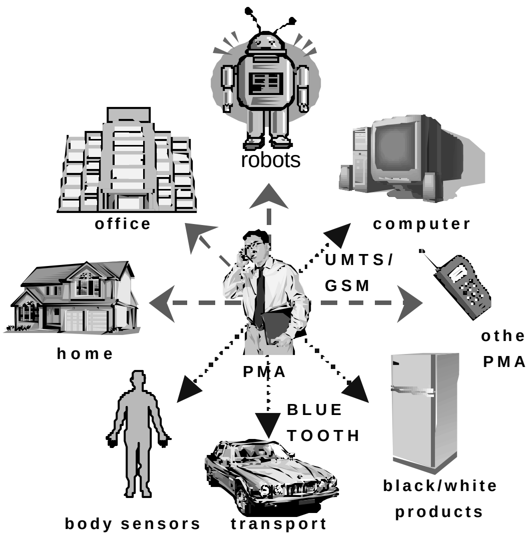

Figure (2): Abstract Communication Model

2. Heterogeneous Multi-Agent Physical Intelligent Environments Application

In previous IB work we have begun to look at the problems associated with a care environment and the ways in which IBs might be able to help [13]. Here we present an application aimed at establishing an agent based care/rehabilitation system in which a collection of building and robot agents cooperate to care for human occupants as show in Figure (3). The building agents handle control of all building services (e.g. heat, light entertainment, etc) maximising energy efficiency by being attached to various sensors and controllers. Davidsson has shown that such systems can save as much as 40% of the buildings energy without any loss of comfort or safety [8]. Mobile robots communicate with building agents so as to operate more effectively (e.g. command doors to open, locate themselves, responding to remote requests, etc) providing services such as delivery of meals, medicine etc. Wearables allow remote monitoring of the patient and can be set to signal any worrying deviation from expected values.

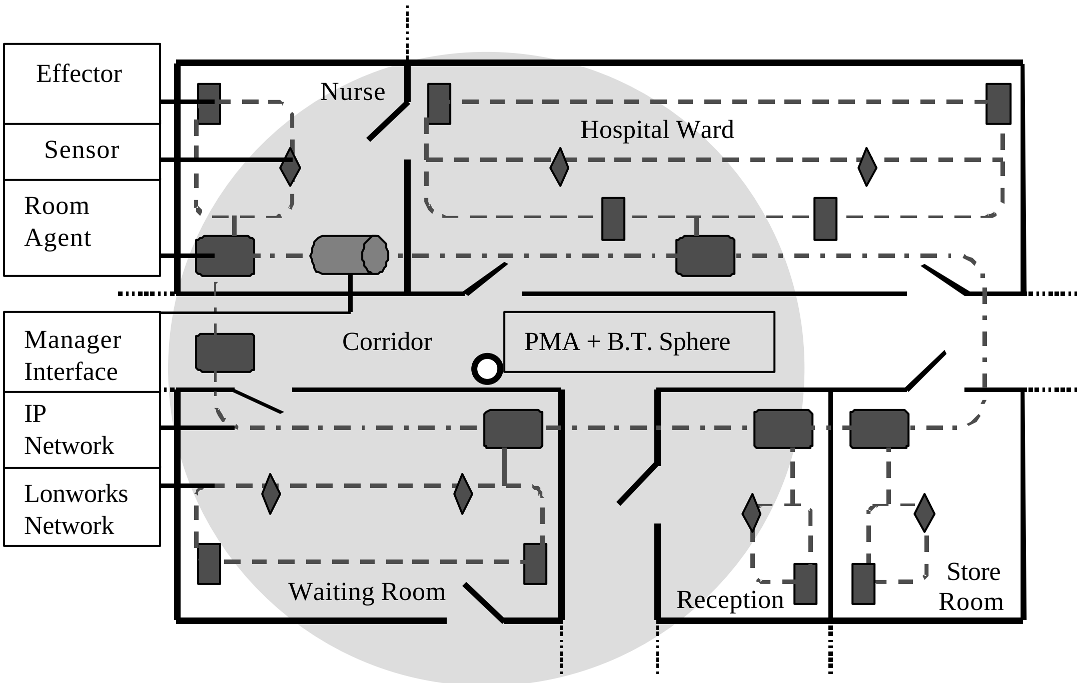

Figure (3): Example Intelligent Environment Infrastructure

Learning in intelligent buildings mostly takes the form of identifying cause-effect relationships based on building occupant actions in response to changes in the environment. A detailed discussion of the challenges involved is presented by the authors elsewhere [5]. In simple terms, the learning mechanism needs to be able to interpret these cause-effect relationships based on the combinational states and temporal sequences of numerous input vectors. The agent uses these relationships as part of a mechanism to serve the occupant's needs by taking pre-emptive actions. For efficient and robust learning, it is necessary to have mechanisms to identify, combine or resolve similar or contradictory events. IB based learning is focused around the actions of people. Buildings are, largely, occupied by people who for a variety of reasons (e.g. time, interest, skills, etc) would not wish, or be able to cope with much interaction with the building systems. Thus in general, learning should as far as possible, be non-intrusive and transparent to the occupants (i.e. requiring minimal involvement from the occupants). IB agents are sensor rich and it is difficult to be prescriptive about which sensor parameter set would lead to the most effective learning of any particular action. Thus, to maximise the opportunity for the agent to find an optimum input vector set, whilst containing the processing overloads, the ideal agent would be able to learn to focus on a sub-set of the most relevant inputs. Thus, minimally constrained approaches, that maximise agent learning opportunities, are favoured. Finally, the computationally compact nature of agents in IB (e.g. limited memory) has an effect that permeates all aspects of learning such as altering the extent of particularisation versus generalisation or the granularity of similarity clustering. One particular challenge here that needs to be addressed is how an agent can identify a change in the person's behaviour that might signal a need for specific forms of help available via mobile robots.

3. The Robots

We have built a number of robots ranging from small desktop vehicles to large outdoor diesel powered agricultural vehicles.

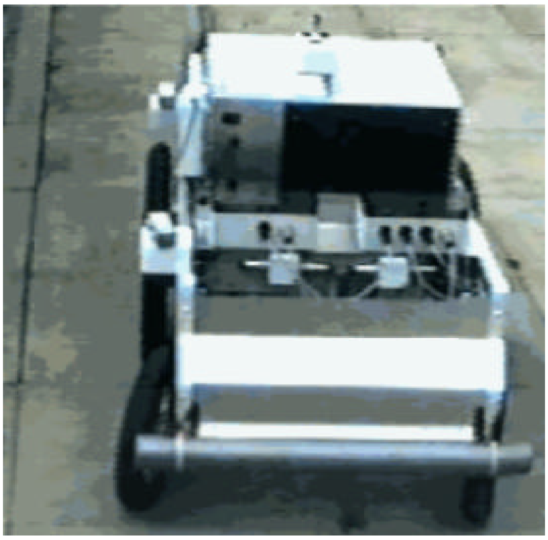

Figure (4): Wheelchair Robot

Of these robots, the wheelchair robot shown in Figure (4) is the one we are using inside our IB environment.

3.1. Agents: Technical Details

The internal agents that make up the mobile robots and the intelligent-buildings are based on the same principle: a double fuzzy logic and genetic learning mechanism. Details of this are published elsewhere [11, 5], but a figure showing the main parts is reproduced below in Figure (5).

In general terms, the architecture utilises fuzzy logic and genetic system principles, the fundamentals of which are widely known and thus are not reproduced here. The high-level operation of the control scheme belongs to a "school" labelled "behaviour based control architecture" pioneered by Rodney Brooks of MIT in the late 80's [3]. In this approach a number of concurrent behaviours (mechanisms to attain goal or maintain a state) are active (sensing environment, effecting machine) to a degree determined by the relationship between the machine and environment. At this macro-level, the novelty lies in our unique combination of fuzzy-based *behaviours* and *behaviour-integration* and a genetic-based "*Associative Experience Engine*" (AEE) (the latter itself containing various novel

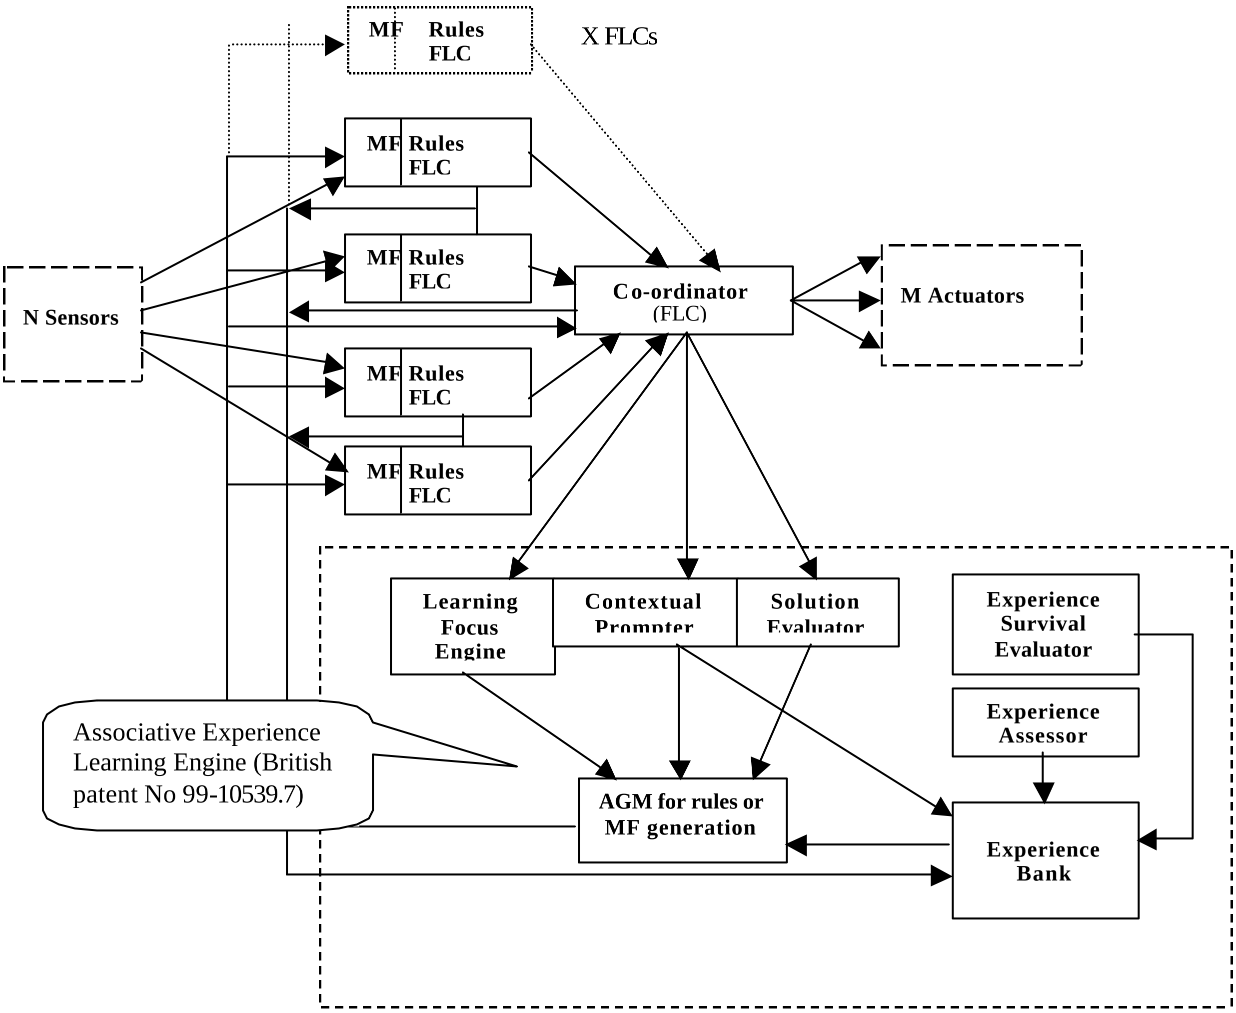

Figure (5): Architectural Overview of Associative Experience Learning Engine (British patent No 99-10539.7).

genetic mechanisms). The diagram provides an overview of both the novel macro and micro aspects of the agent. Behaviours are represented by parallel Fuzzy Logic Controllers (FLC). Each FLC has two parameters that can be modified; these are the *Rule Base* (RB) of the behaviour and its *Membership Functions* (MF). The behaviours receive their inputs from sensors. The output of each FLC is then fed to the actuators via the *Co-ordinator*, which weights its effect. In our agricultural agent, we employ four FLCs namely: Obstacle Avoidance (OA), Left Edge Following (LF), Right Edge Following (RF) and Goal Seeking (GS). When the system fails to have the desired response (e.g. deviating largely when following a crop edge or colliding with an obstacle), the learning cycle begins. Learning depends on the *Learning Focus* that is supplied by the *Coordinator* (the fuzzy engine which weights contributions to the outputs). When the *Learning Focus* learns the MF for individual behaviours, the MF for these behaviours is learnt in isolation. When the *Learning Focus* is learning an individual rule base of a behaviour, then each rule base of the behaviours is learnt alone. When the *Learning Focus* is adapting the coordinated behaviours online, then the algorithm will adapt different rules in the different behaviours in response to the environment. The system recalls similar experiences by checking the stored experiences in the *Experience Bank*. The robot tests different solutions from the *Experience Bank* by transferring the most recent experiences, which are stored in a queue. If these experiences show success than they are stored in the FLC and thereby avoid generating new solution for our system. The Experience Assessor assigns each experience solution a fitness value. When the *Experience Bank* is full, we have to delete some experiences. To assist with this the *Experience Survival Evaluator* determines which rules are removed according to their importance (as set by the Experience Assessor). When past experiences did not solve the situation we use the best-fit experience to reduce the search space by pointing to a better starting point. We then fire an *Adaptive Genetic Mechanism (AGM)* using adaptive learning parameters (except when learning behaviour with immediate reinforcement, we use an optimum mutation parameter) to speed the search for new solutions. The AGM is constrained to produce new solutions within certain ranges as defined by the *Contextual Prompter*. This is supplied by sensors and defined by the coordinator according to the *learning focus* in order to avoid the AGM searching options where solutions are unlikely to be found. By using these mechanisms, we narrow the AGM search space massively thus improving its efficiency. After generating new solutions (either rules or MFs) the system tests the new solution and gives it fitness through the *Solution Evaluator*. The AGM provides new options until a satisfactory solution is achieved. The system can be viewed as a double hierarchy system. In this, both fuzzy behaviours and the online learning mechanism can be seen to be hierarchies. In the case of the latter, at the higher level we have a population of queued solutions stored in the Experience Bank. If any

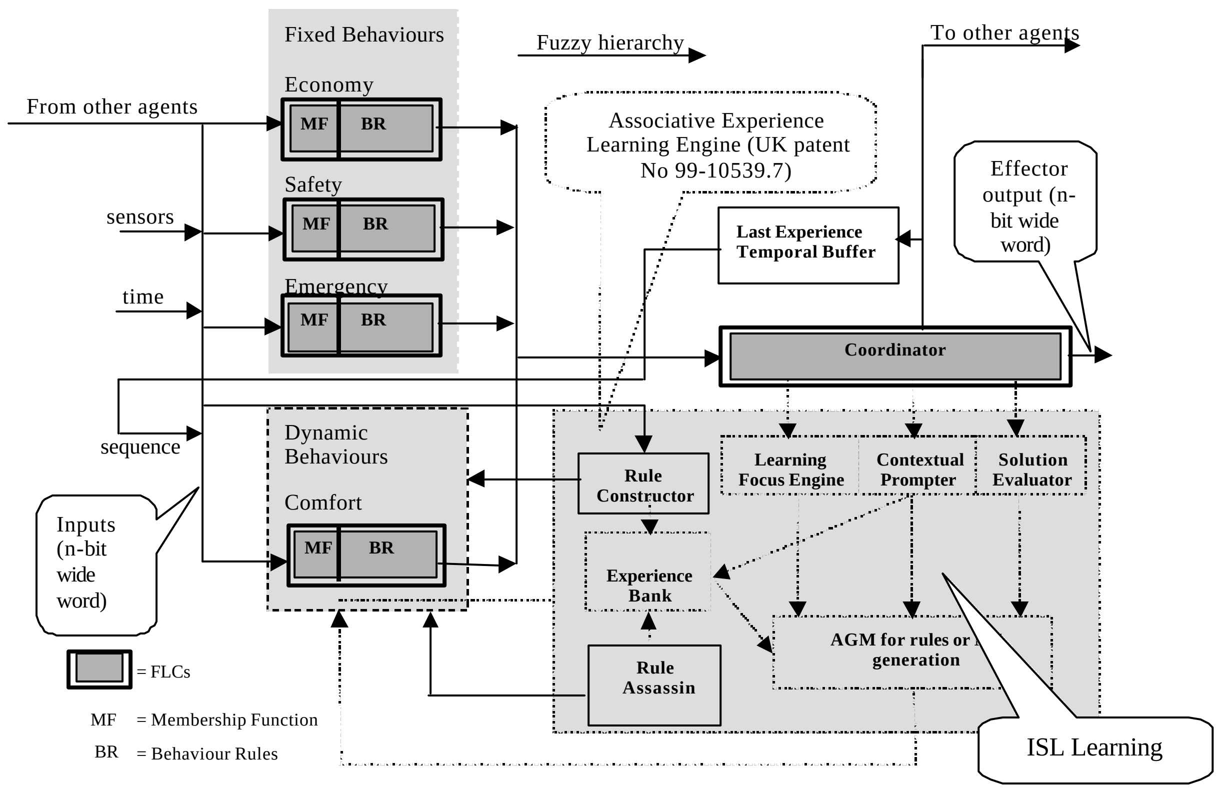

Figure (6): Modified AEE Embedded-Agent Architecture (ISL Learning) of these stored experiences leads to a solution then the search ends, if not then each of these experiences is assigned a fitness value. The fittest experience is used as a starting position to the lower level AGM that is used to generate a new solution.

The current IB agent is built on the same architecture as the mobile robot agent apart from the subtle alterations to the learning engine to provide it with a non-intrusive learning mechanism as shown in Figure (6). We have shown elsewhere that learning by forced-error methods as with the mobile robot would be both frustrating and incompatible with the IB learning paradigm. Thus, we have developed a non-intrusive learning technique for our agent that we refer to as *Incremental Synchronous Learning* (ISL). In simple terms this works as follows: when an occupant changes an effector setting manually, the system responds by immediately carrying out the action, setting the building to the requested state, generating a new rule based on that instance and initiating a new learning sequence. In this case, the learning sequence is equivalent to one iteration of forced-error learning in our mobile robot agent. At this point any further action is suspended until there is another interaction with the occupant. That is, *there are no forced interactions* with the occupant but rather the *occupant's spontaneous interactions* trigger a simple learning process. Thus learning is made unobtrusive by spreading the iterations over an extended period using the natural interactions of the user with the system. For example, in previous work [5] we performed experiments using an AEE based IB agent, configured as a temperature controller, based on a 68000 Motorola processor shown in Figure (1). The agent has a built-in "economy behaviour" (eg minimises heat low & ventilation in vacated rooms); "safety behaviour" (eg prevents temperature going below zero degrees which would result in pipes freezing) and "comfort behaviour" (eg maintains the room environment to the users liking). The agent was tested under various conditions such as multiple occupancy, different temperatures, human activity and times of day. The agent was able to generate a satisfactory rule base for different users in an average of 22 trials. Thus, assuming that each day the occupant makes an adjustment to the system (i.e. one learning trial) the agent would complete a learning cycle in an average of 22 days. We would argue that this is an acceptable time for an agent to learn to particularize its services to a person (given in a manual system the user will always need to command the system, whereas in the agent-assisted system the manual load upon the occupant reduces over time). In addition to providing a non-intrusive learning mechanism, this approach also places *the user in prime control* as it unfailingly and immediately responds to his command.

4. Future Work and Conclusions

We are currently planning a series of larger scale experiments involving a more complex *multi-use* environment made up of many more agents, sensors, effectors and dynamic events. This multi-use space takes the form of an *Intelligent Dormitory* (see figure 7). In our case, the dormitory is a student bed sitting room based upon standard student accommodation on the Essex University campus (being university all the necessary expertise and equipment is readily available). This test-bed will allow us to monitor and eventually control the environment based upon the activity of the room's occupier – a student researcher.

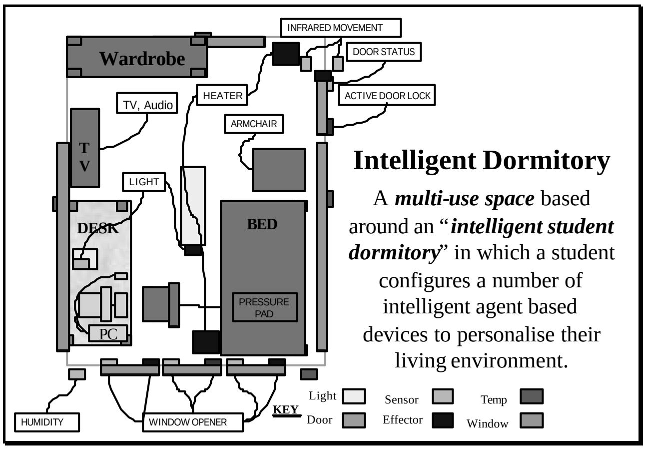

Figure 7 – Intelligent Student Dormitory

In parallel we are carrying out research into a novel position sensor and the use of fuzzy-based voice recognition and agent communication (the latter in partnership with the University of Otago).

Since writing this paper we have been given two new grants to take aspects of this work forward. The first under the EU's "Disappearing Computer Programme" (in Partnership with CTI Patras & NMRC Cork) concerns the placement of intelligence within a wider set of passive commodities (eg cups, pencils, cloth etc), the second under the UK-Korean Scientific fund concerns the integration of robots and intelligent-buildings in a rehabilitation environment (in partnership KAIST, Korea). Thus development of a multi heterogeneous agent *intelligent inhabited environment* is well under way and we look forward to report results on these projects in the near future.

5. Acknowledgements

We are pleased to acknowledge Malcolm Lear and Robin Dowling for their help building the experimental systems, together with Sue Sharples, Gillian Kearney, and Filiz Cayci for many stimulating discussions on embedded-agent issues.

6. References

[1] Genetic-Fuzzy Controller, UK No 99 10539.7, 7th May 1999

[2] A Bonarini, "Comparing Reinforcement Learning Algorithms Applied to Crisp and Fuzzy Learning

Classifier systems", *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 52-60, Orlando, Florida, July 1999.

[3] R Brooks, "Intelligence Without Representation", *Artificial Intelligence*, No. 47, pp139-159, 1991

[4] V Callaghan, G Clarke and A Pounds-Cornish, "Buildings As Intelligent Autonomous Systems: A Model for Integrating Personal and Building Agents", *Proc. of 6th International Conference on Intelligent Autonomous Systems*, Venice, Italy; July 25 - 27, 2000.

[5] V Callaghan, G Clarke, M Colley and H Hagras "Computing DAI Architecture for Intelligent *Studies in Fuzziness and Soft Computing on Soft Computing Agents, Springer-Verlag*, (currently due for publication in July 2001)

[6] F Cayci, V Callaghan and G Clarke, "DIBAL - A Distributed Intelligent Building Agent Language", *Proc. of 6th International Conference on Information Systems Analysis and Synthesis*, Orlando, Florida, July 2000.

[7] H Hagras, V Callaghan and M Colley, "A Fuzzy-Genetic Based Embedded-Agent Approach to Learning and Control in Agricultural Autonomous Vehicles", *IEEE International Conference on Robotics and Automation*, pp. 1005-1010, Detroit-U.S.A, May 1999.

[8] P Davidsson "Energy Saving and Value Added Services; Controlling Intelligent-Buildings Using a Multi-Agent System Approach" in *DA/DSM Europe DistribuTECH*, PennWell, 1998.

[9] H Hagras, V Callaghan and M Colley, "Online Learning of Fuzzy Behaviours using Genetic Algorithms & Real-Time Interaction with the Environment", *IEEE International Conference on Fuzzy Systems*, Seoul-Korea, pp. 668-672, August 1999.

[10] H Hagras, V Callaghan and M Colley, "Online Learning Of Fuzzy Behaviour Co-ordination For Autonomous Agents Using Genetic Algorithms And Real-Time Interaction With The Environment" *IEEE International Conference on Fuzzy Systems*, San Antonio, Texas, USA, 7-10 May 2000.

[11] H Hagras, V Callaghan and M Colley, "On-Line Learning Of The Sensors Fuzzy Membership Functions In Autonomous Mobile Robots", *IEEE International Conference on Robotics and Automation*, San Francisco, April 2000.

[12] J Mendel and L Wang, "Generating Fuzzy Rules by Learning Through Examples", *IEEE Trans. on Systems, Man and Cybernetics*, Vol. 22, pp. 1414-1427, December 1992.

[13] S Sharples, V Callaghan and G Clarke, "A Multi-Agent Architecture for Intelligent Building Sensing and Control" *International Sensor Review Journal*, May 1999.