

Research articles

A multi-agent architecture for intelligent building sensing and control

*Sue Sharples
Vic Callaghan and
Graham Clarke*

The authors

Sue Sharples, Vic Callaghan and Graham Clarke are Researchers in Intelligent Machines at the Department of Computer Science, Essex University, Colchester, UK. Tel: +44 1206 872 682; Fax: +44 1206 872 788; E-mail: vic@essex.ac.uk

Keywords

Healthcare, Intelligent buildings, Multi-sensor systems, Sensors

Abstract

We describe a new approach to intelligent building systems, that utilises an intelligent agent approach to autonomously governing the building environment. We discuss the role of learning in building control systems, and contrast this approach with existing IB solutions. We explain the importance of acquiring information from sensors, rather than relying on pre-programmed models, to determine user needs. We describe how our architecture, consisting of distributed embedded agents, utilises sensory information to learn to perform tasks related to user comfort, energy conservation, safety and monitoring functions. We show how these agents, employing a behaviour-based approach derived from robotics research, are able to continuously learn and adapt to individuals within a building, while always providing a fast, safe response to any situation. Finally, we show how such a system could be used to provide support for older people, or people with disabilities, allowing them greater independence and quality of life.

Electronic access

The research register for this journal is available at <http://www2.mcb.co.uk/mcbrr/sr.asp>

The current issue and full text archive of this journal is available at <http://www.emerald-library.com>

A house is a machine for living in
(Le Corbusier, 1921).

Introduction

The idea of an intelligent building probably conjures up images of helpless occupants struggling to maintain their sanity and freedom, while pitted in a life and death struggle against an all-powerful computer lurking in the basement! Thankfully, the reality is quite different, as buildings that many would describe as intelligent have existed for some time.

The building industry uses the term intelligent, to describe the way the design, construction and management of a building can ensure that the building is flexible and adaptable, and therefore profitable, over its full life-span. A definition which finds favour with many building managers and architects is that "An intelligent building is one that provides a productive cost-effective environment through the optimisation of four basic elements: systems, structures, services, management and the inter-relationship between them" (Robathan, 1989).

Computer scientists, however, have a very different view of intelligence. We are more concerned with giving machines management, analytic and control capabilities that are comparable to intelligent human activity. In the context of a building, a system works by taking inputs from building sensors (light, temperature, passive infra-red, etc.), and using this and other information to control effectors (heaters, lights, electronically-operated windows, etc.). If this system is to be intelligent, an essential feature must be its ability to learn from experience, and hence adapt appropriately. Thus the notion of "autonomous governing" is important, as it implies a system which can adapt and generate its own rules (rather than being restricted to simple automation). We propose a computer science

The authors are pleased to acknowledge the support of Hamilton Lodge Trust (in particular David Heather), Balkerne Gardens Trust (in particular Michael Siggs) and netCam Ltd (www.netcam.ltd.uk), whose building control hardware forms the basis of our experimental systems. The authors also wish to express their appreciation of assistance received from University staff, notably Malcolm Lear, Robin Dowling, Martin Colley and Paul Chernet for assistance in setting up the experimental system.

Sensor Review

Volume 19 · Number 2 · 1999 · pp. 135–140

© MCB University Press · ISSN 0260-2288

definition – “An intelligent building is one that utilises computer technology to autonomously govern the building environment so as to optimise user comfort, energy-consumption, safety and monitoring-functions”. We view intelligent buildings as computer-based systems, akin to robots, gathering information from a variety of sensors, and using embedded intelligent agent techniques to determine appropriate control actions.

Our work is concerned with utilising an intelligent embedded-agent approach (similar to the approach already taken in some areas of mobile robotics), to create an integrated and semi-autonomous building control system. Numerous types of buildings exist to which a system of this type could be applied. For demonstration purposes we have chosen to focus on one specific problem: enabling people who are elderly or have physical or learning disabilities to achieve as great a degree of independence and self-sufficiency as possible. Hence, we are primarily involved with ordinary domestic buildings, as well as small units within residential or nursing homes.

Components of an intelligent building

Intelligent buildings are composed of numerous sensors, effectors and control units interconnected in such a way as to effectively form a machine. In theory, a wide range of sensors and controllers could be utilised. For example, sensors used might include temperature and light-level detectors, movement or occupancy sensors (such as passive IR), pressure pads, and smoke or gas detectors. Less commonly, status sensors (giving information on the current status of, for example, electronically operated windows or household appliances) and tagging systems (to detect the location of specific individuals) could also be used. Devices being controlled by the system, on the other hand, could include heating, lighting, ventilation, alarms, electronically-operated blinds, doors and windows, and standard household appliances (such as kettles or televisions).

For the building to function as an integrated system, the final technological ingredient required is a network. This network needs, ideally, to be real-time, and to have simple device interfaces comparable with the cheap nature of existing building devices such as light switches. This has led to the development of specialist networks described in Table I.

Table I Intelligent building standards

IB standard	Key attributes
X10	Oldest commonly available IB technology allowing limited control of common household control devices, through power line
CEbus	An Electronics Industry Association (EIA) standard which covers devices that communicate through power line wires, low voltage twisted pairs, coax wires, infrared, RF and fibre optics
LonWorks	Popular standard that covers communication media similar to CEbus. The principle focus of LonWorks is a chip known as a “Neuron” chip, which acts as a network node and includes all of the communications hardware, communications protocol plus a fuzzy-control like language
BACnet	An American Society of Heating, Refrigerating and Air-conditioning Engineers (ASHRAE) standard that is modelled on the Open Systems Interconnection (OSI) basic reference model that shields BACnet from obsolescence with respect to networking technologies. Honeywell, a well known building automation supplier, use both BACnet and LonWorks in their products (Newman, 1996)
NEST	Novell has created a standard, Novell Embedded Systems Technology which they aim to be used everywhere where intelligent devices may be useful: offices, cars, homes, etc.
Smart-House	Developed by the National Association of Home Builders (NAHB) for building into new houses (Strassberg, 1995)
CAN	Developed by the German company Bosch for the automotive industry. It is robust and potentially cheap being linked to economies of scale associated with the car industry
EHSA	The European Home Systems Association (EHSA) standard allows connection to a network using any collection of media and thus supports the open systems principle (Boivin and Anguill, 1996)

A historical perspective

Intelligent buildings based on computer technology have been around in one form or another for over 20 years. Perhaps the most significant developments were the introduction to building control systems of embedded processors, dedicated networks and intelligent agent approaches. This view has led us to proposing the following taxonomy for technologically based intelligent-buildings:

- First-generation Intelligent Buildings consist of numerous independent self-regulating (automatic) sub-systems. These sub-systems might be relatively sophisticated (e.g. HVAC or security systems), but they are essentially

disconnected, and operate independently of each other.

- Second-generation Intelligent Buildings are formed when building control systems, such as those described in the previous paragraph, are connected together via a network. By interconnecting them in this way, it becomes possible either to control them remotely (from a building services manager's office), or to facilitate some central scheduling or sequencing (such as securing areas, or turning systems on or off at specific times). Several specialised networks, designed for this purpose, are commercially available and fairly widely used.
- Third-generation Intelligent Buildings have, in addition to the processors and networks of the first two generations, the capability of learning about the building and its occupants, and hence adapting their control behaviour accordingly. This functionality arises from the application of intelligent agent techniques (already widely used in other areas, such as robotics).

Although first- and second-generation Intelligent Building technologies have greatly increased the ease of operation of building control, they still have not given the building any functions that are akin to human intelligence, such as reasoning, learning, or adaptation, that are present in the third generation systems. Only recently have researchers begun facing up to the challenge of giving buildings these third generation capabilities. It is in this area that the work described in this paper lies. BT's recent Telecare project is an example of a third generation intelligent-building project which seeks to "learn" the occupants' living patterns; however, it only uses this information for monitoring (Barnes *et al.*, 1998), rather than control of the building, which is the primary aim of our work.

The challenge of autonomous learning

Our goal is to develop a system that not only performs monitoring and routine tasks, but also actively provides support for occupants within a building (by performing tasks which they are incapable of doing, perhaps due to physical or mental deterioration, for example). It is clear, then, that the main focus must be on these individuals. The system must be able to learn what needs to be done, how, and when, for any particular individual, in such a way that

it does not do anything inappropriate, and does not challenge the individual's independence or authority in any way.

This problem is not deterministic or amenable to modelling, due to the dynamic, complex and unpredictable nature of people and their environment; it can not be solved by classic real-time control or automation either. Even traditional knowledge-based techniques such as Expert Systems are not appropriate, as they would still require detailed knowledge of both the building and its occupants, and would be unable to deal with any situation which had not been "foreseen" by the programmer. Thus, the challenge is to produce an Intelligent Building with a control system that can learn to adapt to individual needs and circumstances, while still being reliable, practical and affordable. In addition, the system needs to be able to overcome all the usual difficulties encountered in real-world situations, such as working with imprecise or incomplete sensor information, imperfect control, and (due to the need for low cost) small, inexpensive hardware.

A behaviour-based solution

Robotic researchers will note that these requirements are very similar to those encountered in mobile robots. It is for this reason that we decided to investigate the applicability of mobile robot control techniques to Intelligent Buildings. In particular, we are investigating the application of behaviour-based techniques, pioneered by researchers such as Brooks, Mataric, and Steels (Steels and Brooks, 1995). This is a relatively new approach to robotics, which has been found to be much more suitable than traditional AI techniques for dealing with embodied systems, in which the information is likely to be partial, unpredictable, and environment-driven, thus preventing an exact or complete solution from being predicted in advance. In this paradigm, control systems are often referred to as "agents". In this context an agent consists of a set of behaviours operating in parallel, each of which decides when to become active, based on the current internal and/or environmental conditions.

The Essex IB Model – a multi-agent distributed architecture

Buildings may be regarded as being made up of rooms of different types. In addition, control

and learning functions can be seen to be based around a room (i.e. our behaviour is often associated with the type of room that we are in, and thus so are our control needs). Most large buildings have a great deal of concurrent human activity distributed widely throughout them. (In particular, in residential and nursing homes, each occupant usually has his own individual bed/sitting room.) Our proposed solution is based on distributed processing and the fact that the physical and logical unit of an Intelligent Building is a single room. Each room contains sensors and output devices, which are monitored and controlled locally by an agent (a small embedded processor). All these agents are connected together via a network, forming a decentralised architecture that enables building-wide collaboration. This architecture is illustrated in Figure 1.

We are therefore dealing with a number of parallel distributed agents, each of which is monitoring a room (or room-like entity, such as a corridor), and responding individually to whatever is occurring there. In this way, each agent is focused on responding as well as possible to the particular needs of the person in the room, rather than finding an efficient way of satisfying the generalised needs of all the people in the building. Of course, there are still some matters that do require communication and co-operation between these distributed agents (e.g. responding to an emergency). The building-wide network allows the agents to selectively share their information when circumstances require, enabling them to make better decisions regarding situations that have a wider impact

on the occupants and building, such as the presence of an intruder, or a fire, for instance. By utilising this decentralised approach, in which most of the control is localised to a particular room, inter-agent communication is minimised, resulting in network bandwidth requirements which are only a fraction of the capacity of most existing building networks.

Inside the room-agent

In order for the room-agent to respond appropriately, it needs knowledge about the environment (i.e. the room itself and the current situation in the room) and about the person or persons currently in the room. We have already explained that it would be almost impossible to create a useful model of these in advance; therefore the room-agent must acquire its knowledge in another way – through its perceptive capabilities, i.e. the sensors in the room.

By gathering information from its sensors over a period of time, the room-agent can notice how a particular person tends to react to particular circumstances, and can then learn to “mimic” or replicate that behaviour itself. Because we are also using sensors to distinguish between different occupants (in the same way as tagging systems are used in industry), the system is able to learn different behaviours for different people. So for example, the system might learn that Person A, who is only partially sighted, prefers a higher level of light than Person B, whose sight is normal. It could then adjust the lighting level appropriately, according to who was in the room at that time (see Figure 2).

Figure 1 The ESSEX IB micro architecture

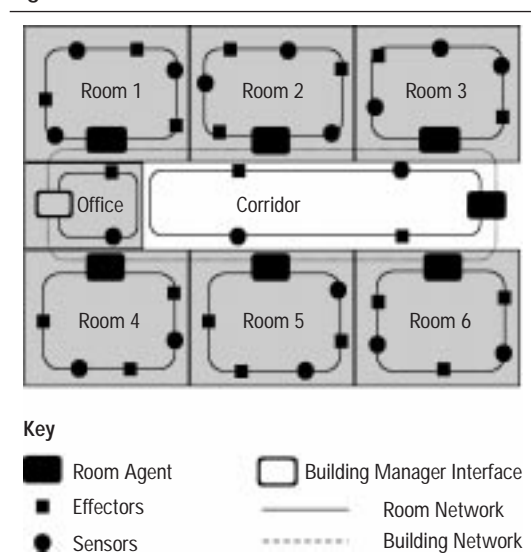



Figure 2 Prototype IB agent (Courtesy of netCam Ltd)



The use of a large number of different types of sensors enables the system to learn a much wider range of behaviours – far exceeding the basic fixed scheduling and conditional behaviours (such as WHEN darkness falls THEN switch the lights on) of second-generation Intelligent Buildings. For example, the system could learn to perform tasks related to security, energy conservation, access control, safety, and comfort, in addition to the more usual control of lighting, heating, and appliances. But the key difference is that, over time, it develops (and continues to adapt) a set of behaviours that are tailored to that particular building and its occupants, by relying on information gathered from sensors instead of from a pre-programmed model.

Of course, there are some situations (such as emergencies) which the agent must always be able to deal with correctly – it must not have to wait to learn these over time. For this reason, we have also included some permanent, fixed basic rules, that ensure the agent always behaves safely and efficiently, and is able to handle emergency situations. For example, an agent controlling a light in a room might contain the following fixed behaviours:

- a safety behaviour, responsible for ensuring the light level in the room is always at a safe level;
- an efficiency behaviour, for ensuring that electricity is not wasted (e.g. that the light is not left on when the room is empty, or when there is sufficient natural light in the room already);
- an emergency behaviour, that controls the light level in an emergency (e.g. it increases the light level if there is a fire, requiring an evacuation of the building, during the night).

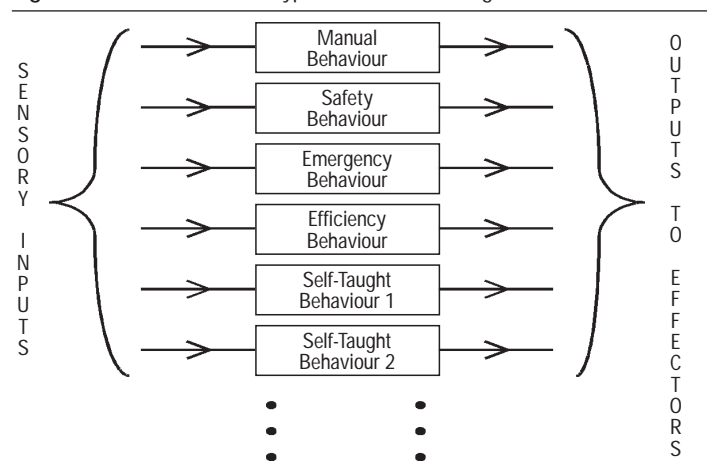
It is evident that the basic behaviour of the system must at least be equivalent to the behaviour of the building if the system were not present. That is, it must react to any orders the occupant gives it (through pressing light switches, turning heating on, etc.), correctly, controlling quantities such as heating and lighting. We refer to this fundamental behaviour as the manual behaviour. However, this behaviour alone is insufficient as a minimum fall-back, which is why we added these extra fixed behaviours. Whatever higher levels of competence are learned, the system always has these minimum levels of performance to fall back on.

The resulting architecture (shown in Figure 3) consists of some simple, fixed behaviours operating in parallel with more sophisticated learned behaviours adapted to different users and environments. The agents do not contain complex modelling or traditional reasoning capabilities (both of which are very processor-intensive). In addition they are reliant on current sensory information. This ensures that the system is able to respond rapidly, in real time, to whatever situation may arise – even if the more sophisticated learned behaviours aren't activated for some reason (either because they can't compute a result in time, or because the system simply hasn't learnt them yet). A reasonable, safe response is therefore always guaranteed, even though it might not always necessarily be the ideal, optimal output.

Current progress and further work

Currently, the agents have been evaluated on a bench-based demonstrator consisting of a “mock” building containing a representative subset of the sensors and effectors to be used in the actual system. (The “mock” building included temperature, light-level, infra-red and occupancy sensors together with domain-specific devices to support care applications. Information from these was used to control heating, lighting, appliances, security and alarms.) The initial results obtained from this test rig have been encouraging. We are currently developing the system to include temporal information, collaborative reasoning between room-agents, and behaviour characterisation, and we are refining the mechanisms needed to generate and maintain the

Figure 3 Behaviours inside a typical device control agent



dynamic behaviours. If resources permit we are also planning to investigate the development of a software front-end tool to elicit knowledge from the building manager, which would enable the fixed rules and behaviours to be tailored to specific users and buildings.

Conclusions

This work is at an early stage, but our hope for the future is that it will (a) make a contribution to the science of intelligent machines; and (b) provide a practical means of applying distributed embedded agents to intelligent-buildings.

We have chosen to evaluate our research by deploying the system in a real care environment. We are therefore collaborating with two local end-user organisations, namely Hamilton Lodge, a Residential Home caring for those with long term disabilities (mainly learning disabilities), and the Balkerne Gardens Trust which has both residential and sheltered housing provision for older people.

Through this evaluation, we hope to demonstrate that third-generation Intelligent Building technology is a viable solution to the increasingly difficult task of improving both the quality and the cost-effectiveness of care delivered to those in need.

References

- Barnes, N.M., Edwards, N.H., Rose, D.A.D. and Garner, P. (1998), "Lifestyle monitoring – technology for supported independence", *IEE Computing & Control Engineering Journal*, August, pp. 12-14.
- Boivin, J.-Y. and Anguill, Y. (1996), "Home automation for the future", *Transm. Distrib. World (USA)*, March, Vol. 48 No. 3, pp. 56-61.
- Newman, M.N. (1996), "Integrating building automation and control products using the BACnet protocol", *ASHRAE Journal*, November, pp. 36-42.
- Robathan, P. (1989), *Intelligent Buildings Guide*, Intelligent Buildings Group and IBC Technical Services Limited.
- Steels, L. and Brooks, R. (Ed.) (1995), "The artificial life route to AI; building embodied situated agents", LEA.
- Strassberg, D. (1995), "Home automation buses – protocols really hit home", *EDN*, Vol. 40 No. 8, pp. 69-80.