DISCLAIMER

This report was prepared as a result of work sponsored by the California Energy Commission (Commission). It does not necessarily represent the views of the Commission, its employees, or the State of California. The Commission, the State of California, its employees, contractors, and subcontractors make no warranty, express or implied, and assume no legal liability for the information in this report, nor does any party represent that the use of this information will not infringe upon privately owned rights. This report has not been approved or disapproved by the Commission nor has the Commission passed upon the accuracy or adequacy of the information in this report.

Published in the Proceedings from the ACEEE 2004 Summer Study on Energy Efficiency in Buildings: Breaking out of the Box, August 22-27, 2004, Asilomar, Pacific Grove, CA. Washington, D.C.: American Council for an Energy-Efficient Economy.

Machine to Machine (M2M) Technology in Demand Responsive Commercial Buildings

August 2004

David S. Watson

Mary Ann Piette

Osman Sezgen

Naoya Motegi

Lawrence Berkeley National Laboratory

Laurie ten Hope

Program Manager, Energy Systems Integration

Sponsored by the California Energy Commission

This work was supported by the California Energy Commission through its Public Interest Energy Research Program and by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State and Community Programs, Office of Building Research and Standards of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

Machine to Machine (M2M) Technology in Demand Responsive Commercial Buildings

*David S. Watson, Mary Ann Piette, Osman Sezgen, and Naoya Motegi
Lawrence Berkeley National Laboratory*

ABSTRACT

Machine to Machine (M2M) is a term used to describe the technologies that enable computers, embedded processors, smart sensors, actuators and mobile devices to communicate with one another, take measurements and make decisions - often without human intervention.

M2M technology was applied to five commercial buildings in a test. The goal was to reduce electric demand when a remote price signal rose above a predetermine price. In this system, a variable price signal was generated from a single source on the Internet and distributed using the meta-language, XML (Extensible Markup Language). Each of five commercial building sites monitored the common price signal and automatically shed site-specific electric loads when the price increased above predetermined thresholds. Other than price signal scheduling, which was set up in advance by the project researchers, the system was designed to operate without human intervention during the two-week test period.

Although the buildings responded to the same price signal, the communication infrastructures used at each building were substantially different. This study provides an overview of the technologies used at each building site, the price generator/server, and each link in between. Network architecture, security, data visualization and site-specific system features are characterized.

The results of the test are discussed, including: functionality at each site, measurement and verification techniques, and feedback from energy managers and building operators. Lessons learned from the test and potential implications for widespread rollout are provided.

Introduction

This paper provides a summary of the control and communications systems evaluated and reported on as part of a larger research report (Piette et al, 2004). The objective of the study was to evaluate the technological performance of Automated Demand Response hardware and software systems in large facilities. The concept in the evaluation was to conduct a test using a fictitious electricity price to trigger a demand-response event over the Internet. Two related papers describe the measurement of the electric demand shedding and the decision making issues with the site energy managers (Motegi et al, 2004 and Shockman, 2004).

The two main drivers for widespread demand responsiveness are the prevention of future electricity crises and the reduction of average electricity prices. Demand response has been identified as an important element of the State of California's Energy Action Plan, which was

developed by the California Energy Commission (CEC), California Public Utilities Commission (CPUC), and Consumer Power and Conservation Financing Authority (CPA). The CEC's 2003 Integrated Energy Policy Report also advocates Demand Response.

A demand responsive building responds to a remote signal to reduce electric demand. This is usually done by altering the behavior of building equipment such as heating ventilating and air conditioning (HVAC) systems and/or lighting systems so as to operate at reduced electrical loads. This reduction is known as "shedding" electric loads. Demand responsiveness and shedding can be accomplished by building operators manually turning off equipment in response to a phone call or other type of alert.

In this paper, we use the term "Automated Demand Response" or "Auto-DR" to describe "Fully-Automated" Demand Response where electric loads are shed automatically based on a remote Internet based price signal. Although the facility operating staff can choose to opt out if desired, Auto-DR systems operate without human intervention.

Previous Research

The California Energy Commission (CEC) and the New York State Energy Research and Development Agency (NYSERDA) have been leaders in the demonstration of demand response programs utilizing enabling technologies. Several studies associated with the California and New York efforts investigated the effectiveness of demand responsive technologies. In California, Nexant was charged with evaluating CEC's Peak Load Reduction Program. The Nexant reports document the performance of all the California funded technology projects including the magnitude of the response and the cost associated with it (Nexant 2001 and Nexant 2002).

In addition to research concerning utility programs, controls, and communications systems, several research studies have examined various topics concerning DR in commercial buildings, including how to operate buildings to maximize demand response and minimize loss of services. Kinney et al reported on weather sensitivity of peak load shedding and power savings from setting up temperatures in buildings to reduce cooling loads (Kinney et al, 2001). This research project also builds on previous LBNL work concerning the features and characteristics of Web-based Energy Information Systems (EIS) for energy efficiency and DR (Motegi et al., 2003).

Project Description

The Automated DR research project took place over approximately two years, beginning with a planning activity in summer, 2002, successful pilot tests in November 2003 and final reporting in March 14, 2004 (Piette, et al 2004). The building sites, including their use, floor area, and equipment loads shed during the Auto DR tests are listed in Table 1.

System Geography

Although all of the Auto-DR pilot sites were in California, the supporting communications infrastructure and several of the developers were distributed throughout North America (see Figure 1).

Table 1. Summary of Sites

	Site A	Site B	Site C	Site D	Site E
Location	Oakland	Concord	Oakland	Palo Alto	Santa Barbara
Use	Supermarket	Office	Office	Pharmaceutical laboratory (Office & Cafeteria)	Library
Floor Area (ft²)	50,000	211,000	978,000	192,000	289,000
Equipment loads shed during test	50% of overhead lighting, Anti-sweat heaters	Supply fan duct static pressure setpoint	Global zone setpoint setup and setback	Constant volume fan shut off	Fan speed reduction, Chilled water valves closed

Figure 1. Geographic Location of Auto-DR Infrastructure

Automated Demand Response System Description

The Automated Demand Response System generated an XML price signal from a single source on the Internet (Figure 2). Each of five commercial building sites monitored the common price signal and automatically shed site-specific electric loads when the price increased beyond predetermined thresholds. Other than price signal scheduling, which was set up in advance by the project researchers, the system was designed to operate without

human intervention during two one week pilot periods. The test process followed these steps:

1. LBNL defined the price vs. time schedule and sent it to the price server.
2. The current price was published on the server.
3. Clients requested the latest price from the server every few minutes.
4. Business logic determined actions based on price.
5. EMCS (energy management control system) carried out shed commands based on logic.

Figure 2. Auto-DR Network Communications Sequence

Web Services / XML

The infrastructure of the Auto-DR system is based on a set of technologies known as Web services. Web services have emerged as an important new type of application used in creating distributed computing solutions over the Internet. Properly designed Web services are completely independent of computer platform (i.e., Microsoft, Linux, Unix, Mac, etc.). The following analogy helps to describe Web services: Web pages are for people to view information on the Internet, Web services are for computers to view and share information on the Internet. Since human intervention is not required, this technology is sometimes referred to as “Machine-to-Machine” or “M2M”. M2M is a superset of technologies that includes some XML/Web services based systems.

XML is a “meta-language” — a language for describing other languages —that allows design of customized markup languages for different types of documents on the Web (Flynn, 2003). It allows designers to create their own customized tags, enabling the definition, transmission, validation, and interpretation of data between applications and between organizations (Webopedia, 2004). Standard communication protocols (TCP/IP, HTTP and SOAP) are used on the Internet and LAN/WANs (local area network/wide area network) to transfer XML messages across the network.

Price Server

The server publishes the current (fictitious) price for electricity (\$/kWh). Each of the participating sites has a Web services client that polls the server to find out the latest price.

Price Scheduling Software

A scheduling client software application was provided to researchers at the LBNL. The password protected scheduling client was used to set up the price vs. time profile in the price server. The price profile could be set up hours, days or weeks in advance, however the building operators were not given any prior knowledge of upcoming fictitious price increases planned by researchers.

Electric Price Signal and Test Description

Figure 3 shows the fictitious price signal that was in effect on the afternoon of November 19, 2003. During the rest of that day, the price remained at \$0.10/kWh.

Figure 3. Price Signal on November 19, 2003

Polling-Client Price Verification

The price server included a feature that verified and time-stamped client receipt of correct pricing information. This feature was implemented by requiring that each building-generated

request for the latest price included the current price (from the client's perspective) and a time stamp. All pricing data were stored in a database. Although the intent of the client price verification feature was to verify client receipt of server generated pricing, there was another unforeseen benefit as well. When pre-pilot testing began, we could see which sites were polling to the server as they each came online. Even after each system was online, there were cases where clients would stop polling for known or unknown reasons. Program managers and system administrators were able to manually observe these intermittent losses of communication problems and make phone calls to resolve the problems.

Web Services – Clients

The polling client is the software application that checks (polls) the Web services server to get the latest price data. The polling client resides on a computer managed by the building operators (or their representatives) for each site. In the pilot test, each client polled the server at a user-defined frequency of once every 1 to 5 minutes.

Controls and Communications Upgrades

In order to add Auto-DR functionality to each pilot site, some upgrades and modification to the controls and communications systems were required. The upgrades were built to work in conjunction with the existing EMCS and EIS (energy information system) remote monitoring and control infrastructure in place at each site. For this project, custom software was written for each site, including: price polling client, business logic, and site-specific EMCS modifications.

Results

Upon sensing an increase in the electricity price above a specified value, all five participating sites initiated pre-determined shed strategies. As an example, the results of site D are shown below in Figure 4. Savings are determined by comparing actual metered demand on the day of the shed with a calculated normal (non-shed) baseline (Motegi et al, 2004). The vertical lines show boundaries of the price range.

Figure 4. Site D Electricity Use, November 19, 2003

The average load shed is shown in figure 5. Each bar represents the average shed over one hour of the three hour elevated price test. The electricity price during the first, second and third hours was \$0.30/kWh, \$0.75/kWh and \$0.30/kWh, respectively (Motegi et al, 2004).

Figure 5. Average shed (kW and W/ft²) during the 3-hour test on November 19, 2003

Lessons Learned from each Site

This section identifies the unique attributes of each participating Auto-DR facility. Controls and communications infrastructures and shed strategies are discussed for each system.

Site A (Supermarket) System Overview

The site A building telemetry data system is split between two systems. The EMCS (excluding electric power monitoring) is accessed via a dial-up modem. EIS data is available via any Web browser through the EIS Web site. The segregated nature of the EMCS and EIS make it a burdensome task for the facility operator to change a temperature setpoint or lighting schedule and then observe the effect on electric usage.

However, the integration between the enterprise networks and control networks is rather tight. The corporate WAN (Wide Area Network) is used to communicate between the business logic/polling client and the on-site Internet Protocol Input/Output (IP I/O) relay device. The enterprise network is also used for mission critical point of sale data communications within the nationwide organization. The fact that the energy data are shared and communicated over the mission critical enterprise network indicates a high level of collaboration and trust between the site A energy managers and other department managers involved with the core business of the organization.

The shed strategy was not objectionable to the store managers or patrons. Although the transition between 100% overhead lighting to 50% was noticeable, there were no complaints. The reduction of overhead lighting appeared to make the other light sources in the store, such as case lights, seem more intense. There is no evidence that the freezer doors fogged up during the shed, even though the setpoint of the anti-sweat heaters was reduced. If the transition of overhead lights to 50% were gradual (e.g., through use of dimmable ballasts) the entire shed would probably not be noticeable.

Site B (Bank Office) System Overview

Integration between enterprise networks and control networks at this site is tight. The site B corporate WAN is used to communicate across the country to the on-site gateway. This network is also used for mission critical financial data communications within the site B organization. Like site A, the fact that energy and HVAC (heating, ventilation, and air conditioning) data are shared over the mission critical enterprise network indicates a high level of collaboration and trust between site B's energy managers, IT (information technology) security managers and other department managers involved with the core business of the organization. The use of highly secure and reliable hardware VPN (virtual private network) routers and the use of a co-location site to host the polling client and business logic computers are indications that system availability and security are high priorities.

With regard to the shed strategy employed at this site, there is no evidence that a modest reduction in duct static pressure for short durations caused any negative comfort effects to the occupants during the test. However, as shown in the measured data, the extent of the electric demand shed is negligible. If this strategy were extended so as to produce significant sized electric sheds, the method may pose some fundamental drawbacks. When the duct static pressure is reduced below the minimum required by the terminal boxes in VAV (variable air volume) systems, airflow is reduced in the zones. But the reduction is not shared evenly between all the boxes. The zones of greatest demand are the ones that are

starved for air most immediately and most severely. In the building used in the Auto-DR pilot, the potential problem is exacerbated by the lack of sensors. Fan airflow is not measured and only nine “representative” zone temperature sensors are available for the entire 211,000 ft² building. There were not enough sensors to estimate the effect that reductions in airflow would have on occupants.

When the third party energy management company (company X) takes action to reduce energy at its connected sites, it uses a centralized control paradigm. While demand response systems are inherently centralized (signals to shed loads are generated in a one-to-many relationship), centralized control for day-to-day operation is less common.

In most control system markets (commercial buildings, industrial controls, etc.) there has been a trend for several decades toward decentralized control. In decentralized control, the control logic is moved (physically) as close to the sensors and final control elements (e.g., relays, valves, etc.) as possible. Decentralized control systems have traditionally been less costly, more flexible and more robust. However, in the IT community, there has been a movement in certain areas toward hosted solutions, application service providers and other centralized solutions. Ubiquitous Internet connectivity and other advances in IT technology have made these systems less costly, more flexible and more robust for certain applications.

The company X system alternates between centralized and decentralized paradigms on a minute-by-minute basis. One minute a fan system maintains a setpoint entered by onsite building operators. In the next minute, a neural network algorithm may define the setpoint from over 3,000 miles away.

Site C (Government Office) System Overview

The enterprise and EMCS infrastructures used to enable Auto-DR at this site are linked together in a long series of serial components and communication links. The prototype system was assembled at low cost using spare parts. With so many links, it is not surprising that there were communication failures due to an unexplained equipment lock-up during the first test. To make the system more robust, a review of the components and architecture should be conducted.

The second test was quite successful, as communications were functional from end-to-end. The shed strategy produced an electric shed about as large as the other four sites combined. Because the temperature setpoint reset was at the zone level, comfort for each occupant could be maintained within the revised, relaxed constraints. To implement this strategy, it was necessary to revise the software parameters and some logic in each of the 1,050 VAV terminal box controllers. For most EMCS systems, the labor required to make these revisions would be substantial (one to three weeks). In this building, the process had been somewhat automated by previous system upgrades. This allowed EMCS reprogramming for Auto-DR to be conducted in about three hours.

Site D (Offices and Cafeteria) System Overview

A third party software framework (company Y) ties together three different EMCS protocols at site D in a seamless fashion. The Web interface provides operators with

compete monitoring and control capability from anywhere on the campus. It was relatively straightforward to interface the Auto-DR polling client and associated business logic to the system. The most challenging part of the project was setting up the “extra” computer outside of the Site D firewall and establishing communications to devices inside of the secure corporate network.

Site E (University Library) System Overview

Remote monitoring and control of the EMCS and EIS was available over the Internet prior to the Auto-DR pilot. However, at the time of the test, remote control of the EMCS was not available. The software gateway between the enterprise network and the EMCS network lost remote control functionality during an “upgrade” of company Zs third party server software. To meet the test schedule of the Auto-DR pilot, an IP I/O relay was added to allow the Auto-DR business logic to initiate the control functions such as initiating sheds. The shed strategy proved to be very effective. The books and other thermal mass in the library buildings acted as a thermal “flywheel” to help keep the space comfortable during the shed periods. In addition, the shed strategy reduced airflow without shutting off fans completely. The coastal climate of the site helped provide a temperate airflow even when the cooling and heating valves were closed.

State of the Art in Automated Demand Responsive Systems

By evaluating the systems demonstrated in the November 2003 Auto-DR pilot, along with other existing technologies found in the EMCSs, EIS and the IT Industries, state-of-art Auto-DR systems can be envisioned. The five participating sites all successfully met the functionality criteria of the pilot (under tight schedules and limited budgets). However, a truly state-of-the-art system would use the “best of the best” components, systems and strategies from end to end. Such a system would be designed from scratch to meet a very specific set of requirements. The “best” system would meet or exceed the requirements at the lowest installed cost. Characteristics of state-of-the-art Auto-DR systems are discussed in this section.

Flexible Designs for the Future

Today’s state-of-the-art Auto-DR technology could be applied in many different ways, depending on the scenarios and use cases that they are designed to satisfy. As the scenarios, use cases, and driving forces behind Auto-DR become better defined, systems will be designed and deployed accordingly. Since these design criteria are likely to remain in flux, Auto-DR system flexibility and future-proofing are a very high priority.

Features

Customers should have numerous options about how they can participate in Auto-DR programs. For any given motivating force that drives customers to consider Auto-DR (i.e.,

price), each will have different circumstances under which they will want to participate. Customer flexibility must be built into any state-of-the-art Auto-DR system. They should have the ability to use custom business logic that is applicable to their own operations. Some may choose to allow remote real-time control (for its extra value) while others may want some advanced warning (via pagers, cell phones, etc.) and the ability to opt out, if desired. Other important features in state-of-the-art Auto-DR systems include real-time two-way communications to the final control and monitoring elements (e.g., sensors), high security, and high system availability.

Leveraging Trends in Technology

The lower the installed cost of state-of-the-art Auto-DR systems, the sooner they will find their way into mainstream use. One of the most important ways to keep costs low is to leverage existing trends in technology. The use of existing IT technology in Auto-DR systems wherever possible is an important way of keeping costs low. The public Internet and private corporate LAN/WANs are ideal platforms for Auto-DR controls and communications due to their ubiquity, especially in large commercial buildings. In addition to the availability of networks, the performance of IT equipment (e.g., routers, firewalls, etc.) continues to improve and prices for this equipment continue to drop.

Enterprise, EMCS & EIS Integration

Another way to obtain high system performance and keep the system costs low is through increased integration within the building. Since energy data from EISs is simply another type of measured data, EISs and EMCSs should share the same networks so as to maximize system performance and functionality and minimize cost. In addition to eliminating a redundant EIS network, other aspects of the system are also unified through this approach. Use of an integrated EMCS/EIS database and associated archiving and visualization tools increases user functionality while reducing cost. The ability to change setpoints for HVAC equipment and observe and analyze the effect on electric consumption from the same Human Machine Interface (HMI) is an important enhancement to both the EMCS and EIS.

State-of-the-art Auto-DR systems should also have tight integration between the EMCS/EISs network and enterprise networks within buildings. Once the integrated controls and communications infrastructures are in place, many applications in addition to Auto-DR are enabled. Auto-DR is a subset of a larger class of functionalities known as “telemetry”. Telemetry is defined as the process by which remote sensors or actuators are controlled and/or monitored. Some other telemetry applications include: energy management, energy forecasting, aggregation, access control and regulatory record keeping.

The network architecture of a state-of-the-art Auto-DR system would tend to be flatter than most of the sites in the November 2003 pilot. A flat architecture is one in which there are a minimum number of layers of control networks and protocols between the HMI and the final control and monitoring elements. The most robust and least costly systems should have no more than one enterprise network protocol and one control network protocol.

Open Standards

For flexibility and future-proofing as well as the option to choose “best of breed” products, state-of-the-art Auto-DR systems should use open standards wherever possible. Unlike proprietary systems, truly open systems are interoperable. In other words, a device from one company will easily and naturally reside on a network with products from other companies. Most products in enterprise networks are interoperable. They all communicate using the TCP/IP protocol and can be set-up and managed using common network management tools. TCP/IP is clearly the worldwide protocol of choice for LAN/WAN, Internet and enterprise networks.

There are several open standards control networks including BACnet (ASHRAE Standard 135-2001) and LonTalk (ANSI/EIA/CEA 709.1). Several database formats have become de facto open standards as well. Although the use of the meta-language XML is becoming a standard framework for communicating over enterprise networks and the Internet, XML alone does not define data formats that could be used to convey measured building or energy data. Standards of this type are being developed by oBIX, ASHRAE and other organizations.

With only two network protocols in the state-of-the-art Auto-DR system, only one type of gateway is required for translation/abstraction between them. An embedded gateway device that conforms with IT industry standards for reliability, security, and network management should be used.

Shed Strategies

State-of-the-art shed strategies should be designed to minimized discomfort, inconvenience and loss of revenue to the participating sites. Shed strategies should be devised by customers to meet their needs. In general, shed strategies that use fine granularity closed loop control are less likely to negatively impact building occupants. Ideally, sheds would vary commensurately with a variable shed signal. Transitions would be fast enough to be effective, but slow enough to minimize the notice of building occupants.

In addition to HVAC control strategies, lighting and switch-able plug loads should be considered for sheds as well. By increasing the controlled load to the point where it approaches the whole building load, each load type (HVAC, lighting, etc.) would need to shed a smaller amount in order to achieve a given shed target for the whole building.

Future Directions

Although more common in the industrial controls marketplace, TCP/IP protocol is used to create control networks. In these systems, traditional open control protocols such as BACnet and LonTalk are eliminated all together. TCP/IP could be used in an end to end integrated enterprise, EMCS/EIS system. This trend is likely to gain momentum once the next generation of Internet Protocol, IPv6 is implemented. To achieve greater flexibility,

increased control granularity and lower costs, increased use of wireless devices in Auto-DR , EMCS and EISs is likely to occur.

Acknowledgements

The authors are grateful for the extensive support from Ron Hofmann, a consultant to the California Energy Commission, Laurie ten Hope (CEC), Gaymond Yee (CIEE), George Hernandez (James Ruben Associates), and Karen Herter (LBNL), Joe Desmond and Nicolas Kardas (Infotility), and numerous individuals from the participated sites. This project was supported by the California Energy Commission's Public Interest Energy Research Program and by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State and Community Programs of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

References

Flynn, P. (Editor). 2003. "The XML FAQ" Maintained on behalf of the World Wide Web Consortium's XML Special Interest Group v. 3.01 (2003-01-14).

Kinney, S., M.A. Piette, L. Gu, and P. Haves. 2001. Demand Relief and Weather Sensitivity in Large California Commercial Office Buildings, Proceedings of the 2001 International Conference for Enhanced Building Operations, LBNL Report 48285

Motegi, Naoya, Mary Ann Piette, David S. Watson, and Osman Sezgen Lawrence Berkeley National Laboratory, "Measurement and Evaluation Techniques for Automated Demand Response Demonstration", 2004 ACEEE Summer Study on Energy Efficiency in Buildings. LBNL Report #55086.

Motegi, Naoya, Mary Ann Piette, Satkerter Kinny and Karen Herter. 2003. "Web-based Energy Information Systems for Energy Management and Demand Response in Commercial Buildings." High Performance Commercial Building Systems, PIER Program. LBNL#-52510. <http://buildings.lbl.gov/hpcbs/Pubs.html>

Nexant. 2002. AB 29X and SB 5X Program Evaluation. First Quarter Report (January 1 to March 1,2002). Submitted to California Energy Commission. Contract #400-00-070.

Nexant. 2001. AB 979, AB 29X and SB 5X Peak Load Reduction Programs. December 2001 Annual Report. Submitted to California Energy Commission and the California State Legislature.

Piette, Mary Ann, David S. Watson, Osman Sezgen, Naoya Motegi, Christine Shockman and Ron Hofmann. Forthcoming in 2004. "Development and Evaluation of Fully Automated Demand Response in Large Facilities." LBNL Report #55085.

Shockman, Christine. 2004. "Market Transformation Lessons Learned from an Automated Demand Response Test in the Summer and Fall of 2003." *Proceedings of the ACEEE 2004 Summer Study on Energy Efficiency in Buildings*. American Council for an Energy-Efficiency Economy.

Webopedia. 2004. <http://www.webopedia.com/TERM/X/XML.html> Viewed on Mar. 15, 2004.