Pre-Publication version of chapter in in a book called 'Intelligent Buildings' published by ICE Publishing, August 2013,
Paperbound ISBN 978-0-7277-5734-0, Electronic (ebook) ISBN 978-0-7277-5792-0

Intelligent Buildings, Second edition
Edited by Derek Clements-Croome

Part 2 Chapters 5

Intelligent Environments

Victor Callaghan, University of Essex (vic@essex.ac.uk)

ABSTRACT
Although the physical architecture is without doubt the most striking aspect of a building, from a more technical prospective it is but one of many building architectures, some of which are hidden to the human eye. One such invisible architecture is the information technology that is integrated into a building to support knowledge-based activities and to automate building services. During the last half century there have been huge advances in the use of IT to manage and control building services moving from simple service automation through to today's so-called intelligent environments. In this chapter I will explore the nature of intelligent environments, explaining what they are and how they work. In particular, the chapter examines the autonomy continuum and presents an architecture that allows building occupants to vary the intelligence level of a building. Finally, I report on studies, conducted in the Essex iSpace, that have explored the attitudes of building occupants to intelligent technology and use these to discuss some of the consequences for designers of intelligent environments.

Keywords: Intelligent-Buildings, Embedded-Agents, Adjustable-Autonomy, HCI, Socio-Technical Research.

1. Intelligent environments

Intelligent Environments are everyday settings (eg buildings, vehicles, clothing etc) that are equipped with advanced networked computer based systems, whereby their coordinated activity is orchestrated by so-called intelligent agents with the aim of enabling better or new lifestyles for people. For example, such technology can lead to design of living environments that are more comfortable, usable, productive, secure, caring, social, entertaining or energy efficient. One example of such an intelligent environment is an intelligent building.

2. Facets of Intelligence

Since this chapter is discussing intelligent environments, it is important at the outset to understand what is meant by the term intelligent. Seemingly, intelligence is an intrinsic property of most life-forms and, as such, it would seem to be a term that most people would understand and be able to define. However, it turns out that intelligence can mean different things to different people, varying between differing contexts and applications. For example in the building industry, commonly, the term is used in a holistic way that seeks to capture all the phases of a building's lifespan from design, through construction to management by using methods that ensure that the building is *flexible* and *adaptable*, and therefore *fit for purpose* and *profitable*, over its full life. As Chen stated "*The Lifespan of buildings is composed of a series of interlocking processes starting from initial architectural and structural design through to actual construction, and then to maintenance and control as well as to the eventual demolition or renovation of the building*" (Chen et-al 2006). As if to emphasis the intelligence aspect, there are a variety of metrics developed that measure the "intelligence" of a building in its various phases of life, such as BREEAM (Building Research Establishment Environmental Assessment Method), LEED (Leadership in Energy and Environmental Design), CASBEE (comprehensive assessment system for building environmental efficiency), BIQ (Building Intelligence Quotient), IBI (Intelligent Building Index). Most of these metrics measure performance, which includes aspects such as health, safety, productivity, energy efficiency, environmental impact, life cycle cost and marketability. Performance benefits generally lie in economy and flexibility to meet the working and sustainability needs. In these respects an intelligent building achieves and maintains optimum *performance* by automatically responding and adapting to the operational environment (climate, occupancy, type of use, services) and user requirements (occupant, owner, developer and agent); facilitating speedy and cost-effective adaptation to changes in user requirements, e.g. space reconfiguration etc; use of the best materials, concepts and systems to meet the needs of the owner, occupants and the community.

In stark contrast to the building industry, computer scientists have an entirely different understanding of what intelligence is, considering it to be related to the human thought process. This view stems from the founding fathers of computer science (eg Von-Neumann) who created computers largely as a means to explore human intelligence, later spawning disciplines such as artificial intelligence and computational intelligence [Muhlenbein 2009]. Thus, in this view, an intelligent-building is seen as one that contains the type of governance processes that are commonly associated with needing human thought, principally *reasoning*, *planning* and *learning*. In this definition the reference to human thought is critical and can be seen as creating computational process that are akin to a person acting on another's behalf (ie an agent) to monitor, analyse, plan and learn how to control a building. In this way (assuming the person is intelligent!), the process that mimics someone is regarded as an *intelligent agent*. Thus, from a computer science perspective, intelligent agents are the basic building blocks of intelligent environments. Most agents are embedded into controllers or other appliances and so are more frequently referred to as embedded-agents. A somewhat more formal definition of an Intelligent Environment is one *where the functionality of the environment is derived from networks of computer based artefacts which <u>sense</u> user behaviour and "purposefully" <u>coordinate</u> their actions to <u>effect</u> higher level meta functionality required by the users* (Callaghan 2003). Before I leave this topic is it perhaps worth highlighting the difference between automation and intelligence. In simple terms automation can be regarded as a controller that executes pre-programmed rules continuously whereas an intelligent system is one that is *self-governing*. By self-governing I mean a system that is capable of generating its own rules (or laws) much in the same way a governments of countries do. Laws (or rules) are generally enacted by reasoning and learning, which have direct parallels in embedded agents. Later is this chapter I will revisit some of these issues and explore how simple embedded agents are designed.

3. The Changing Nature of Building Appliances

When considering the role of intelligence in buildings and other environments it is also important to understand how building appliances may evolve in the future. For example there is a credible school of research that is arguing that future appliances will no longer be monolithic in nature (eg HVACs, TVs etc) but will take more distributed or decomposed, forms. For example companies such as British Telecom have been exploring scenarios whereby buildings are equipped with a basic set of IT services such as video displays, audio transducers, media streamers, digital processors, raw sensors, effectors and interaction devices which can be dynamically inter-connected to make regular building appliances such as televisions, security systems etc). This approach goes by various names such as "*virtual appliances*" or "*soft-appliances*" and is regarded as a highly disruptive technology [Chin 2009]. Thus, for this approach, buildings may be provided with a basic set of IT services and the commissioning process would involve connecting them together to form virtual-appliances that mirror current appliances such as HVACS, building access, security or telephones etc. Moreover, as these virtual appliances are simply created by interconnecting network services, it is possible for certain functionalities (or virtual appliances) to be created by the building occupants [Chin 2010]. Thus, should this approach come to fruition, it promises to radically transform and disrupt current practice and expand the role of building intelligence, as will be evident in the following section.

4. The Intelligence Continuum

In computer science, the nature and extent of intelligence is generally fixed at design. Over the life of the agent its performance may improve as it learns how to model the task better (ie acquires more data about the task) but the agent's quota of intelligence and, more importantly, the autonomy it enjoys, remains fixed. A natural question to ask is why is an agent's autonomy fixed? Viewed another way one could ask, "*why can't an occupant of a building vary the amount and type of assistance to receive from intelligent technology*". There are various reasons that a building's occupants may want to vary the intelligence or the amount of autonomy of their building control systems. For example, depending on a person's mental or physical state (that may vary according to mood, age, health, ability etc) they may prefer more or less assistance from technology. Another argument for people wanting to be able to manage the level of assistance from technology is that people are intrinsically creative beings and too much automation can undermine this pleasurable aspect of life. For example, some people take pleasure in designing the interior finish and furnishings of their homes and, even if this could be automated, would prefer to do the design themselves. Configuring virtual appliances can also be seen as a creative process, (akin to home DIY) that some people may enjoy and not others. Therefore, in such circumstances, having some say over the amount of autonomy in a building can be advantageous. Another reason people may want to vary the extent of intelligence is that technology cannot accurately predict a person's intentions, and when it makes mistakes it can be annoying. The amount of mistakes a predictive agent makes is linked to both the type of task (how repetitive is it), the person's persona (how spontaneous they are) and the nature of the sensing. Given that sensors are unlikely to be able to accurately read people's minds and that an intrinsic aspect of being human is "free will" (non-determinism) it is likely the predictive performance of agents will continue to vary according to the situation, resulting in advantages to the user in being able to reduce or increase the use of intelligence to match the circumstances. Finally, various surveys have concluded that an obstacle to deploying more intelligent solutions is that people are fearful of too much intelligence and have a strong desire to remain in control [Ball 2011]. Being able to affect the level of autonomy is one-way people can exercise control. Thus, at times people may wish to exercise more explicit control over the design of a building's functionality and intelligence or autonomy should not be presented as a binary option (ie either it's there or not).

Looking at this from the viewpoint of a simple building management system, if a person is doing simple repetitive tasks, that an agent can quickly model, then its pre-emptive actions can work well (eg managing heating for fuel efficiency) but if the occupant's behaviour is irregular or complex (eg a one-off, or spontaneous tasks) then the occupant may find the agent's actions in conflict with their needs and benefit from reducing the agents control of the environment. In all these cases the ability to vary intelligence can be desirable. To vary the amount of assistance received from an intelligent environment could be seen as moving along a continuum which at one extreme the user does all the work, where at the other extreme the intelligent agent does all the work; see figure 1.

Figure 1: The Intelligence Continuum

The Essex work has investigated how to create such an adjustable intelligence (or autonomy) control, that is much like a thermostat control in a building that can be adjusted from no-autonomy (fully manual, no agent pre-emption) to full-autonomy (max agent pre-emption). It does this for each building sub-system thereby enabling the building occupant to exercise fine grained control over what aspects of the buildings functionality are automated in a way that is akin to a music mixing desk. I will elaborate on this later in the chapter and describe some user findings, as part of the case study.

5. A simple embedded-agent architecture

As was explained in section 2 of this chapter, embedded-agents (intelligent-agents) are the basic building blocks of intelligent environments. Thus, an important question is how do they work, and how can they be designed? The challenges facing the design of an embedded agent for intelligent buildings are significant. First, there is the choice of a centralised or decentralised computational approach. Historically, systems have been centralised, as they are logically simpler to design. However centralised systems suffer from well-known shortcomings such as single point failures which can bring down an entire building. Also, centralised architectures are not readily scalable, as the processor is of a fixed size and routing back connections for sensors and actuators is more difficult. On the other hand distributed architectures are more difficult to design but are more scalable and malleable to the building structure. Fortunately, in the behaviour architecture I will describe here, creating a distributed architecture is relatively simple.
Second, considering the agents themselves, they are essentially real-time controllers receiving vast amounts of sensor data that is noisy and relatively sparse. Also, attached to these agents are effectors which can be electro-mechanical in nature and thus, prone to malfunction. In addition, the embedded processors are relatively computationally small, compared to a centralised system, and so there is an additional challenge of developing computational intelligence schemes that fit these small resources. Artificial intelligence has a particular computational problem, as it normally uses at least one layer of abstraction (a model of the world) that has proven both difficult to keep current and is computationally demanding.
A third issue is that environments are rarely occupied by just one person meaning that any controlling agents, in addition to coping with individuals, need to model and manage multiple-occupancy.

Thus, when considering all these issues, the challenge to create an architecture that can work within all these constraints is considerable. Fortunately, despite the apparent overwhelming complexity of creating a disturbed intelligent agent architecture that is capable of operating in such a challenging environment, it turns out there is a remarkably simple solution. The field of mobile robots had faced very similar challenges and had developed a number of solutions, the most relevant being behaviour based architectures popularised by Rodney Brookes at MIT [Brooks 1986 & 1991]. The principle is simple in that, instead of operating on a data abstraction layer, as most artificial intelligence had previously done, Brookes dispensed with the usual abstracted model preferring to operate directly on the real world coining the phrase, "t*he world is its own best model*". In addition, he proposed a horizontally partitioned architecture where the overall task was decomposed into a collection of sub-tasks (called behaviours) each providing some independent sensor-to-effector control. By arranging these behaviours in this way, the architecture maintained a number of concurrent processes (behaviours) that were competing for control of the system. The interaction between these behaviours (what permutation of behaviours was in control) provides another level of adaptability and gave rise to a property called emergent behaviour. In short, this was equivalent to reasoning and planning in more traditional AI, as it can be shown that it solves the same problems. Figure 2 shows a hierarchical fuzzy-logic implementation of the Brookes' Behaviour Based Architecture.

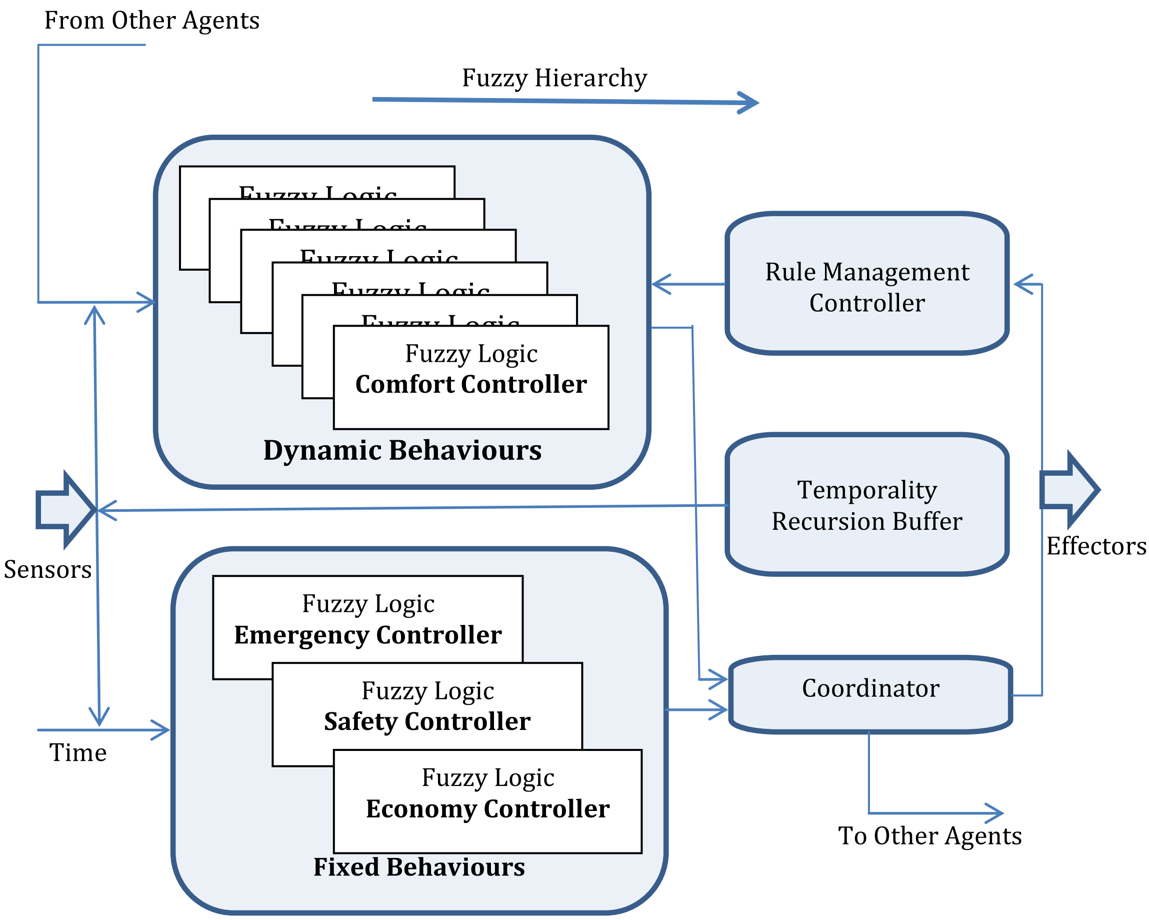

Figure 2: A Fuzzy Logic Implementation of a Behaviour Based Agent

In this agent the behaviours are divided into fixed and dynamic behaviours. The fixed behaviours reflect "must happen" conditions permanently set by the building's stakeholders, which might include government 'health and safety' rules or security needs. The dynamic rules are ones that are learnt on-the-fly by monitoring the building's occupant's habitual behaviour and creating a rule set that matches the user's needs. The agent works by switching between behaviours using an arbiter (coordinator); this decides what behaviour is active based on the current context (the sensed environment). Thus, the coordinator acts as a type of value-added sensor, codifying the context of all its sensed inputs into a single word describing the state of its observed world. It is this fact that leads to a mechanism for creating a simplified distributed agent coordination mechanism, in which other agents are just regarded as pseudo-sensors (but of added-value because of the additional processing) that by being connected to other agents inputs provide a "semantic free" distributed coordinated mechanism referred to as an *Agent Semiotic Language* (ASL). The use of agent semiotic schemes to simply distributed agent processing is a significant boon to the use of a behaviour-based approaches for intelligent environment management. In this model, temporality (basing future decisions on past experience) is catered for by the classic state machine principle of including a feedback loop that links current decisions (time n) to past decisions (time n-1) and by the principle of recursion to all previous states [Lewin 1987]. Concerning the issue of multiple occupancy, there are various ways to address this issue and the approach adopted in this agent design is the so called "*corporate identity method*" in which groups of people can be viewed as a single persona that an agent models and manages. As the name suggests the inspiration for this approach came from the nature of companies that, while being composed of multiple people, are regarded as equivalent to a single person in law. This methodology is discussed in more detail in other papers by the author [Callaghan 2000, 2002 & 2004]. Finally, a problem with such ad-hoc interconnection schemes is that they are prone to cyclic instabilities (eg live-locks) producing symptoms such as unwanted flashing lights. Put simply, the cause of this erroneous behaviour can be traced to closed loops in which the action of a given agent is based on the action of another that is in-turn dependent on the former (ie the interdependence of agents in a multi-agent system). In practice these interactive loops occur across numerous agents and complex routes (both spatially and temporally) which conspire to mask and complicate their identification and eradication. In fact, although the symptoms had been observed and reported in various intelligent environment projects, the relatively embryonic state of this field meant that until recently little was known about their cause or cure. It has been found that this erroneous behaviour can be eradicated by breaking the loops but this must be done with due regard to the disabling effect on the overall system's functionality. A more detailed discussion is presented in other papers by the author [Zamudio 2009].

The fuzzy logic aspects of the agents operation are highly mathematical, and readers interested in that theory are referred to earlier papers by the author that provides a detailed mathematical explanation of these agents [Callaghan 2002 & 2004]. Regarding the behaviour-based architecture, in the world of robotics, this process is easily understood via its correspondence to the physical world. Thus, for example, when a moving robot is endowed with a few simple switched behaviours (either on/off) such as (1) obstacle avoidance and (2) goal-seeking, plus some priority scheme, the robot could be understood to function as follows; if the robot is not near an obstacle and not at its goal, that behaviour will be off and the goal seeking behaviour will be on and it will head towards its goal; however, should it encounter an obstacle, that behaviour will become dominant, switching on and muting the goal seeking behaviour. Through this interplay of behaviours it's possible to argue that a mobile robot can solve difficult problems such as navigating to a goal through a field of obstacles. Of course this is just a simplified explanation to convey the principle of behaviour-based architectures and much more sophisticated arrangements exist. In this architecture each of the behaviours is a simple rule-based process such as "*if obstacle to front, reverse*" etc. The extraordinary aspect of behaviour-based architectures is that whilst they are composed of collections of extremely simple rule-based interacting processes, they solve very difficult problems that hitherto required large and sophisticated artificial intelligence (AI). The lightness of the rules means it can run in real-time on a small processor. The advent of behaviour-based architectures broke the impasse that had existed in the ability of artificial intelligence to control robots, and was a significant breakthrough.

The advance with regard to intelligent buildings was the observation that robots and buildings are logically identical which is perhaps best captured by the phrase coined by the author that *"a building is a robot we live inside"* [Callaghan 1999 & 2000]. This is a parody on the rather more famous quote by the well-known 20th century Swiss architect of the modern movement, 'Le Corbusier', who is reported to have stated *"the house is a machine for living in"* [Le Corbusier 1923]. In the case of the authors inspiration, it came from the fact he founded the Essex University mobile robots laboratory (the Brooker Lab) and was working on a robot which had the appearance of a square box containing various sensors and actuators managed by an intervening processor; the similarity to the room (that the robot was in) was striking as it was another square box that contained various sensors and actuators managed by an intervening processors; thus this thought gave rise to the notion robots and buildings have important similarities. In addition to this empirical connection, a more reasoned explanation is that buildings actually move through a 'data space' in much the same way as a robot moves through a physical space.

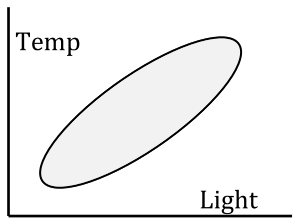

Figure 3: A simplified sMap

Figure 3 depicts this using an abstraction called an sMap (sensory map). The example sMap shows the correlation between 'temperature' and 'light' in a building and depicts the situation that as the sun rises, the light level and temperature rise, but as sun sets, then the light level falls and temperature falls, but lags due to thermal absorption. Clearly other entities such as windows, doors or HVACs can 'enter' this space and effectively move around it by changing the temperature-light balance (in the same ways as a mobile robot moves around its space). Thus, in this abstracted space, similar objects exist such as an open window causing a dynamic "temperature object" to move and obstruct the movement of the climate control system (akin to a mobile robot). Some of the thermal characteristics of the building represent more fixed objects. In this way the same theory is applicable to problem-solving in both the mobile robot and building domains. The final piece of the jigsaw connecting robotic control principles to buildings was identifying the fixed behaviours for a behaviour-based intelligent building architecture which were identified as being *goal-seeking*, *safety*, *manual control* and *comfort* [Sharples 1999]. Other combinations of behaviour are possible; the criteria used to generate these was to first identify the "*must be guaranteed behaviours*" and assign these to be fixed behaviours and then leave the less critical ones to be either leant or emerge from the behaviour interplay. Learning is introduced through the use of dynamic behaviours in which the agent learns new rules though a mechanism dubbed as "evidential learning" [Sharples 1999]. Briefly, this works by recording the sensor values of all devices whenever a significant event occurs (eg a user changing a building control setting); in this way rules are assembled. Whilst the basic principle is remarkably simple, to make the agent more robust the architecture is usually refined in various ways such as by the addition of fuzzy logic, learning inertia (a mechanism to ensure only significant rules are learnt) and rule a management system (see figure 3). Learning inertia can be more formally represented as follows:

$$w_1 \cdot li(x_1)\frac{\Delta x_1}{\Delta t_1} + w_2 \cdot li(x_2)\frac{\Delta x_2}{\Delta t_2} + \cdots\cdots\cdots w_i \cdot li(x_i)\frac{\Delta x_i}{\Delta t_i}$$

In this the number of occurrences of a behaviour pattern or cycle occurring during a learning phase is modelled by a simple differential function li(x) where 'x' is the complete set of observed occurrences during the learning phase, 'i' is the number of learning phases with different durations, Δt_i corresponds to an individual learning phase duration, Δx_i is the number of occurrences observed within the respective learning phase duration, and x_i is the set of observed occurrences within the respective learning phase duration. A constant 'w_i' weights how much the individual differentials (representing different learning phase durations) affect the learning process. By setting Δt_i to a fixed value each differential term can capture behaviours with differing cyclic periodicies (eg hourly, daily, weekly, annually etc). Alternatively, by setting Δx_i to a fixed value each differential term can capture behaviours with differing occurrences (eg 1, 2, 3 etc). In the example presented here, the learning system adopts the latter approach using an occurrence parameter (learning inertia) of 3 as a minimum figure to trigger learning. Clearly this equation may be applied in various ways to design an agent to manage differing kinds of human and environment behaviour [Ball 2010]. Finally, most current agents sense the parameters associated with the physical environment, using changes in these to trigger adaptation of the environment. Given that the environment is being adapted to mirror the changing needs of occupants, there is interest in sensing more personal properties of the occupants such as their mood or physical state as a means to improve the accuracy of agent managed control of the environment. Other researchers are interested in collecting such personal data as a way of evaluating intelligent environment concepts (eg automatically collecting emotional response to experimental aspects of intelligent environment design). This is a complex field that would merit a chapter in its own right and interested readers are referred to a more detailed exposure on this by the author in another publication [Leon 2010]. Likewise, more detailed discussion on the embedded-agent operation is given in other papers by the author [Callaghan 2002 & 2004].

6. End User Programming

The behaviour based agent described in the previous section works effectively, taking approximately two days to learn about 200 rules that seem to characterise a typical user in a single room [Callaghan 2004]. However the architecture, as presented, is effectively a fixed full-autonomy system. Other work has explored the opposite extreme, end user programming, where the building's occupant is in complete control, programming all the functionality of the building. The issue with involving users in programming a building's system is twofold. First users are, in general, not technologist and are usually unable to make use of the type of programming tools that scientist and engineers use. Second, there are aspects of a building system that need to be kept away from the general occupant for either safety or other stakeholder needs (eg the occupant may not be the owner, or the duly authorised manager as buildings can have multiple stakeholders). To solve the first challenge the general solution is to provide the occupant with a set of familiar graphical or physical objects that have a metaphorical relationship to the programming possibilities of the building or intelligent environment. For example, one popular metaphor is a jigsaw puzzle, where the building occupant is presented with a collection of pieces they can recombine into a number of differing pictures, each picture being a particular programmed building functionality. Usually these pieces are directly analogous to programming constructs. Thus, the jigsaw pieces are a bridge between the real building and the underlying computer system [Humble 2003]. At Essex we have adopted a *programming-by-example* mechanism. In this approach the building occupant puts the system into a learning mode and then simply demonstrates the desired behaviour to the system. So, for example, if the owner of a home-cinema room wanted to program the room to "*on the receipt of an incoming telephone call, pause the movie, raise the lights and divert the call to the AV system*", they would simply demonstrate this to the system by first putting it in a learning mode then, use their mobile-phone to call the house, on hearing the ring they would then manually raise the lights and switch the call to the AV system. The learning mode would be terminated and the result stored as a portable ontology based description called a MAp (Meta-appliance/application). The system would then remember this, so if the same context reoccurred, than that action would be replayed. The system goes beyond a simple macro arrangement, as the information is not sequence dependent and encodes the task using ontology to make portable soft objects that can be carried between environments by users (or even traded) [Chin 2010]

6. Adjustable Autonomy Agents

Having described how agents at either end of the autonomy scale can operate the question is, how can these approaches be combined to produce an adjustable autonomy system that allows a user to find a "sweet spot" between the extremes of "being controlled by the system" and "being in control of the system"*?* To achieve this there are a number of options that range from varying how many agents are active in a system, through schemes for switching agents between discrete autonomy levels to creating controllable learning mechanisms. At Essex we have explored the latter two approaches. The first approach involves a switched discrete system that allows the user to select one of four discrete autonomy states, namely:

1. **Full autonomy**: the agent learns from the user's behaviour, automatically creates/maintains rules as the agent deems it necessary.
2. **High autonomy**: the agent learns rules from the user's behaviour which can only become active when confirmed by the user (agent teamwork).
3. **Low autonomy**: the user creates/maintains rules assisted by the agent presenting suggestions (agent teamwork)
4. **No autonomy**: the user creates/maintains rules with no assistance from the agent.

The second approach is based on managing the learning mechanism of a behaviour-based agent. The general principle is there are two sets of behaviours (rules sets) one active and the other potentially active (see Figure 4). Each rule (or rule set) has a 'usefulness' parameter; a numeric quantification of how frequent & accurate a rule has proved.

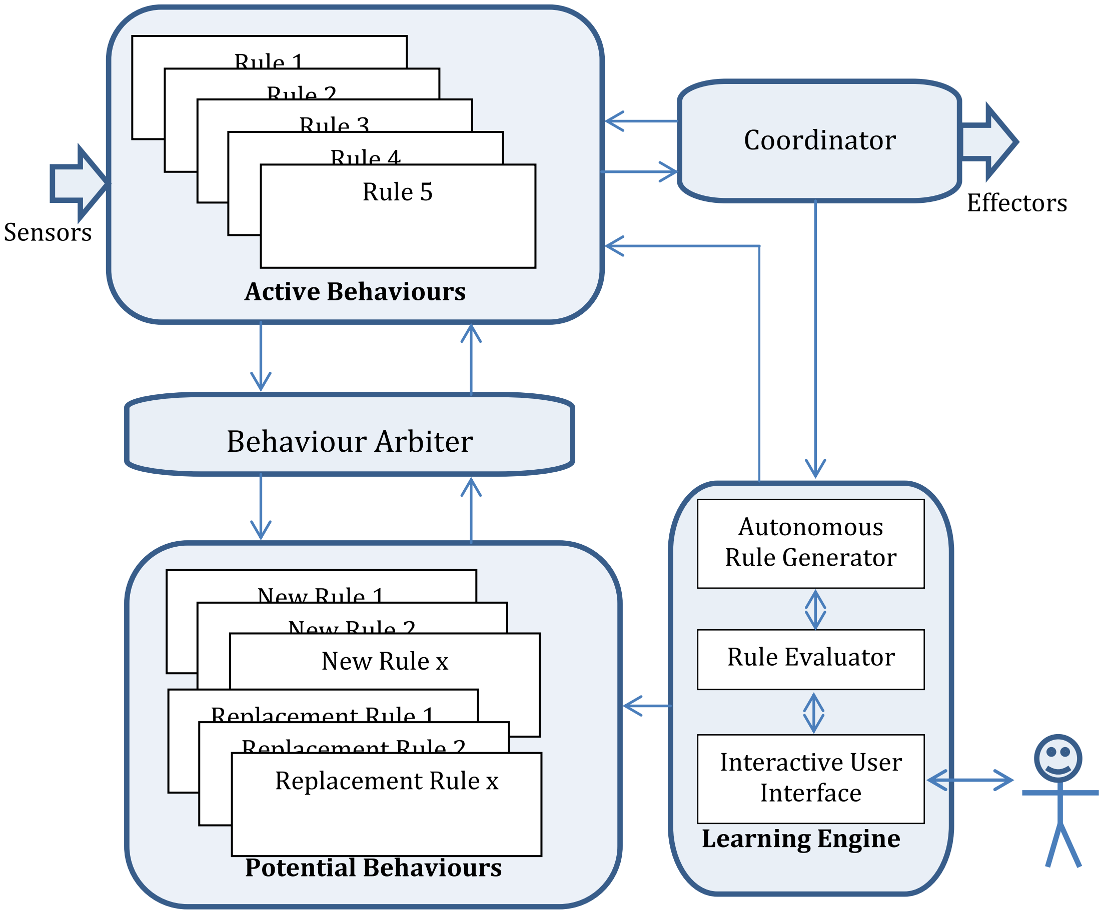

Figure 4 Adjustable-autonomy Behaviour-Based Agent

Adjustable autonomy is achieved through comparing this to an adjustable threshold that determines whether a particular rule can be active or not. The learning equation presented earlier in section 5 encapsulates the generality of the operational functionality available to this architecture. Thus for example, a simplistic approach might be to set Δt_i to 24 hours, w to 1 (and w in other terms to 0) and use Δx_i as an adjustable autonomy variable. Clearly the choice of parameters is a key issue in the design of adjustable autonomy systems and readers wishing to have a more detailed insight to the issues are referred to [Ball 2010].

7. Intelligent Environments and People

One of the incidental benefits of using an adjustable autonomy agent is that in addition to providing an end user with a more effective means to control their intelligent environment, it can also be used to assess user attitudes towards intelligence and autonomy. By giving a number of users the ability to vary the level of autonomy for each function in their environment, statistics may be gathered on people's attitude to the use of autonomy in intelligent environments. Understanding users concerns relating to intelligent environment technology is important for companies wishing to overcome market barriers and for users to get systems they actually need. Various studies have been conducted on users attitudes towards building based technology, the main ones being the University of California's study of *attitudes to smart home technologies* [Venkatesh 2001], a study by the Samsung Corp and American Institute for Research on *smart home requirements in USA & South Korea* [Chung 2003], the University of Copenhagen on context-awareness [Barkhuus & Dey 2003], the Fraunhofer Institute, Philips Research and France Telecom *study of cross cultural expectations of smart homes (in multiple European countries)* [Röcker 2004], Tampere University Hypermedia laboratory *study of expectations of digital homes* [Mäyrä 2006], Goteborg University's *study of attitudes to smart homes* [Montano 2006], Carnegie-Mellon University *investigation of the type of control of digital homes* [Davidoff 2006], the University of Munich's *evaluation of interaction with technology in digital homes* [Rukzio 2006], the University of Essex *study of user control issues in smart home* [Chin 2008], and investigation of *perceptions of autonomy* [Ball 2011].

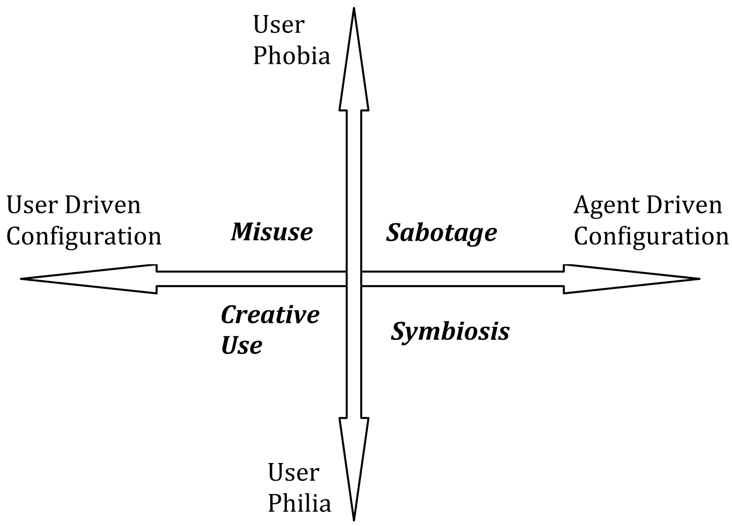

Figure 5: The 3C User Attitudes to Intelligent Environments model

A common finding of these studies was that users felt maintaining control was of paramount concern. Additional issues included adaptability, personalization, privacy and trust of intelligent environment technology. The net result of an aggregation of these concerns was that, in the extreme, users were either attracted to or repelled from these systems. These attitudes are depicted in figure 5, which summarises the main reactions to different levels of agent autonomy [Callaghan 2009].

Each quadrant represents an extreme type of usage that may be encountered as a system moves between being exclusively autonomous or end-user driven. Users have different views of technology and this diagram enables feelings of phobia (fear) or philia (love) to be depicted. Ideally the design of an intelligent environment technology should aim to avoid misuse and sabotage of the system and engender creative use or symbiosis between the user and system. In the next section I will present a case study that provides a more detailed and evidence-based insights to users' views of intelligent technology in everyday living environments.

8. A Case Study; The Essex iSpace

In this case study I describe the Essex iSpace, a purpose built experimental intelligent environment in the form of a two bed-roomed apartment, see figure 6.

Figure 6: The Essex iSpace

This chapter also reports on the use of the Essex discrete model of adjustable agent autonomy which provided users with 4 switchable settings; '*full-autonomy*' in which the agent learnt from the user's behaviour and automatically created and activated rules as the agent deemed necessary; "*high-autonomy*" where the agent learnt rules from the user's behaviour which only become active when confirmed by the user (a strong form of agent teamwork); "*low-autonomy"* where the user created and activated rules assisted by the agent presenting suggestions (a weaker form agent teamwork)*; "no-autonomy"* where the user created and activated rules with no assistance from the agent. The adjustable agent was built and deployed in the University of Essex iSpace. The aim of the study was to gain an understanding of people's opinions relating to the use of autonomous-agents in intelligent environments. Twenty participants completed three short tasks using the adjustable autonomy agent and were asked provide feedback on their experiences. The participants interacted with the system using an Apple iPad, which provided a Rule Creator, Rule Viewer, Autonomy Settings screen and Room Control, see figure 7. The first task involved creating a simple scheme for managing the opening/closure of the curtains, the light levels and the air conditioning and AV (eg TV) settings. The second task extended the complexity by introducing correlated state conditions, including time. So, for example, a person might set a condition that the lights would not be turned on during daylight. The third task added additional dynamics to the conditional mix for example, creating settings relating to the users activity or location. After an initial instruction phase, where participants were given a chance to try out the system, the participants completed these three tasks, followed by a debriefing interview. The interview took a semi-structured approach [O'Leary 2004], covering issues such as, what autonomy levels they chose and why, exploring how their choices/preferences changed over the course of the task, if any autonomy levels raised concerns and understanding whether their choices were time or function dependent? The participants were equally split between male and female, aged between 20 and 45 and mostly with university level backgrounds. On average, 40% use computers over 40 hours per week, 40%

used computers between 20 and 40 hours per week and 20% of the participants used computers between 2 and 20 hours per week. 80% of the participants had no experience of computer programming and 75% of the participants had not heard of a smart homes or intelligent environments.

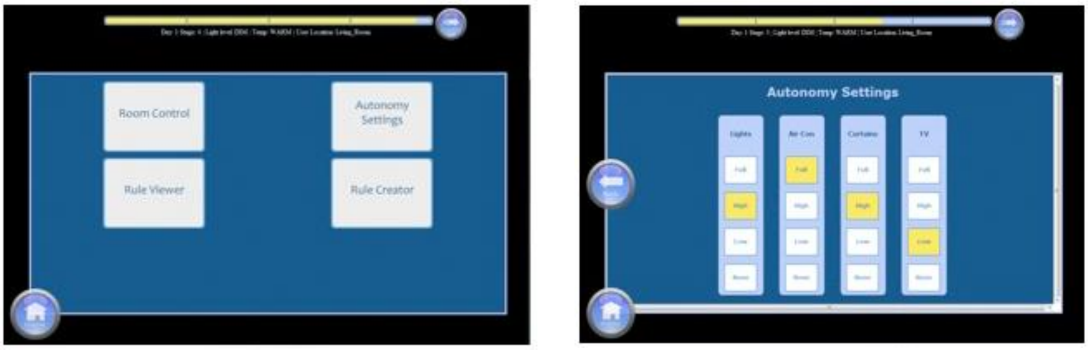

Figure 7: Autonomy setup screens on Apple iPad

The results produced some interesting findings which are summarised in. Figure 8. Superficially the results are intuitive in that, the more "*personal*" a function was, the more the participants needed direct control over it whereas the more "*shared*" a function was, the less control they required. Thus, for example, participants wanted explicit control of their entertainment system but were happy to delegate HVAC control to an agent. However discussions with the participants exposed that peoples reasoning can be more complex with some of the participants displaying a risk versus benefits calculation of their decisions to use any particular function.

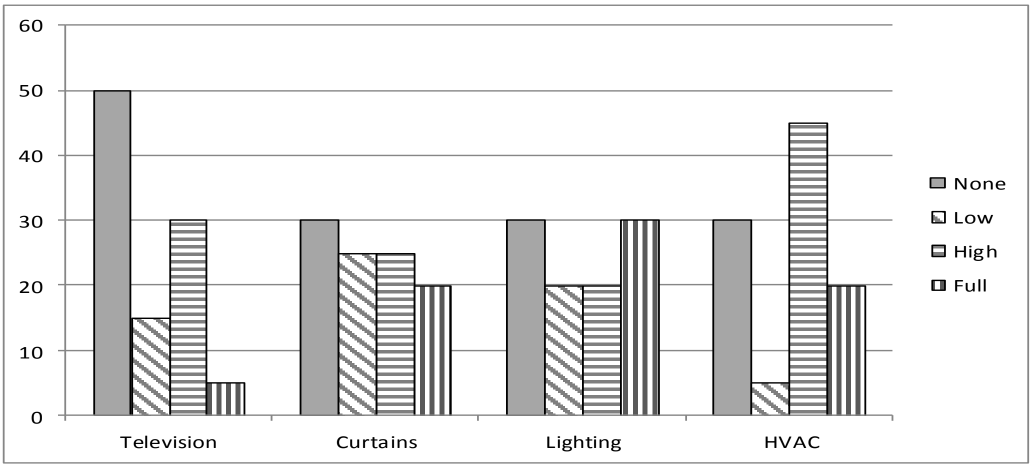

Figure 8: Participants autonomy preferences for different functions

As explained earlier in the chapter, the technology is not perfect and is error prone. The cost of errors varies from being just a mild irritation, in the case of the temperature being slightly wrong, to severely annoying where the agent made a wrong choice of music. It is clear that each of these domains represents a different level of difficulty to the agent making correct predictions relating to the users needs, so whilst it can deal with climate control issues relatively accurately, its ability to deal with the human psycho for media taste issues is beyond its capacity and therefore the performance of an agent-controlled climate change system will inevitably be much better than that of an agent-controlled AV control system, which in-turn influences people's opinions. That said, the findings represent the current state-of-the-art in these technologies and therefore, have implications for people designing today's intelligent environments and buildings. In terms of understanding the broader question of peoples overall views on autonomy in intelligent environments, the study revealed that people prefer to be in control, rather than to be controlled, which is consistent with all the surveys reported earlier in this chapter. From the Essex work, and the various studies conducted by other organisations, it is clear that an intelligent building designer need to be very thoughtful about where and how autonomy is included in buildings, if this technology is to be successful and of genuine use to people. Clearly this is a complex topic and such a short section cannot adequately discuss the issues; thus, interested readers are referred to our other papers that describe the architecture and evaluations is much greater detail [Ball 2012].

9. Summary

In this chapter I described what an intelligent environment is, and likened it to being *"a robot people live inside*". This analogy has allowed agent design techniques taken from the field of mobile robots, notably behaviour-based design, to be applied to the design of intelligent buildings and other ambient intelligent systems. I described how behaviour-based design allows intelligent embedded agents to be built using very simple principles involving interacting sets of rule-based processes in which the reasoning and planning arises both from explicit execution of the rules, and also the interaction between the rule sets; so-called emergent behaviour. In addition to discussing how such simple agents (just a few lines of code per behaviour) can solve complex problems I described how the scheme also simplifies the task of coordination in distributed agent architectures which is important, as distributed architectures bring scalability and reliability advantages to the implementation of intelligent environments and buildings. Likewise I explained how agents can deal with multi-occupancy using the "*corporate identity method*". I pointed out a frequently observed, but poorly understood phenomenon that causes erroneous behaviour in systems of distributed intelligent agents, namely cyclic instability and explained how designers of intelligent buildings can overcome this problem. This chapter also raised the issue of the need to consider social-technical issues as part of the design of an intelligent environment and I presented a model, the 3C socio-technical framework that lubricates discussion about these interdisciplinary topics. In connection with this model I described research on technologies at the two extremes of what is termed the Intelligence Continuum; end-user programming and autonomous agents. I then introduced a test-bed for intelligent environments research, the Essex iSpace. As part of describing a case study, I presented the concept of adjustable autonomy and described an implementation that allowed users to explore other points on the intelligence continuum. I briefly described some related user evaluations, which revealed that building occupants have complex views on the use of agent technology in intelligent environment and buildings. The findings confirmed earlier surveys that suggested users like to be in control, rather than to be controlled which is an issue that should be borne in mind when designing intelligent buildings and environments. Finally, although adjustable autonomy was introduced as a means to explore user concerns, it is clear that it forms an interesting option for designing future agents and intelligent environments and, given user attitudes to such technology varies so much, it allows each individual to select the balance that suits them best. It is my hope in writing this chapter that readers will be motivated to create technology for intelligent environments that gives people more rather than less control over their environments.

Acknowledgement: The roots of this chapter stretch back to 1997 and a chance conversation at a postgraduate welcoming party with Dr Graham Clarke which led to a proposition to combine concepts from building architecture and artificial intelligence into a new research area that became the Essex University Intelligent Environments Group. Since then the group has attracted talented staff, bright students and significant research funding that has produced some truly fascinating and useful research. Much of the value of this chapter is down to the valuable and inspiring work of this group; my thanks to all of you.

References

[Barkhuus 03] Barkhuus, L, Vallgårda, A, "*Smart Home in Your Pocket*", Adjunct Proceedings of UbiComp 2003 (2003) 165-166

[Ball 11] Ball M, Callaghan V, "*Perceptions of Autonomy: A Survey of Users' Opinions Towards Autonomy in Intelligent Environments",* Intelligent Environments 2011 (IE'11), Nottingham 27-29th July 2011

[Ball 10] Ball M, Callaghan V, Gardner M, "*An Adjustable Autonomy Agent for Intelligent Environments*", 6th International Conference on Intelligent Environments (IE'10), Kuala Lumpur, Malaysia. 19-21st of July 2010.

[Brooks 91] Brooks R, "*Intelligence Without Reason*", Computers and Thought, IJCAI-91, April 1991

[Brooks 91a] Brooks R, "*Intelligence Without Representation*", Artificial Intelligence, 47, 1991, 139-160.

[Brooks 86] Brooks R, "*A Robust Layered Control System for a Mobile Robot*", IEEE Journal of Robotics and Automation, RA-2, April, 1986, 14-23.

[Callaghan 09] Callaghan V, Clarke G, Chin J, "*Some Socio-Technical Aspects Of Intelligent Buildings and Pervasive Computing Research*", Intelligent Buildings International Journal, Earthscan, Vol 1 No 1, Jan 2009, ISSN 1750-8975.

[Callaghan 08] Callaghan V, Clarke G, Chin J, "*Some Socio-Technical Aspects Of Intelligent Buildings and Pervasive Computing Research*", Intelligent Buildings International Journal, Vol 1 No 1, 2008

[Callaghan 04] Callaghan V, Clark G, Colley M, Hagras H Chin J, Doctor F, "*Intelligent Inhabited Environments*", BT Technology Journal , Vol.22, No.3 . Klywer Academic Publishers, Dordrecht, Netherlands, July 2004

[Callaghan 02] Callaghan V, Colley M, Clarke G, Hagras H, "*A Soft-Computing based Distributed Artificial Intelligence Architecture for Intelligent Buildings*", in "Studies in Fuzziness & Soft Computing", Springer-Verlag, Vol 75, pp. 117-145, 2002

[Callaghan 00] Callaghan V, Clarke G, Pounds-Cornish A, Sharples S, "*Buildings as Intelligent Autonomous Systems: A Model for Integrating Personal and Building Agent*s", 6th International Conference on Intelligent Autonomous Systems (IAS'06), Venice, Italy, July 25-27 2000

[Callaghan 99] Callaghan V, Clarke G, Sharples S, '*Buildings as Intelligent Machines*', IEE Colloquium (Control Division) on Intelligent Buildings, Savoy Place, London, 17th November 1999

[Chen 06] Chen Z, Clements-Croome D, Hong J, Li H, Xu Q "*A Review of Quantitative*

Approaches to Intelligent Building Assessment" Proceedings 6th Conference for Enhanced Building Operations, Shenzhen, China, November 6 - 9, 2006

{Chin 10] Chin J, Callaghan V, Clarke G, "*End-user Customisation of Intelligent Environments*" Handbook of Ambient Intelligence and Smart Environments, Springer, 2010, Spring, pp 371-407, ISBN 978-0-387-93807-3

[Chin 09] Chin J, Callaghan V, Clarke G, "*Soft-appliances: A vision for user created networked appliances in digital homes*", Journal of Ambient Intelligence & Smart Environments, IOS Press, Netherlands, 2009

[Chin 08] Chin J, Callaghan V, Clarke G, "*A Programming-By-Example Approach To Customising Digital Homes*", Intelligent Environments 2008, Seattle, 21-22 July 2008.

[Chung 03] Chung KH, Oh KS, Lee CH, Park JH, Kim S, Kim SH, Loring B, Hass C, "*A User-Centric Approach to Designing Home Network Devices*", CHI '03, 2003, pp648 – 649

[Davidoff 06] Davidoff S, Lee MK, Yiu C, Zimmerman J, Dey A, "*Principles of Smart Home Control*", Proc of Ubicomp 2006, pp19-34, Springer-Verlag, 2006

[Humble 03] Humble J, "*Playing with the Bits, User-Configuration of Ubiquitous Domestic Environments*", UbiComp 2003, Springer-Verlag, (2003), pp 256-263

[Le Corbusier 23] Le Corbusier "*Towards A New Architecture*" (Vers Une Architecture), The Architectural Press, London, 1923

[Leon 10] Leon E, Clarke G, Callaghan V, Doctor F "*Affect-aware behaviour modelling and control inside an intelligent environment*", Pervasive & Mobile Computing, Elsevier B.V, Vol 6, Issue 4, 2010.

[O'Leary 04] O'Leary, Z., "*The Essential Guide to Doing Research*", Sage Publications Ltd, pp.164 (2004)

[Lewin 87] Lewin D, "*Design of Logic Systems*" Van Nostrand Reinhold, Wokingham, 1987, ISBN 10: 0442306067

[Mäyrä 06] Mäyrä F, Soronen A, Vanhala J, Mikkonen J, Zakrzewski M, Koskinen I, Kuusela K, "*Probing a Proactive Home: Challenges in Researching and Designing Everyday Smart Environments*", Human Technology Journal, Volume 2 (2), October 2006, 158-186

[Muhlenbein 09] Muhlenbein H "*Computational Intelligence: The Legacy of Alan Turing and John von Neumann*" Chapter 1, in Computational Intelligence; Collaboration, Fusion & Emergence, Springer-Verlag 2009, ISBN 978-3-642-01798-8

[Montano 06] Montano C, Lundmark M, Mähr W, "*Control vs Convenience: Critical Factors of Smart Homes*", 2nd Scandinavian Student Interaction Design Research Conference, Göteborg, Sweden, 25-26 February 2006.

[Röcker 04] Röcker C, Janse M, Portolan N, Streitz N, "*User Requirements for Intelligent Home Environments: A Scenario-Driven Approach and Empirical Cross-Cultural Study*", Joint sOc-EUSAI conference, 2004, 111-116

[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, pp.135-140, May 1999

[Venkatesh 01] Venkatesh A, "*The Home of the Future: An Ethnographic Study of New Information Technologies in the Home*", (Project NOAH), Advances in Consumer Research XXVIII, Association for Consumer Research, 2001, 88-96.

[Zamudio 09] Victor Zamudio and Victor Callaghan , "*Understanding and Avoiding Interaction Based Instability in Pervasive Computing Environments*", Journal of Pervasive Computing and Communications, Vol. 5 Issue: 2, pp.163 – 186, 2009