James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

The Tailored Fabric of Intelligent Environments

James Dooley, Hani Hagras, Vic Callaghan, and Martin Henson

Intelligent Environments Group,
The University of Essex,
COLCHESTER,
Essex CO4 3SQ,
UK

Abstract The traditional Internet of Things (IoT) vision states that passive, everyday objects are uniquely identified through some computer-readable means such as barcodes or RFID so that electronic systems can identify them. The identity is then used to retrieve a virtual representation for the object - a source of information that forms the basis for context awareness, decision making or action invocatoin. It was envisioned that every object in the world could be tagged and that the Internet could provide the network across which these *"things"* could be active (resolved, interacted, etc.). In this chapter, we describe how this vision converges with the vision for Intelligent Environments (IEs) as Ubiquitous Computing deployments that are endowed with an Ambient Intelligence. In particular we see the marriage of passive-objects from IoT and active-objects from IE as symbiotic if real-world deployment can ever be achieved - it is from these objects that the fabric of IEs will be woven.

1 Introduction

As the vision for an Internet of Things (IoT) becomes closer to reality, the number of objects that are deployed in the real world with a digital presence increases towards a massive scale. Familiar objects that already exist around us in the spaces we occupy will be given a digital-identity and possibly endowed with computational and communication capabilities. They, along with new and novel objects that include the virtual, will be interconnected and reflected by a digital presence. Collectively, these objects form part of the IoT – a massive distributed system that requires infrastructure to enable operation, discovery and management while simultaneously protecting scope, security and privacy. With such a rich and diverse landscape of information, there arises the necessity for standard middleware and novel Artificial

James Dooley · Hani Hagras · Vic Callaghan · Martin Henson
University of Essex, Wivenhoe Park, Colchester, Essex, U.K. CO4 3SQ.
e-mail: {jpdool, hani, vic, hensm}@essex.ac.uk

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environment*s", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

Intelligence (AI) to be used in dealing with / operating such a body of knowledge that results.

This chapter introduces the concepts that outline how active and passive objects are interconnected and unified in the real world. Collectively they form the fabric from which Intelligent Environments (IEs) are constructed and provide a layered support for intelligent agents and software applications to operate atop. The formation of an agent population within an IE results in an emergent and collective Ambient Intelligence (AmI) that exists as a product of interaction, cooperation and even competition among intelligent agents. The model scales-up to form a view of the world as a set of interconnected IEs between which human users and their subservient agents may roam. The purpose of this effort is to facilitate a better quality of life and continuity-of-experience as perceived by human inhabitants through environmental adaptation.

Herein, we describe convergence between the IoT and IE research fields. Although the two are distinct in their vision, the real world will be deployed by a hybrid of both - this chapter identifies how the envisioned Large-Scale Intelligent Environments (LSIEs) infrastructure can support the digital identity that is given to passive objects by mixed technologies such as RFID and barcodes. Conversely, the IoT vision enhances the operation of IEs by enabling passive objects to exist along-side active objects (which have embedded computation and communication capabilities).

Sect. 2 introduces the more significant IoT and IE literature that has led us to the present state-of-the-art. This includes a description of the four *"living-labs"* that we have constructed on the University of Essex campus. Sect. 3 then presents a view of the world that scales from individual passive / active objects, through the IEs they occupy and up to clouds as virtual collections of IEs. Sect. 4 describes the significant requirement on enabling technology - such as middleware for the interconnection of entities (Sect. 4.1) and agents, applications and virtual appliances for intelligent operation (Sect. 4.2). The material discussed is then illustrated through the use of a case study in Sect. 5. The chapter is then summarised and concluded with some remarks on future work in Sect. 6.

2 State of the Art

In 1991, Mark Weisers seminal work [1] described a grand *"Ubiquitous Computing"* (UC) vision for our future in which computer technology becomes transparently embedded in the world around us. This takes a user-centric approach in which the cognitive load of technology on people is reduced by making technology recede into the background of our lives, beyond human perception. This is in contrast to the modern day model of people staring *"awkwardly"* at a desktop screen - interacting on the terms of technology. The emergence of mobile-computing can be seen as a stepping stone between the two where users are always connected through the device they carry with them. In the years since, UC has flourished and stimulated a

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environment*s", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

great many works that span the entire spectrum of technology and society. The related fields of UC, IE, IoT, AmI and *Pervasive Computing* all offer variations around the same theme with focus on specific parts of the problem space. In the following two sub-sections, we focus on the more significant works in the IE and IoT areas.

The convergence of these two areas is object-centric and relies on the availability of common infrastructure. The real world will likely be the result of work that evolves from both IE and IoT fields and so the convergence of the two must be achieved. The problems of infrastructure and sensor availability that afflict the IoT vision are implicitly solved by the IE vision, whilst the IE need for a mix of passive and active objects is augmented by the IoT body of knowledge.

2.1 Internet of Things

Circa 1999, the IoT concept was suggested as a means to connect the internet to the physical world through the large-scale deployment of sensors. The intended purpose was to remove the dependance of humans by machines in acquiring information and to allow the information to be directly sampled from the real-world. With this proposition, the *"Auto-ID Centre"* was established. The purpose of the Auto-ID Centre was focused towards the investigation into Radio-Frequency IDentification (RFID) technology so that everyday objects could be given an Electronic Product Code (EPC) to aid supply-chain management [2] [3]. The work carried out under this Centre surpassed the standardisation of RFID and also investigated other associated problems, such as the specification of a common description language used to describe objects, processes and environments [4].

The *"Cooltown"* project explored the possibilities of linking every object in the world to a web-presence [5]. The work made use of several contact and non-contact sensing technologies and was motivated to link the physical and virtual worlds towards a mobile computing vision. Of note, this work examined other works in the field and classified the nature of links between a physical and virtual object by application [6] :

- **Physical Browsing:** The association of digital documents with physical objects - *users designate entities that interest them, and thereby obtain documents (pages) about them.*
- **Content Repositories:** The association of some digital content with physical objects - *so that users may transfer the content to one another or move it from place to place by passing the corresponding object around.*
- **Copy-and-Paste:** The temporary association of content with a physical *"clipboard"* object - *so that users can copy content from a source and paste it to a sink.*
- **Communication Points:** The association of communication medium with physical objects - *so that users who encounter the same object can communicate (for example using bulletin boards, email, voicemail, etc.).*

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

- **Physical Icons:** The association of actions with physical objects - *so that users can invoke actions such as turning a light on, by manipulating a physical object (in a similar fashion as selecting an icon on a desktop PC will invoke some action)*.

The authors of [6] also describe *Physical-Icons* as inputs to computational functions where the physical object can be mapped to a virtual entity. We would however argue that this deserves its own classification:

- **Object Reflection:** The association of a virtual *entity* with a physical object - *so that a user can identify an object, for the system to manipulate in some way through its virtual counterpart (for example identifying a camera will then provide a user-interface for the virtual functions of the camera, such as "view photos")*. The identification of users can also be considered as part of this classification.

Various technologies have been used to achieve the tagging of objects so that they can be identified uniquely, reliably and quickly by electronic systems [7]: Visual Object-recognition [8] [9], barcodes [10] [11], 2D barcodes (such as QR codes) [12], Infrared beacons [13] and badges [14], Ultrasound [15] [16], RFID [17], Ubisense (an RF-based realtime location tracking technology) [18], Wi-Fi [19], etc.

The salient point of the original IoT work is to identify *"passive"* objects (objects which have no computational or communication capabilities beyond that required to identify themselves) and to then do something based on that identity. It is worth noting that the tagged objects are passive - they do not actively do anything themselves, but rather the system that identifies them will carry out some action based on that recognition.

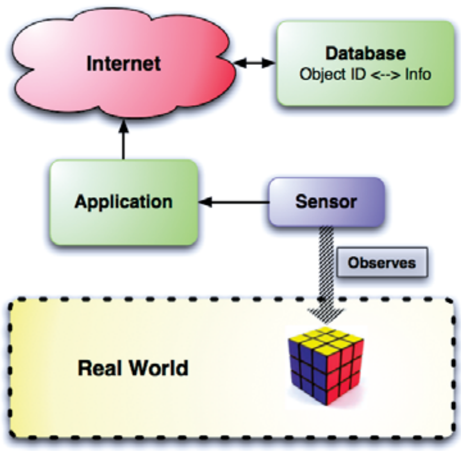

Fig. 1 The typical IoT architecture.

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environment*s", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

Fig. 1 shows the typical architecture taken: *an application uses some sensing apparatus to observe an object in the real world. A database is then queried over the internet (or some other network) to resolve necessary information from the identity that has been gained from the real-world object. The application can then use this information to achieve its functional operations*.

As the work towards an IoT progresses with works such as [20], the vision evolves and converges with that of the IE field. In particular the inclusion of not just passive objects, but active ones as well is increasingly popular. This now causes confusion as to what the *"things"* in the IoT actually are [21] - *"are they sensors, are they devices, or are they passive objects?"*

2.2 Intelligent Environments

As a multi-disciplinary research area, there are a huge variety of topics into which IE researchers delve. Consequently, this results in diverse approaches being taken to construct *"living labs"* in which the research is conducted. For example, the Cisco *"Internet House"* was constructed on a full building scale, but its purpose was to showcase a home with always-on Internet connectivity and appliance automation (where the home and its appliances could be controlled over the internet). Similarly, the Philips *"HomeLab"* [22] was a fully functional apartment whose purpose was aimed more at user experience evaluation through the use of monitoring technologies (such as cameras and microphones). The greater extent of technology deployment in the MIT *"Placelab"* also took place in a dedicated apartment scale space and focused on the space construction, technology deployment and user experience. The Stanford *"iRoom"* [23] and National Institute for Science and Technology (NIST) *"smart space"* [24] have investigated the deployment of ubiquitous computing in the office / meeting room context. The Fraunhofer inHaus-Center run two labs called the *"SmartHome"* and *"SmartBuilding"* for research into many different areas of innovation including user interfaces [25], an area of research also investigated by the *"iRoom"* at LIMSI [26]. Facilities such as the Duke University *"Smart Home"* have been used primarily for student projects, while the more recent emergence of community-lead *"hackerspaces"* around the world have promoted public participation in technology-oriented projects. Being rich in interconnected computing devices, sensors and actuators, these *"technology rich"* environments are the precursors to the IE - lacking only a quality of intelligence that is achieved through the deployment of suitably endowed software, such as intelligent embedded-agents.

Recognising this disparity, researchers have deployed intelligence into numerous spaces. At the University of Colorado, the *"Adaptive Home"* used a centralised neural-network based controller that monitored approximately 75 sensors (light, temperature, sound, motion, door/window state, etc.) and then took appropriate action on related actuators in the home [27]. Over the lifetime of this lab, many experiments were conducted and results published regularly. Such a rich publication history also exists for the Georgia Institute of Technology *"Aware Home"* that ex-

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

plores a huge diversity of subject areas including sociological applications such as assisted-living and home-care [28]. The *"PEIS home"* at the Orebro University further extends the capability of environment manipulation that lies within control of software intelligence by deploying and integrating mobile robots into its infrastructure [29]]. Elegantly, some labs (such as the iRoom at the German University in Cairo [30] and the MavHome at Washington State University [31]) are used to experiment with populations of software agents that provide the ambient intelligence (this is especially interesting when considering emergent behaviour from populations of agents that compete or collaborate).

At the turn of the century, when technology became cheap, small and abundant, there was a renewed energy in the field of UC. Works such as [32] were stimulated and the *"disappearing computer"* [33] was being chased. Among the fray of projects that we spawned, the e-Gadgets (extrovert gadgets) project was started and focused on the creation of pro-active *"Intelligent Artefacts"* [34]. In support of this, the Intelligent Dormitory (also known as the *"iDorm"*) was constructed as a test-bed that mimicked a single room student accommodation where individuals could stay for short periods of time (1-2 weeks). Within this seemingly normal place, heterogeneous technology was embedded and interconnected to form a grid computing deployment [35]. The iDorm identified and motivated continued work into the various challenges of UC, such as the Pervasive Home Environment Network (PHEN) project that continued to investigate the middleware and end-user interaction challenges [36].

Fig. 2 The University of Essex iSpace (living-area).

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

A range of devices, technologies and networks were used and almost every aspect of the space could be monitored or controlled by the software agents that constantly executed and evolved. This has been previously and comprehensively described over years of publication, such as in [35] [37] [38]. In combination with desktop PCs and hand-held devices; motion, pressure, temperature and light sensors sampled the world, blinds could be opened / closed, lamps and lights could be dimmed / switched, doors could be unlocked, heaters / coolers could be controlled, etc. From this work, the *"iSpace"* (a fully functional apartment, shown in Fig. 2 discussed more in Sect. 5) and the *"iClassroom"* [39] were evolved.

The salient point to note about the IE field is that objects are considered *"active"* - they are envisioned to have embedded systems and communications capabilities within them and so are able to perform tasks themselves. The combination of infrastructure and a population of these active-objects results in a complex and dynamic distributed system - one that intelligent agents are envisioned to operate, thus giving spaces an AmI quality . The overall IE resource is itself intended to be adaptable to changes in context and user preference through software agents that not only sense the real world, but also act upon it through actuators (see Fig. 3)

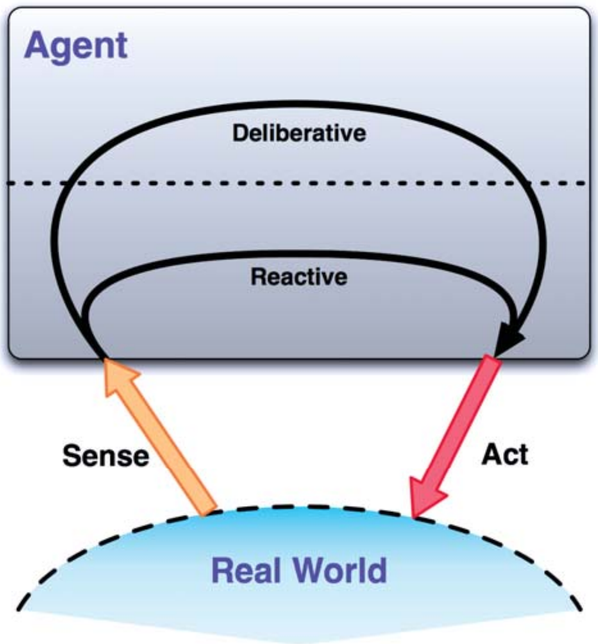

Fig. 3 The sense / act cycle that software agents conduct.

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

3 The World View

Beyond the test-beds of research and proof-of-concept works; the realisation of the UC / IoT / IE visions in the real world rely on operation at large-scale. To date this is something that the IoT field has accomplished very well and that the IE field is only just beginning to venture into [40] [41]. From a top-down perspective, our future world can be seen as a set of geographically distributed IE *"Spaces"* - interconnected by the internet (Fig. 4). Users roam through the physical world entering and exiting these IEs in a transient nature - each user having a distinct *"role"* in each and carrying with them a digital profile that contains their data, agents, preferences, etc. [42] [43]. Users can then be considered to have a history of occupancy within a subset of the IE superset (i.e. through the life of a user, that user will have visited some of all the Spaces that exist within the world).

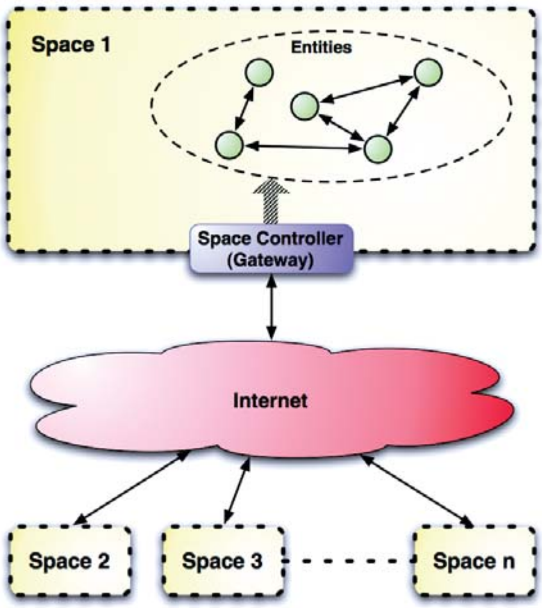

Fig. 4 An architecture for a *"world-of-spaces"*.

Within this model, there are two distinct architectures that come together: the inter-Space and intra-Space . The inter-Space architecture is large-scale and formed by the interconnection of Spaces over a large network such as the global internet. The intra-Space architecture is concerned with how a space is composed from its constituent devices and entities. We use the concept of an abstract *"Entity"* [44] to describe the uniquely identifiable digital-presence for an object of some form (such as a sensor, actuator, file, process, user, place, etc.) regardless if it is real, virtual,

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson “*The Tailored Fabric of Intelligent Environments*“, In Volume 460 of the book series ‘Studies in Computational Intelligence’, special edition on ‘ Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence’ pp 321-344, ISBN 9783642349515, 2013

logical or otherwise. Entities reside on physical devices and are grouped into sets that are published together by a *“Peer”* to a Space (Fig. 5).

In its simplest description; a Space (S) is a virtual machine (VM) that is distributed over a set of interconnected Peers that communicate through a network using secure middleware [45] . This virtual machine abstraction provides the conceptualisation that a space is centralised (with the associated advantages of management and security), even though it is indeed distributed (with the inherent properties of scalability and robustness). Several independent and isolated Spaces can exist across the same set of network-connected devices.

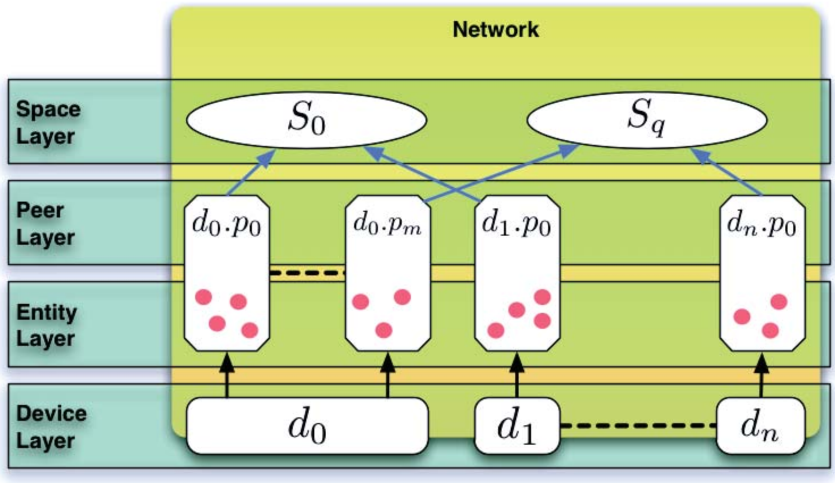

Fig. 5 Construction of a Space from its fabric (*“devices”*, *“peers”* and *“entities”*).

The Space Controller (SC; Fig. 4) represents the convergence of the inter-Space and intra-Space architectures. Itself an Entity; the SC acts as a gateway between locality and the wider large-scale. It also has the responsibility to manage the Space, its identity, members, etc. Through this gateway, Entities within a Space can be securely and safely accessed from outside. This model represents the convergence of IE and IoT towards a structured *“internet-of-entities”* .

4 Enabling Technology for IoT and IE convergence

A UC deployment is implicitly a distributed system - it relies on the interconnection of many computing devices and their software components across a network. Fig. 6 shows this and also shows the enabling technologies that are deployed across those devices. Of note, the middleware component (Sect. 4.1) must exist on every

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson “*The Tailored Fabric of Intelligent Environments*“, In Volume 460 of the book series ‘Studies in Computational Intelligence’, special edition on ‘ Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence’ pp 321-344, ISBN 9783642349515, 2013

device that wishes to participate in the distributed system. This provides network transparency for the software that is deployed on top of the middleware - allowing higher level components to operate. Virtual Appliances can be formed, while Agents , Applications and entities can communicate to achieve behaviour and functionality within an IE (Sect. 4.2).

In the remainder of this section, these enabling technologies are presented.

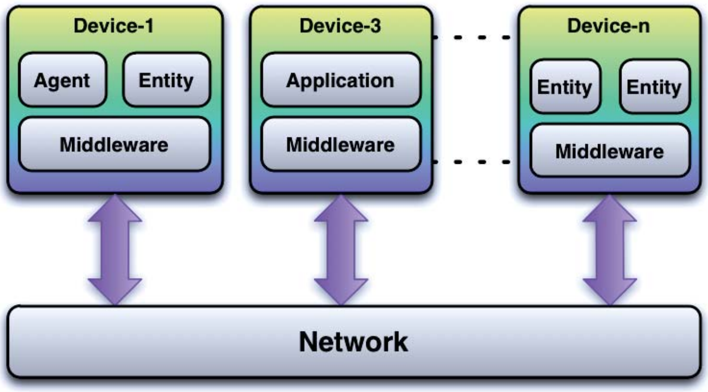

Fig. 6 A distributed system involving many devices, middleware, agents, applications and entities.

4.1 Middleware

Middleware provides a common functionality to higher level software such as agents and applications while abstracting the underlying implementation. It is the enabling technology that permits the processes on a single computing device to be rendered in a wider distributed computing environment - local resources can be exported and remote resources imported. As discussed in Sect. 3, the realisation of a converged IoT / IE reality depends on two kinds of architecture - the inter-Space and the intra-Space. While they must both support the same functionality (such as Entity discovery, interaction, eventing / subscription, etc.) they require slightly different approaches that are tailored to the conditions under which they must operate. The Space-Controller represents a convergence of these two approaches. At an inter-Space level, functionality is required to connect between spaces across the internet, while at the intra-Space level the emphasis is on the interconnection of Entities on a local network. The purpose of these architectures is to provide an end-to-end support for the interconnection of communicating entities that may reside in separate Spaces on a global scale.

Many approaches have been investigated for middleware that operates at the intra-Space level, [44] [46] but the core functionality that has been evolved here

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

has not been scaled-up to an inter-Space level (although the proposition of this has been suggested [47], it is still an open and exciting area of investigation):

- **"Entity Discovery":** The ability to discover a previously unknown entity given some search parameters - this is a particularly difficult thing to achieve in a distributed system that is subject to any real entropy. On a large scale (such as searching web-pages on the internet for content) this is usually achieved by centralised *"Search Engines"*, while on the intra-Space level it is realistic to use distributed search requests through broadcast / multicast messaging.
- **"Entity Resolving":** The ability to resolve an entity identity to the current location of the entity so that further interaction may occur (by routing messages to it). On the WWW, a URL acts as both a page identity and location, but in the IE / IoT vision, entities may be mobile and move from location to location - identity and location should therefore be de-coupled.
- **"Interaction":** The ability to send messages to an entity and receive responses.
- **"Eventing":** The ability to subscribe to an entity so that it may send asynchronous messages back (this is in contrast to polling that is inefficient, particularly on a large scale).

Technically, the scope of approaches used to achieve functional middleware varies, but the two most common are Remote Procedure Call (RPC) and Message-Oriented Middleware (MOM). These are very distinct in their approach - the former treats remote objects like local ones and presents software with a proxy of some form upon which procedures can be invoked as if it were a local resource, while the latter achieves communications by routing messages between entities. More recently, the concept of leveraging web technologies (HTTP, SOAP, etc.) and applying them in a service oriented fashion has attracted much attention. Although most of the work in this area still focuses on the use of larger, more powerful desktop / server hardware, the performance limitations are plainly seen when attempting to apply the same techniques with embedded systems that are less capable of processing the comparatively large message sizes that are typically encoded in XML documents (despite some more recent work in the past few years towards overcoming this limitation [48] [49]).

It has become a popular practice to use these underlying technologies as simply a transport mechanism and expose an Object model to the higher levels through an API. This approach provides a more convenient / usable middleware (that can sometimes be swapped out for alternative middleware) for higher level software to utilise. As part of this, the middleware layer will typically also incorporate extra features to aid with reliability and quality-of-service such as automatic failure detection and selection of new candidates. In some cases, the higher level API actually dictates the underlying model and results in what has come to be known as Object-Oriented Middleware (OOM). The early Object Request Broker (ORB) approach [50] is quintessential of this kind of middleware and attempted to provide an implementation independent specification (through the Interface Description Language - IDL) from its inception. This approach is still popular and has resulted in many flavours of OOM, such as [51], while forming the basis for further investiga-

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

tions such as the emergence of Reflective middleware [52] that makes assurances regarding the fidelity between an Object and its remote representations on the client side.

It is well-understood that there is a necessity for entity identity that is unique across space and time in a distributed system. Active-objects pro-actively present their own identity to the Space in the form of an Entity , but passive-objects have no way to achieve this and rely on the infrastructure to carry out the correct actions following identification. Passive objects must therefore be resolved from their ID by using logic / knowledge that exists either within an application / agent, a Space, a user-profile or some other entity in the wider world. The effect of resolving a passive-object identity can vary depending on what that object ID links to and the context in which it is used (see Sect. 2.1: Physical-Browsing, Content-Repositories, Copy-and-Paste, Communication-Points, Physical-Icons, Object-Reflection). For example; an RFID tag that is linked to a user ID - when the tag is identified by an access-control agent (that identifies users at a door and controls the door lock), the agent will seek to establish if that user has permission to enter and either unlock the door or provide some feedback to the contrary. However, the same tag linked to the same user in a different context will have a different effect - for example if the tag is identified to a coffee table, then that table may undergo some adaptation such as display artwork / messages. Likewise, an RFID tag that is linked to a song / album and identified by the same table may display the artwork for that music and begin playback via a media control agent. This is further explored in Sect. 5.

4.2 Agents, Applications and Virtual Appliances

Across the large-scale of deployment that is the world of Spaces, there is a very large scope for *"things-that-do"* as consumers of existing information and producers of synthetic information. The purpose of software that falls under any of these categories is to achieve some functionality - i.e. to do something. The variation among them is due to *how* that something is done:

- **Applications:** These are pieces of software that are designed to achieve some specific function and are generally developed to operate in the same way as traditional distributed system software - interacting with distributed entities across a network. For example, a digital photo-frame that loads image entities across the network and shows them sequentially on a display. This kind of application may also expose some interface that allows other things-that-do to manipulate its behaviour (for example, to pause on the current photo or flick through to the next photo).
- **Agents:** These are somewhat more complex than standard applications; agents are embedded with some form of AI or computational intelligence and are characterised by being pro-active, that is they do not simply react to user control, but actively operate independant of it. Some, but not all, will also have a capability to learn from experience and self-adapt behaviour / structure. While they can oper-

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson “*The Tailored Fabric of Intelligent Environment*s“, In Volume 460 of the book series ‘Studies in Computational Intelligence’, special edition on ‘ Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence’ pp 321-344, ISBN 9783642349515, 2013

ate independantly, there is also a huge scope for populations of interacting agents that cooperate and compete. Agents have been used extensively in the IE field, where they are given the ability to interact with the real world through sensors and acuators (see Fig. 3).

- **Virtual Appliances:** Virtual appliances can be constructed at runtime by linking together several component entities [53]. For example, a music-player can be constructed by linking a data source (for example an mp3 entity) to a decoder and then to one or more speakers. By assigning input and output *“ports”* to entities, simple graph-theory can be applied to construct a great number of appliances from the same set of component entities - it is the flow on information between them that achieves functionality. The appliances can be constructed, modified and deconstructed in real-time by simply linking / unlinking their IO ports. Hence, applications are recombinant [54].

All three kinds of things-that-do are portable across Spaces and can travel with a user from Space-to-Space. They do however all rely on middleware functionality to resolve component dependencies from those entities that are available at runtime (a process known as *“Runtime-Discovery”*). This concept can be extended to improve reliability / robustness to component failure by swapping out components for *“better”* ones as and when they are found.

While virtual appliances are essentially instantiated by the interconnection of entities, they are inherently bound to intra-Space deployment. Applications and agents can, however, reside *“in the cloud”* and peer into Spaces by interacting with entities that are accessible through the Space-Controller. This is especially useful considering that user-profiles will have some presence and dependance upon the cloud to facilitate the migration of digital assets from Space-to-Space.

A mixed population of things-that-do within an environment gives the user an experience that has a variable level of autonomy and transparency. This eases the cognitive load on the user by hiding away some decision-making and operation whilst making others overt. Filtering of user-direction makes the increasingly technological world more tractable without removing the sense of control that users need in order to be accepting of UC. In particular, as a user moves from Space-to-Space they experience a continuity of experience as the environment is adapted to the preferences of the user. This gives an impression that there is a collective and coordinated AmI at work on behalf of IE occupants.

5 Case Study: the Essex iSpace

The *iSpace* is a purpose-built, fully functional apartment that resides within the School of Computer Science and Electronic Engineering (CSEE) at the University of Essex, UK (see Fig. 2). Its layout consists of a main living area, kitchen, bedroom, study-room, bathroom and control room. A false ceiling and false walls provide additional space to hide technology such as sensors, actuators and computational devices from small embedded systems up-to full desktop PCs. As a UC deployment

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson “*The Tailored Fabric of Intelligent Environments*“, In Volume 460 of the book series ‘Studies in Computational Intelligence’, special edition on ‘ Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence’ pp 321-344, ISBN 9783642349515, 2013

the iSpace is equipped with numerous sensors that sample the various phenomena of the real world, actuators that manipulate aspects of the environment, computational devices that run software and a firewall protected network that interconnects the entire resource. The architecture of this deployment is shown in Fig. 7, where the Ethernet / WiFi backbone can easily be seen as the convergence of many devices, some of which act as gateways into specific technologies. UPnP is deployed as a middleware that homogenises the heterogeneous and distributed UC resources, thus providing a consistent and accessible view of the network for software agents and applications. At the time of writing (March 2012), there are over 100 UPnP devices deployed within the iSpace network, each representing an entity of some form (logical, virtual, real, etc.). The availability of dynamic-discovery, event-subscription and action-invocation within this *living-lab* allows loosely-coupled agents and applications to interact directly with every entity on the network - UPnP device / service types define common interfaces; sensors can produce asynchronous events and actions can be invoked to achieve some function (such as turning on a light). While a few implementations have been used, the UPnP functionality is primarily achieved through the use of a Java based library called Youpi that was developed as part of the Atraco project . This library has been released as open-source and is used in both living-labs (such as the LIMSI iRoom) and commercial products (such as those offered by inAccess networks).

Although each gateway device in the iSpace is unique in its configuration, there are two main types. The first configuration uses a Java based OSGi framework that provides component management. Bundles of functionality (including the middleware) are deployed in this framework and can be done so dynamically during runtime. In this configuration a single runtime exists on the device and capabilities are added by installing bundles of functionality. This is easy to manage and efficient in operation due to the fact that only one instance of the middleware is running per device (and so only one middleware runtime needs to communicate over the network). In the second configuration, each component of functionality is wrapped in its own application. This requires more effort to manage and is less efficient from the perspective of the middleware, but is necessary where multiple components need to be deployed to a device and each component has been developed using different languages / tools. For example, our Windows PC-3 has two applications deployed on it (as shown in Fig. 7):

1. A Java application that advertises functionality to control Curtains and also provides a management GUI. This application operates its own instance of the middleware, has the control logic in-built and communicates via a RS-232 serial connection to the curtain actuators.
2. A C++ application that advertises functionality for the Ubisense real-time location tracking system (RTLS). This application operates its own instance of the middleware that wraps the installed software system (a Windows based application) through a C++ API.

While some of the bundles / applications wrap functionality by communicating directly with microcontrollers (as is the case for the curtains, lights, Phidget sensors

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

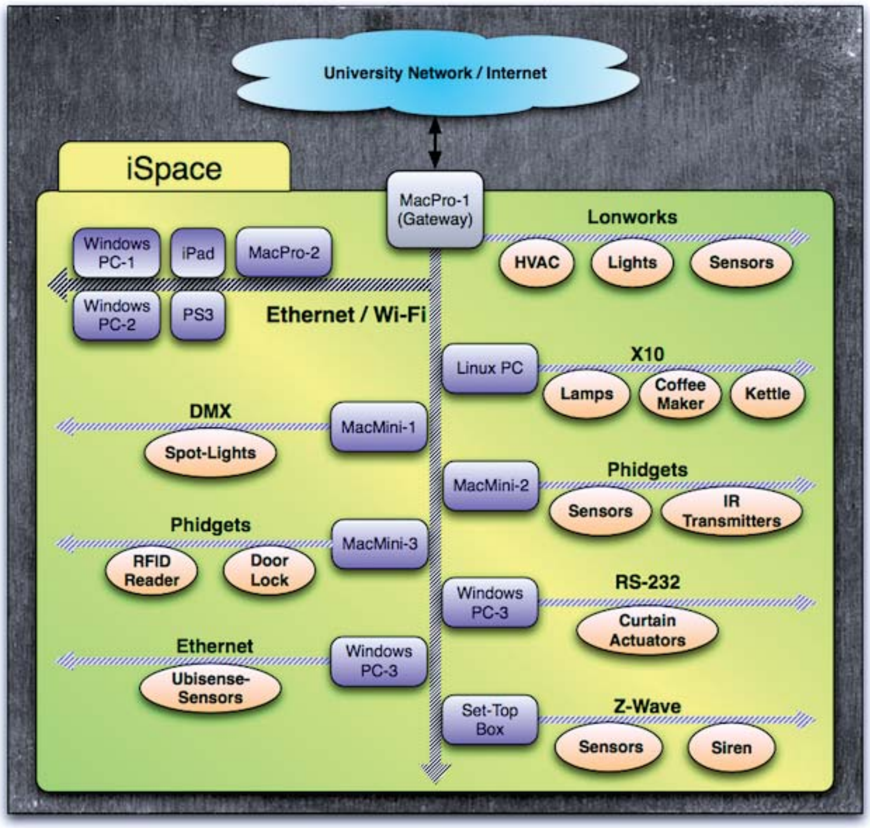

Fig. 7 Architecture of technology deployment within the iSpace.

/ actuators and X10 devices) others exploit programming APIs of other software packages (such as the Ubisense RTLS and the Lonworks sensors / actuators), or have the exposed functionality in-built (such as is the case for the HTML5 based user interfaces and media repositories). The technical details regarding each implementation is beyond the scope of this chapter and could be realistically achieved using several approaches - what should be noted is the functionality they provide and the way in which they can be utilised. For example it is important to understand that a light can be advertised, described, manipulated and inspected through its software representation that is made available to agents / applications that are distributed across the network. In this specific example, three separate lighting technologies are used in parallel throughout the iSpace - but to a software consumer, a light of each type is indistinguishable from a light of any other type as the interfaces they implement are the same. The heterogeneity of the numerous components in the iSpace is homogenised through the middleware, allowing a single consistent model to be used.

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

The iSpace is an excellent experimental facility for multi-disciplinary research , this is especially useful across the spectrum of UC investigations where there is a symbiotic relationship between computer-science and social-science; two of the demonstrations from its portfolio are described below to holistically illustrate the concepts introduced through this chapter: *"FollowMe"* and *"HotSpot"* . Both of these demonstrators make use of RFID technology to recognise user-initiated events - a form of HCI that permits explicit user control. The experimental setup is the same for both the demonstrators and is shown in Fig. 8 below.

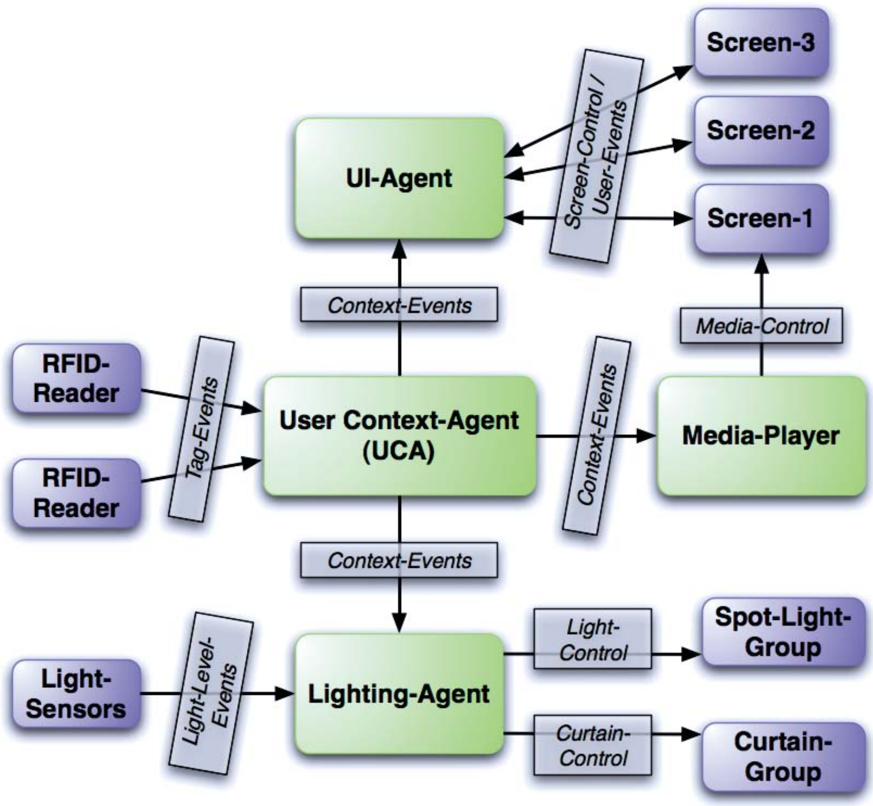

Fig. 8 Experimental setup of entities within the iSpace living-area for both FollowMe and HotSpot.

In particular, a multi-agent approach is adopted and the deployment consists of the following entities that are used to achieve desired operation:

1. **Real Entities:**

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

a. **Spot-Lights:** Eight dimmable spot lights are embedded in the ceiling of the living / kitchen area. Each is individually represented by a single UPnP device and can have its state (on / off) and intensity (0-100) controlled.
b. **Light-Level Sensors:** Light sensors embedded in the walls and ceiling of the iSpace provide localised measurements for light-levels and can be used together to build a picture of the overall lighting conditions in the space. Values from each sensor can be retrieved through action invocation or by subscribing to the UPnP device for asynchronous event notifications.
c. **Curtains:** Two windows are equipped with motorised curtain controls. Each window is represented as a single UPnP device and can be in one of two states: OPEN or CLOSED.
d. **RFID-Readers:** Two RFID readers are deployed - embedded into the furniture of two contextual zones (markers indicate where a user must *"tag-in"*).
e. **Screens:** Of the six screens available, three are used within these demonstrators. Each shows a full-screen HTML5 web-browser that is connected to the UI agent (each screen provides a unique ID to the UI agent through the URL that it GETs). The HTML5 web-sockets feature is used to maintain a bidirectional link with the UI agent - this allows events to flow asynchronously in both directions.
   i. **Screen-1:** A 40" LCD TV with a touch-sensitive overlay. This screen is used as a user interface.
   ii. **Screen-2:** A 40" wall-mounted Plasma screen used as an ambient media display.
   iii. **Screen-3:** A table-top projection (top-down LCD projector onto kitchen table) that provides user interface through a wireless trackpad.

2. **Virtual Entities:**

a. **Light-Group:** The eight *"real"* spot-lights can be addressed / controlled through a single UPnP device that represents them as a virtual group. Virtual light groups also exist for each of the contextual zones, but are not utilised in FollowMe / HotSpot.
b. **Curtain-Group:** In a similar way to the Light-Group; the Curtain-Group provides a single and convenient representation for the two curtain devices to be treated as one.

3. **Logical Entities:**

a. **User Context-Agent (UCA):** This software agent has knowledge of fixed RFID reader locations *a priori*. It also has a database of known-users, each with an associated identity, RFID-tag and profile. A subscription to the two UPnP RFID readers allows RFID events to be monitored (the location of the event can be inferred by using the ID of the source RFID-Reader). A UPnP interface allows other agents to access user profile information and to subscribe for user context changes that are initiated when a user *"tags-in"* to a zone. Non-user RFID tags can be registered within a user-profile to generate

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

specific events to subscribers (this is discussed further in the FollowMe and Hotspot sub-sections).

b. **Lighting-Agent:** This Fuzzy-Task Agent (FTA) [55] controls the spot-lights to achieve lighting adaptation in response to context events from the UCA. A subscription to light-level sensors provides feedback from the environment and a fuzzy membership function is used in conjunction with the learned user preferences for light-levels.
c. **UI-Agent:** This software agent has knowledge of fixed screen locations *a priori* and provides a HTML5 web-server for each screen to connect to (each screen provides a unique screen ID when it connects). This agent subscribes to the UCA and will modify the content of each individual screen when a user context change event is received.
d. **Media-Player:** A UPnP media player is used to render audio and video on demand. When active it occupies full-screen on Screen-1 (replacing other active screen content such as the environment UI). For simplicity, the audio is simply output through the TV speakers, but could be direted to some other UPnP audio renderer if desired.

Why are there virtual groupings for lights / curtains? There are two reasons for this: *firstly*, it is more convenient and robust to develop software that deals with one remote resource than many. *Secondly*, Action invocation over a network using middleware has an inherent problem - it incurs a time overhead. More specifically, the current open-source *"Youpi"* implementation used requires $\sim$100ms to invoke an action on a UPnP device. And so, when an agent / application wants to achieve something like *turn on a bunch of lights*, there is a perceivable delay between the first and last light illuminating. Using a singly addressable group removes this problem as only one action invocation is required and so the individual lights respond together.

5.1 FollowMe

In this demonstrator, the user-interface for a specific occupant will migrate from screen-to-screen as he / she roams through the iSpace - thus it follows the user [46]. This relies on knowledge of screen locations (which are fixed) and user location (which is dynamic).

When a user touches an RFID-tag (such as the one shown in Fig. 9) onto a reader, the UCA attempts to match the identity to one in its database. If the tag matches one registered as a user-context-tag, then a *context-change* event is distributed to all subscribed listeners. This event consists of the user-ID, a timestamp and a location-ID (inferred from the RFID-reader-ID). If the user is not already *logged-in* to the space, the listeners will utilise the user profile (available from the UCA) to configure certain aspects of the environment. The Lighting-Agent sets the appropriate light level and curtain state, the media player stops any currently playing media (and

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

may start some background music if the user-profile specifies this preference), the UI-Agent transfers the UI to the screen closest to the user location (Screen-1 or Screen-3) and then sets the artwork on Screen-2 to the user preference. A textual message is also popped-up on Screen-2 so that the user is informed of what just happened.

Fig. 9 RFID-tag attached to the keys of a user.

As the user roams through the iSpace, they can touch-in to other locations - this prompts the UCA to generate a new *context-change* event to all subscribed listeners. When the UI-Agent receives this event it will migrate the UI from whichever screen it is currently on and transfers it to the screen at the new user location.

5.2 HotSpot

In this demonstrator, a user can explicitly express some wish to the iSpace by placing an RFID-tagged object onto a reader. Fig. 10 shows three kinds of tagged objects - a document, some DVDs and a toy duck. The effect of each object on the space is specified in the user profile - it should be noted that a single object can therefore have a different effect depending on which user is currently logged-in. For the purpose of this discussion, we will present one of the authors profiles (physical document maps to digital document, DVDs map to movies and toy duck maps to music).

When a tagged object is placed on the reader, the UCA attempts to match the identity to one in the current user profile. Described below are the effects of three kinds of entity that are linked:

1. **Document:** A *context-change* event is generated by the UCA to indicate the user-activity is *"reading"*. The Lighting-Agent adapts the lights and curtains to the user preference for this task (for example all lights to full brightness and curtains

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson “*The Tailored Fabric of Intelligent Environments*“, In Volume 460 of the book series ‘Studies in Computational Intelligence’, special edition on ‘ Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence’ pp 321-344, ISBN 9783642349515, 2013

Fig. 10 RFID-tagged objects: a) a document, b) two DVDs, c) a toy duck.

open). The UI-Agent displays a digital form of the document (PDF) on Screen-2 and the Media-Player stops any current video / music.

2. **Movie:** A *context-change* event is generated by the UCA to indicate the user-activity is *“watching-movie”*.The Lighting-Agent adapts the lights and curtains to the user preference for this task (for example all lights to low brightness and curtains closed). The UI-Agent displays coverart on Screen-2 while the Media-Player goes full-screen on Screen-1 and then begins playback of the movie (streamed from a URL source over the network).
3. **Music:** Here the Genre of the linked music playlist is examined by the UCA in order to further specify the generation of a *context-change* event. As a result, the UCA notifies its listeners to indicate a *“relaxing-with-music”* activity.The Lighting-Agent adapts the lights and curtains to the user preference for this task (for example all lights to low brightness and curtains closed). The UI-Agent displays coverart on Screen-2 while the Media-Player goes full-screen on Screen-1 and then begins music playback (streamed from a URL source over the network).

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environment*s", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

6 Conclusions and Future Challenges

In this chapter we have discussed the convergence of IoT and IE approaches into a workable model that caters to both active and passive objects. A case study of the University of Essex iSpace is provided in which two demonstrators are described that illustrate the discussed convergence.

The abstraction of all things as entities allow software populations to make use of common functionality in order to reason with and manipulate a vast array of *"things"*. And so, the entities that form the fabric of an IE can be tailored to the preferences of a user. This is facilitated by middleware that renders a homogeneous distributed system from heterogeneous *"things"*. To ease the need on humans for direction and orchestration, intelligent agents collectively form an AmI that interacts with the real world through the use of sensors and actuators - providing a quality of intelligence that achieves autonomy.

The grand vision we have is for a world of Spaces, where each space constitutes an IE. People will be able to roam from Space-to-Space and enjoy a continuity of experience. There is still a lot of work to be done to achieve this. In particular, the inter-Space and intra-Space relationships need to be integrated and aligned to allow universal and global interoperability. The work towards realising AmI must also make breakthroughs, particularly in the support of multiple-users from its current proof-of-concept state in which single user scenarios are the norm. Security and privacy must also see a vast improvement before widespread adoption is made; perhaps the grandest challenge of all however, is to address the legal and societal boundaries to acceptance. The world is certainly becoming more accepting of technology in society - mobile phones, set-top boxes, tablet computers, etc. are already pervasive and have broken down the *"digital divide"* that once excluded certain groups of people from adoption. However, the more exciting current trend is the emergence of a global community that includes hobbyists and professionals alike that are in particular taking advantage of cheap and freely-available electronics such as [56] [57] [58] to realise new and novel creations that contribute to the IoT - an exemplar of technology in society.

What next could we envision once we have this world of Spaces ? perhaps Spaces that are structurally reconfigurable - a challenge more for architects and engineers within this, a multi-disciplinary field.

Acknowledgements We gratefully acknowledge the support of HEFCE, the EPSRC, EU Framework Programmes, King Abdulaziz University, and the numerous other sources of funding that the Intelligent Environments Group has attracted over the past 10 years.

References

1. Mark Weiser. The computer for the twenty-first century. *Scientific American*, pages 94–104, September 1991.

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

2. David Brock. The networked physical world, proposals for engineering the next generation of computing, commerce and automatic-identification. *Auto-ID Centre Whitepaper*, January 2001.
3. Paolo Magrassi. A world of smart objects: The role of auto identification technologies. *Strategic Analysis Report, Gartner*, 2001.
4. David Brock. The physical markup language. *Auto-ID Centre Whitepaper*, February 2001.
5. T. Kindberg, J. Barton, J. Morgan, G. Becker, D. Caswell, P. Debaty, G. Gopal, M. Frid, V. Krishnan, H. Morris, J. Schettino, B. Serra, and M. Spasojevic. People, places, things: Web presence for the real world. In *Mobile Computing Systems and Applications, 2000 Third IEEE Workshop on.*, pages 19 –28, 2000.
6. John Barton and Tim Kindberg. The challenges and opportunities of integrating the physical world and networked systems. Technical report, HPL Technical report HPL-2001-18, 2001.
7. Roy Want and Daniel M. Russell. Ubiquitous electronic tagging. *IEEE Distributed Systems Online*, 1(2), February 2000.
8. Till Quack, Herbert Bay, and Luc Van Gool. Object recognition for the internet of things. In *Proceedings of the 1st international conference on The internet of things*, IOT'08, pages 230–246, Berlin, Heidelberg, 2008. Springer-Verlag.
9. J. Rekimoto and Y. Ayatsuka. Cybercode: Designing augmented reality environments with visual tags. *Proceedings of the ACM Designing Augmented Reality Environments*, 2000.
10. Robert Adelmann, Marc Langheinrich, and Christian Floerkemeier. Toolkit for bar code recognition and resolving on camera phones - jump starting the internet of things. In *Proceedings of Workshop Mobile and Embedded Interactive Systems (MEIS06) at Informatik 2006*, pages 366–373, 2006.
11. David Jarrett Collins and Nancy Nasuti Whipple. *Using Bar Codes – why it's taking over*. Data Capture Institute, 1990.
12. T. Falas and H. Kashani. Two-dimensional bar-code decoding with camera-equipped mobile phones. *Fifth Annual IEEE International Conference Pervasive Computing and Communications Workshops (PerCom)*, pages 597 — 600, 2007.
13. Tim Kindberg and John Barton. A web-based nomadic computing system. *Computer Networks, Special Edition on Pervasive Computing*, 35(4):443–456, 2001.
14. V. Falcao J. Gibbons R. Want, A. Hopper. The active badge location system. *ACM Transaction on Information Systems*, 10(1):91–102, 1992.
15. Andy Ward, Alan Jones, and Andy Hopper. A new location technique for the active office. *IEEE Personal Communications*, 4(5):42–47, 1997.
16. Hari Balakrishnan and Nissanka Bodhi Priyantha. The cricket indoor location system: Experience and status. *Proceedings of the 2003 Workshop on Location-Aware Computing, Held in conjunction with UbiComp 2003*, pages 7 – 9, 2003.
17. E. Welbourne, L. Battle, G. Cole, K. Gould, K. Rector, S. Raymer, M. Balazinska, and G. Borriello. Building the internet of things using rfid: The rfid ecosystem experience. *Internet Computing, IEEE*, 13(3):48 –55, may-june 2009.
18. Jay Cadman. Deploying commercial location-aware systems. *Proceedings of the 2003 Workshop on Location-Aware Computing, Held in conjunction with UbiComp 2003*, pages 4 – 6, 2003.
19. T. Garcia-Valverde, A. Garcia-Sola, H. Hagras, J. Dooley, V. Callaghan, and J.A Botia. A fuzzy logic based system for indoor localisation using wifi in ambient intelligent environments. *To appear in IEEE Transactions on Fuzzy Systems*, 2012.
20. Jan S. Rellermeyer, Michael Duller, Ken Gilmer, Damianos Maragkos, Dimitrios Papageorgiou, and Gustavo Alonso. The software fabric for the internet of things. In *Proceedings of the 1st international conference on The internet of things*, IOT'08, pages 87–104, Berlin, Heidelberg, 2008. Springer-Verlag.
21. Stephan Haller. The things in the internet of things. *Procceedings of the Internet of Things Conference 2010, Tokyo, Japan*, 2010.
22. B. de Ruyter and E. Aarts. Ambient intelligence: visualizing the future. *In Proceedings of the Working Conference on Advanced Visual interfaces (AVI '04)*, pages 203 – 208, 2004.

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environment*s", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

23. Brad Johanson, O Fox, and Terry Winograd. The interactive workspaces project: Experiences with ubiquitous computing rooms. *IEEE Pervasive Computing*, 1:67 – 74, 2002.
24. Hamchi I. Degré S. Diduch L. Rose T. Fiscus J. Fillinger, A. and V. Stanford. Middleware and metrology for the pervasive future. *IEEE Pervasive Computing Mobile 8and Ubiquitous Systems*, 8(3):74 – 83, 2009.
25. C. Ressel, J. Ziegler, and E. Naroska. An approach towards personalized user interfaces for ambient intelligent home environments. *2nd IET International Conference on Intelligent Environments, 2006. (IE'06)*, 1:247 – 255, 2006.
26. Y. Bellik and C. Jacquet. From the intelligent room to ambient intelligence. *1st Digiteo Annual Forum, Poster n°2008-01*, 2008.
27. M. C. Mozer. The neural network house: An environment that adapts to its inhabitants. *In M. Coen (Ed.), Proceedings of the American Association for Artificial Intelligence Spring Symposium on Intelligent Environments*, pages 110 – 114, 1998.
28. C. D. Kidd, R. Orr, G. D. Abowd, C. G. Atkeson, I. A. Essa, B. MacIntyre, E. D. Mynatt, T. Starner, and W. Newstetter. The aware home: A living laboratory for ubiquitous computing research. *In Proceedings of the Second international Workshop on Cooperative Buildings, integrating information, Organization, and Architecture*, 1670 (Lecture Notes In Computer Science):191 – 198, 1999.
29. M. Broxvall, B.S. Seo, and W.Y. Kwon. The peis kernel: A middleware for ubiquitous robotics. *Proceedings of the IROS-07 Workshop on Ubiquitous Robotic Space Design and Applications.*, 2007.
30. B. El-Desouky and H. Hagras. An adaptive type-2 fuzzy logic based agent for multi-occupant ambient intelligent environments. *In proceedings of the 5th International Conference on Intelligent Environments (IE'09)*, 2009.
31. D. J. Cook, M. Youngblood, E. Heierman, K. Gopalratnam, S. Rao, A. Litvin, and F. Khawaja. Mavhome: An agent-based smart home. *In proceedings of the IEEE International Conference on Pervasive Computing and Communications*, pages 521 – 524, 2003.
32. Friedemann Mattern, Manuel Ortega Cantero, and Jesús Lorés Vidal. Ubiquitous computing – the trend towards the computerization and networking of all things. *Upgrade*, 2(5), October 2001.
33. N. Streitz and P Nixon. The disappearing computer. *Communications of the ACM*, 48(3):32–35, 2005.
34. A. Kameas, I. Mavrommati, and P. Markopoulos. Computing in tangible: using artifacts as components of ambient intelligence environments. *Ambient Intelligence: The evolution of Technology, Communication and Cognition*, pages 121 – 142, 2004.
35. A. Pounds-Cornish and A. Holmes. The idorm - a practical deployment of grid technology. *Second IEEE International Symposium on Cluster Computing and the Grid (CCGRID'02)*, 2002.
36. R. Limb, S. Armitage, J. Chin, R. Kalawsky, V. Callaghan, P. Bull, M. Colley, and H. Hagras. User interaction in a shared information space - a pervasive environment for the home. In *Proceedings of the IEE Workshop on Perspectives in Pervasive Computing*, October 2005.
37. H. Hagras, V. Callaghan, M. Colley, G. Clarke, A. Pounds-Cornish, and H. Duman. Creating an ambient-intelligence environment using embedded agents. *IEEE Intelligent Systems*, 19(6):12 – 20, 2004.
38. F. Rivera-Illingworth, V. Callaghan, and H. Hagras. A neural network agent based approach to activity detection in ami environments. *Proceedings of the first IEE International Workshop on Intelligent Environments 2005 (IE'05).*, 2005.
39. James Dooley, Hani Hagras Vic Callaghan, Michael Gardner, Mohammed Ghanbaria, and Daniyal AlGhazzawi. The intelligent classroom : Beyond four walls. *Intelligent Campus Workshop (IC'11) held at the 7th IEEE Intelligent Environments Conference (IE11)*, 2011.
40. James Dooley, Matthew Ball, and Mohammed Al-Mulla. Beyond four walls: Towards large-scale intelligent environments. *In proceedings of the first Workshop On Large Scale Intelligent Environments (WOLSIE), col-located with the 8th International conference on Intelligent Envrironments 2012 (IE'12).*, 2012.

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson “*The Tailored Fabric of Intelligent Environment*s“, In Volume 460 of the book series ‘Studies in Computational Intelligence’, special edition on ‘ Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence’ pp 321-344, ISBN 9783642349515, 2013

41. D. J. Cook and S. Das. Pervasive computing at scale: Transforming the state of the art. *Journal of Pervasive and Mobile Computing*, 8(1):22 – 35, 2012.
42. Laurent Beslay and Hannu Hakala. Digital territory: Bubbles. *In European Visions for the Knowledge Age: A Quest for New Horizons in the Information Society.*, 2007.
43. Luke Whittington, James Dooley, Martin Henson, and Abdullah AL-Malaise AL-Ghamdi. Towards followme user profiles for macro intelligent environments. *In proceedings of the first Workshop On Large Scale Intelligent Environments (WOLSIE), col-located with the 8th International conference on Intelligent Envrironments 2012 (IE'12).*, 2012.
44. James Dooley. *An Information Centric Architecture for Large Scale Description and Discovery of Ubiquitous Computing Objects*. PhD thesis, University of Essex, 2011.
45. James Dooley, Martin Henson, Vic Callaghan, Hani Hagras, Daniyal AlGhazzawi, Areej Malibari, Mohammed AlHaddad, and Abdullah AL-Malaise AL-Ghamdi. A formal model for space based ubiquitous computing. *Intelligent Environments Conference 2011 (IE'11)*, 2011.
46. James Dooley, Christian Wagner, Hani Hagras, and Gaten Pruvost. Followme: The persistent gui. *1st International Workshop on Situated Computing for Pervasive Environments (SCOPE'11) at the 6th International Symposium on Parallel Computing in Electrical Engineering (PARELEC'11)*, 2011.
47. Yichao Jin, Ruchuan Wang, Haiping Huang, and Lijuan Sun. Agent-oriented architecture for ubiquitous computing in smart hyperspace. *Wireless Sensor Network*, 2(74 – 84), 2010.
48. Francois Jammes, Antoine Mensch, and Harm Smit. Service-oriented device communications using the devices profile for web services. *Proceedings of the 3rd international workshop on Middleware for Pervasive and Ad-hoc Computing (MPAC)*, 2005.
49. E. Zeeb, A. Bobek, H. Bohn, and F. Golatowski. Service-oriented architectures for embedded systems using devices profile for web services. *21st International Conference on Advanced Information Networking and Applications Workshops (AINAW)*, 1:956–963, May 2007.
50. W. Emmerich. *OMG/CORBA: An Object-Oriented Middleware*, chapter in Encyclopedia of Software Engineering, John J. Marciniak (ed), pages 902–907. John Wiley and Sons, 2002.
51. M. Henning. A new approach to object-oriented middleware. *IEEE Internet Computing,*, 8(1):66 – 75, jan-feb 2004.
52. Fabio Kon, Fabio Costa, Gordon Blair, and Roy H. Campbell. The case for reflective middleware. *Communications of the ACM*, 45(6):33–38, June 2002.
53. J. Chin, V. Callaghan, and G. Clarke. Soft-appliances: A vision for user created networked appliances in digital homes. *Journal of Ambient Intelligence and Smart Environments (JAISE)*, 2009.
54. W. Keith Edwards, M.W. Newman, and J.Z. Sedivy. The case for recombinant computing (technical report csl-01-1). Technical report, Xerox Palo Alto Research Center, 2001.
55. Aysenur Bilgin, James Dooley, Luke Whittington, Hani Hagras, Martin Henson, Christian Wagner, Areej Malibari, Abdullah AlGhamdi, Mohammed AlHaddad, and Daniyal AlGhazzawi. Dynamic profile-selection for zslices based type-2 fuzzy agents controlling multi-user ambient intelligent environments. *Proceedings of the 2012 FUZZ-IEEE Conference*, 2012.
56. V. Callaghan. Buzz-boarding; practical support for teaching computing, based on the internet-of-things. In *1st Annual Conference on the Aiming for Excellence in STEM Learning and Teaching*, April 2012.
57. Eben Upton and Gareth Halfacree. *Raspberry Pi User Guide*. John Wiley and Sons, 2012.
58. Leah Buechley, Mike Eisenberg, Jaime Catchen, and Ali Crockett. The lilypad arduino: using computational textiles to investigate engagement, aesthetics, and diversity in computer science education. In *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, CHI '08, pages 423–432. ACM, 2008.

James Dooley, Hani Hagras, Victor Callaghan, Martin Henson "*The Tailored Fabric of Intelligent Environments*", In Volume 460 of the book series 'Studies in Computational Intelligence', special edition on ' Internet of Things and Inter-cooperative Computational Technologies for Collective Intelligence' pp 321-344, ISBN 9783642349515, 2013

Index