Smart Home Area Management

David Lillis, Tadhg O'Sullivan, Thomas Holz, Conor Muldoon, Michael J. O'Grady and Gregory M. P. O'Hare
CLARITY: Centre for Sensor Web Technologies
School of Computer Science and Informatics
University College Dublin, Ireland

ABSTRACT

Autonomically managing energy within the home is a formidable challenge as any solution needs to interoperate with a decidedly heterogeneous network of sensors and appliances, not just in terms of technologies and protocols but also by managing smart as well as "dumb" appliances. Furthermore, as studies have shown that simply providing energy usage feedback to homeowners is inadequate in realising long-term behavioural change, autonomic energy management has the potential to deliver concrete and lasting energy savings without the need for user interventions. However, this necessitates that such interventions be performed in an intelligent and context-aware fashion, all the while taking into account system as well as user constraints and preferences. Thus this chapter proposes the augmentation of home area networks with autonomic computing capabilities. Such networks seek to support opportunistic decision-making pertaining to the effective energy management within the home by seamlessly integrating a range of off-the-shelf sensor technologies with a software infrastructure for deliberation, activation and visualisation.

INTRODUCTION

The home area network (HAN) market is set to reach $3 billion by 2014, with 30 million households containing nearly 50 million smart home and energy management devices (On World, 2010). However, this opportunity can only be availed of if real and tangible energy savings are achieved. In fact, over 80% of household owners are willing to pay for energy management services if they achieve savings of 30% or more (ON World, 2010).

Achieving savings of such a magnitude is a challenging problem. As numerous studies have shown, merely displaying energy consumption leads only to minimal savings (e.g. McKerracher & Torriti, 2012). Incentives, usually in the form of monetary savings, can have some effect on usage (Faruqui et al., 2009; Lui et al., 2010), but this may still not lead to lasting behavioural change, with users regressing to old consumption patterns over the course of several months (Hazas et al., 2011; Kluckner et al., 2013). Autonomically and intelligently managing energy consumption in the home, on the other hand, has the potential to deliver energy savings that are completely independent of occupant behaviour and, as such, are not susceptible to the weaknesses outlined above.

To enable autonomic energy management, there is a need for a pervasive sensing and networking configuration within the home. This configuration, referred to as an Autonomic Home Area Network Infrastructure (AUTHENTIC), must be capable of supporting opportunistic decision-making pertaining to residential energy management by seamlessly and effortlessly integrating several key enabling technologies. These include HAN technologies, physical sensing devices and a mechanism for sensing contextual data outside the home. Going forward, it will also be essential to provide a mechanism for integration with the smart grid.

The aim of this chapter is to present the design and implementation of the AUTHENTIC architecture. It first describes the overall system architecture, along with brief discussions of its constituent modules and their functionalities. Following this, particular focus is paid to the SIXTH middleware, which is a fundamental component of the system, acting as a conduit for sensor data and actuation commands so as to become the central hub of the overall architecture. The reasoning module of the system is based on the multi-agent paradigm and is discussed in detail in the next section. The effectiveness of the system is then illustrated via a simple use case scenario. A summary of related approaches to the problem is then outlined, before finally presenting conclusions and some ideas for further work.

SYSTEM ARCHITECTURE

The AUTHENTIC system follows a component-based software engineering approach adhering to the OSGi framework and is made up of five different modules (see also Figure 1):

- The communications module, which provides a unified interface to all sensors and actuators in the HAN;
- The semantic module, which infers situational context from low-level sensor data;
- The deductive module, which realises intelligent decision-making using the SIXTH middleware together with a multi-agent systems approach;
- The appliance scheduling module, which employs a constraint-based reasoning engine to schedule appliances based on different user preferences;
- The AUTHENTIC graphical user interface (GUI), which allows the occupant to interact with all facets of the system.

The following sections provide an overview over each of these modules, before a more in-depth look is taken at some of the enabling technologies of the deductive component, which is the focus of this paper.

Figure 1. AUTHENTIC System Architecture

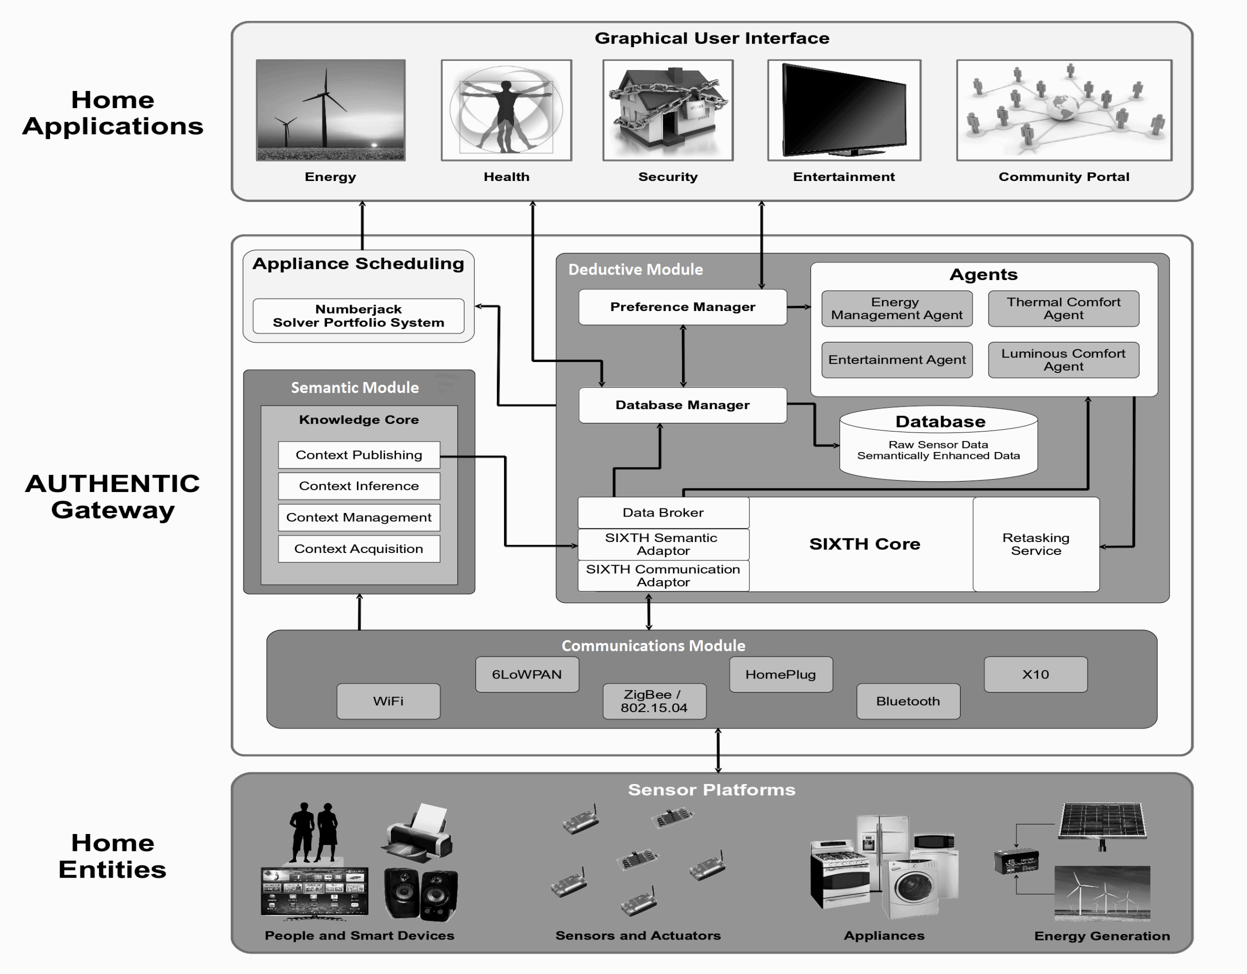

Communications Module

The purpose of the communications module is to hide the inherent heterogeneity of the HAN infrastructure and to provide a unified interface to the other components of the architecture. It uses the publish/subscribe pattern to allow components to register their interest in device messages. When devices (such as appliances, actuators or sensors) are added to the network, or already existing devices sense or actuate, the communications module triggers an event message that is translated into a system-wide common format (JSON) and sent to subscribed components. All requests (e.g. command and get/set state variables) issued to devices within the HAN are similarly translated from JSON into the native format and protocol of the target device. The module currently supports communication via WiFi, Bluetooth, ZigBee, 6LoWPAN and X10. The use of JSON as a common message format means that this communications module can operate in a distributed fashion, for example in situations where remote monitoring of the HAN is desired.

Semantic Module

The semantic module aims to infer contextual information by classifying low-level sensor data gathered from the communications layer and to transform said information into a machine-readable format to be used by the deductive module. The Web Ontology Language (OWL) (McGuinness & Van Harmelen, 2004) is used to automate logical reasoning in order to infer new contextual knowledge from the information gathered. Examples of context inference include whether a room is occupied (e.g. based on data from a passive infrared sensor) or what activity an inhabitant is currently engaged in (e.g. that the inhabitant is making tea based on the kettle being switched on). Once an inference has been made, it is then passed on to the semantic adaptor in the deductive module.

Deductive Module

The deductive module is primarily concerned with decision-making and data management. In order to gather the information about the system upon which reasoning is based, it is necessary to connect to the semantic and communications modules, the latter of which necessitates two-way communication in order to allow actuation commands to be sent to devices also. This is achieved by deploying the SIXTH sensor middleware within the deductive module. SIXTH, which is described in more detail below, enables components within the deductive module to subscribe only to data that is of interest via the SIXTH data broker.

For persistent storage of data, a database manager is deployed within the deductive module. The recorded data is currently used for providing historical usage accounts (e.g. for energy usage reports) but can in future enable the learning of energy usage patterns (Ruzzelli et al., 2010; Heierman & Cook, 2003). The database also facilitates the storage of persistent user preferences, which are set and made available via a preference manager component.

The key decision-making component within the architecture is a multi-agent system based on Agent Factory Micro Edition (AFME) (Muldoon et al., 2006; O'Hare et al., 2012b). Agents are responsible for particular aspects of home management (e.g. adjusting heating or light settings) and do so by combining data received through the SIXTH middleware with user preferences that are set via the user interface. The multi-agent reasoning aspect of the architecture is described in more detail below.

Graphical User Interface

The AUTHENTIC graphical user interface is an application (currently developed for Android smartphones and tablets) that allows users to interact with all aspects of the HAN. Through the GUI, users can set or alter their preferences, such as temperature levels or light levels. Current and historical data associated with devices within the HAN (e.g. the current energy usage of a specific appliance or the overall energy used within a specific room) can also be reviewed. Finally, the GUI provides access to the appliance scheduling feature described in the next section.

Appliance Scheduling

Managing energy usage within the home involves a variety of stakeholders, each of which can have multiple, possibly conflicting, goals. For example, utility providers might want to reduce a household's total energy consumption in order to balance the overall load on the grid. On the other hand, a household owner might want to minimise overall energy cost but without compromising comfort in the home. Appliance scheduling aims to optimise how devices function within the HAN based on time preferences (scheduling windows), energy costs at the time of appliance usage and the energy consumption of schedulable appliances. Users are provided, through the GUI, with detailed schedules for all appliances with the option of overriding the schedules at their own discretion. The primary purpose of said schedules is to ensure that appliances operate in an efficient manner with regard to energy usage and cost while staying within the constraints set by the user.

SIXTH MIDDLEWARE

SIXTH is a Java-based sensor middleware that is capable of sourcing data from a variety of sources (O'Hare et al., 2012a). This incorporates both physical sensing apparatus and also other programmatically-accessible data sources (accessed via *cyber-sensors*). SIXTH features an *adaptor* abstraction that allows sources to be accessed in a consistent manner. Applications running on SIXTH can access this data in a source-agnostic fashion, while also gaining the ability to actuate sensing devices through a *retasking* service. This allows devices to be enabled or disabled, or for their behaviour to be changed (e.g. to change the sensing frequency).

For the purposes of the AUTHENTIC architecture, a number of adaptors were developed in order to interact with the various architectural components of the system (as outlined in Figure 1). These include: a *communications adaptor* (to interact with physical devices) and a *semantic adaptor* (to gain semantically-enhanced data from the semantic module). Both adaptors are outlined in the following sections.

Communications Adaptor

The communications adaptor acts as the gateway between the SIXTH middleware and the communications module. This adaptor allows sensor and status data to be injected into the architecture in a consistent manner. It also allows devices to be actuated through a uniform interface that is agnostic of device type or the underlying communications protocols used. This is done through the SIXTH retasking service. In a typical SIXTH deployment (e.g. those envisaged in Carr et al., 2012; O'Grady et al., 2013), a separate SIXTH adaptor is developed for each type of device desired. However, the architecture illustrated in Figure 1 is intended to demonstrate that the additional features of SIXTH (e.g. the data broker, described below) are equally applicable when some existing effort has been invested in developing a module that is capable of interacting with devices.

In the present architecture, the communications module exposes a number of APIs relating to HAN devices and sensors using a variety of communications protocols. In addition to allowing agents and other components to invoke these APIs, the communications adaptor also adds a level of abstraction to facilitate the reasoning process. For example, one commonly-used home automation device is a dimmer switch. Dimmer switches based on the X10 protocol are typically one-way devices, in that they can receive actuation commands but do not report on their own status. An additional challenge is that commands are available to increase or reduce the light intensity level by a set delta amount, rather than setting the light to a particular level (using *bright* and *dim* commands respectively). Thus in order to set lights to a specific desired level, it is necessary to record the current light level in software, since the device cannot report its current level. The communications adaptor provides an abstraction that allows agents (or other applications) to change light levels to specific values, translating these to the appropriate *bright* or *dim* commands, based on previously recorded light levels, to pass to the communications module.

Semantic Adaptor

The semantic adaptor offers a service similar to the communications adaptor, feeding data from the semantic layer through the SIXTH middleware. This also allows the agents and other applications to gain access to semantic data by registering their interest with the data broker. Unlike the communications adaptor, this is a one-way process, since no additional configuration or retasking of the semantic module is required.

Data Broker

The data broker is the conduit through which data is routed through the SIXTH middleware. All data from the communications and semantic adaptors are sent via the SIXTH core to the data broker for further dissemination. The task of the data broker is to forward only relevant data to interested parties, rather than flooding the system with superfluous and irrelevant data notifications.

Rather than subscribing to a type of message, SIXTH supports the registration of queries with the data broker through a simple Java interface. The registration of a query is associated with the component that created it, so that data received by the data broker that is a match to any registered queries is forwarded to the component associated with a particular query. There is no restriction to the number of queries that an individual component may register.

For example, a heating agent may register a query whereby it receives a notification only if a room's temperature falls below (or exceeds) a threshold temperature, in which case it can switch the heating on (or off). Under the simpler publish/subscribe mechanism, the agent would get constant updates from the room's temperature sensor rather than only when it is actually relevant to the agent's decision-making process.

This approach ensures the extensibility of the SIXTH middleware, allowing a diverse range of applications to be deployed on the HAN. Within the architecture outlined in Figure 1, the database manager registers to receive all data flowing through the system so that it can be recorded. Agents within the deliberative module can register to receive only data that is relevant to their particular tasks. Further applications that reside outside the architecture itself are facilitated in the same way.

Database Manager

A MySQL database is used to store information relevant to the HAN. This includes descriptions of all devices and appliances in the network, any data associated with said devices and appliances (e.g. location, sensor readings, state), semantic information and user preferences set through the GUI. The database manager is an OSGi service that acts as a database front-end for a number of components within the HAN that require database access.

One of the more interesting features of the database manager from a user's perspective is the ability to provide dynamic real-time and historical data from devices within the HAN. Users can specify devices they wish to monitor and the relevant information is extracted from the HAN database (e.g. device type, measurement type, measured value and measurement timestamp). This is then displayed in a dynamic (in the case of real-time data) or static graph, as illustrated in Figure 2. Through the AUTHENTIC GUI, users can also request an energy report, which is also provided by the database manager. It presents the user with the energy usage and cost of devices within the system between time periods specified by the user. Both of these features allow users to better understand how energy is being used within their home environment and better manage both their energy consumption and energy cost (Eßer, Kamper, Franke, Möst, & Rentz, 2007).

Figure 2. Dynamic Sensor Graphing

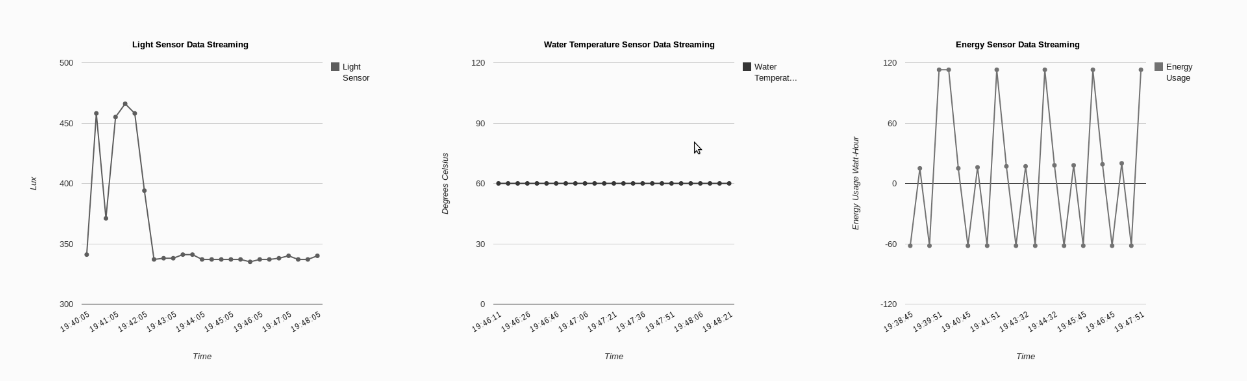

MULTI-AGENT REASONING

The multi-agent reasoning functionality with the HAN has been implemented using the Agent Factory Micro Edition (AFME) agent platform. AFME is a minimised-footprint intelligent agent platform based on the Agent Factory agent development framework (Collier et al. 2003), but designed for use with the Java Micro Edition (JME) Constrained Limited Device Configuration (CLDC). Although, primarily intended for highly constrained devices, applications developed for JME CLDC can also be used on desktop machines running Java Standard or Enterprise Edition. The agents developed for the HAN, at present, are deployed on the HAN Gateway, but it is envisaged that some of these will in future be deployed directly on the sensor or actuation devices to enable inter-device collaboration.

AFME is concerned with the development of computationally reflective agents. Computational reflection is a technique that enables a system to maintain meta-information about itself and to use this information to determine its behaviour. In the agent community parlance, this meta-information is commonly referred to as an agent's belief set or an agent's model of the world. Many intelligent agent platforms, including AFME, draw from folk psychological concepts, such as those identified by the philosopher Daniel Dennett (Dennett, 1987). Specifically, those types of agent platforms employ the notion of the intentional stance as a tool for modelling complex systems through the attribution of mental attitudes, such as beliefs and goals, to agents so as to explain and predict behaviour. According to Dennett, there are three different strategies we use when confronted with an object or system, namely the physical stance, the design stance, and the intentional stance.
To predict the behaviour of an entity according to the physical stance, we use information about its physical constitution along with information about the laws of physics. The physical stance, for example, is employed when we predict the path of a ball in flight.

With the design stance, we assume that the entity in question has been designed for a particular purpose and our predictions are based on the idea that the entity will behave as designed. When someone turns on an electric fan, they predict that it will behave in a certain manner i.e. the fan will cool down the room. They do not need to know anything about the physical constitution of the fan to make the prediction. Predictions made from the design stance are based on two assumptions that the entity is designed for the purpose that the user thinks it to be designed for and that it will perform as designed without malfunctioning. This does not mean that the design stance is always used for entities that have been designed. The physical stance could be used to predict what would happen to the fan if it were knocked onto the floor or if it malfunctioned, but in most cases there is no need to go to a lower level of granularity.

We can often improve our predictions of the design stance by adopting the intentional stance. When making predictions from this stance, we interpret the behaviour of an entity by treating it as a rational agent whose behaviour is governed by mental attitudes. The intentional stance is adopted where it is useful to do so. This is often the case in situations whereby we do not fully understand the design of the system, for example, when considering living organisms. It is less useful when we do understand the inner workings of a particular system. Suppose, for instance, we apply the intentional stance to a doorbell, i.e. we imagine that it is a rational agent that reasons about its beliefs and desires and intends to alert us when someone is at the door. This is not particularly useful because we can understand the functionality of a doorbell in simpler physical or mechanical terms. In contrast, suppose we wish to explain or predict the behaviour of a person or complex computer system. In such cases, it is necessary to form a higher level of abstraction if we do not fully understand their inner workings or design. The intentional stance can be applied to anything, but it is more practical to use when it leads to simpler descriptions than would otherwise be available.

The behaviour of agents in AFME is represented using declarative antecedent-consequence rules that determine the conditions under which commitments are adopted and actions are performed. To facilitate this, the conditions are matched against belief sets (meta-information maintained by agents) at periodic points throughout execution.

The HAN multi-agent architecture comprises a set of agents, a set of actuators, a set of perceptors, and the HAN service, which is a class that enables agents to interact with the SIXTH middleware. At present, within the HAN, there are six agents: the Luminosity Comfort Agent, the Thermal Comfort Agent, the Security Agent, the Entertainment Agent, the Energy Management Agent, and the Schedule Agent. The Thermal Comfort Agent ensures that temperature levels in the room match the user preferences. The Luminosity Comfort Agent acts in a similar manner, but with regard to light. The Security Agent informs the user when a security event is triggered, while the Entertainment Agent controls the application behaviour when the user is sitting down and watching television. The Energy Management Agent proactively heats water in anticipation of user behaviour, while the Schedule Agent informs the user when water has been wasted and asks the user to change their scheduling preferences.

Perceptors were developed to enable agents to receive information in relation to sensor data, user preferences, deductions arising from the semantic module, and the user's schedule. This information is obtained through the AUTHENTIC service, which is a SIXTH receiver that registers with the SIXTH data broker and preference manager when the application begins to operate. Actuators were developed to enable agents to send commands to control devices in the HAN. The actuators also use the AUTHENTIC service in delivering this functionality. In particular, the actuators make use of the SIXTH retasking service.

USE CASE SCENARIO

In order to demonstrate the functionality of the HAN, the system was deployed in a physical environment and a use case scenario was envisioned to illustrate the effectiveness of the HAN system in regards to internal stimuli. A number of appliances, sensors and actuators were deployed in room (see Figure 3) in order to simulate a living room.

When an individual enters the room, the passive infrared (PIR) sensor fires and an occupancy message is sent to the deductive module. The middleware then retrieves the user-set light and temperature levels from the preference manager. The lights in the room are turned on and set to the preferred level. If the current temperature of the room is below the user's preferred level, the heating is also turned on until the room temperature reaches said level. If the heating is turned on, the user is notified of this through the GUI.

If the user sits on the couch, the pressure sensor sends a message through the communications module and the smart TV is turned on. At the same time, the lights in the room are dimmed to the user's preferred level for this activity. If the user stands, the smart TV turns off and the lights return to their previous luminescence levels.

When the window in the room is opened or closed, the contact sensor fires and sends a stated changed message through the communications module to the deductive module. If the state of the sensor is open then the user is notified through the GUI that the window has been opened.

Figure 3. Sensors involved in the use case scenario: (1) HAN, (2) PIR sensor, (3) temperature & luminescence sensor, (4) radiator, (5) smart TV with smart plug, (6) couch with pressure sensor, (7) lamp with dimmer switch and (8) window with contact sensor.

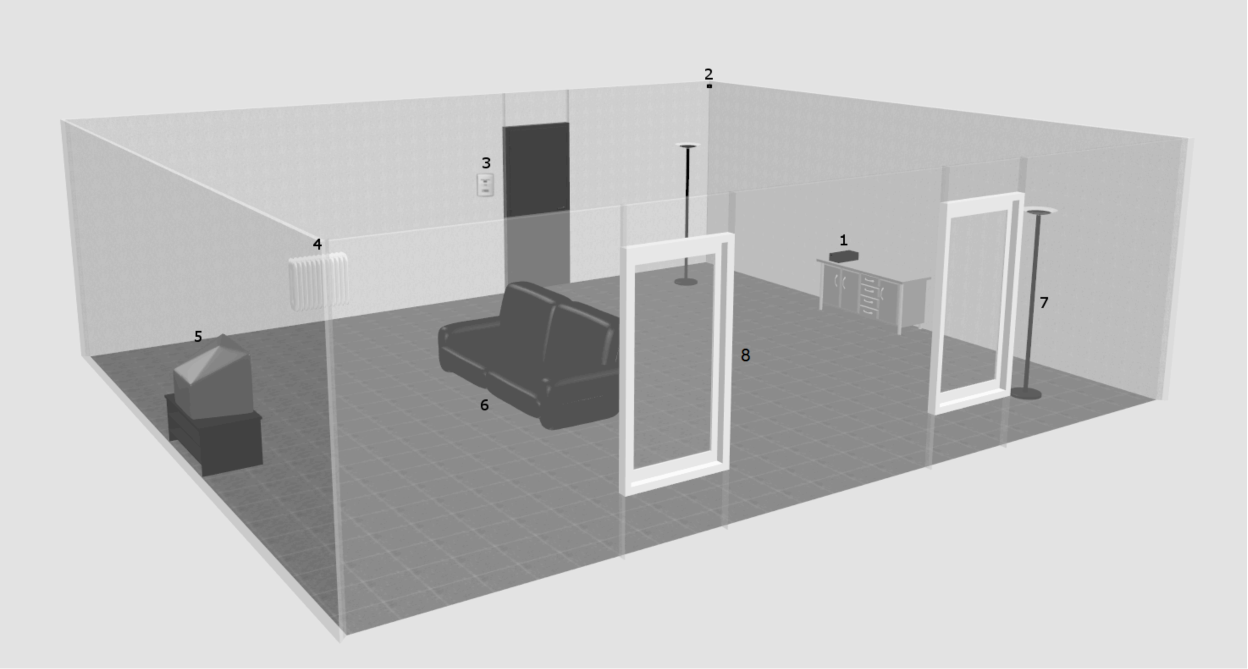

RELATED WORK

Current HAN automation research is often limited to one type of technology, most often ZigBee or WiFi (e.g. Gill et al., 2009; Han & Lim, 2010), or power line communication (Cook et al., 2006, Son et al., 2010). However, as Parikh and colleagues (2010, see also Güngör et al. 2011) note, different technologies offer different benefits and drawbacks. Therefore, heterogeneity must be assumed to be inherent to any HAN, which is why AUTHENTIC was deliberately designed to handle a multitude of technologies, protocols and data formats by abstracting from these details and providing a unified interface to higher-level services.

Furthermore, most research on home energy management is focused on providing feedback about energy usage to consumers rather than proactively managing energy consumption. Jahn and colleagues (2010), for example, developed a smart home system on top of the OSGi-based Hydra middleware (Eisenhauer et al., 2010). Similar to the communications module in AUTHENTIC, Hydra provides a unified interface to heterogeneous networks of embedded devices irrespective of communication protocol. The developed system harnesses smart plugs to measure energy consumption on a per-device basis and displays this information *in situ* via an augmented reality app on the home owner's smart phone whenever its camera is pointed at a device. Apart from the restriction to smart plugs as the sole type of sensor deployed, the only example of home automation within the system involves the home owner setting a maximum energy price, below which the (virtual) washing machine is switched on. In stark contrast, AUTHENTIC incorporates a range of sensors (e.g. occupancy, temperature, luminosity and humidity sensors) and controllable devices (e.g. smart plugs, controllable thermostats and dimmer switches) to allow autonomic energy management of the whole home.

Zhang and colleagues (2005) employ OSGi and agents to develop a so-called "control system architecture for smart homes", but treat an individual home as the smallest unit of interest within the wider smart grid rather than actually specifying the internal operation an individual smart home. Further, Zhang and colleagues assert that "[t]he Home Gateway should be designed to have a high availability and deliver services with a specified level of determinism. These requirements necessitate remote management and monitoring Home Gateway services and home devices." AUTHENTIC, instead, puts full control of the internal HAN management in the hands of the user. All data remains local and all autonomic decision-making is carried out in line with user-specified preferences and constraints. We believe this degree of user control to be vital in alleviating legitimate privacy concerns and engender trust within the user (Lui et al., 2010; Krishnamurti et al., 2011).

Zhao and colleagues (2010) present a conceptual framework for energy management in both residential and commercial buildings, which also takes a multi-agent approach. However, while their proposed framework shows obvious parallels to the AUTHENTIC architecture, it also takes only a very high-level approach to the problem, for example by assuming "[a]n underlying layer of communication infrastructure [...] to be in place that interconnects the control systems with the controlled entities." This layer forms an integral part of AUTHENTIC in the shape of the communications module. Furthermore, Zhao and colleagues assume the individual agents to be in charge of collecting data as well as controlling appliances. AUTHENTIC, instead, amalgamates all the data, including contextual information provided by the semantic module, in one place (the SIXTH data broker) and then disperses only information directly relevant to a specific agent.

Similar to AUTHENTIC's appliance scheduling facility, Abras and colleagues (2010) present a multi-agent home automation system that has individual agents negotiate execution schedules on behalf of the appliances they represent (e.g. heater, washing machine) in order to minimise overall energy consumption. Agents take into account predicted power consumption as well as user comfort and are mapped one to one onto an appliance. By contrast, scheduling of appliances in AUTHENTIC is handled by a constraint-based reasoning system, whereas agents, rather than controlling access to a singular appliance, deliver services that include potentially multiple appliances and, more importantly, are based on intelligent decisions borne out of sensor data gathered from within and outside the home. This approach also allows for the easy integration of additional services not related to energy management but still dependent on the HAN infrastructure (e.g. security, assisted living and entertainment).

Closest to the spirit of AUTHENTIC is probably MavHome (Cook et al., 2006), an agent-based smart home. MavHome exhibits a similar architecture to AUTHENTIC, with a decision layer (the equivalent of the deductive module), an information layer (semantic module), a communication layer (communications module) and a physical layer (home entities). However, rather than constituting an overall architecture for the system, every agent in MavHome is realised as an entity consisting of these four layers. The physical layer, corresponding to the perceptors and actuators in an AFME agent, represents physical sensors and appliances as well as potentially other agents. The communications layer is responsible for connecting the agent's higher levels to the physical layers, while the information layer generates contextual knowledge from raw sensor data. The decision layer, finally, selects which actions to perform based on this knowledge.

Separating these layers into distinct individual modules, rather than instantiating them on every agent deployed in the system, allows agents in AUTHENTIC to gather information from multiple sensor sources in a centralised fashion through the SIXTH data broker and its query interface, thereby minimising the amount of information each agent has to process. Furthermore, abstracting communication technologies through a dedicated module implies that adding a new protocol involves only an update to the one module rather than every single agent. Likewise, devices added to the system, as well as inferences added to the semantic module, are automatically available to all agents, ensuring AUTHENTIC's extensibility.

In summary, while the constituent technologies and approaches of the AUTHENTIC system might have been evident in the prior literature, this project has succeeded in delivering a complete and integrated solution that is uniquely adapted to the delivering autonomic intelligent energy control in a smart home environment.

CONCLUSION

Achieving effective energy management within the home is an objective shared by home occupants and government agencies alike. How best to meet this objective is still open to question, however. The approach proposed in this chapter advocates the augmentation of HAN technologies with a physical sensing infrastructure. Such an infrastructure is low-cost, as it harnesses off-the shelf components and can avail itself of, and co-exist, with a pre-existing HAN. Though smart grid technologies are in their infancy in terms of consumer adoption, harnessing any services enabled by such grids is essential going forward.

Many challenges must be overcome before homeowners will adopt such technologies at mass scale. These are not confined solely to system issues; rather those of privacy, security and trust must also be addressed. Introducing systems, such as those envisaged by AUTHENTIC, radically challenge perceptions and expectations of how energy management is delivered and experienced within the home. Should these be perceived in a negative fashion, adoption will be slower and insufficient. Many efforts to date have been technology-centric. If effective energy management within homes is to be realised, a more human-centric approach must be enabled, and delivered.

FUTURE WORK

One of the primary aims of the AUTHENTIC project going forward is to enable integration with the smart grid, for example, to allow agents to intelligently negotiate tariffs with utility providers in order to feed micro-generated power into the grid, or to adapt HAN-internal consumption to enable utility providers to balance load and reduce peak usage over a large number of customers (Niyato et al., 2011; Lui et al., 2010).

With future plans of integrating machine learning technologies into the current system, the HAN will not only be able to react via a set of preconceived rules but also to learn based on users' activities (Ruzzelli et al., 2010; Heierman & Cook, 2003). It is hoped that this added layer of intelligence will be of great benefit to future users of the system.

Within the AUTHENTIC project, the HAN is viewed as a key enabling technology that should be able to accommodate the coexistence of multiple applications. These may consist of a variety of services, for example, ambient assisted living (Korhonen et al., 2013), home security and home entertainment (Messer et al., 2006), alongside the presented energy management functions.

AUTHENTIC thus seeks to future-proof such requirements within its design by framing it within the key drivers of cost, ease of deployment, ease of use, adherence to standards and the need for it to be an ambient technology offering.

ACKNOWLEDGEMENTS

The support of Enterprise Ireland under grant TC20121002A is gratefully acknowledged. In addition this work is supported by Science Foundation Ireland under grant 07/CE/11147.

REFERENCES

Abras, S., Ploix, S., Pesty, S., & Jacomino, M. (2008). A multi-agent home automation system for power management *Informatics in Control Automation and Robotics* (pp. 59-68): Springer.

Carr, D., O'Grady, M. J., O'Hare, G. M., & Collier, R. (2013). SIXTH: A Middleware for Supporting Ubiquitous Sensing in Personal Health Monitoring *Wireless Mobile Communication and Healthcare* (pp. 421-428): Springer.

Choi, T.-S., Ko, K.-R., Park, S.-C., Jang, Y.-S., Yoon, Y.-T., & Im, S.-K. (2009). *Analysis of energy savings using smart metering system and IHD (in-home display).* Paper presented at the Transmission & Distribution Conference & Exposition: Asia and Pacific, 2009.

Collier, R., O'Hare, G., Lowen, T., & Rooney, C. (2003). Beyond prototyping in the factory of agents *Multi-Agent Systems and Applications III* (pp. 383-393): Springer.

Cook, D. J., Youngblood, M., & Das, S. K. (2006). A multi-agent approach to controlling a smart environment *Designing smart homes* (pp. 165-182): Springer.

Dennett, D. C. (1987). *The intentional stance*: The MIT press.

Eisenhauer, M., Rosengren, P., & Antolin, P. (2010). Hydra: A development platform for integrating wireless devices and sensors into ambient intelligence systems *The Internet of Things* (pp. 367-373): Springer.

Eßer, A., Kamper, A., Franke, M., Möst, D., & Rentz, O. (2007). Scheduling of electrical household appliances with price signals *Operations Research Proceedings 2006* (pp. 253-258): Springer.

Faruqui, A., Sergici, S., & Sharif, A. (2010). The impact of informational feedback on energy consumption—A survey of the experimental evidence. *Energy, 35*(4), 1598-1608.

Gill, K., Yang, S.-H., Yao, F., & Lu, X. (2009). A ZigBee-based home automation system. *Consumer Electronics, IEEE Transactions on, 55*(2), 422-430.

Gungor, V. C., Sahin, D., Kocak, T., Ergut, S., Buccella, C., Cecati, C., & Hancke, G. P. (2011). Smart grid technologies: communication technologies and standards. *Industrial informatics, IEEE transactions on, 7*(4), 529-539.

Han, D.-M., & Lim, J.-H. (2010). Smart home energy management system using IEEE 802.15. 4 and ZigBee. *Consumer Electronics, IEEE Transactions on, 56*(3), 1403-1410.

Hazas, M., Friday, A., & Scott, J. (2011). Look back before leaping forward: Four decades of domestic energy inquiry. *Pervasive Computing, IEEE, 10*(1), 13-19.

Heierman III, E. O., & Cook, D. J. (2003). *Improving home automation by discovering regularly occurring device usage patterns.* Paper presented at the Data Mining, 2003. ICDM 2003. Third IEEE International Conference on.

Hussain, S., Schaffner, S., & Moseychuck, D. (2009). *Applications of wireless sensor networks and RFID in a smart home environment.* Paper presented at the Communication Networks and Services Research Conference, 2009. CNSR'09. Seventh Annual.

Jahn, M., Jentsch, M., Prause, C. R., Pramudianto, F., Al-Akkad, A., & Reiners, R. (2010). *The energy aware smart home.* Paper presented at the Future Information Technology (FutureTech), 2010 5th International Conference on.

JSON. JavaScript Object Notation. Retrieved 24/11, 2013, from http://www.json.org/

Kluckner, P. M., Weiss, A., Schrammel, J., & Tscheligi, M. (2013). Exploring Persuasion in the Home: Results of a Long-Term Study on Energy Consumption Behavior *Ambient Intelligence* (pp. 150-165): Springer.

Korhonen, I., Parkka, J., & Van Gils, M. (2003). Health monitoring in the home of the future. *Engineering in Medicine and Biology Magazine, IEEE, 22*(3), 66-73.

Krishnamurti, T., Schwartz, D., Davis, A., Fischhoff, B., de Bruin, W. B., Lave, L., & Wang, J. (2012). Preparing for smart grid technologies: A behavioral decision research approach to understanding consumer expectations about smart meters. *Energy Policy, 41*, 790-797.

Lui, T. J., Stirling, W., & Marcy, H. O. (2010). Get smart. *Power and Energy Magazine, IEEE, 8*(3), 66-78.

McGuinness, D. L., & Van Harmelen, F. (2004). OWL web ontology language overview. *W3C recommendation, 10*(2004-03), 10.

McKerracher, C., & Torriti, J. (2012). Energy consumption feedback in perspective: integrating Australian data to meta-analyses on in-home displays. *Energy Efficiency*, 1-19.

Messer, A., Kunjithapatham, A., Sheshagiri, M., Song, H., Kumar, P., Nguyen, P., & Yi, K. H. (2006). *InterPlay: a middleware for seamless device integration and task orchestration in a networked home.* Paper presented at the Pervasive Computing and Communications, 2006. PerCom 2006. Fourth Annual IEEE International Conference on.

Muldoon, C., O'Hare, G. M., Collier, R., & O'Grady, M. J. (2006). Agent factory micro edition: A framework for ambient applications *Computational Science–ICCS 2006* (pp. 727-734): Springer.

Niyato, D., Xiao, L., & Wang, P. (2011). Machine-to-machine communications for home energy management system in smart grid. *Communications Magazine, IEEE, 49*(4), 53-59.

O'Grady, M. J., Murdoch, O., Kroon, B., Lillis, D., Carr, D., Collier, R. W., & O'Hare, G. M. (2013). *Pervasive Sensing: Addressing the Heterogeneity Problem.* Paper presented at the Journal of Physics: Conference Series.

O'Hare, G. M. P., Collier, R. W., Dragone, M., O'Grady, M. J., Muldoon, C., & Montoya, A. D. J. (2012b). Embedding Agents within Ambient Intelligent Applications. In T. Bosse (Ed.), *Agents and Ambient Intelligence* (pp. 119-135): IOS Press.

O'Hare, G. M. P., Muldoon, C., O'Grady, M. J., Collier, R. W., Murdoch, O., & Carr, D. (2012a). Sensor web interaction. *International Journal on Artificial Intelligence Tools, 21*(02).

ON World (2010). The Smart Energy Home Market to Reach $3 Billion in 2014. Retrieved from http://www.onworld.com/news/newssmartenergy.htm

Parikh, P. P., Kanabar, M. G., & Sidhu, T. S. (2010). *Opportunities and challenges of wireless communication technologies for smart grid applications.* Paper presented at the Power and Energy Society General Meeting, 2010 IEEE.

Ruzzelli, A. G., Nicolas, C., Schoofs, A., & O'Hare, G. M. (2010). *Real-time recognition and profiling of appliances through a single electricity sensor.* Paper presented at the Sensor Mesh and Ad Hoc Communications and Networks (SECON), 2010 7th Annual IEEE Communications Society Conference on.

Son, Y.-S., Pulkkinen, T., Moon, K.-D., & Kim, C. (2010). Home energy management system based on power line communication. *Consumer Electronics, IEEE Transactions on, 56*(3), 1380-1386.

Zhang, H., Wang, F.-Y., & Ai, Y. (2005). *An OSGi and agent based control system architecture for smart home.* Paper presented at the Networking, Sensing and Control, 2005. Proceedings. 2005 IEEE.

Zhao, P., Simões, M. G., & Suryanarayanan, S. (2010). *A conceptual scheme for cyber-physical systems based energy management in building structures.* Paper presented at the 9th IEEE/IAS International Conference on Industry Applications (INDUSCON), 2010.