Citation for published version:

Padget, J, Riat, H, Forchhammer, B, Warnier, M, Brazier, FMT & Natarajan, S 2011, An agent-based infrastructure for energy profile capture and management. in *2011 International Conference on Networking, Sensing and Control, ICNSC 2011*. pp. 50-55, 2011 International Conference on Networking, Sensing and Control, ICNSC 2011, Delft, Netherlands, 11/04/11. <https://doi.org/10.1109/ICNSC.2011.5874876>

DOI:

[10.1109/ICNSC.2011.5874876](https://doi.org/10.1109/ICNSC.2011.5874876)

Publication date:

2011

Document Version

Peer reviewed version

[Link to publication](#)

© 2011 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

An Agent-Based Infrastructure for Energy Profile Capture and Management

Julian Padget, Harpreet Riat, Benedikt Forchhammer, Martijn Warnier, Frances M.T. Brazier and Sukumar Natarajan

Abstract—Accurately monitoring changing energy usage patterns in households is a first requirement for more efficient and eco-friendly energy management. Such data is essential to the establishment of the Smart Grid, but at this stage, domestic data collection devices are still in development and monitoring-enabled domestic appliances are rare, so that any experimental software framework must be flexible and adaptable both in respect of sensor sources and developer and user requirements. These considerations have been the drivers behind the distributed agent-based platform this paper proposes. It provides: (i) a generic sensor interface that can be specialised for new devices as required, while insulating the rest of the platform from such changes, (ii) persistent unstructured (RDF) data storage, permitting both semantic annotation and semantic-based queries, independent of data sources, and (iii) a flexible, dynamic browser interface, that allows for remote configuration of the sensor platform and accessibility via a wide range of devices. Two small case studies show the utility of the approach.

I. INTRODUCTION

Energy has become a major focus of governmental organizations in recent years. Energy resources are limited and with the emergence of the Smart Grid [1], [2], [3] energy infrastructures are changing. More and more information about energy usage and transfer is becoming available at different locations and levels of granularity. Information that can be used for numerous applications such as smart energy routing, utilization and (micro) production. 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, two drivers for reducing energy consumption are:

- 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

J.A. Padget, H. Riat and B. Forchhammer are with the Department of Computer Science, University of Bath, Claverton Down, Bath, United Kingdom jap@cs.bath.ac.uk

M.E. Warnier and F.M.T. Brazier are with the Faculty of Technology, Policy and Management, Delft University of Technology, Delft, the Netherlands m.e.warnier@tudelft.nl, f.m.brazier@tudelft.nl

S. Natarajan is with the Department of Architecture and Civil Engineering, University of Bath, Claverton Down, Bath, United Kingdom s.natarajan@bath.ac.uk

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 optimization of micro-generation technologies to complement network load [5], [6], appliance to electricity network communication to allow Energy Service Companies (ESCOs) to help schedule appliance operation [7] and smart metering to exploit occupant energy use awareness benefits [8], [9]. 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 [10]. Broadly, home automation excluded, these efforts are driven by the desire to improve energy conservation and energy security through reducing ‘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 households to shift energy demand [11], [12], [13], [14] to different time periods in response to finer-grained time-dependent pricing (as part of the move to supply-led rather than demand-led generation). Virtual powers stations [15], 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 presented study 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 Plogg¹ sensors. 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 remain 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. This architecture provides a 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.

II. ENERGY PROFILING

As set out above, researchers in both demand-side management (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 developed architecture 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

Our technical solution for these requirements is, building on the framework outlined in [17], to use: (i) an agent plat-

¹The Plogg is a particular example of a plug-in appliance energy monitor [16] other similar devices are available.

Fig. 1. Monitoring architecture

form 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.

III. MONITORING ARCHITECTURE

The monitoring architecture introduced in this paper is an extension and refactoring of the monitoring architecture put forward in [17]. The framework is implemented as an agent-based application that runs on the AgentScape [18] platform.

A. Overview

With reference to Fig. 1, the monitoring architecture comprises three sub-systems:

- 1) **Collection:** Ploggs [16] (see section III-B) 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 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.

B. 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 Fig. 1).

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.

C. 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. 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 Fig. 1 agents access individual sensors through a generic sensor service [17]. 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 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 attraction of this approach is the flexibility afforded by the RDF triple structure and the fact that a triple store naturally accommodates semantic annotation. See for example the records in Figure 2, recorded from one particular Plogg, where each record comprises three elements: (i) (subject) the unique agent handle that is the domain of the relation (ii) (predicate) a `urn` that identifies the relation by means of an element from a simple ontology, and (iii) (object) a value in the range of the relation

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 Fig. 1 may provide such functionality. However, because monitors and devices are (at present) separate objects, it is not possible to know for certain whether a particular monitor is delivering data about a particular appliance.

The monitoring architecture provides an abstraction over physical sensors called virtual sensors, which may be assigned to one or more physical sensors. See Figure 3, where the sensor is associated with an appliance as part of the platform’s support for naming and grouping. The indirection afforded by this mechanism means that down-stream software components can refer to virtual sensors making them independent of up-stream changes in physical sensors. For example, if the Plogg by which the coffee machine is connected stops working, it can be replaced and one only has to reconfigure, the virtual sensor and down-stream code is not affected. Virtual sensors and the presentation agents acquire their data from the triple store by making queries using the SPARQL [19] language, which is essentially a development of SQL for querying collections of RDF triples.

```

3  bb1...761ab6#93 urn:cumulative_watts_con 310941.0
...
7  bb1...761ab6#93 urn:logger_agent 8689E1225B5F5FE1...
...
9  bb1...761ab6#93 urn:plogg_id 0021ED000004503E
...
14 bb1...761ab6#93 urn:sensor_id ae954447-18ad-4fb5-...
15 bb1...761ab6#93 urn:submission_time 1273157624885
...
17 bb1...761ab6#93 urn:type PloggDataset
...
20 bb1...761ab6#93 urn:watts 0.0

```

Fig. 2. Actual Sensor RDF

```

1 32e...9c2#158 urn:last_updated 1273163917000
2 32e...9c2#158 urn:name CoffeeMachine
3 32e...9c2#158 urn:sensor_owner 4562de51-9bb1-4561-...
4 32e...9c2#158 urn:type Sensor
5 32e...9c2#158 urn:uuid ae954447-18ad-4fb5-9095-...

```

Fig. 3. Virtual Sensor RDF

Fig. 4. Web view of energy display

D. Presentation

A presentation agent is used to display the (aggregated) information from the data store. Different types of presentation agent can be used. Fig. 1 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 [20], [21] for details.
- **Web service:** The presentation agent uses AgentScape's WS-Gateway [22] 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.

The dynamic web page approach accesses the platform via a RESTlet interface. The web interface allows for the creation and naming of new virtual sensors, associating them with one or more ploggs, and then constructing a simple display for each virtual sensor (see Figure 4).

E. Architecture Implementation

The architecture presented in the section has been implemented on 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 ongoing work.

In the current architecture, sensor agents that collect the data do not communicate directly with aggregator or presentation agents. The latter agent types acquire data through the RDF data store, making this a potential communications bottleneck. For most application 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. The architecture also supports distributed deployment as agent can communicate with a (possible distributed) data store via the same message passing mechanism. As a result data store and agents do not have to run on the same platform.

Consequently, care must be taken to store the data in the centralized RDF store for further (non-real time) aggregation and presentation purposes. More details about platform implementation are given in [17], [20].

IV. DEPLOYMENT

Section III explains that the collector component employs a wireless interface to the Ploggs and a software component that downloads records from the Ploggs at some frequency, from seconds to days to weeks. The on-line collector then sends the records to the sensor platform for processing, storage and presentation, while the off-line version just saves the records for subsequent processing. The platform components in both scenarios have been trialled.

A. Off-line

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, a set of Plogg sensors was installed in 4 households in Bath. They ran 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 I).

Occupants were asked to connect sensors to some typical household appliances, see Table II with reference to the households in Table I. The Ploggs were configured to collect cumulative 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. Figure 5 illustrate the kind of data that can be collected with this approach. It provides an example of monitored data for one appliance: the energy consumption of the television across the four household types. The graph shows that all TVs are off during the night, but some consume more energy (on standby) than

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

Fig. 5. Monitored hourly average energy consumed by TV in 4 households, see TABLE I for legend.

others (not on standby and off). Most households watch some TV during the afternoon and evening, with the exception of the household with children who do not seem to watch TV at all. Also note that The two person middle age household without children have the TV on during most of the day.

Clearly, with such a small sample and a subset of possible household types, the data and its analysis is not representative of the wider population. 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, 19 usable data sets were collected 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.

B. On-line

To test feasibility of the the platform it was trialled via a short-term deployment in student housing on the University of Bath campus in Spring 2010, with live feedback through

a web interface (see Figure 4). Again using ploggs, three appliances were instrumented (toaster, microwave, kettle) in three communal student kitchens, and as a reference, the Computer Science department kitchen. Here the objective was behavioural response to different types of displays, such as the digital one shown in Figure 4 and “ambient” visualizations such as colour changes and informative icons. This work will be the subject of a future paper.

C. Deployment Issues

The current sensor architecture provides end-to-end functionality in the form of real-time collection, processing and presentation of data. However, 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 substantially more or less than 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 IVA reports, the Ploggs themselves are probably not reliable enough. Small embeddable devices, such as the Sun SPOT², are probably more suitable, but availability seems limited. The Arduino-based OpenEnergyMonitor³ is reliable and easy to use, but implies considerable packaging overhead for anything beyond small-scale deployment.

V. DISCUSSION AND FUTURE WORK

Two further case studies are in planning:

- 1) Utilisation of the platform in conjunction with a whole building management system for a large new academic building (4 West, University of Bath), where there are significant numbers of embedded sensors both in the physical structure as well as the control systems.
- 2) A larger and longer term study follow-up of the on-line trial reported above in conjunction with a building management systems using a recently completed student housing block (Woodland Court, University of Bath).

We have also developed a first version of a mechanism to switch Ploggs off remotely. This functionality can be

²<http://www.sunspotworld.com/>, accessed 20100923

³openenergymonitor.org, accessed 20100923

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 [11], [13]. The ‘processing’ part of the architecture introduced in Section III can be extended with another agent type (effector) that can control energy consumption by switching thermostatically controlled appliances, such as fridges or ac-units, off or on in a coordinated manner thereby shifting energy consumption [23], [12] 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 [24] and security [25] form other areas that have to be considered in future work. Privacy in particular forms an obvious challenge in this context.

VI. CONCLUSIONS

This paper presented a framework for the collection, processing and presentation of sensor data in general [17] applied to energy consumption by household appliances in particular. Main features of the platform are (i) connection with a RDF database, which permits ontological annotation of sensor feeds and ontology-based querying (ii) virtual sensors, which can act as indirection, grouping and aggregator components, decoupling down-stream analysis from up-stream collection (iii) utilisation of a distributed agent platform, which offers scope for genuinely distributed data collection and analysis Overall, we believe this demonstrates the flexibility and scope for future development that the use of agents has to offer in the domain of energy monitoring. In the near future, this can be extended to energy management, by closing the loop and selectively turning devices on and off subject to policy and user behaviour, as part of the same framework.

VII. ACKNOWLEDGMENTS

This work is partially supported by NGI project “Self-Managed Dynamic Institutions in Power Grids: Sharing the Cost of Reliability” 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).

REFERENCES

- [1] C. Gellings, *The Smart Grid: Enabling Energy Efficiency and Demand Response*. Fairmont Press, 2009.
- [2] C. Shu-yong, S. Shu-fang, L. Lan-xin, and S. Jie, “Survey on Smart Grid Technology,” *Power System Technology*, vol. 33, no. 8, pp. 1–7, 2009.
- [3] R. Duan and G. Deconinck, “Future electricity market interoperability of a multi-agent model of the Smart Grid,” in *Networking, Sensing and Control (ICNSC), 2010 International Conference on*. IEEE, 2010, pp. 625–630.
- [4] F. Haldi and D. Robinson, “Interactions with window openings by office occupants,” *Building and Environment*, vol. 44, no. 12, pp. 2378–2395, 2009.
- [5] A. Peacock and M. Newborough, “Impact of micro-combined heat-and-power systems on energy flows in the uk electricity supply industry,” *Energy*, vol. 31, no. 12, pp. 1804 – 1818, 2006.
- [6] 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*, vol. 84, no. 4, pp. 389 – 403, 2007.
- [7] M. Newborough and S. D. Probert, “Intelligent automatic electrical-load management for networks of major domestic appliances,” *Applied Energy*, vol. 37, no. 2, pp. 151 – 168, 1990.
- [8] W. Anderson and V. White, “Exploring consumer preferences for home energy display functionality,” Centre for Sustainable Energy, Tech. Rep., 2009.
- [9] A. Faruqui, S. Sergici, and A. Sharif, “The impact of informational feedback on energy consumption—a survey of the experimental evidence,” *Energy*, vol. In Press, Corrected Proof, pp. –, 2009.
- [10] G. Wood and M. Newborough, “Energy-use information transfer for intelligent homes: Enabling energy conservation with central and local displays,” *Energy and Buildings*, vol. 39, no. 4, pp. 495 – 503, 2007.
- [11] 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).
- [12] 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*, vol. 24, no. 2, pp. 285 – 295, 2009.
- [13] G. Strbac, “Demand side management: Benefits and challenges,” *Energy Policy*, vol. 36, no. 12, pp. 4419–4426, December 2008.
- [14] J. Joo and M. Ilic, “Adaptive load management (ALM) in electric power systems,” in *Networking, Sensing and Control (ICNSC), 2010 International Conference on*. IEEE, 2010, pp. 637–642.
- [15] 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.
- [16] “Plogg, wireless energy management,” <http://www.plogginternational.com>, last retrieved March 8, 2010.
- [17] 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.
- [18] 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)*, ser. Lecture Notes in Computer Science, vol. 3732. Copenhagen, Denmark: Springer, June 2004, pp. 675–679.
- [19] E. Prud’hommeaux and A. Seaborne, “W3C Recommendation: SPARQL query language for RDF,” <http://www.w3.org/TR/rdf-sparql-query/>, January 2008.
- [20] H. S. Riat, “Multi agent based energy monitoring system,” Master’s thesis, University of Bath, 2009.
- [21] B. Forchhammer, “Data analysis and presentation tools for heterogeneous sensor networks,” Master’s thesis, University of Bath, 2010.
- [22] 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.
- [23] 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*, vol. 0, pp. 265–271, 2009.
- [24] P. McDaniel and S. McLaughlin, “Security and privacy challenges in the smart grid,” *IEEE Security and Privacy*, vol. 7, no. 3, pp. 75–77, 2009.
- [25] L. AlAbdulkarim and Z. Lukszo, “Information security implementation difficulties in critical infrastructures: Smart metering case,” in *Networking, Sensing and Control (ICNSC), 2010 International Conference on*. IEEE, 2010, pp. 715–720.