An Agent-Based Infrastructure for Energy Profile Capture and Management

Julian Padget
Harpreet Riat
Dept. of Computer Science
University of Bath
Claverton Down, Bath, UK
jap@cs.bath.ac.uk

Martijn Warnier
Frances M.T. Brazier
Dept. of Technology, Policy and Management
Delft University of Technology
Delft, NL
m.e.warnier@tudelft.nl
f.m.brazier@tudelft.nl

Sukumar Natarajan
Dept. of Architecture and Civil Engineering
University of Bath
Claverton Down, Bath, UK
s.natarajan@bath.ac.uk

ABSTRACT

Accurately and dynamically monitoring energy usage patterns in households forms a first requirement for more efficient and eco-friendly energy management in the future. Monitored energy usage data can be used by power systems engineering—to inform demand-side management systems in the near future term—and in architecture/civil engineering—where it can be used to carry out long-term studies across populations and sectors to estimate future demand and to evaluate prospective (social) policies. This paper presents an agent based prototype of an architecture to meet these needs. The proposed system remains flexible to new functional requirements and adaptable to new edge devices for data collection, as well as offering the potential to 'close the loop' and permit remote control of power supplies to individual appliances. Some preliminary analyses of data collected is used to illustrate what may be possible in the longer term.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Management, Measurement

Keywords

multi-agents systems, energy monitoring, sensor networks

1. INTRODUCTION

Energy has become a major pre-occupation of governments in recent years. Sustainable energy usage by reduction of energy consumption, primarily through more efficient utilization, forms a particular focus. Citizens initially observe the issue through the effect on household energy bills, but there is also rising awareness of both the need for greater efficiency and concerns about energy security at both individual and national levels.

Thus, there are two drivers for reducing energy consumption:

1. At the micro-level, individuals are concerned about day-to-day costs, as well as convenience and individual comfort

2. At the macro-level, regional or national authorities are concerned about total carbon footprint and its evolution over decade and longer time-frames, as well as (i) availability and affordability of energy and (ii) the security of supply in the short and long term.

Examples in the energy domain where both micro and macro levels are of importance include the improved understanding of occupant operation of window systems [7], optimization of micro-generation technologies to complement network load [16, 27], appliance to electricity network communication to allow Energy Service Companies (ESCOs) to help schedule appliance operation [12] and smart metering to exploit occupant energy use awareness benefits [2, 6]. Centralized or remote operation of home appliances has also emerged under the banner of 'home automation' as the number and complexity of appliances in the home has increased [26]. Broadly, home automation excluded, these efforts are driven by the desire to improve energy conservation and energy security through reducing what we might term 'occupant related losses'. In many cases, these efforts also align well with carbon emission reduction targets.

A key starting point to address each of these perspectives, is the collection of data about actual household consumption. Such data might then be used *on-line* as part of a control system for individual household appliance energy management. On-line monitoring of energy usage also allows the households to shift energy demand [19, 23, 24] to different time periods in response to finer-grained time-dependent pricing (as part of the move to supply-led rather then demand-led generation). Virtual powers stations [13], where, in a small geographical region, households store over-capacity generated by solar and wind energy and sell this back to the market, forms another area where on-line monitoring is likely to be beneficial.

Additionally, *off-line* monitoring can be utilized to carry out forward simulation of energy requirements and policy analysis using empirical data on consumption across domestic, commercial and demographic populations.

Our aim is to develop the means to collect that data and enable each of the uses identified above through an open scalable agent-based architecture. The study presented here examines the potential for a highly disaggregated energy use monitoring and feedback system for home electricity consumption. Unlike most smart metering solutions, this system can be used for:

1. Collection and display of use information from an arbitrary number of appliances either directly or as groups connected via extension sockets or wall sockets using the Plogg[1] sen-

[1]The Plogg is a particular example of a plug-in appliance energy monitor [18]—see Figure 1—other similar devices are available.

sors. This means that users, and potentially researchers, have an unprecedented level of detail about electricity usage by minute, hour, day, month, year, further disaggregated by end-use and location within the building. Further development will allow arbitrary re-aggregation into meaningful groups that make sense to individual users.

2. Deployment in both new and existing buildings, since the wireless energy monitors connect to standard household sockets using standardized communication protocols such as Zigbee and Bluetooth.

3. Manage an arbitrary number of sensors so users / researchers are not limited by cost so long as adequate measures are taken to ensure that at least one sensor is within range of the Zigbee/ethernet bridge and the remainder close enough to each other to form a mesh network. More careful positioning is necessary in the case of Bluetooth.

Thus, the main contributions of this paper are a proof of concept implementation of an agent-based architecture for the real-time collection of energy-use data and an illustration of the kind of off-line analysis that is feasible once such data sets are available. We believe this architecture provides a sound practical basis for both the live monitoring that is necessary for the various on-line applications identified above as well as the collection of long-term data needed for vertical studies and policy analysis.

The remainder of this paper is laid out as follows: in section 2 we set out the needs of our stakeholders in electrical engineering and in architecture, followed in section 3 by a detailed description of the architecture that has been developed and its constituent components. In section 4, we sketch a possible distributed deployment that is realizable using our architecture, while section 5 provides an illustrative analysis of period data collected from a pilot deployment of the energy monitoring devices. Sections 6 and 7 discuss short and medium term future work and present our conclusions.

2. ENERGY PROFILING

As set out above, researchers in both demand-side management (as we term the collection of on-line functions) and housing policy (collectively off-line data analysis) need comprehensive data sets against which, respectively they can experiment with market mechanisms or explore the potential long-term effect of policy initiatives to affect individual behavior. This translates not only to a means for the collection of long-term data sets from which population characteristics can be extracted, but also to a need for a much shorter feedback loop to examine the effectiveness of economic and physical control mechanisms.

The architecture we have developed is intended to satisfy both these requirements, but while it offers a high degree of flexibility in terms of the analysis that can be carried out, the agents are strictly on the "inside" and the interfaces are presented in terms of widely accepted HTTP protocols.

Stated in more neutral terms, the domain requirements are to:

R1: Capture appliance specific data over extended periods

R2: Present data sets for analysis not defined at the time of capture

R3: Allow for technology change in data capture and device control

R4: Provide low-overhead integration with current and future networking facilities

Figure 1: Two Ploggs, one for European and one for UK power sockets [18]

Our technical solution for these requirements is, building on the framework outlined in [8], to use: (i) an agent platform to provide separation of concerns, distribution, loose-coupling and scope in the future for institutionally-directed component behavior—this may be particularly relevant for the control aspect identified above—and (ii) a semantically annotated, unstructured data representation, that, while computationally more expensive to process than a conventional relational database, offers complete flexibility in respect of future analysis requirements A more detailed account of the architecture and its use comes in the following sections.

3. MONITORING ARCHITECTURE

The monitoring architecture introduced in this paper is an extension and refactoring of the monitoring architecture put forward in [8]. The framework is implemented as an agent-based application that runs on the AgentScape [14] platform.

3.1 Overview

With reference to Figure 2, the monitoring architecture comprises three sub-systems:

1. **Collection:** Ploggs [18] (see section 3.2) are used to monitor energy usage from individual power sockets. A number of Ploggs are deployed in a single household and connected through a Zigbee mesh network. A collector component is used to gather the data stored on the different Plogg sensors. The diagram shows just one collection network, but there could be many, in which case there would normally be one sensor agent for each collector.

2. **Processing:** The collector component posts the collected data to a servlet that forwards it to a sensor agent on the Agentscape platform. This sensor agent will (pre)process the data into RDF format that is subsequently forwarded to a (semantics-enabled) data store. An aggregation agent can (optionally) access, process and store the aggregated data again in the data store. The diagram also shows only one processing network with a single database; however this could also be replicated and federated to support large-scale geographical deployments.

3. **Presentation:** The (aggregated) energy usage data can be presented in different formats. The task of a presentation agents—and there can be many of these, each providing different perspectives on the data—is to process selected data

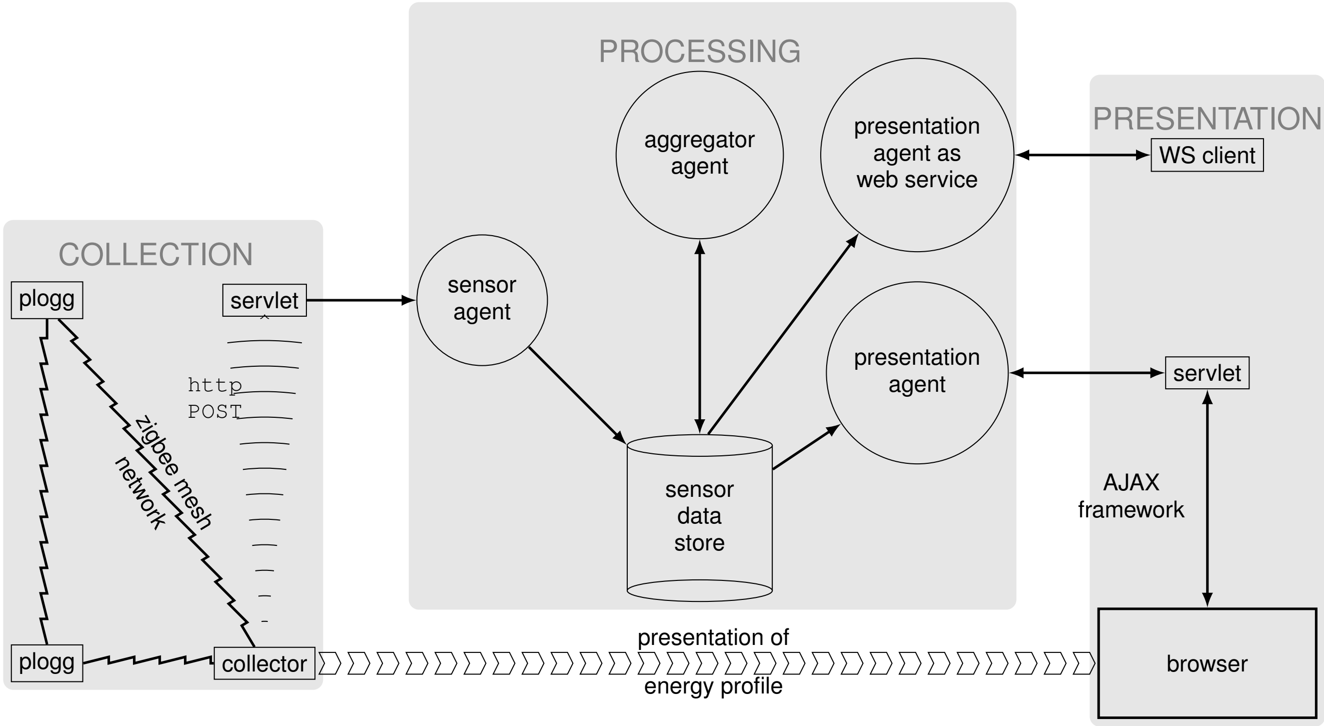

Figure 2: Monitoring architecture

from the database in order to compute its particular view on the current situation. Thus, it can publish a web-service for further usage by other applications or provide a dedicated ajax-based web page that delivers live updates on (aggregated) energy usage for a specified collection of energy monitors (by household, by usage, by area, etc., by means of the (semantic) annotation on the entries in the database).

Note that from a typical user's point of view, the processing component is completely hidden. Users should effectively only be aware of the sensors and the presentation of the sensor data.

3.2 Collection

The energy usage over time of an electrical device is monitored using 'Ploggs'. These devices measure some fourteen parameters, but the most important ones for our present needs being the current (live) energy consumption in Watts and the cumulative energy consumption (since the Plogg was first plugged in) in kWh. The device can store a user-defined selection of the parameters in the plogg, at a frequency also set by the user, in the range of once every tens of minutes to once every second.

Each Plogg only has a small internal memory for data collection (64Kb). Accordingly, if the data stored on the Plogg is not retrieved sufficiently frequently, the earliest measurements will be overwritten. For example, consider the following two scenarios:

1. **Offline:** Ploggs are deployed (stand-alone) in a household. Energy usage is monitored and stored on the Plogg every ten minutes. After a suitable period, say a month, the Ploggs are retrieved, and the contents of the internal memory is downloaded and stored in a database.

2. **Online:** Ploggs are deployed in a household. Energy usage is monitored and stored on the Plogg every second. Every 5 seconds the stored data is accessed by a collector. Live energy usage is displayed in a web browser (see Figure 2).

The *collector* serves as a customizable software layer on top of the hardware device that is used to access the sensors and that sends the collected data to the sensor architecture. Individual Ploggs either communicate directly with the collector via the Bluetooth protocol or they can form a mesh network using the Zigbee protocol and communicate with the collector as a group. Collectors for both Plogg types have been implemented.

3.3 Processing

The AgentScape platform supports agents as autonomous processes. A uniform middleware layer provides an agent run-time that is available for several heterogeneous platforms. Within AgentScape, *agents* are active entities that reside within *locations*, and *services* are external software systems accessed by agents hosted by the AgentScape middleware. Agents in AgentScape can communicate with other agents and can access services. Agents may migrate from one location to another.

Agentscape defines a 'location' as a collection of hosts that typically run at the same physical site, for example a household or an organization, see Figure 3. AgentScape is a middleware and has been designed for modularity, extensibility and scalability. This makes it well-suited to the implementation of a distributed sensor infrastructure.

In the sensor architecture from Figure 2 agents can access individual sensors through a generic sensor service [8]. This sensor service provides an abstraction mechanism for implementing interfaces for different (hardware) sensor types. Sensors are individually accessed on a per URI basis. After the agent provides the service with the URI of the sensor, an interface belonging to the specific sensor type is returned. This latter interface forms a

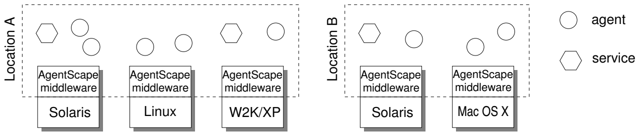

Figure 3: Two AgentScape locations. One with three and one with two hosts.

specialized version of the generic interface provided by the sensor service.

Data collected from a specific sensor instance can be filtered and processed by the sensor agent. This data can be in the form of a continuous stream or discrete (polled) data. Consequently, the processed data can be used directly, or be stored in RDF format in a database. The JRDF package[10] that provides an API to a triple store is currently used to implement this. The attraction of this approach is flexibility afforded by the RDF triple structure and the fact that a triple store naturally accommodates semantic annotation.

The sensor data store can be queried by agents who in turn can aggregate the information stored in one or more data stores. Such an agent could, for example, calculate the energy usage of a complete household or the total energy consumed by all the televisions in a town, etc. The aggregator agent from Figure 2 may provide such functionality, but of course, this depends on whether a particlar monitor is connected to a particular device and is unverifiable in practice with present technology.

3.4 Presentation

A presentation agent is used to display the (aggregated) information from the data store. Different types of presentation agent can be used. Figure 2 shows two examples:

- Dynamic web page: The presentation agent forwards the sensor data to a servlet. The Ajax framework is used to display the (aggregate) energy usage of a (collection of) household(s) continuously. See [21] for more details.

- Web service: The presentation agent uses AgentScape's WS-Gateway [15] service to publish the sensor data as a web service. Other applications, for example a web application targeted at mobile phones, could access and display the energy usage patterns of households.

These are two possible examples of presentation mechanisms. Dedicated agents can be developed and (re)used to display sensor data, as desired, independently of the rest of the architecture.

3.5 Architecture Implementation

The architecture presented in the section has been implemented on top of the AgentScape platform. End-to-end functionality in the form of real-time collection, processing and presentation of data—from ploggs to browser—is currently working, though the presentation of data is still simplistic, and is the subject of current work.

In the current architecture, sensor agents that collect the sensor data do not communicate directly with aggregator or presentation agents. The latter agent types access the sensor data through the RDF data store, making this a potential communications bottleneck. For most applications types this does not matter as aggregated data, for example, does not need to be presented in real-time. However, if real-time presentation is required then agents can circumvent the RDF store and communicate directly with each other, using (AgentScape's) asynchronous message passing. In this case, the sensor agent parses the (plogg) sensor data, encapsulates it in a (Java) object and sends the object to the presentation agent, that can access the object directly and display the (real-time) data. Note that, in this instance, the semantic information stored in the RDF store is not used. Alternatively, a local sensor store *cache* can be used that can be accessed by a presentation agent. The local cache circumvents the bottleneck identified above.

In both cases, additional care needs to be taken to store the data in the centralized RDF store for further (non-real time) aggregation and presentation purposes. More details about the implementation of the sensor architecture can be found in [8, 21].

4. ONLINE MONITORING

AgentScape has been developed for scalability and deployment in open, insecure environments. This has some obvious advantages when developing a monitoring architecture that intended to be widely deployed and scalable.

In principle, the Ploggs can be used as stand-alone monitoring devices, as described in scenario 1 in Section 3.2. However, offline monitoring is not always possible, nor is (offline) automated aggregation using data from multiple households. The sensor architecture introduced in the previous section allows both on-line monitoring and automated aggregation. This section describes a typical usage scenario of the energy monitoring sensor infrastructure.

4.1 Sensor Deployment

Figure 4 shows a possible deployment scenario of the energy monitoring framework from the previous section. There are three monitored households in the scenario. All three are located in Bath (at location 1, 2 and 3). A couple of Ploggs are located at location 1 together with an Ethernet Access Point (EAP) for Zigbee which communicates with the Ploggs and forwards the monitored data via a regular Internet connection to a sensor agent (agents are represented by small circles in Figure 4). In this case the sensor agent runs at another household in Bath (location 2). Ploggs are also deployed here that gather information of the current household. A third location in Bath also runs the AgentScape middleware and monitors energy consumption at this household. The two AgentScape locations are connected with each other and with a third AgentScape instance that runs in Delft. Here the RDF database is located that collects all the monitored data. A presentation agent accesses the RDF database to present a live feed of the monitored data which can be viewed with any web-browser, for example in Toronto.

4.2 Deployment Issues

The current sensor architecture provides end-to-end functionality in the form of real-time collection, processing and presentation of data. However, while the scenario from Figure 4 can be realized, a number of deployment issues remain.

In the current implementation there is only one (centralized) database. A more realistic, and more scalable, option would be to use a distributed database. Alternatively, each group of households (a city block, for example) could have its own database. A centralized database containing aggregated can be added to such a scenario. In this case the information in the centralized (aggregated) database can be used to find global trends while the local databases can provide more insight into local energy usage patterns. Ideally, aggregated databases can be added at multiple levels (single households, neighborhood, city, region etc.). This can also provide valuable insight into the energy usage of different regions, make it easier to compare single households, cities or regions and allow to identify both over and under energy consumers, i.e., households that use

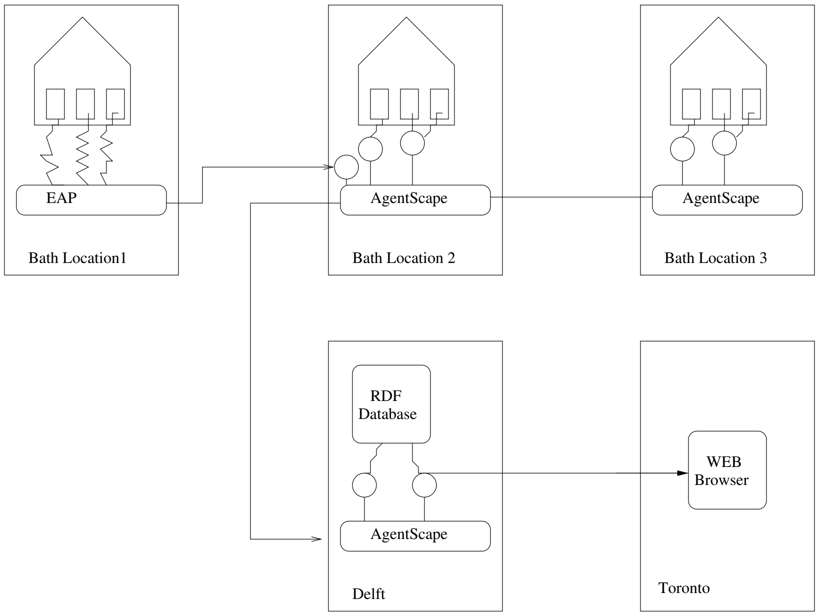

Figure 4: Illustration of a possible deployment scenario of the monitoring architecture

substantially more or less then the average consumption of a similar household.

Multiple (possibly aggregated) databases can also help in case of network or hardware failures. To limit the impact of failures as much as possible local caches can be used to store recently generated sensor data. These caches can potentially be on the devices (the Ploggs) themselves, though as Section 5 shows, the Ploggs themselves are probably not reliable enough. Small embeddable devices, such as the Sun SPOT[2], are probably more suitable for this. This will be further explored in follow-up case studies.

5. ANALYZING HOUSEHOLD ENERGY CONSUMPTION

The motivation for this section stems from studies and investigations carried out in the UK, but similar reports will exist for many other countries.

5.1 Motivation

The UK government is committed to making large cuts in carbon emissions to mitigate the impact of climate change. This is coupled with programs for improving energy efficiency and reducing energy usage to increase energy security. Buildings in the UK are currently responsible for around half of the UK's CO_2 emissions [17] and constitute a major focus for both mitigation and adaptation strategies. A significant part of the effort to reduce buildings-related carbon emissions has been directed at energy efficiency. However, energy use as determined by occupant thermal comfort, changes in ownership of appliances and hours of use are poorly understood and out of date: the last major report on domestic appliance ownership and efficiency was the DECADE report in 1998 [5]. This report significantly alters, by a factor of two, the prevailing assumptions regarding lighting and appliance energy, spurring a review of the UK Building Research Establishment's BREDEM[3] model on which almost all current UK national domestic energy calculations are based [1].

Consequently, it is projected that *every household in the UK* will need to reduce its carbon emissions by at least 80% by 2050 from 1990 levels in order to meet the government's long term emission reduction target. While technological options are expected to deliver significant cuts, individual household *energy behaviors* will have a major impact on actual emissions. The Inter-governmental Panel on Climate Change (IPCC) fourth assessment report on mitigation states that "occupant behavior, culture and consumer choice and use of technologies are also major determinants of energy use in buildings and play a fundamental role in determining CO_2 emissions". However, although the impact of occupant behavior is widely regarded as important, the IPCC report also recognizes that there is limited evidence to support it [9, p.389].

5.2 Realization

In section 3, we explained that the collector component may be realized as either a Zigbee or Bluetooth wireless interface to the Ploggs and a software component that downloads records from the Ploggs at some frequency, from seconds to days to weeks, even. In its on-line incarnation, the collector software component then sends the records to the sensor architecture for processing, storage and presentation, while the off-line version just saves the records for subsequent processing.

In order to prototype the data analysis phase, while the development of the collection, processing and presentation framework was in progress, we needed some data sets recording actual appliance usage. Consequently, we installed a set of Plogg sensors in 4 households in Bath and ran them in off-line mode for a period of 4 weeks in July 2009. The households comprised a sample range of family structures and ages (See Table 1).

Occupants were asked to connect sensors to some typical household appliances, as identified in Table 2 with reference to the households in Table 1. The Ploggs were configured to collect cumulative

[2] `http://www.sunspotworld.com/`

[3] Building Research Establishment Domestic Energy Model: a model that estimates energy uses in a property based on, among other variables, dwelling size, occupancy and heating system installation.

Household Code	Household Structure	Age Bracket
1P-0C-Y	One-person household	Young
2P-0C-M	Two-person household with no children	Middle-Aged
2P-2C-M	Two-person household with children	Middle-Aged
2P-0C-O	Two-person household	Old

Table 1: Household Structure + Age matrix for test case energy monitoring where occupant ages (in years) are represented as Young $<= 35 <$ **Middle Age** $<= 60 <$ **Old**

Household Code	PC	WM	KT	TV	MW	FZ
1P-0C-Y	x	x		x		
2P-0C-M	x		x	x	x	
2P-2C-M	x	x	x	x		
2P-0C-O	x			x	x	x

Table 2: Appliance-sensor installation matrix for households in Table 1, where PC = Personal Computer + peripherals, WM = Washing Machine, KT = Electric Kettle, TV = Television, MW = Microwave, FZ = Freezer

kWh—power consumption—at 10 minute intervals; this data set being small enough to fit in the 64K of on-Plogg memory for the duration of the deployment, but recording sufficient data to reveal useful information. In a larger scale study, this would allow us to examine patterns of electricity use across households.

Some practical issues were encountered in the deployment, relating to the design and size of the sensor which did not accommodate all types of wall socket/switch combinations. In some cases, the sensors could not be installed because it obscured the switch and so the appliance could not be operated easily by occupants who choose to switch the appliance off at the socket—note: there is no switch in the Plogg itself. In other cases, the depth of the sensor reduced available space so that the appliance could no longer be plugged in. In a larger sample, these issues can be mitigated through the use of a small extension socket or board. It may also be possible to re-house the sensor in tighter, custom built housing, but the bottom-line physical constraint is the diameter of the ring around the live connection, which is of the order of 2cm.

5.3 Analysis of Monitored Data

Clearly, with such a small sample and also only a subset of possible household types, the data and its analysis cannot be representative of the wider population. However, our purpose at this stage was a feasibility study to evaluate (i) the reliability of the Ploggs themselves (ii) the usefulness of the data that could be collected, and (iii) the kinds of analysis that were subsequently possible. In respect of reliability, we collected 19 usable data sets from 22 deployed devices. One unit failed completely and two others reverted to the default configuration of collecting data on all the parameters permitted every minute, thus filling (and wrapping around) the on-Plogg memory. This was a significant improvement on the first (4 week) deployment during which about half the Ploggs reverted to the default configuration. This behavior is attributed to the battery that maintains the on-Plogg memory having insufficient charge to retain the configuration between setup and deployment. Subsequently all Ploggs were left plugged in for 24hrs prior to configuration for the second deployment. Clearly, this presents a risk and a practical problem for larger scale studies as well as justifying the time spent on the pilot study.

The sample data sets do provide us with data on PC and TV usage for all four households. Consequently, we present an illustrative analysis of these two devices to show the *type of data* such monitoring in larger, more representative, samples may deliver and their potential benefits.

5.4 Energy used by PCs

PC energy use varies by the rated power consumption of the main unit (motherboard and graphic cards etc.) and peripherals (display, printers etc) and the activity level of the machine. PCs that are not on stand-by will use more power than ones that are, and those doing high-end gaming or other intensive tasks still more. With modern PCs, typical rated power consumption may be anywhere from 60W on a low-end machine (typically laptops) to a few hundred watts on a high end machine in addition to energy used by peripherals such as monitors, printers etc. Therefore, in order to accurately estimate domestic energy use from PCs, the actual consumption needs to be logged. In our sample, PCs and peripherals were 'on' an average of 50% of the time within the 4 week monitored period and consumed between 3.4kWh (2P-0C-M household) to 36.9kWh (2P-0C-O household). If this were taken as a typical month, annual energy consumption could vary from 40.8kWh to 442.8kWh through the use of PCs alone. This translates to a difference in the annual electricity bill of between £4 and £40 (at 10p per kWh) or carbon dioxide emissions of 21.9kgCO_2 and 237.8kgCO_2 (at 0.537 kgCO_2 /kWh [20]). Formally, Green House Gas (GHG) emissions are reported in tons of Carbon (tC) for comparability and in this case, these ranges would be 0.006tC and 0.065tC [4]. These figures are 1.5% and 16.3% of the *estimated* average carbon emissions from a UK dwelling through the use of lights and appliances (0.4tC, Domestic Energy Fact File 2006 [22]).

We emphasize 'estimated', because these data are derived for the Domestic Energy Fact File through a model that estimates end-uses based on the BREDEM model and not actual measurements. Current proposals for revisions to the BREDEM model's empirical data refer to the DECADE study in 1994 which was the last major study to examine domestic energy use from lights and appliances [5]. It is noteworthy that the DECADE data classes energy use from PCs in the 'miscellaneous' category as much of the recent proliferation in home computing and the effect of the Internet were not well known at that time. As much UK national policy is driven by calculations based on the BREDEM model, updating the model with more accurate monitored energy use data is essential.

In addition, such data may also tell us how long appliances are left 'on' without active use and therefore the potential value of occupant feedback specifically for certain classes of appliances. For example, a PC left 'on' may have a built in display that counts the cost of energy wasted and could display this for both the current session and historic sessions. Wood and Newborough [25] found that device level Energy Consumption Indicators (ECI) for electric cookers allowed 7 out of 10 households in a group to achieve greater than 10% savings compared to a control group (max of 39%). These savings were generally better than the groups with (i) an information pack and no ECI, and (ii) an information pack and an ECI. If such savings could be realized for a broad range of domestic and commercial appliances, significant carbon savings might be possible.

5.5 Energy used by TVs

Energy used in TVs varied from 0.8kWh (2P-2C-M) to 48.7kWh (2P-0C-M) for the 4 households over the monitored period. As with

[4] The atomic weight of Carbon is 12 and Oxygen is 16 so every ton of CO_2 is equivalent to 12/44 tons of carbon.

the case for PCs, this is a wide range and could represent between 0.4% and 21.4% of the estimated carbon emissions for an average UK dwelling through the use of lights and appliances [22].

If these data were representative of the housing stock there would be a significant impact on estimated carbon savings. For example, BREDEM currently estimates energy use for lights and appliances based on dwelling size and the number of occupants. However, it does not take into account the age of the occupants or the type of household. Therefore, our estimates of carbon savings do not account for the fact that different types of households *may* use energy differently. Although we emphasize 'may', there is already some evidence to suggest that age and household structure have an impact on energy use [4, 3]. This means that should the distribution of household structure change in the UK—for example, due to a demographic shift—emissions targets might easily be missed. Currently, we have no method for modeling by how much those targets may be missed, because we do not have robust data on exactly what effect occupant age or household structure has on energy use. The monitoring framework and analysis set out here provide one route for obtaining such data: using objective energy-use monitoring in combination with detailed survey information on household age, structure, income etc. Such data will also allow the optimization and tailoring of energy saving information to fit occupant types - for example, through improved smart meter displays.

6. DISCUSSION AND FUTURE WORK

The presentation of data is currently simplistic, because our first priority was to achieve end-to-end functionality in the form of real-time collection, processing and presentation of data. With the components of this chain in place, we are currently developing a range of analysis and rendering tools for the construction and visualization of aggregating sensors via a web browser interface or published via a web service (see Figure 2). Several further case studies are in planning:

1. a medium-term deployment in student housing on the University of Bath campus in the beginning of 2010
2. a larger and longer term domestic deployment in Bath 2010
3. an investigation of power consumption profiles of numerically controlled machinery in intelligent manufacturing that will involve collecting data on three-phase power supplies

We have also developed a first version of a mechanism to switch Ploggs off remotely. This functionality can be made available to humans via the browser interface, even permitting remote switch off through mobile phones. Additionally, aggregator agents building up a whole household picture could be capable of identifying situations where disconnection of a device from the energy supply may be appropriate and this action can now be achieved through this mechanism. The challenge here lies in the decision-making procedure making the right choice at the right time.

Another promising area for future work includes demand-side energy management systems [19, 24]. The 'processing' part of the architecture introduced in Section 3 can be extended with another agent type (an effector) that changes the energy consumption in a household by switching thermostatically controlled appliances, such as fridges or ac-units, off or on in a coordinated manner thereby shifting energy consumption [11, 23] and removing (global) peaks in energy consumption. An (analysis) agent can be used to process the (monitored) data in the database to make plans for the effector agents to change the energy consumption in a positive manner. The generic design of the monitoring framework is ideal for this kind of extension, and allows experimentation with different types of analysis and effector agents.

Privacy and security form other areas that have to be considered in future work. Privacy in particular forms a obvious challenge in this context.

7. CONCLUSIONS

We have presented a framework for the collection, processing and presentation of sensor data in general [8] applied to energy consumption by household appliances in particular and sketched a preliminary analysis of the data that such energy monitors are able to collect. What we believe this demonstrates is the flexibility and scope for future development that the use of agents has to offer in the domain of energy monitoring and in the future, energy management, since it is also possible to close the loop and selectively turn devices on and off through the same framework.

Acknowledgments: This work is partially supported by the NLnet Foundation (`http://www.nlnet.nl`) and the Engineering and Physical Sciences Research Council UK ("Creativity in Systems Design: The Bath Interactive Ideas Factory", EPSRC Bridging the Gaps initiative, EP/E018122/1).

8. REFERENCES

[1] B. Anderson, P. Chapman, N. Cutland, C. Dickson, S. Doran, G. Henderson, J. Henderson, P. Iles, L. Kosminat, and L. Shorrock. Bredem-8 - model description: 2001 update, 2002.

[2] W. Anderson and V. White. Exploring consumer preferences for home energy display functionality. Technical report, Centre for Sustainable Energy, 2009.

[3] V. Beillan, E. Cayre, A. Pouget, D. Sanna, D. Novakov, U. Rochard, and R. Trotignon. Habitenergie. Success factors in the construction of very low-energy housing: the weight of stakeholder relationships and of household practices - a survey in three european countries. In *ECEEE 2007 Summer Study: Saving Energy - Just do it!*, pages 849–858, 2007.

[4] BRE. Estimates of hot water consumption from the 1998 EFUS. implications for the modelling of fuel poverty in england. `http://www.berr.gov.uk/files/file16568.pdf`, 2005.

[5] ECI. *2MtC (DECADE)*. Environmental Change Institute, Oxford, 1997.

[6] A. Faruqui, S. Sergici, and A. Sharif. The impact of informational feedback on energy consumption–a survey of the experimental evidence. *Energy*, In Press, Corrected Proof:–, 2009.

[7] F. Haldi and D. Robinson. Interactions with window openings by office occupants. *Building and Environment*, 44(12):2378–2395, 2009.

[8] T. Harman, J. Padget, and M. Warnier. A multi-layered semantics-ready sensor architecture. In *Proceedings of the Third International Workshop on Agent Technology for Sensor Networks (ATSN-09)*, 2009.

[9] IPCC. Climate Change 2007: Mitigation. Contribution of Working Group III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change [Metz, B., Davidson, O., Bosch, P., Dave, R., Meyer L. (eds.)]. Technical report, IPCC, 2007.

[10] Jrdf. `jrdf.sourceforge.net`, retrieved 20090215. Java Resource Description Framework API for graphs, object persistence and querying.

[11] J. Li, G. James, and G. Poulton. Set-points based optimal multi-agent coordination for controlling distributed energy loads. *Self-Adaptive and Self-Organizing Systems, International Conference on*, 0:265–271, 2009.

[12] M. Newborough and S. D. Probert. Intelligent automatic electrical-load management for networks of major domestic appliances. *Applied Energy*, 37(2):151 – 168, 1990.

[13] E. Ogston and F. M. T. Brazier. Apportionment of control in virtual power stations. In *In the proceedings of the international conference on infrastructure systems and services 2009: Developing 21st Century Infrastructure Networks*, 2009.

[14] B. J. Overeinder and F. M. T. Brazier. Scalable middleware environment for agent-based internet applications. In *Proceedings of the Workshop on State-of-the-Art in Scientific Computing (PARA'04)*, volume 3732 of *Lecture Notes in Computer Science*, pages 675–679, Copenhagen, Denmark, June 2004. Springer.

[15] B. J. Overeinder, P. D. Verkaik, and F. M. T. Brazier. Web service access management for integration with agent systems. In *Proceedings of the 23rd Annual ACM Symposium on Applied Computing (SAC)*. ACM, March 2008.

[16] A. Peacock and M. Newborough. Impact of micro-combined heat-and-power systems on energy flows in the uk electricity supply industry. *Energy*, 31(12):1804 – 1818, 2006.

[17] D. Pearce. The Social and Economic Value of Construction. Technical report, nCRISP, London, 2004.

[18] Plogg, wireless enegery management. `http://www.plogginternational.com`, last retreived March 8, 2010.

[19] E. Pournaras, M. Warnier, and F. M. T. Brazier. Local Agent-based Self-stabilisation in Global Resource Utilisation. *International Journal of Autonomic Computing (IJAC)*, 2010. (to appear).

[20] M. T. Programme. BNXS01: Carbon dioxide emission factors for UK energy use. `http://efficient-products.defra.gov.uk/spm/download/document/id/785`, 2009.

[21] H. S. Riat. Multi agent based energy monitoring system. Master's thesis, University of Bath, 2009.

[22] L. Shorrock and J. Utley. Domestic energy fact file 2008. Technical report, Building Research Establishment, Garston, Watford., 2008.

[23] M. Stadler, W. Krause, M. Sonnenschein, and U. Vogel. Modelling and evaluation of control schemes for enhancing load shift of electricity demand for cooling devices. *Environmental Modelling & Software*, 24(2):285 – 295, 2009.

[24] G. Strbac. Demand side management: Benefits and challenges. *Energy Policy*, 36(12):4419–4426, December 2008.

[25] G. Wood and M. Newborough. Dynamic energy-consumption indicators for domestic appliances: environment, behaviour and design. *Energy and Buildings*, 35(8):821 – 841, 2003.

[26] G. Wood and M. Newborough. Energy-use information transfer for intelligent homes: Enabling energy conservation with central and local displays. *Energy and Buildings*, 39(4):495 – 503, 2007.

[27] A. Wright and S. Firth. The nature of domestic electricity-loads and effects of time averaging on statistics and on-site generation calculations. *Applied Energy*, 84(4):389 – 403, 2007.