

An Agent-based Domestic Electricity Consumption Advisory System

Gregory M. P. O'Hare
CLARITY: Centre for Sensor
Web Technologies
University College Dublin
gregory.ohare@ucd.ie

Mauro Dragone
CLARITY: Centre for Sensor
Web Technologies
University College Dublin
mauro.dragone@ucd.ie

Antonio Ruzzelli
CLARITY: Centre for Sensor
Web Technologies
University College Dublin
ruzzelli@ucd.ie

Conor Muldoon
Computing Laboratory
Oxford University
Oxford, UK
conor.muldoon@
comlab.co.uk

Richard Tynan
CLARITY: Centre for Sensor
Web Technologies
University College Dublin
richard.tynan@ucd.ie

Michael J. O'Grady
CLARITY: Centre for Sensor
Web Technologies
University College Dublin
michael.j.ogrady@ucd.ie

ABSTRACT

This paper introduces an agent-based domestic electricity consumption advisory system. It reflects upon the difficulties of realizing the ubiquitous sensing vision which underpins such systems. It advocates the need for an effective middleware which will support the evolution of heterogeneous, distributed, collaborative intelligent sensing artifacts. To this end, it introduces the SIXTH Middleware.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems

General Terms

Design

Keywords

Energy consumption, Intelligent agents

1. INTRODUCTION

In a world of ever depleting natural resources and increasing CO₂ emissions various initiatives have begun in the Information Technology sector that address the need for Green Technologies. This is typified by initiatives like IBM's Big Green and HP's Green Business Technology Initiative. Big Green is a 1 billion dollar IBM offering to dramatically increase the efficiency of IBM products and services, which has diversified into a green IBM agenda. It is universally recognized that carbon emissions need to be significantly reduced. Indeed carbon taxes will introduce new carbon credit trading centres and emerging business models. Within this paper, we do not address the corporate demand for intelligent green technologies. Such technologies assist in sensing carbon emissions and advise as to remedial process interventions that reduce emission rates and identify causal consequences of such process interventions. Rather we address the challenging domestic market where unit cost to the householder presents a challenging cost benefit equation.

1.1 Advising the Consumer in an ever Changing Market

At present, within Europe, a varying number of tariff bands exist in member states for electricity charging. In Ireland, there are currently two, but this is soon to be expanded to three. In other member states, there are five and soon to be seven. In this rapidly changing landscape, it is becoming increasingly difficult for domestic consumers to understand charging models and more specifically how to remedy household behaviour so as to achieve reductions in charges. Recent standardization efforts have generated an increasing trend towards the integration of sensor systems in Building Automation Systems. 802.15.4, ZigBee, and 6LowPan represent key enabling technologies that allow connecting low-cost sensing and monitoring units and that gather energy consumption information in real-time.

In order to achieve this, it is necessary to recognize the activation of individual appliances and to sense the differential energy consumption contributed by that specific appliance. One such system entitled REAR, developed within Clarity, seeks to do exactly this by profiling individual appliances. This is based on a portfolio of parameters including: real power, power factor, RMS current, RMS voltage, peak current, and peak voltage.

2. ENERGY MONITORING

Broadly speaking, two approaches exist in the sensing and monitoring of household electricity consumption.

1. The deployment of plug in domestic socket based sensors that record the switching on and off of a given socket and capture the sensed portfolio described above for the given appliance;
2. The deployment of a single sensor at the fuse board that records the aggregated energy consumption at a given instance;

The energy monitoring system discussed in this paper explores both approaches. The later lends itself to a non-intrusive, cost effective, centralised approach to the problem of appliance recognition. The former meanwhile lends itself

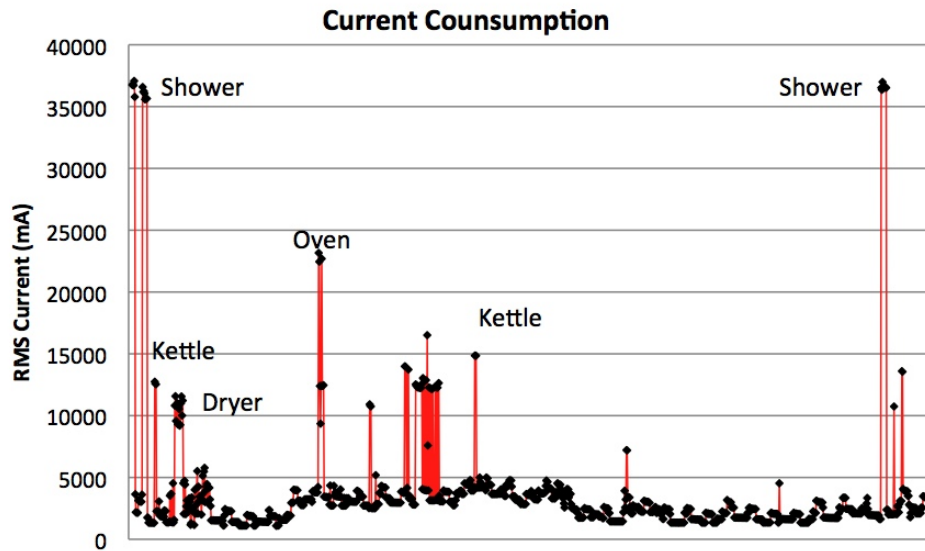


Figure 1: Domestic Energy Consumption

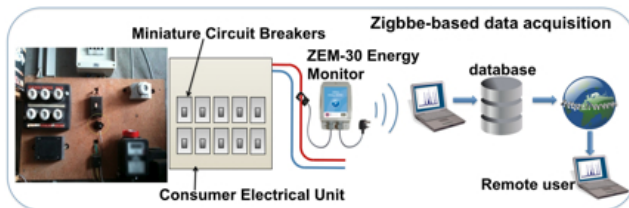


Figure 2: ZigBee-based Data acquisition System

to a more distributed multi-agent systems solution that is more expensive in terms of sensor deployment. The current incarnation of REAR uses machine learning techniques to determine unique appliance signatures¹.

Figure 2 depicts an off-the-shelf data acquisition system based on ZigBee. It consists of an Episensor ZEM-30 energy monitor clipped around the live wire of the consumer electrical unit. The energy monitor transmits periodic energy data to a local gateway, which then forwards the information to a remote database.

A stepping-stone for users to reduce electricity expenditure is to provide a visual interface that displays electricity consumption in real-time. This is achieved through the electrical appliance activity recognition system. This information enables the electricity bill to be decomposed in terms of device energy consumption; it will enable energy savings should the users decide to change their behaviour based on the recorded data. Figure 1 shows typical domestic energy consumption over a certain period of time with some appliance activity annotation. The main aim of this research is to attribute appliance names to each of the energy spikes in real time. It is interesting to note how certain appliances stand out and can immediately be identified whilst other appliances have similar or low energy consumption patterns and

¹In this paper, the signatures are unique in terms of household usage patterns of a particular appliance type rather than uniqueness in terms of make and model.

require some background intelligence to differentiate them.

Some initial experiments based on Episensor energy monitors have demonstrated that an appliance generates a unique set of parameters based on the way it draws current from the mains. For example, Figure 3 shows the signature pattern that is generated by an electric oven over a few minutes of operation. The oven generates a saw tooth waveform pattern when it reaches the temperature chosen by the user and it starts switching on and off periodically due to the thermostat. The idea is to combine this information together with other data collected from a smart meter and use it to create *unique appliance signatures* to recognize when and what type of appliances are turned ON/OFF in real-time. Again, in this situation, by unique appliance signature, we mean unique in terms of appliance type (e.g. oven rather than kettle) and unique in terms of user or household usage patterns. Information related to appliances, such as power consumption, is recorded by the system and stored in a database. REAR uses a neural network that is trained on data downloaded from the database to identify appliance signatures. The advantage of using neural networks is that they have a high tolerance to noise and can classify patterns on which they have not been a priori trained. Figure 4 shows a system architecture of the REAR. REAR includes a data acquisition module, a reasoning module, and a user interface that provides user with information related to appliance usage.

The user interface of the system enables the user to click on each active appliance and observe its overall cost. The real-time information can be used not only for energy consumption estimation, but also for real-time energy peak levelling, user activity recognition, and faulty appliance detection.

Within the current system, three agents cooperate in the identification of a given appliance. A monitoring agent recognizes the activation of a new appliance and senses and extracts the required parameters from the background energy landscape and stores these appropriately. An appliance recognition agent uses similarity metrics to compare a given

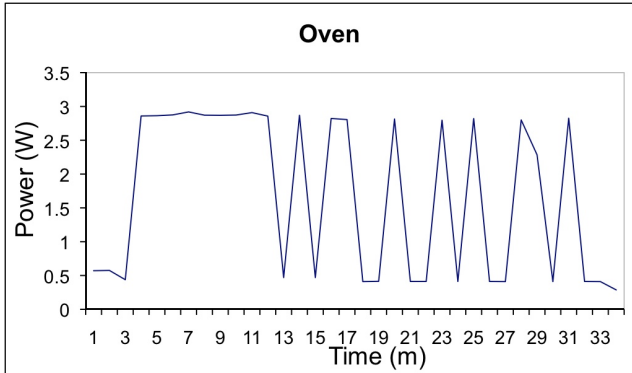


Figure 3: Appliance signature for an electric Oven

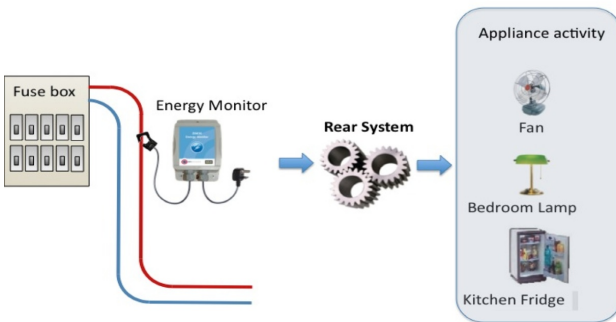


Figure 4: REAR: Real time Appliance Recognition

Level	Requirements
BASIC (1)	Data Access / Discovery / Routing
EXTENDED (2)	Messaging / Monitoring / Notification
SERVICE (3)	SIXTH Service Registry / Services
AGENT (4)	Agent Enabled

Figure 5: A SIXTH Sensing Device Classification Hierarchy

appliance signature with a library of known appliances. A consumption recommender system, which based upon profiles of consumption, advises as to ways in which more timely activation of appliances could effect cost savings.

The second approach envisages a ubiquitous sensed infrastructure on a variety of sensing devices which in addition to monitoring individual sockets may monitor ambient temperature door closures and so forth within the home, and based upon this more complete data set, advise as to behaviour change in electricity consumption. To effectively achieve agent controlled sensor networks, places some considerable demands [3] on the system and in the remainder of this paper we introduce a sensor based middleware that supports this vision.

3. AN AGENT BASED ARCHITECTURE

The architecture that we adopt in this work partitions the network according to the capabilities of the devices (See Figure 5). At the top level, devices with the most resources co-operate using the SIXTH middleware and intelligent agents to deliver system level adaptation. At the intermediate level are devices capable of supporting agents but not the middleware. Finally, are devices which operate using a basic embedded Operating System, such as TinyOS or Contiki.

3.1 SIXTH Middleware

Current middleware offerings for sensor networks tend to encompass systems that can aid application development to any system between the application and the operating system of the sensor. While this broad definition includes many solutions, there are a number of common goals: to facilitate application development and to provide the capability of deploying a sensor network with little in depth knowledge of WSNs. The SIXTH Middleware, as illustrated in Figure 6, is in keeping with the existing goals of WSN middleware but also provides additional capabilities not currently present: modularity, flexibility, reusability, openness, extensibility, universality, and multiple abstractions.

One of the primary goals of SIXTH is to facilitate the creation of reusable components capable of intelligent behaviour, which can be brought together to form one or more applications. The SIXTH middleware is therefore the collection of components along with their integration framework, which in our case is OSGi. On a macro level this means that the application uses the middleware components and each component is part of many applications. Given the

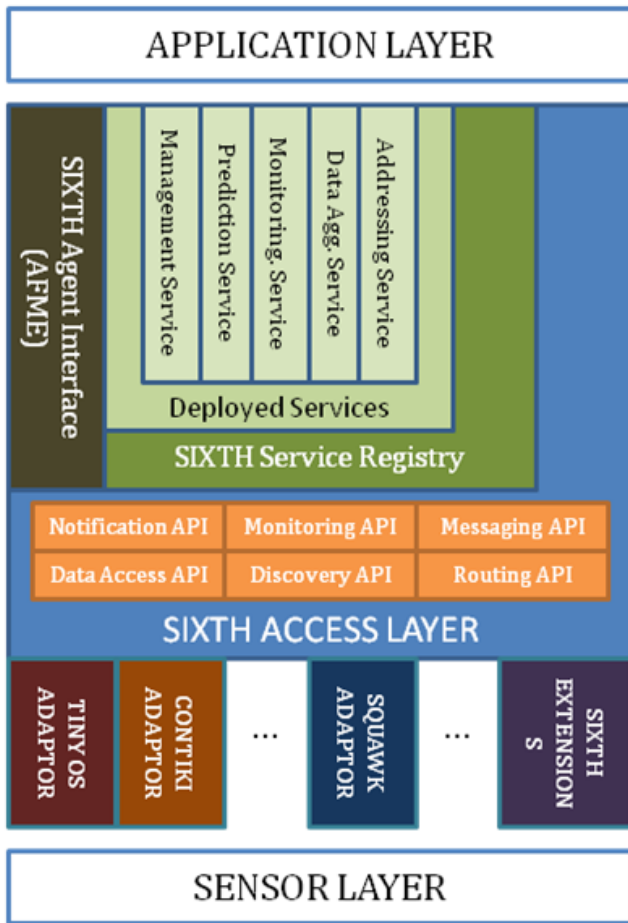


Figure 6: The SIXTH Middleware Architecture

uniqueness of the sensor devices, it is important that many applications coexist on the same sensor infrastructure and we achieve this through the use of multiple agents.

3.2 Agent Capable Devices

When creating applications for devices with sufficient processing capabilities, such as the Sun SPOT, SIXTH provides support for the development and deployment of intelligent agents. Specifically, it supports the construction of Agent Factory Micro Edition (AFME) agents [2]. AFME represents a minimised footprint agent platform that facilitates the execution of intelligent agents on resource constrained devices. Broadly speaking, it conforms to the semantics of the Agent Factory Agent Programming Language [1]. Agents adopt beliefs, which represent their model of the world, and use these beliefs to determine a set of commitments. Commitments represent an intended course of action and will be revised at various points throughout execution as circumstances change.

AFME incorporates functionality for intelligent scheduling, dynamic role adoption, and resource bounded reasoning. These capabilities provide agents with the capability to dynamically alter their behaviour and computational overhead in accordance with quality of service (QoS) requirements, such as application lifetime requirements. For instance, when power is running low on nodes, agents will

reduce their computational overhead by executing their reasoning algorithm less frequently. Additionally, using the resource-bounded reasoning capabilities, agents will drop commitments with a low priority or utility value.

3.3 Leaf Nodes

For data collection and network configuration SIXTH defines a standardized API. System specific implementations of this API exist, including ports for TinyOS. For instance, the TinyOS implementation utilizes the multi-hop collection tree protocol (CTP) and the low power listening (LPL) access control defined within TinyOS. In each case, the low-level communication (e.g. the serial interface) is encapsulated by an OSGi component implementing the SIXTH API, which can be used to interact with all the motes communicating with a gateway from any Java-enabled device connected to it. By simply introducing component's configuration variables mapping to SIXTH's messages, it is possible to extend the actuator/perceptor capabilities of the AFME agents. Specifically, by changing the value of these variables, agents can send SIXTH request packets from the gateway to the network. Each request packet can be broadcast to the entire network or unicast to individual nodes. Common requests allow the instructing of individual nodes as to which data to transmit and furthermore to set important operative parameters, such as sleep duty cycle and modality of the data transmission, for instance, with commitment rules such as the following:

```
Belief(temp(?moteX, ?T)) & Belief(greaterThan(?T,25))
=>
COMMIT(configure(gateway-component, duty_cycle(20)))
```

In particular, each mote can be instructed to send data periodically, at a given sampling rate, or on an event basis. For the latter, special query requests are sent, each consisting of a subject, a verb, a value, and an epoch, concatenated with logic connectors. The subject allows for the selection of packet fields like node ID, sensor reading or link quality. The verb represents the verifying condition while the epoch denotes the duration of the event before the sentence is validated. The Logic connectors interlace sentences to formulate composite queries such as “if a node is within certain xy coordinate and temperature is greater than z then report”. After the composite query is injected into the network, the component sets a countdown timer and waits for notification from the nodes before raising the relevant event to the agent in charge of the component. Sensor nodes are provided with an NesC interpreter module to decode the sentence and activate the enquired mode.

4. CONCLUSIONS

This paper introduced on-going research that embraces an agent-based metaphor in the monitoring and the subsequent provision of advice to domestic electricity consumers. It explores the vision of ubiquitous sensing and how intelligent agents offer a candidate technology for the effective management of wireless sensor networks. To this end, the SIXTH Middleware is introduced as a catalyst to support the graceful expansion of the network through the support of sensor hosted agents.

5. ACKNOWLEDGMENTS

This work is supported by Science Foundation Ireland under grant 07/CE/I1147

6. REFERENCES

- [1] R. W. Collier. Agent Factory: A Framework for the Engineering of Agent-Oriented Applications. *Ph.D. Thesis*, 2001.
- [2] C. Muldoon, G. M. P. O. Hare, R. W. Collier, and M. J. O. Grady. Agent Factory Micro Edition: A Framework for Ambient Applications. In *Intelligent Agents in Computing Systems*, volume 3993 of *Lecture Notes in Computer Science*, pages 727–734, Reading, UK, 28-31 May 2006. Springer.
- [3] R. Tynan, C. Muldoon, G. O. Hare, and M. O. Grady. Coordinated Intelligent Power Management and the Heterogeneous Sensing Coverage Problem. *The Computer Journal*, page bxp109, 2009.