

Creating an Ambient-Intelligence Environment Using Embedded Agents

Hani Hagra, Victor Callaghan, Martin Colley, Graham Clarke, Anthony Pounds-Cornish, Hakan Duman

Department of Computer Science, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK.

Abstract— We outline an approach based on using embedded agents for realising ambient environmental intelligence within a ubiquitous computing environment. We describe a practical ubiquitous computing test-bed environment, the Essex intelligent Dormitory (iDorm), which combines a number of heterogeneous computational artefacts and networks. We show novel results where a user occupied the iDorm for five and half days. During this period the iDorm was monitored and controlled by an intelligent and adaptive embedded agent using our fuzzy logic based Incremental Synchronous Learning (ISL) approach. The agent operates in a life-long learning mode realising the ambient intelligence vision in which it learns the user's behaviour and adapts to his needs and controls the iDorm according to the user's preferences in a personalised and non-intrusive way.

Keywords: Fuzzy Logic, Learning, Embedded Agents, Ambient Intelligence, Ubiquitous Computing, Pervasive Computing, Intelligent Environments

Ambient Intelligence in Ubiquitous Computing Environments

Ubiquitous computing is a paradigm where technology becomes virtually invisible. Instead of having a desktop or laptop machine, the technology we use is embedded in our environment [12]. Recent estimates show that of the order of 8 billion microprocessors that were produced in 2001, only 2% of which went into PCs, the rest went into embedded computer devices most people wouldn't recognize as computers [3]. The ever decreasing price and size of microelectronics devices opens the opportunity for embedded-computers to move beyond the bounds of familiar electronic goods into new areas such as being integrated into clothing fabric, building materials, etc [6]. As embedded-computers get smaller, and become integrated into non-computing artefacts, they become pervasive and effectively disappear. The addition of a communication capability to such artefacts and the use of pervasive networking allow such artefacts to be associated together and remotely accessed in both familiar and novel arrangements to make highly personalized systems. This new and exciting computing paradigm promises to revolutionize the way we interact with computers, services, and the surrounding physical spaces, yielding higher productivity and a more seamless interaction between users and computing services, however the challenge is how to manage, program or direct such systems. The vision of ambient intelligence can help to address this challenge.

Ambient intelligence refers to an exciting new information technology paradigm where people are empowered through a digital environment that is aware of their presence and context, and is sensitive, adaptive and responsive to their needs [6]. Ambient intelligence improves the quality of life by creating the desired environment conditions and functionality via intelligent, personalised interconnected systems and services. Ambient intelligence environments are characterised by their ubiquity, transparency and intelligence [6]. Ubiquitous because the user is surrounded by a multitude of inter-connected embedded systems, transparent because the computing equipment appears invisible to the user as it is seamlessly integrated into the background. Intelligent because the system is able to recognise the people that live in it and is able to program itself to meet their needs by learning from their behaviour [6].

For the vision of ambient intelligence to be realised, people must be able to use and configure the computer-based artefacts and systems present in their ubiquitous computing environments in a seamless, unobtrusive and non intrusive way without being cognitively overloaded [6]. If the user has to program each device and work out how to connect them together to achieve the required functionality, this vision may never be realised. The complexity associated with the number, varieties and uses of computer based artefacts leads to the need to design a system that would allow intelligence to disappear into the infrastructure of active spaces (e.g. buildings, shopping malls, theatres, domestic environments, etc)

November/December 2004 (Vol. 19, No. 6) pp. 12-20

[3], automatically learning to carry out everyday tasks based upon the users' habitual behaviour. One approach to achieve the vision of ambient intelligence in such ubiquitous computing environments is to use embedded agents to control such environments according to the user needs and preferences.

The work proposed by this paper focuses on the development of learning and adaptation techniques for embedded agents seeking to provide an online, life-long, personalised learning of anticipatory adaptive control to realise the ambient intelligence vision in ubiquitous computing environments. We will use the Essex intelligent dormitory (iDorm) as a test bed for this work and an exemplar of this approach.

Intelligent Embedded Agents

Embedded intelligence can be regarded as the inclusion of some capacity for reasoning, planning and learning in an artefact. Embedded-computers that contain such an intelligent capability are normally referred to as "*embedded-agents*" [3]. Each embedded agent is an autonomous entity and it is common for such embedded-agents (as intrinsic parts of "*intelligent artefacts*") to have a network connection thereby facilitating communication and cooperation with the other embedded agents to form multi embedded-agent systems [3].

An embedded agent in the form of a mobile robotic agent can learn and adapt, online, its navigation behaviours [7]. However we will concentrate on embedded agents present in ubiquitous computing environments that will try to realise the vision of ambient intelligence. Each embedded agent is connected to sensors and effectors comprising a ubiquitous computing environment. The agent uses our fuzzy logic based Incremental Synchronous Learning (ISL) system to learn and predict the needs of the user to automatically adjust the agent controller based on a wide set of parameters [2] in a non-intrusive and invisible way, which is one of the requirements for ambient intelligence [6]. There is thus a need to modify effectors for environmental variables (such as heat and light, etc) on the basis of a *complex multi dimensional input vector*. An added control difficulty is that people are essentially non-deterministic and highly individual. As the embedded agents are located on small embedded computers with limited processor and memory abilities, therefore any learning and adaptation system should be able to deal with these computational limitations.

Most automation systems, involving a minimum of intelligence, utilize mechanisms that generalize actions across a population e.g. set temperature or loudness that is the average of many peoples' needs. However to achieve the ambient intelligence vision, we argue that any type of intelligence applied to personal artefacts and spaces needs to *particularize* to the individual [3]. Further, we argue that it is essential that any agent serving a person should always and immediately carry out any requested action i.e. people are always in control, subject to overriding safety considerations, to achieve the responsive property implied in the ambient intelligence vision [6]. The embedded-agent learning technique we have adopted is characterized by its ability to particularize its actions to individuals *and* immediately execute the user commands. Our embedded agent will be tested in its ability to control the iDorm, which is a ubiquitous computing test bed with a big number of embedded sensors and actuators and a heterogeneous network set up.

Intelligent Inhabited Environments and Intelligent Buildings

Intelligent Inhabited Environments (IIE) are spaces such as cars, shopping malls, homes and even our own bodies that will allow these environments to respond "thoughtfully" to our needs. Such environments would consist of a multitude of, possibly disconnected active spaces providing ubiquitous access to system resources according to the current situation of the user. Such environments promise a future where computation will be freely available everywhere, in a similar fashion to the current availability of batteries and power sockets. These intelligent environments will personalize themselves in response to our presence and behaviour. Precursors to such environments can be found now in intelligent buildings [3].

A typical 'container' environment for ubiquitous computing is an intelligent building which could be a house or office. The heterogeneity, dynamism and context-awareness in a building make it a good choice to explore design challenges within the range of ubiquitous systems. We view intelligent buildings as computer-based systems, gathering information from a variety of sensors (and other computers), and using intelligent embedded-agents to determine the appropriate control actions of the various devices [3,7,8]. In controlling such systems one is faced with the imprecision of sensors, the large number of information sources, lack of adequate models of many of the processes and the non-deterministic aspects of human behaviour. It is important for the intelligent building that embedded agents are able to continuously learn and adapt to the needs of the individuals within it, whilst always providing a safe and timely response to any situation [8].

There are a growing number of research projects concerned with applying intelligent agents to IIE and intelligent buildings. In Sweden, Davidsson [5] utilised multi-agent principles to control building services. These agents are based on

November/December 2004 (Vol. 19, No. 6) pp. 12-20

the Artificial Intelligence (AI) thread that decomposes systems by function rather than behaviour, their work does not address issues such as occupant based learning. In Colorado, Mozer [10] uses a soft computing approach - neural networks - focusing solely on the intelligent control of lighting within a building. Mozer's system, implemented in a building with a real occupant, achieved a significant energy reduction, although this is sometimes at the expense of the occupant's comfort. Work at MIT on the HAL project [4] concentrated on making the room responsive to the occupant by adding intelligent sensors to the user interface. Context Aware systems such as that of the Aware Home work at Georgia Tech [1] represent a large body of current research effort but differ to our work in that they are more concerned with time independent context rather than temporal history or learning which is a central issue in our work. There are other high profile intelligent environments projects such as the Microsoft Smart House, BT's Tele-care and Cisco's Internet Home [11]. However most of these industrial projects, including home automation technologies, like Lonworks and X10 are geared toward using networks and remote access with some smart control (mostly simple automation) with sparse use of AI and little emphasis on learning and adaptation to the user's behaviour.

The iDorm – A Testbed for Ubiquitous Computing and Ambient Intelligence

We have chosen the Essex Intelligent Dormitory (iDorm) pictured in Figure (1-a) to be a demonstrator and test-bed for IIE and ubiquitous computing environments. We will use the iDorm to test our intelligent learning and adaptation mechanisms needed by the embedded agents for the realisation of the ambient intelligence vision in ubiquitous computing environments. Being an intelligent dormitory it is a multi-use space (i.e. contains areas with differing activities such as sleeping, working, entertaining, etc) and can be compared in function to other living or workspaces such as a one-room apartment for elderly or disabled people, an intelligent hotel room or an office. The iDorm contains the normal mix of furniture found in a study/bedroom allowing the user to live comfortably. The furniture (most of which are fitted with embedded sensors that provide data for the network) includes a bed, a work desk, a bedside cabinet, a wardrobe and a PC-based work and multimedia entertainment system. The PC contains most office type programs to support work and the entertainment support includes audio entertainment (e.g. playing music CDs and radio stations using Dolby 5.1 surround sound) as well as video services (e.g. television, DVDs, etc).

In order to make the iDorm as responsive as possible to the needs of the occupant it is fitted with an array of embedded sensors (e.g. temperature, occupancy, humidity, light level sensors, etc) and effectors (e.g. door actuators, heaters, blinds, etc) [9]. Amongst the many interfaces, we have produced a virtual reality system (VRML) shown in Figure (1-b) that marries the Virtual Reality Modelling Language with a Java interface controlling the iDorm. It provides the user with a visualisation tool showing the current state of the iDorm and allows direct control of the various effectors in the room.




Figure (1): a) Photo of iDorm. b) The iDorm VRML interface.

November/December 2004 (Vol. 19, No. 6) pp. 12-20

Although the iDorm looks like any other room, above the ceiling and behind the walls hide a multitude of different networks and networked embedded devices. In building the iDorm, we have installed devices that reside on several different types of networks. As such access to the devices needs to be managed, gateways between the different networks can be regarded as critical components in such systems, combining appropriate granularity with security. Currently the iDorm is based around three networks: Lonworks, 1-Wire (TINI) and IP, providing the diverse infrastructure present in ubiquitous computing environments and allowing the development of network independent solutions [9].

Lonworks is Echelon's proprietary network and encompasses a protocol for building automation. There are many commercially available sensors and actuators for this system. The physical network installed in the iDorm is the Lonworks TP/FP10 network. The gateway to the IP network is provided by Echelon's iLON 1000 web server. This allows the states and values of sensors and actuators to be read or altered via a standard web browser using HTML forms. The majority of the sensors and effectors inside the iDorm are connected via a Lonworks network.

The 1-Wire network, developed by Dallas semiconductor was designed for simple devices to be connected over short distances. It offers a wide range of commercial devices including small temperature sensors, weather stations, ID buttons and switches. The 1-Wire network is connected to a Tiny Internet Interface board (TINI board), which runs an embedded web server serving the status of the networked devices using a Java servlet. The servlet collects data from the devices on the network and responds to HTTP requests.

The IP network forms a backbone to interconnect all networks and other devices like the Multi-media PC (MMPC). The MMPC will be the main focus for work and entertainment in the iDorm. Again the MMPC uses the HTTP protocol to display its information as a web page.

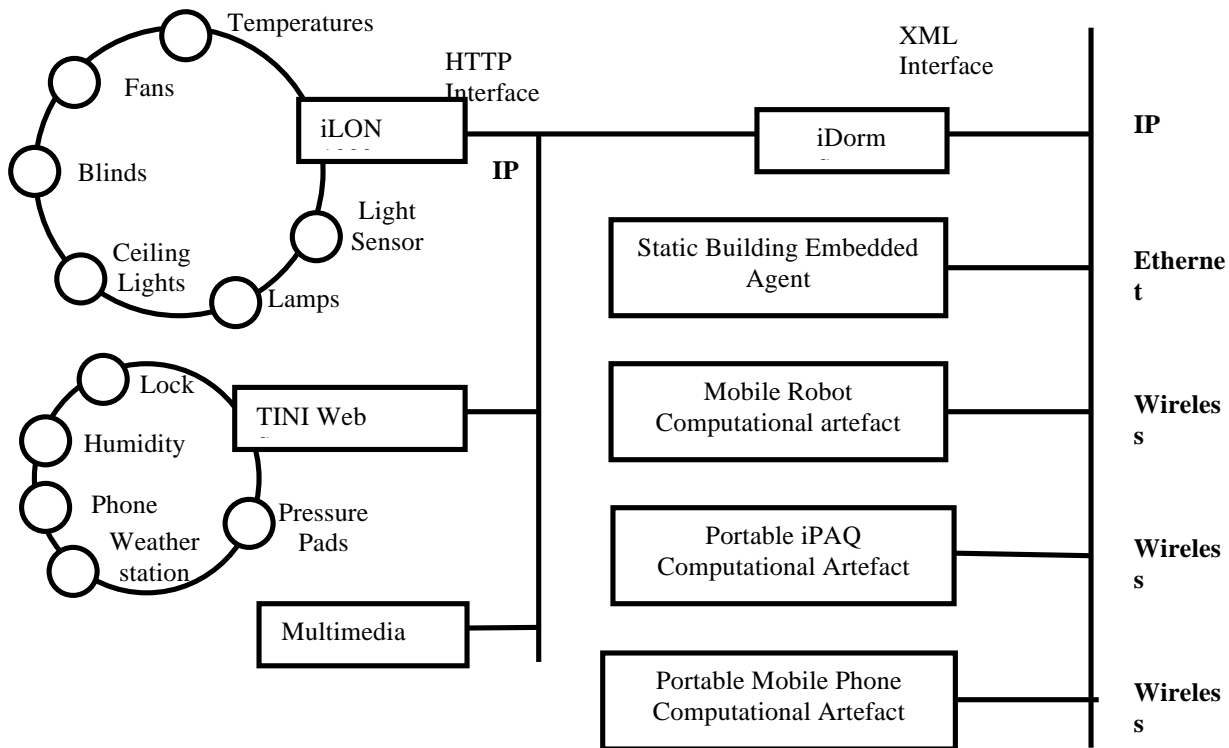


Figure (2): The logical network infrastructure in the iDorm.

The iDorm's gateway server is a practical implementation of a HTTP server acting as a gateway to each of the room's sub networks. This illustrates the concept that by using a *hierarchy of gateways* it would be possible to create a *scalable architecture* across such heterogeneous networks in IIE and ubiquitous computing environments [9]. The iDorm gateway server allows a standard interface to all of the room's sub networks by exchanging XML formatted queries with the entire principal computing components, which overcomes many of the practical problems of mixing networks. This gateway server will allow the system to operate over any standard network such as EIBus, Bluetooth and Lonworks and could

November/December 2004 (Vol. 19, No. 6) pp. 12-20

readily be developed to include 'Plug N Play' allowing devices to be automatically discovered and configured using intelligent mechanisms [9]. In addition, it is clear such a gateway is an ideal point to implement security and data mining associated with the sub network. Figure (2) shows a logical network infrastructure in the iDorm.

The iDorm Embedded Computational Artefacts

The iDorm has three types of embedded computational artefacts connected to the network infrastructure. Some of these devices contain agents.

The first type is a *physically static* computational artefact closely associated with the building. In our case this artefact contains an agent and thus is termed the iDorm Embedded Agent which is an example of the embedded agents targeted by this paper that try to realise the vision of ambient intelligence in ubiquitous computing environments. The iDorm agent receives the iDorm sensor values through the network infrastructure and contains the fuzzy Incremental Synchronous Learning system to learn the user's behaviour and compute the appropriate control actions and send them to iDorm effectors across the network. The iDorm Embedded Agent is shown in Figure (3-a) and is based on 68000 Motorola processor with 4 Mbytes of RAM, an Ethernet network connection and runs VxWorks Real Time Operating System (RTOS). The sensors and actuators available to the iDorm Agent are as follows:

The agent accesses eleven environmental parameters (some, such as entertainment, being parameters on multi-function appliances):

- Time of the day measured by a clock connected to the 1- Wire network
- Inside room light level measured by indoor light sensor connected to the Lonworks network
- Outside outdoor lighting level measured by an external weather station connected to the 1-Wire network
- Inside room temperature measured by sensors connected to the Lonworks and the 1-Wire networks
- Outside outdoor room temperature measured by external weather station connected to the 1- wire network
- Whether the user is using his audio entertainment on the computer – sensed by custom code publishing the activity on the IP network
- Whether the user is lying or sitting on the bed or not, measured by pressure pads connected to the 1-Wire network
- Whether the user is sitting on the desk chair or not, measured by a pressure pad connected via a low power wireless connection to the the 1-Wire network
- Whether the window is opened or closed measured by a reed switch connected to the 1-Wire network
- Whether the user is working or not, sensed by custom code publishing the activity on the IP network
- Whether the user is using video entertainment on the computer - either a TV program (via WinTV) or a DVD using the Winamp program sensed by custom code publishing the activity on the IP network

The agent controls nine effectors, which are attached to the Lonworks network:

- Fan Heater
- Fan Cooler
- A dimmable spot light above the Door
- A dimmable spot light above the Wardrobe
- A dimmable spot light above the Computer
- A dimmable spot light above the Bed
- A Desk Lamp
- A Bedside Lamp
- Automatic blind status (i.e. open/closed and angle)

The room is also supplied with other sensors such as a smoke detector, a humidity sensor, activity sensors and telephone sensor (to sense whether the phone is on or off the hook) as well as a camera to be able to monitor what happens inside the iDorm. It is possible to follow (and control) the activities inside the iDorm, via a live video link over the Internet.

The second type of the embedded computational artefacts is a *physically mobile* computational artefact. This takes the form of a service robot and contains an agent and thus termed robotic agent. The robotic agent can learn and adapt, online, the robot navigation behaviours [7] (which is different to iDorm embedded agent that seeks to realise ambient intelligence). The robot prototype used in the iDorm is shown in Figure (3-b). The robot can be regarded as a servant-

November/December 2004 (Vol. 19, No. 6) pp. 12-20

gadget with the aim of delivering various objects of interest to the user of the iDorm such as food, drink and medicine. The mobile robotic agent has a rich set of sensors (9 ultrasound, 2 bumpers and an IR beacon receiver) and actuators (wheels). It uses 68040 Motorola processors and runs VxWorks Real Time Operating System (RTOS). The robot is equipped with essential behaviours for navigation, such as obstacle avoidance, goal-seeking and edge-following. These behaviours are combined and co-ordinated with a fuzzy coordination module so the robot can reach a desired location whilst avoiding obstacles. The robot's location is passed to and processed as an additional input by the *static* iDorm embedded agent that controls the iDorm. In the experimental set up we use a simplified system in which the robot can go to two locations identified by infrared beacons to pick up objects. After picking up an object the robot can deliver it to the user and then go to its charging station, which is identified by another infrared beacon. The robotic agent sends information about its location to the iDorm agent and it takes destination instructions from that agent depending on the user's previous behaviour. For example the robot might have learned to go and fetch a newspaper from a specific location whenever it is delivered in the morning.

The communication between the *static* iDorm embedded agent and the *mobile* robotic agent is implemented via a wireless link. Communication is established by initiating a request from the iDorm embedded agent to the mobile robotic agent server. Once the request has been sent the server passes it to the robotic agent to carry out the task and informs the iDorm embedded agent of the robot's current status. If the task is in progress or not completely finished then the server sends a message indicating that the job is not complete. Every time the iDorm embedded agent wants to send out a new request, it waits until the previously requested job has been successfully completed.

The third type of the embedded computational artefacts is a *physically portable* computational artefact. Typically these take the form of wearable technology that can monitor and control the iDorm wirelessly. The handheld iPAQ shown in Figure (3-c) contains a standard Java process that can access and control the iDorm directly, this forms a type of "remote control" interface that would be particularly suitable to elderly and disabled users. Because the iPAQ supports Bluetooth wireless networking, it was possible to adjust the environment from anywhere inside and nearby outside the room. It is also possible to interact with the iDorm through mobile phones as the iDorm central server can also support the WML language. Figure (3-d) shows the mobile phone WAP interface which is a simple extension of the web interface. It is possible for such portable devices to contain agents but this remains one of our longer-term aims.

The learning mechanism within the iDorm embedded agent is designed to learn behaviours relating to different individuals. In order to achieve this the embedded agent needs to be able to distinguish between users of the environment. This is achieved by using an active lock, designed and built by our research team, based on Dallas Semiconductors 1-Wire protocol. Each user of the environment is given an electronic key, about the size of a penny. This is mounted onto a key fob and contains a unique identification number inside its 2-kilobyte memory. The Unique ID Number of the user is passed to the iDorm embedded agent so that it may retrieve and update previous rules learnt about that user.




Figure (3): The iDorm embedded computational artefacts a) Static iDorm Embedded Agent. b) Mobile Service Robot. c) Portable iPAQ interface. d) Portable mobile phone interface

Fuzzy Incremental Synchronous Learning Technique

In our work, learning is achieved through interaction with the actual environment and we call this *online* learning as adaptive behaviours cannot be considered as a product of an agent in isolation from the world, but can only emerge from strong coupling of the agent and its environment [8].

The Incremental Synchronous Learning (ISL) architecture is shown in Figure (4). The ISL forms the learning engine within the iDorm embedded agent and is the subject of British patent 99-10539.7. The ISL system aims to provide life-long learning and adapts by adding, modifying or deleting rules. It is memory based in that the system can use its previous experiences (held as rules) to narrow down the search space and speed up learning. The embedded agent is an augmented behaviour based architecture, which uses a set of parallel Fuzzy Logic Controllers, each forming a behaviour. We have used the Fuzzy Logic Controller (FLC) approach as it has proven to be very useful where the processes are too complex for analysis by conventional quantitative techniques or when the available sources of information are interpreted qualitatively, imprecisely or uncertainly [7,8], which is the case of embedded agents operating in IIE and ubiquitous computing environments. For embedded agents, the number of inputs and outputs are usually large and the desired control behaviours are complex. However, by using a *hierarchical* assembly of fuzzy controllers, the number of rules required can be significantly reduced [7,8]. In general we divide the behaviours available to the iDorm embedded agent into fixed and dynamic sets, where the dynamic behaviours are learnt from the person's behaviour and the fixed behaviours are pre-programmed. These latter behaviours need to be predefined because they cannot easily be learnt e.g. the temperature at which water pipes freeze. The fixed behaviours include a safety behaviour, an emergency behaviour and economy behaviour. The *Safety behaviour* ensures that the environmental conditions are always at a safe level. The *Emergency behaviour*, which in the case of a fire alarm or another emergency, might open the emergency doors and switch off the main heating and illumination systems. The *Economy behaviour* ensures that energy is not wasted so that if a room is unoccupied the heating and illumination will be switched to a sensible minimum value. All of the previous behaviours are fixed but adjustable.

Each dynamic FLC (the comfort behaviour in the iDorm case) has one parameter that can be modified which is the *Rule Base* (RB) for each behaviour. Also, at the high level the co-ordination parameters can be learnt [7,8]. Each behaviour uses a FLC using singleton fuzzifier, triangular membership functions, product inference, max-product composition and height defuzzification. The selected techniques were chosen due to their computational simplicity and real-time considerations. The equation that maps the system input to output is given by:

$$Y_t = \frac{\sum_{p=1}^M y_p \prod_{i=1}^G \alpha_{Aip}}{\sum_{p=1}^M \prod_{i=1}^G \alpha_{Aip}} \quad (1)$$

Where M is the total number of rules, y_p is the crisp output for each rule, $\prod \alpha_{Aip}$ is the product of the membership functions for each rule's inputs and G is the number of inputs.

We use a higher-level *Co-ordinator* FLC to combine the preferences of different behaviours into a collective preference (giving a two-level behaviour hierarchy). More information about the fuzzy hierarchical architecture can be found in [7,8]. The ISL works as follows: - when a new user enters the room they are identified by the active key button and the ISL enters an initial *Monitoring mode* where it learns the user's preferences during a non intrusive cycle. In the Experimental set-up we used a period of 30 minutes but in reality this is linked to how quickly and how completely we want the initial rule base. For example in a care home we might want this rule base to be as complete as possible, in a hotel we might want this initialisation period to be small to allow fast learning. The rules and preferences learnt during the Monitoring mode form the basis of the user rules which are reactivated whenever the user reenters the room. During this initialisation period the system monitors the inputs and the user's action and tries to infer rules from the user's behaviour. The user will usually act when a set of environmental conditions (an input vector) is unsatisfactory to him by altering the output vector (he needs to turn a light on or adjust the heating etc). Learning is based on negative reinforcement, as the user will usually request a change to the environment when he is dissatisfied with it.

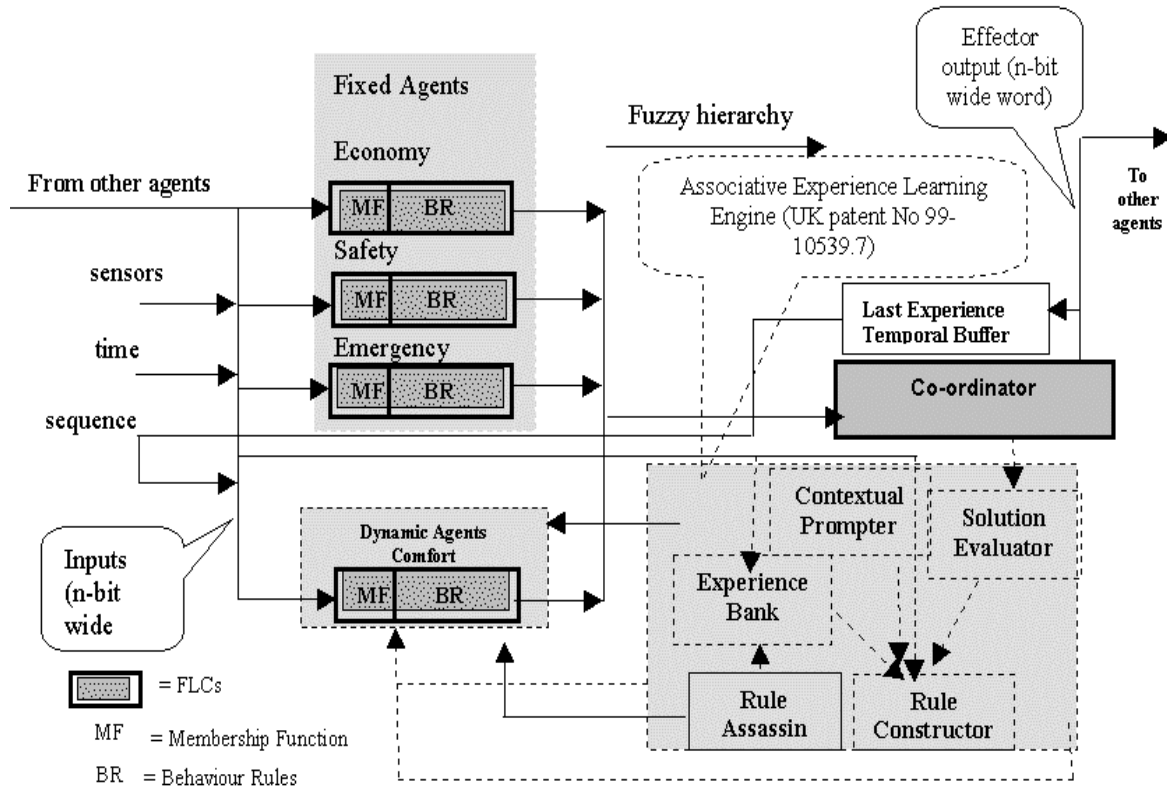


Figure (4): The ISL Architecture

After the monitoring period the ISL enters a *Control mode* in which it uses the rules learnt during the Monitoring mode to guide its control of the room's effectors. Whenever the user behaviour changes, there may be a need to modify, add or delete some of the rules in the rule base. Thus the ISL goes back to the non-intrusive cycle and tries to infer the rule base change to determine the user's preferences in relation to the specific components of the rules that have failed. This is a very short cycle that the user is essentially unaware of and such modifications and adaptation are distributed throughout the lifetime of the use of environment, thus forming a life-long learning phase.

As in the case of classifier systems, in order to preserve system performance the learning mechanism is allowed to replace a subset of the classifiers (the rules in this case). The worst m classifiers are replaced by m new classifiers [2]. In our case we will change all the consequents of the rules whose consequents were unsatisfactory to the user. We find these rules by finding all the rules firing at this situation whose firing strength $\Pi\alpha_i > 0$. We replace these rule consequents by the fuzzy set that has the highest membership of the output membership function. We make this replacement to achieve non-intrusive learning avoiding direct interaction with the user. The learnt consequent fuzzy rule set is guided by the *Contextual prompter*, which uses the sensory input to guide the learning.

During the non-intrusive monitoring and the life-long learning phases the agent encounters many different situations as both the environment and the user's behaviour change, e.g. the agent attempts to discover the rules needed in each situation guided by the occupant's behaviour in response to different temperature and lighting levels inside and outside the room. The learning system consists of learning different episodes; in each situation only small number of rules will be fired. The model to be learnt is small and so is the search space. The accent on local models implies the possibility of learning by focusing at each step on a small part of the search space only, thus reducing interaction among partial solutions. The interaction among local models, due to the intersection of neighbouring fuzzy sets means local learning reflects on global performance [2]. So we can have global results coming from the combination of local models, and

November/December 2004 (Vol. 19, No. 6) pp. 12-20

smooth transition between close models. By doing this we don't need to learn the complete rule base all at once but we learn only the rules needed by the user during the different episodes. There is a significant difference in our method of classifying or managing rules, rather than seeking to extract *generalised* rules we are trying to define *particularised* rules.

After the initial monitoring phase the system tries to match the user derived rules to similar rules stored in the *Experience Bank* that were learnt from other occupiers. The system will choose the rule base that is most similar to the user-monitored actions. By doing this the system is trying to predict the rules that were not fired in the initialisation session thus minimising the learning time as the search starts from the closest rule base rather than starting from a random. This action should be satisfactory for the user as the system starts from a similar rule-base then fine-tunes the rules. Subsequently the agent will operate with rules learnt during the monitoring session plus rules that deal with situations uncovered during the monitoring process, which are ported from the rule base of the most similar user, all the rules that are constructed and added to the system are symbolised by the *Rule Constructor* block in Figure (4). The system then operates in the Control mode with this rule-base until the occupant's behaviour indicates that his needs have altered which is flagged by the *Solution Evaluator* (i.e. the agent is event-driven). The system can then add, modify or delete rules to satisfy the occupant by briefly re-entering the monitoring mode. In this case again the system finds the rules fired and changes their consequent to the action exhibited by the user. It is in this way that the system implements a life long learning strategy.

It should be noted that, as we are dealing with embedded agents with limited computational and memory capabilities, it is very difficult to deal with a large number of rules in the rule base, e.g. for the iDorm comfort behaviours used in our current implementation, the complete set of possible rules for the iDorm is 62208 rules and this would lead to large memory and processor requirements which are not realistic in embedded agents. Therefore we set a limit on the number of stored rules to 450 (in our case, the maximum number the agent can store on the onboard memory without exceeding the memory limit or degrading the real-time performance). Each rule will have a measure of importance according to how frequently this rule is used since it was created. In calculating this *degree of importance* we also include a measure of *most-recent-use*. When the memory limit is reached the *Rule Assassin* retains rules according to the priority of *highest-frequency of use*, followed by most-recently-used. If two rules share the same degree of relative rule frequency recall tie breaking is resolved by a *least-recently-used* rule. This action allows the onboard memory to only store the most efficient and frequently used rules and to reduce degradation of the real time performance of the embedded agent. However the rules chosen for replacement can be stored with the other rules representing the user behaviour in an external hard disk so that they can be recalled when needed.

Experimental Results

We have conducted a novel experiment over five and a half days (132 hours) in which a user occupied the iDorm while the iDorm was under the control of the iDorm embedded agent using the ISL architecture.

The user is identified by his electronic intelligent key, which operates the active lock as shown in Fig (5-b). Throughout the experiment, the user used the wireless iPAQ to monitor and control the iDorm environment whenever they were dissatisfied with the current state of the environment. A history of the user decisions was recorded in a journal. One of our axioms is that "the user is king" by which we mean that an agent always executes the user's instruction immediately, to achieve the responsive property implied in the ambient intelligence vision, unless safety is compromised. This is shown in Figure (5-a) as whenever changes to controls occurred, the iDorm embedded agent received the request, generates a new rule or adjusts a previously learnt rule and allowed the action through. A small parsing tool was written to convert the text file containing the fuzzy rule sets into a human readable format. At the end of the experiment, the rules were converted into this form and examined in two different ways. The first involved comparing the human readable rules with user's journal entries to ensure that the agent had successfully learnt the behaviours the user was intending. The second was to compare the number of rules learnt over time. The embedded agent success can be measured by monitoring how well it matches the environment to the user's demands. If it does this well, the user will intervene less, which will result in less rule generation over time. If it does this badly, the user will intervene more which will result in more rule generation over time.

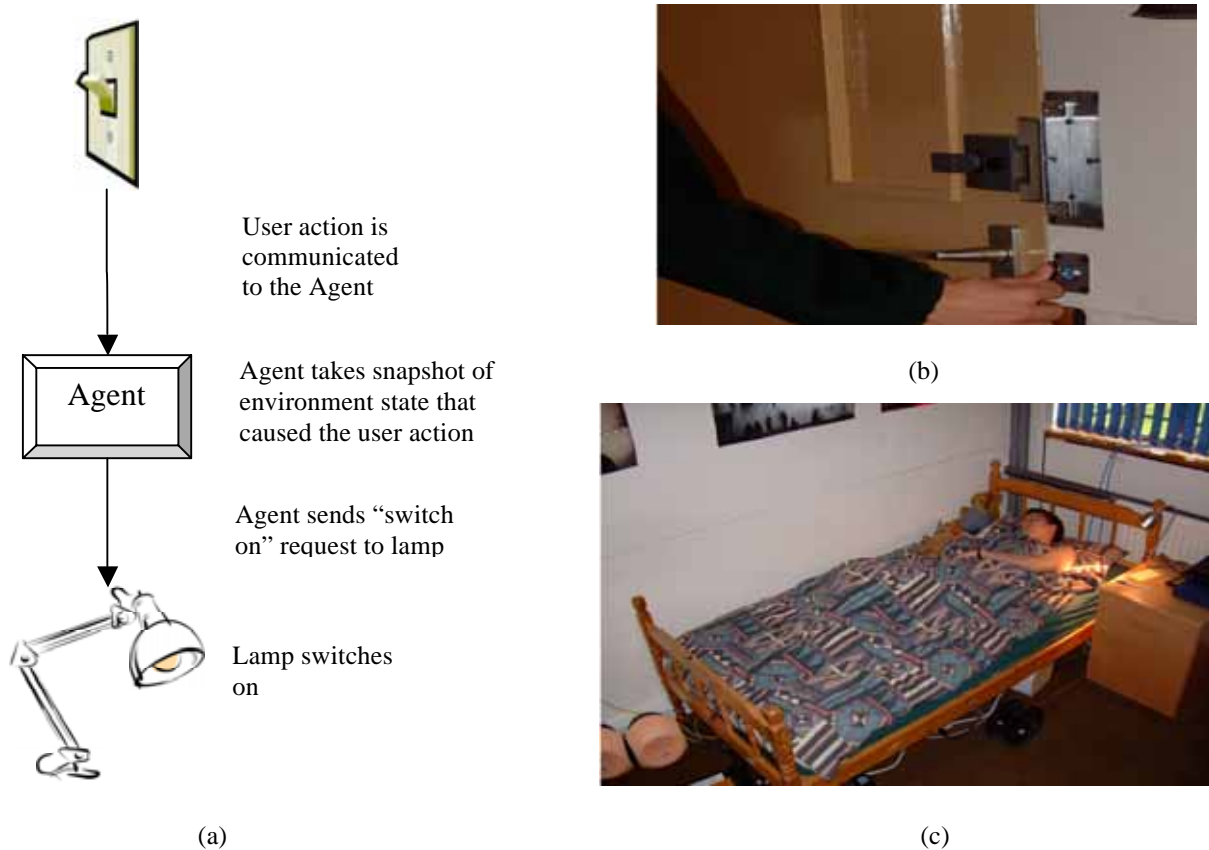


Figure (5): a) The Agent Communication Path. b) Using the active lock. c) The user is sleeping in the bed in the iDorm, which is controlled by the ISL based embedded agent.

The results of the experiment are summarised in Figure (6). The data making up the graph was generated by a logging program, which stored the number of rules learnt by the ISL together with a time stamp. The logging program took a reading every five minutes for the duration of the experiment. There are several sections of the graph worthy of note.

In the first section from *0-24 hours* starting from an initial rule base of zero, it can be seen that a very large number of rules are learnt in a comparatively short period of time. In fact, in the first nine hours of the experiment, the agent learnt 128 rules. This is nearly half the total number of rules learnt by the end of the experiment. These results are consistent with the agent learning and making incorrect decisions for the user in the initial stage. However, as the end of the period is reached it can be seen that the learning rate of the agent (Rules/Time) reduces drastically. This is consistent with the agent's reactions requiring less correction by the user and is therefore consistent with the agent making useful decisions about the environment based on the user's requirements. The latter trend of fewer rules learnt over time is consistent across the whole of the experiment. Hence the level of "comfort" experienced by the user (in relation to the environment state) is high enough for them not to make an environmental change and consequently alter the learning rate.

The second section of the graph from *60-72 hours* shows a sharp increase in the learning rate of the agent. This is explained by the user introducing novel activity into his repertoire of behaviours. It shows that our system, which operates in a life long learning mode, is adaptive to the user needs.

The third section *72-132 hours* shows that in last two days the agent had not generated any new rules. This constitutes a control period where the user was generally satisfied with agent control action that resulted from the embedded agent successfully learning the user behaviour.

The agent has learnt the 280 rules needed to capture the behaviour of this user over the 132 hours experiment, which demonstrated that our system can learn effectively using the ISL and it doesn't need to learn the complete rule base

November/December 2004 (Vol. 19, No. 6) pp. 12-20

(potentially 62208 rules in the case of the iDorm). Figure (6) shows that the agent had to learn less new rules about the user as the experiment progressed. Since this was one of our criteria for measuring the agent’s success the evidence of the continual reduction in the learning rate, leads us to conclude that the agent managed to pick out the pertinent behaviour of the user over time.

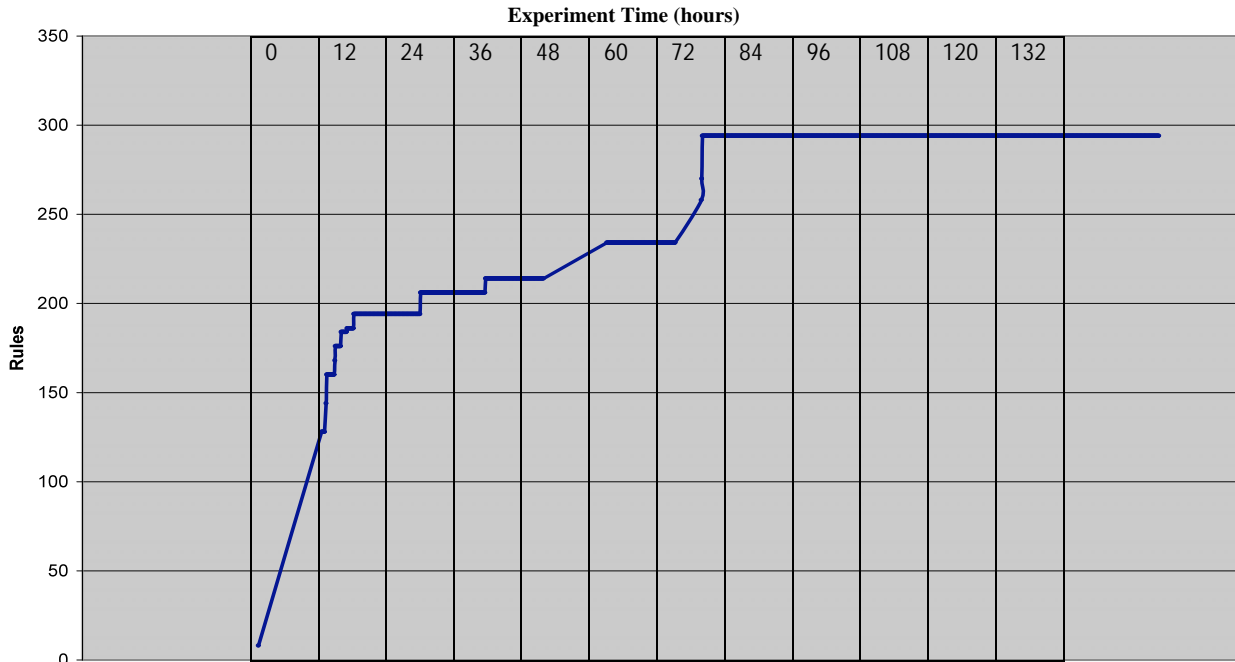


Figure (6): The rules learnt plotted against experiment time

An example of a rule learnt by the agent could be written as follows:

In the evening when it is dark inside and bright outside, temperate inside and warm outside, sitting on the bed, with the window open I switch the heater off, switch the cooler off, switch the door spotlight off, switch the wardrobe spotlight off, switch the computer spotlight off, switch the bed spotlight off, turn off the bed lamp turn on the table lamp and close the blind and send the robot to the bed.

The rule involves the robot attending to the user, who is lying on the bed, responding to a specific input vector. The mobile robot was used as a method of transporting specific objects to the user such as food or a can of drink. These objects were placed on a flat area of the robot that could be used as a tray, retrieved, used and replaced when finished.

In our previous work we tried experimenting with different room users [7] and the role of the Experience Bank was important in reducing the time the ISL takes to learn the user’s behaviour and achieve user satisfaction. This is because it starts learning the user’s behaviour from the best matching behaviour previously recorded rather than starting from random.

These experiments showed how our embedded agent approach can help to realise the vision of ambient intelligence in a ubiquitous computing test bed environment like the iDorm by learning and adapting the user’s habitual behaviour and controlling the iDorm according to the user activity in a responsive and particularised way. The learning was done in a non-intrusive manner without the user aware of the existence of the agent.

Conclusions

We have demonstrated how our embedded agent approach can contribute towards the realisation of the vision of ambient intelligence in ubiquitous computing environments. We demonstrated how we support cognitive disappearance, by making our agent able to learn the user's habitual behaviour so it can effectively programme itself in an adaptive way in a non-intrusive manner. The results suggest that over the experimental period the agent made a significant reduction in the need for intervening by the user. This same capability not only reduces the complexity of use but can bring significant cost and effort savings over the evolving lifetime of products by avoiding expensive programming (and re-programming). In addition it empowers the ordinary users by allowing them to use collections of computer based artefacts to design novel systems to suit their own personal tastes without the need for them to understand the technical complexities or programme the systems (i.e. self programming or learning is done by the agent). The shape of the learning curve in Figure (6) shows the learning (or self-programming) in operation and suggests the agent moved increasingly closer towards the user's environmental preference even though this preference was never static.

We have performed a five and half day experiment using the iDorm and demonstrated that our fuzzy logic based ISL can develop useful particular rules over a short time frame. This technique has the ability to add, delete or modify rules in the dynamic behaviour rule base, which is an essential feature of a life-long learning embedded agent. We have shown that the ISL embedded agent is both immediately reactive and responsive to a user's command (subject to safety restrictions) and is able to particularise itself to the user's behaviour, both of which are an essential feature of any agent that is to be acceptable to people and for the realisation of the ambient intelligence vision. These experiments suggest that surprisingly few rules are required by the agent (i.e. in the order of a few hundred) in order to autonomously create a comfortable environment with diminishing need for user correction. This is important new information as prior to this work it was not known whether a tractable rule-set would emerge and whether embedded architectures with only the order of MBs of memory would be able to host agents for such ubiquitous computing environments.

Future Work

Our current experimental programme includes plans for multi-user and multi-room habitation experiments. We also plan wider deployment of embedded agents (e.g. personal agents inside wearable technology) and experiments on differing agent granularities. We are in the process of building a multi-roomed version of the iDorm, iDorm-2, as a preliminary step towards the construction of a fully functional apartment (iFlat), which will house both visiting researchers and act as a unique ubiquitous computing test-bed. The iDorm-2 is currently being designed for such experimentation from the ground up.

Acknowledgements: We are pleased to acknowledge the funding support from the EU IST Disappearing Computer program and the Korean-UK Scientific Fund programme.

References

- [1] G. Abowd, M. Eibling, G. Hunt G, H. Lei, "Context Aware Computing" IEEE Pervasive Computing, Volume 1, No 3, pp. 22-23, 2002
- [2] 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] V. Callaghan, M. Colley, G. Clarke, H. Hagra, "Embedding Intelligence: Research Issues for Ubiquitous Computing", Proceedings of the Ubiquitous Computing in Domestic Environments conference, Nottingham, September 13 - 14, 2001.
- [4] M. Coen, "Design Principles for Intelligent Environments". In Proceedings of the Fifteenth National Conference on Artificial Intelligence. (AAAI'98). Madison, WI. 1998.
- [5] P. Davidsson "Energy Saving and Value Added Services; Controlling Intelligent-Buildings Using a Multi-Agent System Approach" in DA/DSM Europe DistribuTECH, PennWell, 1998.
- [6] K. Ducatel, M. Bogdanowicz, F. Scapolo, J. Leijten and J.C. Burgelma., "Scenarios for Ambient Intelligence in 2010", (ISTAG 2001 Final Report), IPTS, Seville, 2001.

November/December 2004 (Vol. 19, No. 6) pp. 12-20

- [7] H. Hagra, V. Callaghan, M. Colley, G. Clarke, H. Duman, "Online Learning and Adaptation for Intelligent Embedded Agents Operating in Domestic Environments " In the book entitled Fusion of Soft Computing and Hard Computing for Autonomous Robotic Systems" (Eds.: Zhou, Dario Maravall and Da Ruan), "Studies in Fuzziness and Soft Computing Series , Physica-Verlag, Volume 116, pp. 293-323,November 2002.
- [8] H. Hagra, V. Callaghan, M. Colley, G. Clarke, "A Hierarchical Fuzzy Genetic Multi-Agent Architecture for Intelligent Buildings Learning, Adaptation and Control", The International Journal of Information Sciences, Volume. 150, pp. 33-54, March 2003..
- [9] A. Holmes, H. Duman, A. Pounds-Cornish "The iDorm: Gateway to Heterogeneous Networking Environments", International ITEA Workshop on Virtual Home Environments, 20-21st February 2002. Paderborn, Germany.
- [10] M. Mozer "The Neural Network House: An Environment That Adapts To Its Inhabitants". In Proc of American Association for Artificial Intelligence Spring Symposium on Intelligent Environments, pp. 110-114, AAAI Press, 1998.
- [11] A Sherwin. Internet Home Offers a Life of Virtual Luxury. The Times, pp. 10, 3rd Nov 1999.
- [12] M. Weiser, "Some Computer Science Problems in Ubiquitous Computing," Communications of the ACM, July 1993. (reprinted as "Ubiquitous Computing". Nikkei Electronics; December 6, 1993; pp. 137-143.)