In proceedings of the i3 2001, workshop WS4 on Cognitive Versus Physical Disappearance, Porto, Portugal, April 2001

The Cognitive Disappearance of the Computer: Intelligent Artifacts and Embedded Agents

(or, do artefact based computers need to be intelligent to disappear?)

V Callaghan, M Colley, G Clarke, H Hagras
Intelligent Inhabited Environment Research Group
University of Essex, Colchester, England
vic@essex.ac.uk   cswww.essex.ac.uk/intelligent-buildings

Summary:

In this paper we argue that embedding intelligence into artefacts is an essential step in making the computer cognitively disappear. We do this by explaining the enhanced functionality that embedded-intelligence can provide to everyday products. In particular we describe how intelligence is the key to groups of artefacts learning to work together to achieve higher level, user-determined goals. We outline a scenario based on an "Intelligent Inhabited Environment", being built at Essex that will allow experimentation on cognitive disappearance arising from a networked system of intelligent artefacts. We explain the challenges facing those seeking to develop methods of embedding intelligent into computationally compact and distributed cooperating artefacts. Finally we summarize our arguments as to why "cognitive disappearance" requires intelligent artifacts and describe some of the projects we are working on that address these underlying research issues.

1. Introduction

Today people's personal spaces are increasingly "decorated" by electronic or computer-based artifacts (gadgets) varying from, mobile telephones through CD players to cars and beyond. The variety of computer-based artefacts, and their capabilities, is growing at an unprecedented rate fuelled by advances in microelectronics and Internet technology. Cheap and compact microelectronics means most everyday artifacts (e.g. shoes, cups) are now potential targets of embedded-computers. While ever-pervasive networks will allow such artefacts to be associated, together in limitless and novel ad-hoc arrangements to make highly personalised systems. To realise this dream non-technical users must be shielded from the need to understand or work directly with the technology "hidden" inside such gadgets or gadget worlds; *the computer must disappear!* How can this aim be achieved? In this paper we will seek to show that embedding intelligence into artefacts could provide one viable solution.

2. Embedded-Intelligence and Artefacts

The opening question of a Turing-type test for the disappearing computer - the "Essex Disappearing Computer Test" – might be, "*Which of these machines contains a computer?"* If you can't tell the computer has disappeared! Perhaps, if the computer had truly disappeared the answer would be, "*what machine?*" This is a somewhat light-hearted adaptation of the famous test for machine intelligence posed by Alan Turing, one of the early pioneers of computing and artificial intelligence. Although posed in a humorous vein, there are some useful parallels that can be drawn from this analogy. Essentially the premise is, "*is the user both physically and cognitively aware of the existence of the computer within the machine?*" It might be possible to argue that the computer has physically already disappeared in that numerous computers are already embedded into products such as kitchen goods, entertainment systems, communication devices, security systems, information appliances, transport etc. which are used by non-technical people who are unaware they are using computers (at least in the technical sense). Cognitively, the computer (via the machine interface) remains very evident as, for example, with a washing machine or video recorder, the user is forced to refer to complicated manuals and to use his own *reasoning* and *learning* processes to use the machine successfully.

From the above we have formed the view that, "*the cognitive disappearance of the computer*" is intimately linked to the amount of cognitive processes in the form of reasoning, planning or learning that the user is required to undertake in order to use a particular artefact, or collection of artefacts. Following this, we would argue that if some degree of the reasoning, planning and learning, normally provided by a gadget user, were embedded into the gadget itself, then by that degree the computer would cognitively disappear. Put another way, the proportion of reasoning, planning and learning transferred to the gadget (collectively referred to as "embedded-intelligence") is a "cognitive disappearance " metric! Hence we view *embedded intelligence as an essential property of artifacts for the cognitive disappearance of the computer*.

Our work at Essex University is focused on the development of computationally compact mechanisms of embedding intelligence into artefacts for the development of intelligent inhabited environments. In the remainder of the paper, we discuss the issues involved, the techniques we have developed and describe an exemplary scenario.

3. Disappearance: The AI Challenges

Above we argued that transferring some cognitive load from the users into the artefact was a key element in achieving cognitive disappearance. However, this is far from easy as such "intelligent artefacts" operate in a computationally complex and challenging physical environment which is significantly different to that encountered in more traditional PC programming or AI. Some of the computational challenges associated with creating systems of intelligent-artefacts are discussed below. As a precursor to this discussion we first overview some of the more general issues and terminology.

Embedded intelligence can be regarded as the inclusion of some of the reasoning, planning and learning processes in an artefact that, if a person did it, we would regard as requiring intelligence. An intelligent artefact would normally contain only a minimal amount of "embedded-intelligence"; only sufficient to do the artefact task in question. Embedded-computers that contain such an intelligent capability are normally referred to as "*embedded-agents*" [Callaghan 00]. Intelligent Artefacts would, in effect, contain an embedded-agent. Individually, such an embedded-agent can harness intelligence to undertake such tasks as:

- Enhancing Artefact functionality (enabling the artefact to do more complex tasks)
- Simplifying or automating the user interface (in effect, providing an intelligent assistant)
- Reducing Programming Costs (the system learns its own program rules)

It is now common for such "*embedded-agents*" (as intrinsic parts of "*intelligent artefacts*") to have an Internet connection thereby facilitating multi embedded-agent systems. In a fully distributed multi embedded-agent systems each agent is an autonomous entity cooperating, by means of either structured or ad-hoc associations with its neighbours. Each agent can reason or plan how it might work with those with whom it is currently associated thereby supporting *evolving aims or emerging functionality*. Without autonomous learning and ad-hoc association it is difficult to see how emergent functionality could otherwise be achieved. Because of this we argue that autonomy and intelligence are important attributes for intelligent artifacts if emergent behaviour is going to be possible. It is important to understand that being autonomous and promiscuous (open to making associations with other artifacts) does not imply undirected or unsafe behaviour. Agents can have basic fixed rules built in to them that prevent them taking specified actions deemed unsafe.

An interesting and potentially productive application of intelligent-artefacts arises when they are assembled and operated in synergetic groups. Perhaps artefacts will most commonly find themselves as part of rooms people live in. Rooms are often highly personalised, decorated by artefacts carefully chosen to suit tastes and needs. Rooms can be regarded as the building block of many habitats from cars and offices to homes. Rooms usually have a function (e.g. living, sleeping, driving etc) and the group of artefacts within a room will invariably reflect in part at least this function and well as the characteristics of the person that "decorated" the room with the artefacts.

Most automation systems (which involve a minimum of intelligence) utilise mechanisms that generalise actions (e.g. set temperature or volume that is the average of many people's needs). However, we contend that AI applied to personal artefacts and spaces needs to *particularise* itself to the individual. Further, subject to safety constraints, we contend that it is essential that any agent (or artefacts) serving a person should always and immediately carry out any requested action, no matter how perverse it may appear (i.e. people are always in control, subject to overriding safety considerations). The embedded-agent techniques we will outline are characterised by their ability to particularise their actions to individuals *and* immediately execute command, wherever that is a practical possibility. Elsewhere, the social and commercial issues of future widespread employment of agent based artefacts are more exhaustively discussed [Clarke 00] and related work on applications such as intelligent-buildings are explored [Brooks 97, Callaghan 00, Minar 99, Mozer 98, Davisson 98].

Artefacts that include intelligent agents of the type we describe inherit all these above-mentioned capabilities

3.1 The Issue of Physical Size and Cost

For physical disappearance artefacts will need relatively small low-cost embedded computers (possibly based on application specific micro-electronic fabrication). For example typical specifications might be Cost: £20-£50, Size: $<2^2$cm, Speed: 1-10MHz, Memory: 1-2 MB, I/O: 10-50 I/O channels. Examples of two real devices are shown in figures 1 & 2.

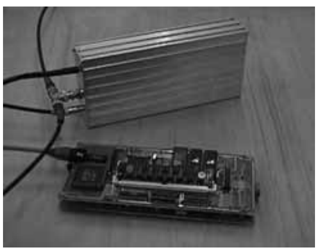

Figure 1 - University of Essex Prototype building services agent

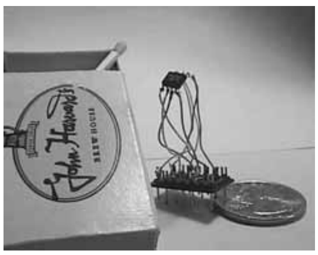

Figure 2 - University of Massachusetts Prototype Embedded-Internet Device

While it is inevitable that the "computing power / cost ratio" will continue to increase (i.e. more mega-everything per dollar), history has shown that functionality will always demand even faster computers. Thus available resources for a given cost always lag behind needs. The classic illustration of this dilemma is the defiance of the hard disk to become extinct despite 30 years of predictions of semiconductor memory becoming cheap and abundant. Of course the prediction that memory will become cheap and abundant has always proved correct but it seems functional demands have outpaced it. The lesson here is that although it is inevitable embedded-computers will become much more powerful, they will always be less powerful than the functionality demanded at that future point!

Traditional artificial intelligence (AI) techniques are well known for being computationally demanding and therefore unsuitable for 'lean' computer architectures. Historically most traditional AI system were developed to run on powerful computers such as workstations, whose specifications are at least 2 orders of magnitude removed from most embedded-computers. In addition traditional AI techniques have proved too fragile to operate real time intelligent machines such as robots. As a result, even implementing simplified traditional AI systems on embedded-computers has proved virtually impossible. However, the authors have techniques from developments of their earlier work in robotics that seem well suited to providing artefact intelligence [Callaghan 01, Hagras 00, Hagras 01] which are discussed later in this paper.

3.2 The Issue of Distribution

In most disappearing computer style scenarios, computer based artefacts are able to form ad-hoc groupings which work together to achieve some higher-level purpose. From an AI viewpoint this raises questions such as:

1. How is AI (agent) functionality and computation distributed (e.g. what is the computational granularity of artefacts, are they computationally and functionally autonomous).
2. How are associations to other artefacts formed and recorded (i.e. does each artefact decide and record its own associations or is this centrally managed and recorded)? Such associations are critical to group coordination, synergy and learning.
3. How are the dynamics of artefact mobility and failure handled (how do artefacts chose between competing services or cope with the removal of a service)?
4. How is group control and contention arbitrated (is there a master artefact in overall charge or is this devolved)?
5. How do artefacts/embedded-agents communicate with each other (what is an appropriate and compact language to support the expression needed for generalised intelligent-artefact communication and cooperation)?

Figure 2 shows a high–level diagrammatic view of a distributed intelligent-artefact architecture that goes some way to address these problems. In this diagram each artefact is responsible for determining which other artefacts (which might include sensors and effectors) to associate with and holding it's own local record (no global record is maintained). The system is initialised with a set of associations deemed the artefacts "sphere of influence" (e.g. these associations may be set using a manual editing tool). In the learning mechanism that is outlined later it will be seen that embedded-agents have the ability to evaluate which of the associations (and associated input stimuli) is important to its event based decision mechanism (discarding those that are not influential). The agent may also look beyond its prescribed

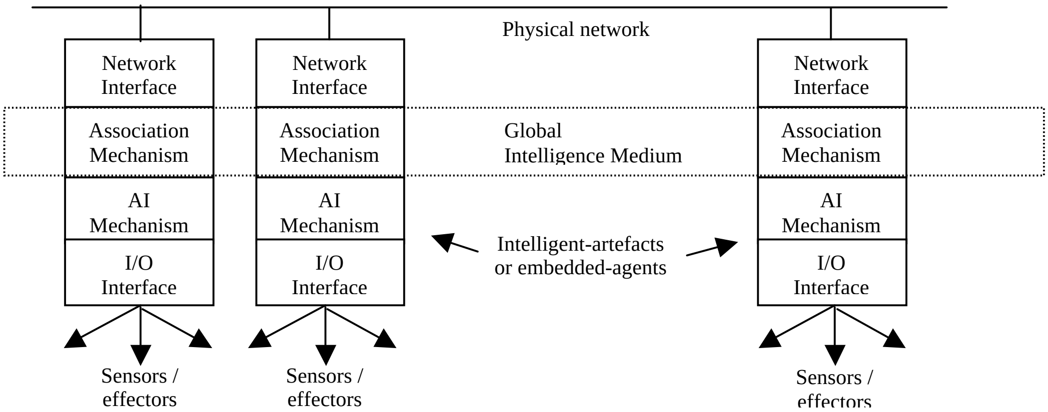

Figure 2 – A Distributed Intelligent Artefact Architecture (with implicit global intelligence)

associations for new associations that might provide input stimuli that improve its decision making process (thus autonomously creating new associations). Through this combined mechanism of association formation and removal, the Global Association medium assumes an implicit and global learning intelligence. Records of association are fully distributed throughout the system with each artefact knowing only about its own associations. These may be interrogated and modified by a manual editor as well as the autonomous self-learning process of each agent.

In order to carry out this sort of co-ordination and communication, intelligent-artefacts need a language to communicate, and to request and provide a variety of services to other artifacts. This needs to operate securely and robustly in a dynamic environment using a minimum of computational resources. Much work has been done on agent communication languages such as KQML, FIPA, Jackal, JafMas etc [Finin 94], the latter being frameworks that utilise JAVA. In a study we have completed we have shown that languages aimed at traditional AI applications are unsuitable in term of their computational demands and functionality. In our related "intelligent-buildings" work we have developed a language, DIBAL, [Cayci 00] which used a tagged hierarchical format to create a highly compact agent language that overcomes many of the problems associated with the more functionally rich traditional agent communication languages such as KQML. It is possible that such a communication language might be adapted for inter-artefact communication based on the architecture described above.

3.3 The Issue of Dimensionality and Temporality

The quality of agent decisions is limited by its knowledge of the world. It gets its knowledge from sensors directly attached to it and other agents (i.e. indirectly from their sensors). Which set of sensor information is sufficient for an agent to make a particular class of decision? Consider a simple heating controller, why does the room's occupant alter the heat value. Is it to do with the current temperature, his current level of activity, what he is wearing, where he is in a room, where he has just been or what? We may decide that it is based upon current temperature and therefore could operate with only one sensor, but later discover that an agent that used only one sensor was not working very effectively. At the other extreme we could decide we should sense 'everything' and then let the agent learn which of these sensed values was important. Clearly in this latter situation the agent would be able to make better-informed decisions and adapt to changing criteria. In addition this problem exposes a central dilemma, what is the best mechanism for selecting relevant sensory sets for agents? Is it the designer or the agents themselves? The problem with the designer is the assumption that people know best what the intelligent agent needs; but is this true? We would argue that it is better to provide a large set of sensory inputs to agents and let them resolve which of the stimuli is important for any given decision wherever possible. Whilst this latter argument may have some appeal it carries with it a penalty, the need to compute using large sensory input vectors. Thus, large sensory sets are an issue for intelligent-artefacts. One solution is the development of mechanisms that allow embedded-agents to “focus” on sub-sets of data relating to specific decisions or circumstances. An additional problem is that of time and sequences. Often the reason an action is taken is not simply related to the current state of the world, but to the sequence of states that led up to the most recent event. Thus, an effective embedded-agent would need to be able to deal with temporality.

In general the foregoing are the most difficult problems to be addressed. In terms of temporality we are beginning to investigate state-machine based methods, whereas for dimensionality our methods largely rely on manual focusing although we are working towards more automated mechanisms based on constraint satisfaction methods.

3.4 The Issue of Non-Determinacy, Intractability and Dynamism

Traditional AI is based around the so-called Sense-Model-Plan-Act (SMPA) architecture. In this there is a presumption that the world the agent acts upon can be abstractly described by either a mathematical model or some form of well-structured representation. In addition, it is usually presumed that the state of the world can be sensed reasonably reliably and compared to the abstract representation so as to reason or plan about the world. This approach works reasonably well for some forms of problem such as chess playing programs where many of these axioms hold true but completely fails in applications such as robotics and other applications that involve an intimate relationship with the physical world. The reason that traditional AI fails in such physical applications has been well described by others [Brooks 91] but a simplified explanation would be that the assumption that the world can be accurately sensed and modelled (the key axiom of SMPA) does not hold. For example, a robot interacting with the world does so via imperfect and sparse sensing, monitoring physical phenomena and people which are either of intractable complexity or are essentially non-deterministic (e.g. the actions of people is well known as defying predication since it is often the outcome of idiosyncratic whims!). In addition, it has proved virtually impossible to adequately represent the world, or to maintain a consistent representation in real-time of a highly dynamic world (e.g. objects and associations changing through deliberate actions or failures), resulting in lose of synchronisation between the model and the real world with associated catastrophic results.

We have shown elsewhere that intelligent-artefacts (containing embedded-agents) are essentially equivalent to robots, experiencing similar problems with sensing, non-determinism, intractability, lose of synchronisation etc [Callaghan 01]. Thus whatever techniques are used to embed intelligence into artefacts will require these issues to be addressed.

Fortunately, robotics has generated a potential solution for this type of problem that works by discarding the abstract model and replacing it by the world itself; a principle most aptly summarised by Rodney Brooks as, “the world is its own best model”. This AI school is known as “new AI” or perhaps more meaningfully “behaviour based AI”. Our earlier work [Callaghan 01, Hagras 00, Hagras 01] was in the field of robotics, which has allowed us to recognise the underlying similarities between robotics and intelligent artefacts.

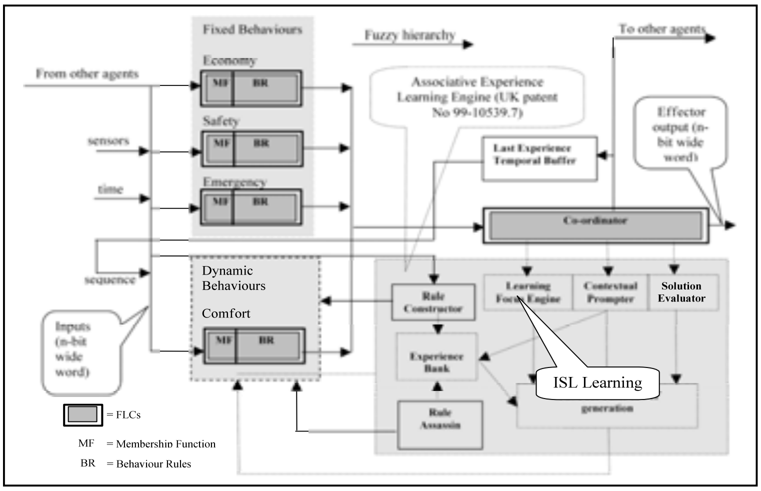

Figure 3 – An Embedded-Agent architecture for Intelligent-Artefacts

In our robot-based embedded-agents, which we have also used within an intelligent building environment, we encode behaviour based architecture principles via hierarchical fuzzy logic in which logic rules (programming) are formed by a novel real-time genetic algorithm. It is not the purpose of this paper to describe this agent mechanism although we include a high-level diagram in figure 3 and refer the interested reader to our other papers which debate these principles in some considerable depth [UK patent 99, Colley 01, Callaghan 01, Hagras 00, Hagras 01]. In simple terms the operation of the agent in an intelligent building scenario is 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 the equivalent of one iteration of forced-error learning in our mobile robot agent. At this point any further action is suspended until there was another interaction with the occupant. That is, there is no forced interaction with the occupant but rather the occupant's spontaneous interactions trigger a simple learning process. Thus, by spreading the iterations over an extended period using the natural interactions of the user with the system learning is made unobtrusive. For example, if we consider a temperature controller, each day the occupant might make an adjustment to the system (i.e. one learning iteration) and complete a learning cycle in, say, 21 days (c.f. our experimental data reported in Callaghan 01 & Hagras 00). We would argue that this is an acceptable time for an agent to learn to particularize its services to a person as, in a manual system, the user will always need to control the system, whereas in the agent-assisted system the manual load upon the occupant should reduce 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 commands.

4. Intelligent Inhabited Environments – An Intelligent Dormitory

In this section we give an overview of an intelligent-artefact space we are constructing at Essex to illustrate the kind of technology involved. More extensive descriptions of the technology involved in these rooms is given elsewhere [Callaghan 01, Colley 01, Hagras 01]. We have chosen a student dormitory (see figure 4) to be a demonstrator and test-bed for some of the techniques involved. The dormitory constitutes a personal space populated by an assortment of personal computer-based artefacts, many of which are to be configured by the occupant. Being a student dormitory it is a multi-use space (i.e. contains areas with differing activities such as sleeping, working, entertaining etc). The occupant of the room (a student) would be free to decorate his room with whatever artefacts he chooses (both computer and non-computer based; passive and active). Because this room is of an experimental nature we are fitting it with a liberal placement of sensors (e.g. temp. sensors, presence detectors, system monitors etc) and effectors (e.g. door actuators, equipment switches etc), which the occupant can also configure and use. Our expectations are that the occupant would chose to decorate his personal space (the room) with a variety of artefacts ranging from building service devices such as heaters to entertainment systems such as CD/TV. A possible scenario is as follows. The student moves into the dormitory, which contains some existing artefacts (mostly connected with the room infrastructure) but brings other more personal artefacts with him. He then runs a configuration program on his PC that allows him to set up associations between sensors and effectors.

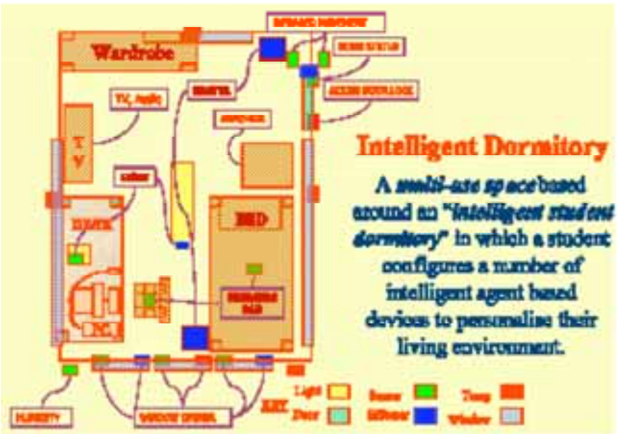

Figure 4 - Intelligent Inhabited Environment

Figure 5 - A World of Interacting Artefacts

To take a mundane example concerning the room's infrastructure, the student might set an association between a light switch immediately inside the door and a number of room lights. In addition he could personalise this space by deciding to associate the same light switch sensor to his radio, so that the radio switches on whenever he enters the room. He then continues until he has associated together all the sensors, effectors and artefacts that interest him. Taking a more speculative example; consider an intelligent-artefact in the form of a pair of shoes with embedded-computers / sensors to measure temperature and pressure both inside and outside of the soles; *magic shoes*! Further, consider the shoes use compression forces from ordinary walking, running etc. to generate electricity and they can communicate with a micro-locality via some form of proximity coupling between the sole and the room surface. By associating these shoes with intelligent artefacts in the room all sorts of possibilities emerge. For instance, the shoes could be associated with the switches within the room and various facilities can be switched on when the occupant is in their vicinity (e.g. on entering the room the lights are automatically switched on). If the shoes were also keeping some information on the spring in the room occupant's step or the level of activity this might be used to help guide a heating system or perhaps notifying a carer of increasing incapacity. By associating the sole of the shoe with an alarm should the occupant be wearing his shoes but the shoes not be in contact with the ground, then it might alert a carer to the possibility that the occupant may have had an accident (e.g. an elderly person who has fallen and cannot get up again). With the inclusion of some minimal AI in the shoes they can provide better quality and higher-level information to the artefacts they are associated with (e.g. they might deduce a likely accident). Clearly this is somewhat speculative but, with a world in which ambient embedded-computing is all-pervasive, it opens up endless possibilities of radically altering people's lives. Having set up a basic artefact association the occupant may then choose to switch the artefacts into an active online learning mode (or leave them as manually set). In general the room and artefacts function as non-agent based systems, interacting with the user through conventional controls (no special embedded-agent controls are necessary and the user is essentially unaware agents exist, or this is anything other than a normal environment). In the active mode artefacts monitor their use, in relation to the state of their local world, programming themselves to satisfy the occupant by doing what he habitually and persistently wants (i.e. not simply learning random whims of a user but rather learning long term persistent requirements). This has been called 'learning inertia' in the embedded-agent research we have undertaken. At the same time as learning habitual and persistent user requirements, the embedded-agents also respond immediately to any command made by the occupant. Thus after some time has passed the intelligent-dormitory may have learnt how to configure and operate the constituent intelligent-artefacts to the benefit of the occupant. This description is not comprehensive in coverage, and clearly speculative in places, but we hope it helps expose some of the issues and gives a feel for they type of operational issues and possibilities involved.

5. Summary / The Future

In his paper we have sought to argue that transferring some cognitive capabilities from people into artefacts was a natural (if not essential!) way to facilitate the disappearance of computers as computers find themselves increasingly embedded into artefacts. We have also argued embedded-intelligence can bring significant cost and effort savings over the evolving lifetime of product by avoiding expensive programming (and re-programming). In particular, if people are to use collections of computer based artefacts to build systems to suit their own personal tastes (which may be unique in some sense) then self programming embedded-agents offer one way of allowing this without incurring an undue skill or time overhead. However, whilst this paper argues strongly that integrating embedded intelligent agents into artefacts is highly beneficial, the paper exposes several significant problems, many of which remain as research challenges. For instance, dealing with the problems of non-determinism, dimensionality and temporality in computationally compact environments are very challenging topics.

We presented an overview of an intelligent inhabited environment in the form of a student dormitory that we plan to use as a test-bed for some of these intelligent artefacts in the eGadgets project (part of the EU Disappearing Computer Programme) and the CareAgent project (part of the Korean-UK Scientific Fund Programme) and we look forward to reporting results from this environment in a future paper. We noted that previous papers from our group have reported on extensive experimental results in similar environment that suggests that embedded-agents can significantly contribute to making effective computer based artefacts in which the computer has cognitively disappeared (to a significant extent)

Acknowledgements: We are pleased to acknowledge the contribution of Malcolm Lear and Robin Dowling (Essex University) for their help building the intelligent dormitory and intelligent-artifacts. We would also like to thank, Anthony Pounds-Cornish, Sue Sharples, Gillian Kearney, and Filiz Cayci for their indirect contributions arising from many stimulating discussions on intelligent-artifact and embedded-agent issues. In addition, thanks are due to Kieran Delaney (NMRC) for the drinks and conversations in a Patras bar that led to some of the unconventional intelligent-artifact ideas!

References

[Brooks 91] Brooks R, "Intelligence Without Representation", Artificial Intelligence 47, pp139-159, 1991

[Brooks 97] Brooks R. "*The Intelligent Room Project*" Proceedings of the Second International Cognitive Technology Conference (CT'97*)*, Aizu, Japan, August 1997

[Callaghan 00] Callaghan V, Clarke, G, Pounds-Cornish A "*Buildings As Intelligent Autonomous Systems: A Model for Integrating Personal and Building Agents*", The 6th International Conference on Intelligent Autonomous Systems (IAS-6), Venice, Italy; July 25 - 27, 2000

[Callaghan 01] Callaghan V, Clarke, G., Colley, M., Hagras, H. "*A Soft-Computing DAI Architecture for Intelligent Buildings*", Journal of Studies in Fuzziness and Soft Computing on Soft Computing Agents, Physica-Verlag-Springer, June, 2001

[Cayci 00] Cayci F, Callaghan V, Clarke G, "*DIBAL - A Distributed Intelligent Building Agent Language*", The 6th International Conference on Information Systems Analysis and Synthesis (ISAS 2000), Orlando, Florida, July 2000

[Clarke 00] Clarke G, Callaghan V, Pounds-Cornish A "Intelligent Habitats and The Future: The Interaction of People, Agents and Environmental Artefacts", 4S/EASST Conference on Technoscience, Citizenship and Culture in the 21st Century , Vienna, 26-28th September 2000

[Colley 01] Colley, M., Clarke, G., Hagras, H, Callaghan V, "Intelligent Inhabited Environments: Cooperative Robotics & Buildings" 32nd International Symposium on Robotics (ISR 2001), Seoul, Korea April 19-21, 2001.

[Davisson 98] P. Davisson "Energy Saving and Value Added Services; Controlling Intelligent-Buildings Using a Multi-Agent System Approach" in DA/DSM Europe DistribuTECH, PennWell, 1998.

[Finin 94] T. Finin, R. Fritzson, D. Mckay, R. McEntire. " KQML: An Information and Knowledge Exchange Protocol, "Knowledge Building and Knowledge Sharing, K. Fuchi and T. Yokoi (Eds), Ohmsha and IOS press, 1994.

[Hagras 00] Hagras H, Callaghan V, Colley M, Clarke G "*A Hierarchical Fuzzy Genetic Agent Architecture for Intelligent Buildings Sensing and Control*", RASC 2000 - International Conference on Recent Advances in Soft Computing June 29 & 30 2000, Leicester, UK

[Hagras 01] Hagras H, Callaghan V, Colley M, Clarke G, "A *Hierarchical Fuzzy Genetic Multi-Agent Architecture for Intelligent Buildings Learning, Adaptation and Control*", International Journal of Information Sciences, August 2001

[Minar 99] Minar N., Gray M., Roup O., Krikorian R., Maes P. "Hive: Distributed Agents for Networking things. MIT Media Labs". ASA/MA August 3, 1999. Cambridge, USA

[Mozer 98] MC 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.

[Sharples 98] Sharples S, Callaghan V, Clarke G, "*The Application of Intelligent Building Techniques to Care Service Provision*", IEE Colloquium on Intelligent Methods in Healthcare and Medicine, York. Oct 1998.

[**Sharples 99]** Sharples S, Callaghan V, Clarke, G "*A Multi-Agent Architecture For Intelligent Building Sensing and Control*", Int'l Sensor Review Journal, Vol. 19. No. 2. May 1999

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